IBM Operations Analytics for z Systems

IBM z Systems Advanced Workload Analysis Reporter (zAware) Guide

Version 3 Release 1

Contents

Figures

Tables

About this information

This information describes IBM® z Advanced Workload Analysis Reporter (IBM zAware) Version 3.1 (V3.1), which consists of an integrated set of applications that monitor software and model normal system behavior. Its pattern recognition techniques identify unexpected messages, providing rapid diagnosis of problems caused by system changes. This early detection helps IT personnel correct problems before they affect system processing.

IBM zAware is a feature of IBM Z Operations Analytics for the following servers:
- An IBM z14 (z14)
- An IBM z13® (z13®) or IBM z13s (z13s)
- An IBM zEnterprise® EC12 (zEC12) or IBM zEnterprise BC12 (zBC12)

Note: Figures included in this document illustrate concepts and are for example purposes. In some cases, they might not be accurate in content, appearance, or specific behavior.

Intended audience

This information is intended for experienced systems programmers and administrators who perform the following tasks:
- Configuring the logical partition (LPAR) on which IBM z Advanced Workload Analysis Reporter (IBM zAware) runs.
- Preparing network and storage resources that IBM zAware requires for operation and securing access to these resources.
- Configuring monitored clients to send data to IBM zAware for analysis.
- Managing the use and operation of IBM zAware, which includes controlling user access to the IBM zAware graphical user interface (GUI).
- Viewing and interpreting analytical data through the IBM zAware GUI and resolving potential problems using IBM Operations Analytics for z Systems®.

Prerequisite and related information

This information describes IBM z Advanced Workload Analysis Reporter (IBM zAware) Version 3.1, which is an element of IBM Z Operations Analytics V3.

For complete prerequisite information, see Chapter 2, “Prerequisites for configuring and using IBM zAware,” on page 13.

Earlier versions of IBM zAware were a feature that was available with specific z Systems servers.

For *System z® Advanced Workload Analysis Reporter (IBM zAware) Guide*, SC27-2623, Version 2.0, see IBM Resource Link® at https://www-01.ibm.com/servers/resourcelink/lib03010.nsf/0/9D5FC2BA5A27447185257DE7004A4347?OpenDocument.

How to use this information

This information provides an overview of IBM zAware, lists the system requirements for its infrastructure and applications, and lists the IT roles and skills required to set up and use it. This information also provides step-by-step instructions, or references to the appropriate hardware or software publications, for systems programmers and administrators who configure and manage IBM zAware or the operating systems that send data to it for analysis.

The following list describes the overall structure and content of this information:

Part 1: Introduction to IBM z Advanced Workload Analysis Reporter (IBM zAware)
Topics in this part describe IBM zAware, explain the benefits of using it, define new terms or concepts, list hardware and software prerequisites, and provide a project plan for configuring and using IBM zAware.

Part 2: Installing or upgrading to the latest IBM zAware
Topics in this part explain the installation instructions that systems programmers must know to be able to install IBM zAware and operate it in IBM Operations Analytics for z Systems.

Part 3: Planning to configure IBM zAware
Topics in this part explain the planning considerations that systems programmers and administrators need to know before they start configuring IBM zAware and its operating environment.

Part 4: Configuring IBM zAware and its monitored clients
Topics in this part provide instructions for configuring the IBM zAware environment, which includes the IBM zAware partition and the monitored clients that send data for analysis. Systems programmers and administrators use these configuration tasks primarily for first-time setup.

Part 5: Managing and using IBM zAware
Topics in this part describe the IBM zAware graphical user interface (GUI) functions for daily operations, which include viewing and analyzing data from monitored clients. Additional topics include management tasks for modifying the IBM zAware configuration or operations.

Part 6: Advanced topics for managing IBM zAware
Topics in this part describe specialized management tasks for IBM zAware, such as modifying IBM zAware models of normal system behavior, planning and setting up IBM zAware for backup and recovery from failures, troubleshooting problems, and reporting problems to IBM.

Appendixes
Appendix topics include Application Programming Interfaces (APIs) that monitoring products can use to extract data from IBM zAware, and a summary of tasks and required IT skills, tools and related information in the z Systems product libraries.

Accessibility features for IBM zAware GUI

Accessibility features help users who have a disability, such as restricted mobility or limited vision, to use information technology products successfully. IBM strives to provide products with usable access for everyone, regardless of age or ability.

The following list includes the major accessibility features in the IBM zAware GUI:
- Keyboard-only operation
- Interfaces that are commonly used by screen readers

Keyboard operations for IBM zAware

This product uses standard operating system navigation keys. You can use keys or key combinations for operations and to open menu actions that are typically done by using mouse actions. You can use the IBM zAware GUI from the keyboard by using the keyboard shortcuts for your browser or screen-reader software. See your browser or screen-reader software **Help** for a list of keyboard shortcuts that it supports.

Interface information

Several pages or display elements in the IBM zAware GUI have more accessible, text-only alternatives:
- The Details pane in the Analysis Heat Map Table provides a graphical and an accessible tabular view of analysis results for one system.

- The Analysis Table provides an accessible view of the analysis results for one or more systems.
- The **Time Line Summary** window provides a text-only format that indicates when a selected message ID was issued during a selected analysis interval. To display the text-only format, position your cursor over the graphic display in the Time Line column on the Interval page, and click to open the **Time Line Summary** window.
- The Summary view provides a text-based version of the dates that are associated with a model of system behavior. The Summary view is the default view when you go to the **Training Sets** > **Manage Model Dates** page.

IBM and accessibility

For more information about the commitment that IBM has to accessibility, see http://www.ibm.com/able/.

Part 1. Introduction to IBM z Advanced Workload Analysis Reporter (IBM zAware)

Topics in this part describe IBM zAware, explain the benefits of using it, define new terms or concepts, list hardware and software prerequisites, and list the IT personnel and required skills for using IBM zAware.

Topics covered in this part are:

Chapter 1. Overview of IBM zAware

Today's complex, integrated data centers require a team of experts to monitor systems for abnormal behavior, and to diagnose and fix anomalies before they result in failures and outages that are visible beyond the data center. These tasks are costly and difficult for many reasons, including the fact that various everyday changes can cause system anomalies. The IBM z Advanced Workload Analysis Reporter (IBM zAware) provides a smart solution for detecting and diagnosing anomalies in z/OS® and Linux on z Systems systems. IBM zAware creates a model of normal system behavior based on prior system data, and uses pattern recognition techniques to identify unexpected messages in current data from the systems that it is monitoring. This analysis of events provides nearly real-time detection of anomalies that you can easily view through a graphical user interface (GUI). You also use the GUI to diagnose the cause of past or current anomalies.

Are your systems behaving badly?

Many everyday activities can introduce system anomalies and initiate failures in complex, integrated data centers; these activities include:

- Increased volume of business activity
- Application modifications to comply with changing regulatory requirements
- Standard operational changes, such as adding or upgrading hardware or software, or changing network configurations.

You can use a combination of existing system management tools to determine whether any of these activities is causing one or more systems to behave abnormally, but none can detect every possible combination of change and failure. Even when you use these tools, you might have to look through message logs to help solve the problem but the sheer volume of messages can make this task a daunting one. For example, a z/OS sysplex might produce more than 4 gigabytes (GB) of message traffic per day for its images and components alone, and application messages can significantly increase that number. More than 40000 unique message IDs are defined for z/OS and the IBM software that runs on z/OS systems.

Modernize detection and diagnosis with IBM zAware

IBM zAware is able to analyze large quantities of message log data. Using prior message log data and mathematical modeling, IBM zAware builds models of normal behavior and uses the appropriate model to compare to current message log data from specific connected systems.

- For z/OS systems, IBM zAware monitors the z/OS operations log (OPERLOG), which contains all messages that are written to the z/OS console, and suppressed messages that are not deleted after passing through any message processing controls at your installation. IBM zAware builds and uses one model for each z/OS system. z/OS system models can be built using both OPERLOG and system log (SYSLOG) data.
- For Linux systems, IBM zAware monitors Linux system log (syslog) messages that are sent through the syslog daemon, such as the open source software utilities rsyslog and syslog_ng. Linux systems are often activated and deactivated frequently to meet various operational needs, such as additional resources, system availability, and system maintenance. For IBM zAware to evaluate message traffic for such systems, they must belong to a model group. IBM zAware builds one model for a group of Linux systems with similar workloads, and uses that model to compare to current syslog data from each system in the group. IBM zAware administrators who manage Linux support determine which Linux systems belong to a particular group, and define the group and its members through the IBM zAware GUI.

IBM zAware detects unusual messages and unusual message patterns that typical monitoring systems miss, as well as unique messages that might indicate system health issues. Its ability to pinpoint deviations in normal system behavior can improve real-time event diagnostics.

IBM zAware automatically manages the creation of the behavioral model and manages the retention of IBM zAware analytical data for each monitored system. The number of monitored systems is limited by the data center resources that are required for collecting and storing data for monitored systems, and for IBM zAware operation.

Through the IBM zAware GUI, you can view analytical data that indicates which system is experiencing deviations in behavior, when the anomaly occurred, and whether the message was issued out of context. Using this information, you can take corrective action for these anomalies before they develop into more visible problems. Early detection and focused diagnosis can help improve time to recovery.

Understanding how IBM zAware calculates and displays anomaly scores

IBM zAware continuously analyzes the current data that connected monitored systems send to it. To produce meaningful analysis results for a monitored system, IBM zAware analyzes the most recent minutes of current data to compare to the model. These minutes are called the *analysis interval*, the length of which varies depending on the type of monitored system.

- For z/OS systems, which typically produce high-volume, consistent message traffic, IBM zAware requires 10 minutes of current data to produce an anomaly score.
- For Linux systems, which tend to produce lower volume, less consistent message traffic, IBM zAware requires 60 minutes of current data to produce an anomaly score.

An *interval anomaly score* indicates the relative difference in behavior of the monitored system, as compared to the model. The IBM zAware server uses unsupervised machine learning and IBM rules to determine anomaly scores for all monitored clients.

- Through *unsupervised machine learning*, the IBM zAware server extracts and organizes message data to build a model of behavior for each z/OS monitored client or each Linux model group. This training process is repeated over time, with the frequency determined by the training interval, which enables the server to update the model with more recent system behavior.
  - Through the training process, the IBM zAware server determines which messages are issued during routine system events, such as starting a batch job or a particular subsystem. For such system events, the server identifies and recognizes groups of messages that are associated with each event. The message groups are called *clusters* and define the normal context for the messages in the cluster.

    When the server detects a specific message that is issued outside of its expected context (that is, without the other messages in the cluster), the server assigns a higher message anomaly score, which is combined with the other message anomaly scores in the interval to assign the interval anomaly score.

    IBM zAware also detects messages that are issued periodically; for example, a message that is issued every 11 minutes. This attribute affects the anomaly score when a periodic message is not issued as expected.
  - Also through the training process, IBM zAware determines the distribution of each unique message ID within a collection of intervals in the message data that is used for training. This distribution influences the interval anomaly score that the IBM zAware server displays for an interval of current data from the client.

  In summary, through the training process, the server learns about expected message patterns, and stores this information as part of the model for a specific client or group. IBM zAware uses this model data to determine interval scores when it analyzes current data that it receives from the client.
- Based on decades of experience, z/OS experts at IBM know which message IDs are likely to indicate potential problems. Message IXC101I, for example, indicates that a system is being removed from a sysplex. For a test system, this removal process could be reflected in the IBM zAware model for this system as a normal, expected behavior pattern. In the analytical data for this test system, you might

expect the server to assign a low anomaly score and light blue color to any intervals that contain such a system removal, when message IXC101I is issued in context.

However, message IXC101I might indicate a potential problem, whether or not it is issued in context. Because the removal of a system from a sysplex warrants further investigation, the IBM zAware server is programmed to assign the highest interval anomaly score to intervals in which message IXC101I is issued. IBM rules for other known messages can alter anomaly scores to a lesser degree.

IBM zAware assigns rules only to messages that z/OS monitored systems issue.

In summary, comparison to the model, context, and IBM rules are key factors that contribute to the interval anomaly scores for systems in the Analysis page display.

To display and record the results of analysis intervals, IBM zAware produces an *analysis snapshot* every 10 minutes for each monitored system. Each analysis snapshot is a point-in-time record of the anomaly score for an analysis interval. For example, the following snapshots:

- For a z/OS system, the snapshot that is recorded at 09:00 UTC represents the analysis score and number of unique messages that are issued by that system from 08:50 to 9:00 UTC. The next snapshot is taken at 09:10 UTC for the analysis interval from 09:00 to 09:10, and so on.
- For a Linux system, the snapshot that is recorded at 09:00 UTC represents the analysis score and number of unique messages that are issued by that system from 08:00 to 9:00 UTC. The next snapshot is taken at 09:10 UTC for the analysis interval from 08:10 to 09:10, and so on. Because of the 60-minute analysis interval, every snapshot for a Linux system overlaps with previous snapshots.

Finding the culprit when a problem occurs

Using the **Analysis** page in the IBM zAware GUI, you can answer key questions to diagnose a system problem. The default presentation of the Analysis page is the Analysis Heat Map view, which is one of several available display formats for the analysis results that IBM zAware produces. Use this display to quickly find which system groups or monitored systems have the highest anomaly scores within a day or hour.

The Analysis Heat Map Table displays the peak anomaly scores per day and per hour for the groups or systems in the IBM zAware topology. In the Interval Anomaly Scores table that is shown in this view, each table cell represents 1 hour; the cell is colored to indicate the highest anomaly score calculated for a monitored system during that hour. The table cell contains the anomaly score itself, which is a link through which you can change the view from group to system, or display more detailed information about a specific system. Table cells of interest are the darkest blue, gold, or orange colors.

Figure 1 on page 6 illustrates the Heat Map view, which can help you answer this question: “Which system is behaving abnormally?”

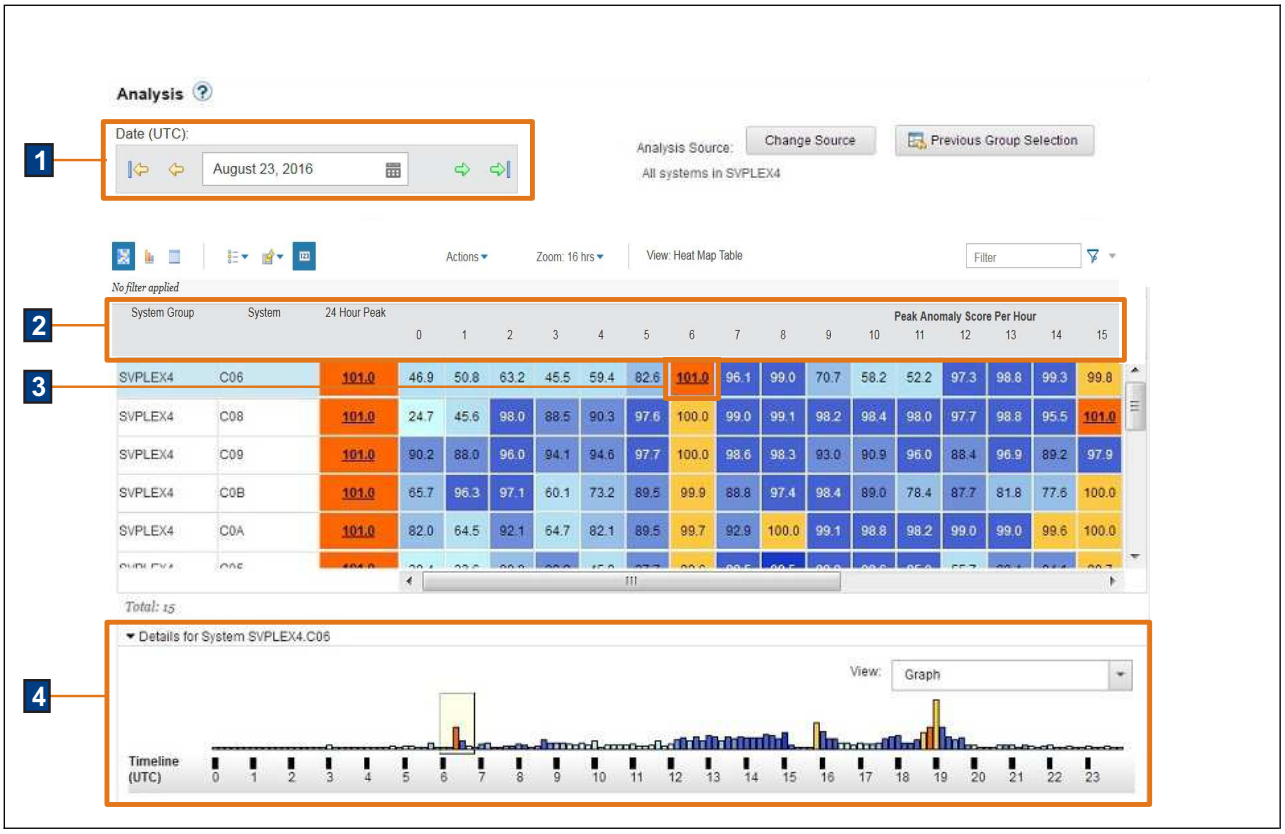

Figure 1. Elements of the Analysis Heat Map view that help identify the problematic monitored system

To begin your search for the problematic monitored system, use the following highlighted controls.

1. Use the Date field to select the date. The **Date** field displays the currently selected date. You can change the current day to any prior date for which IBM zAware has analytic data. You can type the date, select it from the calendar widget, or use the forward and back arrows to change the date.
2. Use the peak score columns in the Interval Anomaly Scores table to quickly find the highest anomaly scores for the selected date. By default, the Analysis Heat Map Table view displays analysis results for all monitored groups. The score in a table cell indicates the top anomaly score from an individual system within the group. The table cell color reflects the anomaly score for the system. A high score indicates unusual message IDs or unusual patterns of message IDs compared to the system model; the table cells that contain higher scores are colored dark blue, gold, or orange.
   - The 24 Hour Peak column indicates the highest anomaly score calculated for an individual system within the 24 hours that constitute the selected date. Rows in the display are sorted in order from the highest to the lowest 24-hour peak anomaly score.
   - The Peak Anomaly Score Per Hour column indicates the highest anomaly score calculated for a monitored system during each hour of the selected date. Each subcolumn indicates the start of each hour of the day according to the 24-hour clock, in Coordinated Universal Time (UTC).
3. To view detailed analysis results for systems within a group, click the table cell with the highest anomaly score. The Analysis Heat Map Table display changes to the system view, which shows one row for each system in the selected group. The top row contains the system with the highest anomaly score; the remaining system rows are sorted in descending order by anomaly score.
4. To view more detailed analysis results for a particular system, click the table cell that contains the highest anomaly score for that system. The Details pane opens to display a bar graph view of analysis results for the selected system, with a transparent rectangle that highlights the selected analysis snapshot. In the Details pane, an analysis snapshot is shown as rectangle.

From the anomaly scores in Figure 1 on page 6, you can determine that C06 is a potential source of the problem because its high anomaly score.

Your next diagnostic question to answer is "When did this system start misbehaving?" To determine when the system began behaving abnormally, switch from the Analysis Heat Map view to the Analysis Graph view, which is shown in Figure 2. The Analysis Graph displays a bar graph for each group or system in the IBM zAware topology. In the Interval Anomaly Scores table that is shown in this view, each rectangle in a bar graph represents the anomaly score that IBM zAware recorded for a monitored system. IBM zAware records an anomaly score every 10 minutes. Rectangles of interest are the darkest blue, gold, or orange colors. Use this display to quickly find the time at which a monitored system exhibited high anomaly scores.

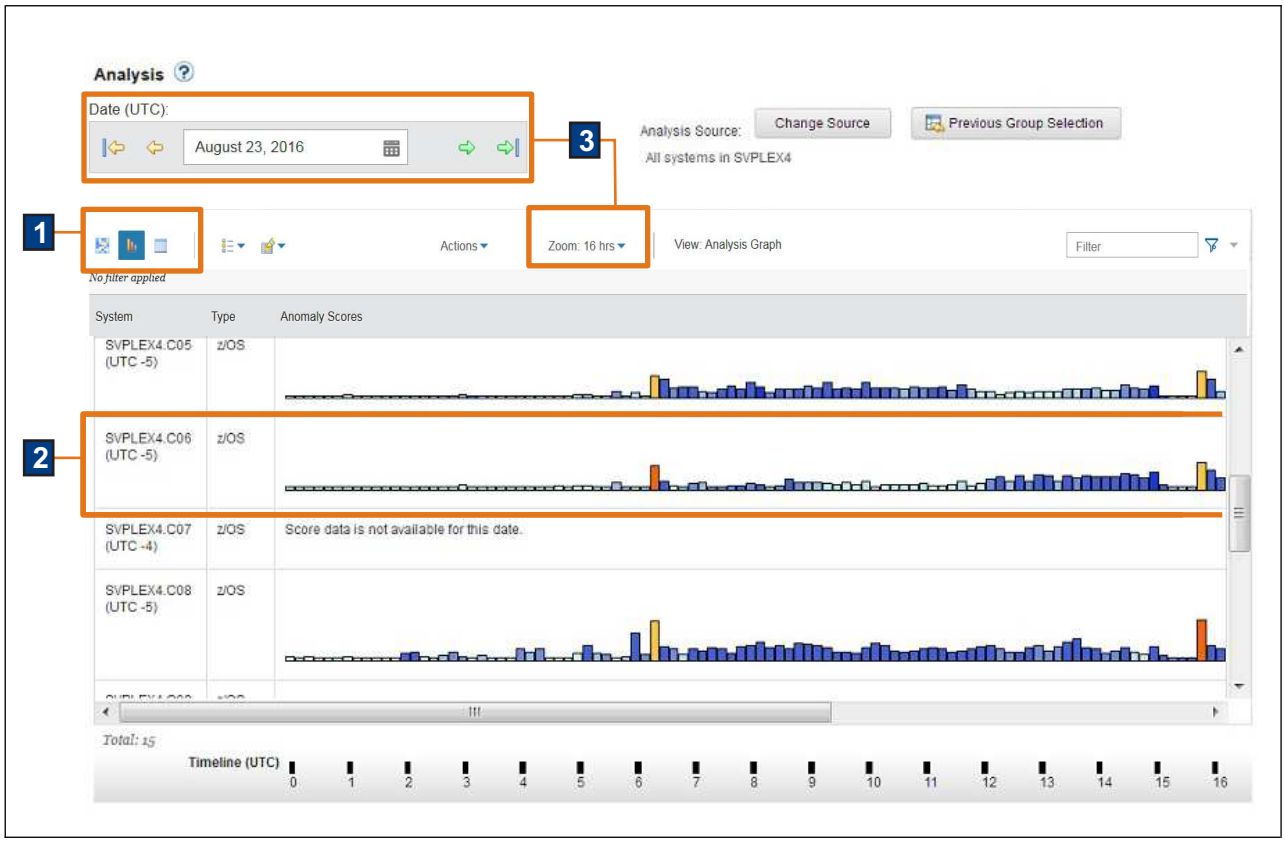

Figure 2. Elements of the Analysis Graph view that help identify when the monitored system began experiencing problems

In Figure 2, the following controls are highlighted:

1. The table toolbar indicates which view is in effect, and also contains controls through which you can change the display. In this figure, the Analysis Graph icon ( ) is selected.
2. Each row, or bar graph, in the Interval Anomaly Scores table shows a timeline of anomaly snapshots for a particular monitored system. Each rectangle in a bar graph represents an analysis snapshot, which is a point-in-time record of the anomaly score for an analysis interval. The rectangle color indicates the anomaly score, and its height is an approximate illustration of the number of unique messages that are issued during the analysis interval. Taller rectangles represent analysis intervals in which a larger number of unique messages were issued.
3. Two controls enable you to change the display to search for the time at which the monitored system began to experience problems.

- The **Date** field displays the currently selected date. You can change the current day to any prior date for which IBM zAware has analytic data. You can type the date, select it from the calendar widget, or use the forward and back arrows to change the date.
- Through **Zoom**, you can control how many hours of the day are displayed in the Analysis page.

As with the Analysis Heat Map view, color is the primary indicator of anomalous behavior, so look for dark blue, gold, or orange rectangles. The height of the rectangle is an extra indicator of unusual activity, with taller rectangles that represent analysis intervals in which a larger number of unique messages were issued. For a quick view of analysis interval results, hover your cursor over any rectangle to display the following information about the analysis snapshot that the rectangle represents.

To pinpoint and diagnose the problem with C06, you might need to ask several questions:
- What messages are unusual?
- How often did the unusual message get issued?
- Are messages issued in context within an expected pattern?
- Is a specific component or application issuing unusual messages?
- When did the message ID first appear?
- Did the message appear when expected?
- Did the message occur at an expected predictable time (for example, every 81 seconds)?

To answer these diagnostic questions, use the Interval page, which is illustrated in Figure 3. You can display the Interval page for any analysis interval by clicking the colored rectangle in the bar graph in the Analysis Graph view.

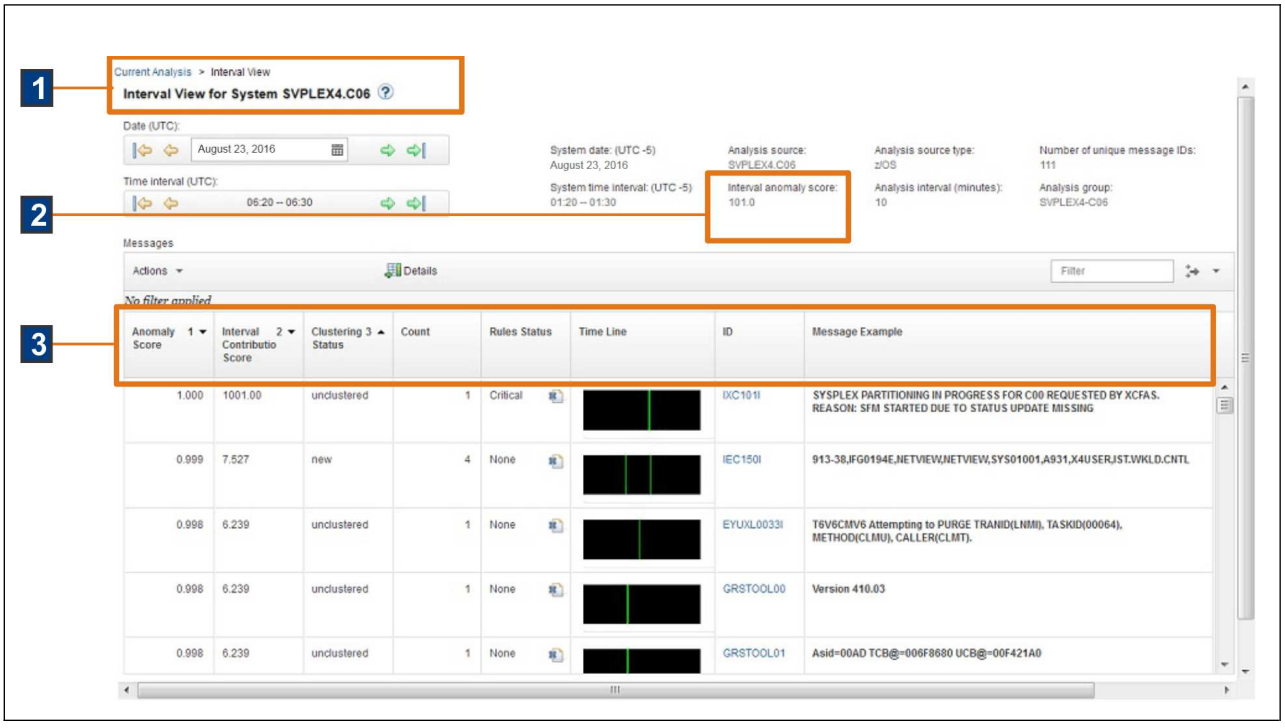

Figure 3. Elements of the Interval page that provide details about unique messages

In Figure 3, the following diagnostic details on the Interval page are highlighted:

1. The system name, the date, and the time interval that you selected from the Analysis page.
2. The “Interval anomaly score” field, which indicates unusual patterns of message IDs within this interval, as compared to the model of normal system behavior. This score determines the color of the rectangle that represents this interval in the bar graph on the Analysis page. Higher scores indicate greater anomaly so intervals with high anomaly scores are more likely to indicate a problem.

3. The Messages table provides diagnostic details about the messages issued during the analysis interval. These details include:

   Anomaly Score
   Indicates the difference in expected behavior for this specific message ID within the analysis interval. The message anomaly score is a combination of the interval contribution score for this message and the rule, if any, that is in effect for this message. Higher scores indicate greater anomaly so messages with high anomaly scores are more likely to indicate a problem. The message anomaly score ranges from 0 through 1.0.

   Interval Contribution Score
   Indicates the relative contribution of this message to the anomaly score for the analysis interval. This interval score is a function of the following analysis results that are reported in the Messages table: Rarity Score, Clustering Status, Appearance Count, and Periodicity Status. Higher scores indicate greater contribution to the interval anomaly score.

   Clustering Status
   Indicates whether this message is part of a cluster, which is an expected pattern or group of messages associated with a routine system event (for example, starting a subsystem or workload). IBM zAware identifies and recognizes these patterns or groups, and the specific messages that constitute a specific cluster. When you analyze data from a monitored client, the server determines whether a specific message is expected to be issued within a specific cluster. A message that is issued out of context (without the other messages in the same cluster) might indicate a problem.

   ID Provides the message identifier.

   - For z/OS system messages, the message identifier is a direct link to the message query in IBM Knowledge Center. In the action menu, next to the message identifier, you can also click **Knowledge Center** or **View Message History** (for example, 

). To display the search results for all occurrences of the z/OS message identifier across all clients that IBM zAware monitors, click **View Message History**, which opens the **Message History** page in a new browser tab.
   - For Linux messages, the message identifier itself is a link that you can click to open the **Message History** page in a new browser tab. The message identifier might be a known Linux system message identifier, or a message identifier that is generated by IBM zAware for its own use. You can also open a browser window and search for the Linux message description by using an internet search engine or Linux repository.

   Example or Summary
   Provides either the full message text for the first occurrence of this message within the analysis interval, or a summary of the common message text that was issued for each occurrence of the same message within the analysis interval. For summaries, only common text is displayed, with asterisks that replace any text that differs between occurrences of this message.

   To control the content displayed in this column, click either **View Message Full Text** or **View Message Summary** from the **Actions** list.

   The message anomaly score is the primary indicator of diagnostic value, so the default organization of the Message table in the Interval page arranges message entries sorted first by **Anomaly Score** in descending order, from highest value to lowest, then by **Interval Contribution Score** in descending order, and then by **Clustering Status**, with new messages listed first.

Monitoring behavior after a change

The Analysis page views and the Interval page are also useful for monitoring system behavior after you make a change to the environment, such as:

- Upgrading operating system, middleware, or application software to new levels

- Modifying system settings or configuration
- Moving a workload to a different system

In such cases, you can use the Analysis page views and the Interval page to determine whether any new unusual messages, or any more messages than you expected, were being issued immediately after the change.

Tracking down a random, intermittent problem

The Analysis page views and the Interval page are also useful for finding the cause of a random, intermittent problem. The analytical data that is available through the IBM zAware GUI can help answer the following diagnostic questions:

- Are new unusual messages issued during periods before the problem was reported, or when the problem was reported?
- Are more messages issued than expected?
- Are messages issued out of context?

Configuring the IBM zAware environment

To reap the rewards of using IBM zAware, you set up a specialized logical partition (LPAR) or IBM z Systems® Secure Service Container partition that is dedicated to running the IBM zAware server. This partition runs on a server, which is also known as a central processor complex (CPC). Figure 4 illustrates the major elements of an IBM zAware configuration with additional monitored clients running on a separate server.

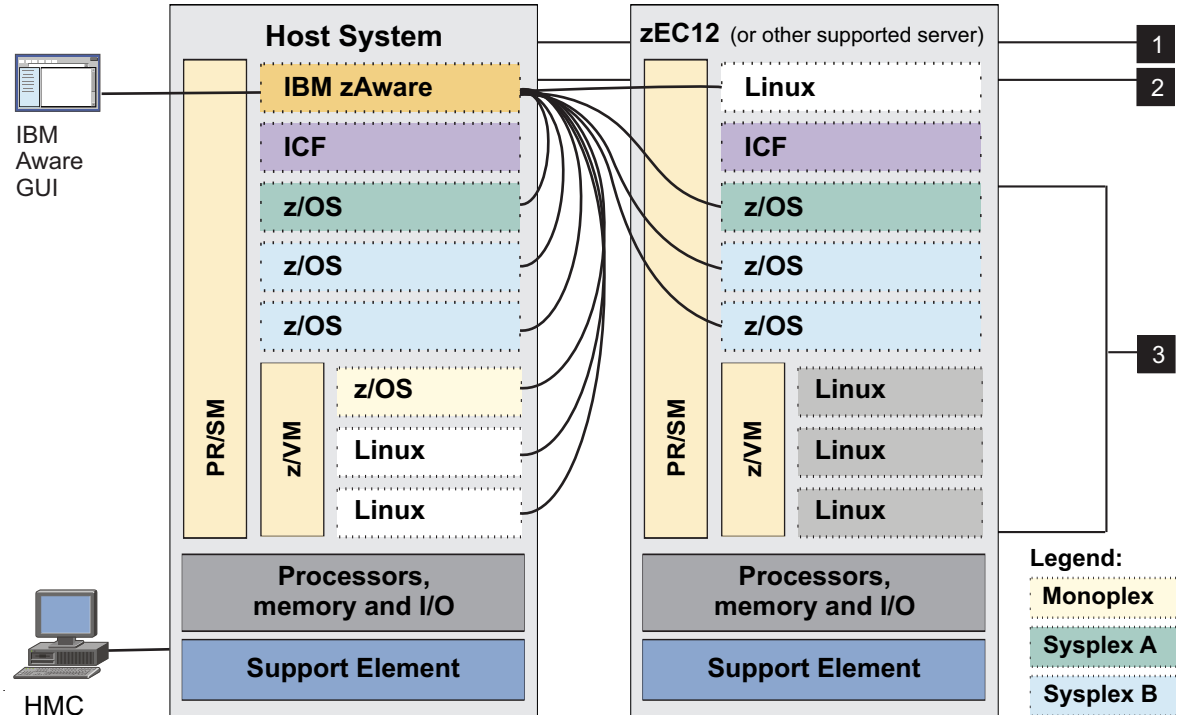

Figure 4. An IBM zAware partition supporting clients in the host system and in one zEC12

1. The server or CPC is called the IBM zAware *host system*. Only the following IBM z Systems (z Systems) servers can be host systems for IBM zAware:
   - An IBM z14 (z14)
   - An IBM z13 (z13) or IBM z13s (z13s)
   - An IBM zEnterprise EC12 (zEC12) or IBM zEnterprise BC12 (zBC12)

2. The specialized LPAR or IBM z Systems Secure Service Container partition is called the IBM zAware *partition*. An instance of IBM zAware that runs in this partition is called the IBM zAware *server*.
3. The z/OS and Linux systems that send message data to the IBM zAware server for analysis are called *monitored clients*.

A single instance of an IBM zAware server can monitor any combination of supported clients. For example, one IBM zAware server can monitor only z/OS systems, only Linux systems, or both z/OS and Linux systems.

The IBM zAware partition and all monitored clients that are sending information to the server running on that partition are collectively known as the IBM zAware *environment*. Monitored clients do not have to run in the same IBM zAware host system that contains the partition. Figure 4 on page 10 illustrates another possible configuration with additional monitored clients running on a separate server.

In Figure 4 on page 10:

- The two z/OS systems in Sysplex A (highlighted in green) are monitored clients sending data to the IBM zAware server that is running in a partition on the host system. Only one z/OS system resides on the host system; the other z/OS system resides on the zEC12.
- Similarly, the four systems in Sysplex B (highlighted in blue) are all monitored clients; two reside on the host system and two reside on the zEC12.
- The z/OS system running as a z/VM® guest, shown on the host system, is also a monitored client that is sending data to the IBM zAware server.
- Two Linux systems, also running as z/VM guests on the host system, are also monitored clients that are sending data to the IBM zAware server.
- The Linux system shown at the top of the zEC12 is also a monitored client, running in its own partition.

Using additional IBM zAware GUI functions

In addition to providing views of analytical data for monitored clients, the IBM zAware GUI provides pages through which you can manage IBM zAware operation. Most of these tasks require the user to have administrator authority to use the GUI.

- Through the **Admin** page, you can set up the **User Profile** to send email alerts when minimum and maximum anomaly thresholds you set are reached.
- Through the **Notifications** page, you can view operational messages that the IBM zAware server issues.
- Through the **Systems** page, you can view information about the monitored clients (systems) that are connected to the IBM zAware server, and manage model group definitions.
- Through **Administration** > **Configuration**, you can perform administrative tasks, such as modifying default values that control the analytics engine, adding or removing storage devices, and managing security mechanisms.
- Through **Administration** > **Training Sets**, you can view information about the generation of IBM zAware models, which are periodically updated by the server. You can also request that IBM zAware build a new model.

Figure 5 on page 12 illustrates how system management products can use the analytical data that is presented in the IBM zAware GUI.

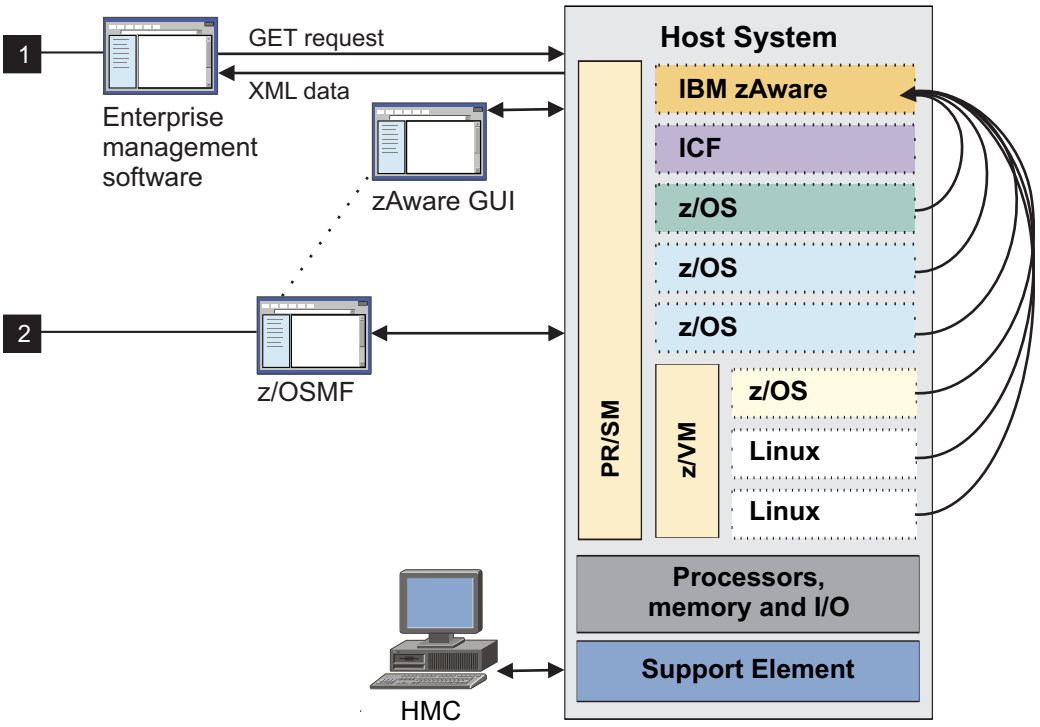

Figure 5. System management products using IBM zAware analytical data

1. Your installation can modify system management products, such as IBM Tivoli® OMEGAMON® for z/OS, to request and receive IBM zAware analytical data in XML format by using the IBM zAware application programming interface (API). This data is equivalent to the information that is available through the Analysis page views and the Interval page in the IBM zAware GUI.

   Starting with Tivoli OMEGAMON on z/OS V5.1.1, IBM zAware data is consolidated with performance and other information to support diagnoses of problems and to include in OMEGAMON XE on z/OS situations. OMEGAMON XE on z/OS provides a workspace through which users can display, manage, and customize IBM zAware data.
2. Your installation also can configure the z/OS Management Facility (z/OSMF) so that users can launch the IBM zAware GUI from the z/OSMF **Links** page.

Chapter 2. Prerequisites for configuring and using IBM zAware

This topic lists the IBM z Systems (z Systems) products for which you can order IBM zAware Version 3.1, and the prerequisites for supported types of monitored clients. The terms for use of IBM zAware are specified in the IBM Customer Agreement, Attachment for the IBM zAware Offering.

IBM zAware Version 3.1 and later is available for the following servers.

- IBM z14 (z14)
- IBM z13 (z13) and IBM z13s (z13s)
- IBM zEnterprise EC12 (zEC12) and IBM zEnterprise BC12 (zBC12)

The server on which IBM zAware runs is called the host system.

IBM zAware Version 3.1 and later runs in an IBM z Systems Secure Service Container on IBM z13 (z13), IBM z13s (z13s), and IBM z14 (z14). For other supported servers, IBM zAware is installed into a Linux type LPAR. Linux is needed for the LPAR type, but there is no need for you to set up Linux.

Additional requirements

The IBM zAware partition that runs on the host or DR system requires the following resources. For more detailed information about network connections, processors, memory, and storage for the IBM zAware partition, see Chapter 8, "Estimating data center resource requirements," on page 49.

- A shared or dedicated Open Systems Adapter (OSA) port, with an IP address that is either dedicated or assigned through Dynamic Host Connection Protocol (DHCP).
  - For OSA-Express4S or later generation features, IBM zAware can use only port 0.
  - With DHCP-type IP addresses, use of a domain name system (DNS) server is required.
- Shared or dedicated Integrated facilities for Linux (IFLs) or central processors (CPs).
- Storage and memory resources that are sufficient to support the IBM zAware server that runs on the partition and the clients that the server monitors.

For optimal performance and operations, configure the IBM zAware partition such that it has access to only those channel path identifiers (CHPIDs), control units, and I/O devices that are required for network connectivity and storage.

Requirements for z/OS systems to be monitored

To become monitored clients of the IBM zAware server, z/OS systems must meet the following requirements.

- The z/OS system must be running on a supported server:
- The z/OS system must be configured as a single-system sysplex (monoplex), a system in a multisystem sysplex, or a member of a Parallel Sysplex®.
- The system must be running a supported release of the z/OS operating system.
- The z/OS system must be using the operations log (OPERLOG) as the hardcopy medium.
- The z/OS system name and sysplex name must uniquely identify the system to be monitored. IBM zAware identifies each monitored client by sysplex and system name, in the format *sysplex_name*.*system_name*; for example: `SYSPLEX1.SYSA`. IBM zAware cannot monitor more than one system with the same sysplex and system name combination.

Requirements for Linux systems to be monitored

To become monitored clients of the IBM zAware server, Linux systems must meet the following requirements.

- A monitored Linux system can run in its own logical partition, or as a z/VM guest, on a supported z Systems server. The z/VM operating system must be a version that is supported for the z Systems server on which it runs. Supported z/VM versions are listed in the Preventative Service Planning (PSP) bucket for the z Systems server.
- Supported servers are:
  - An IBM z14 (z14)
  - An IBM z13 (z13) or IBM z13s (z13s)
  - An IBM zEnterprise EC12 (zEC12) or IBM zEnterprise BC12 (zBC12)

  Although an IBM zAware partition cannot be defined or activated on a host system that has IBM Dynamic Partition Manager (DPM) mode enabled, IBM zAware can monitor Linux systems that run on Dynamic Partition Manager-enabled servers.
- The syslog daemon for the Linux monitored system must be configured to send messages over a plain TCP transport layer to port 2003. The messages must be formatted according to the Internet Engineering Task Force (IETF) syslog protocol RFC 5424, which includes 4-digit years and time zone information. Additionally, each individual message that is transmitted must be preceded by the length of the message; this convention is known as octet framing. IBM zAware supports either rsyslog or syslog-ng as the syslog daemon on the monitored system.
  - The Linux system must correctly, consistently, and uniquely identify itself in the host name portion of the syslog message. IBM zAware interprets different but equivalent host name specifications to be different systems.
  - Each Linux system must be configured to send its syslog directly to the IBM zAware server, without consolidation with other Linux syslogs.
  - When sending syslog messages, the Linux system must provide a correct time stamp, including the Coordinated Universal Time (UTC) offset.
  - For IBM zAware to produce valuable analysis results, the syslog daemon must be configured to send at least the default level of messages, or more. With more message data, IBM zAware can more quickly build a quality model and produce valuable analysis results; message filtering through the syslog daemon has the opposite effect.
- The Linux operating system must be a distribution that was tested for the z Systems server on which it runs. The distributions that support RFC 5424 include:
  - SUSE Linux Enterprise Server (SLES) 10 or later.
  - Red Hat Enterprise Linux (RHEL) 6 or later.
  - Ubuntu 16.04

  For the recommended Linux on z Systems distribution levels and z Systems servers, see the IBM tested operating systems at this URL: www.ibm.com/systems/z/os/linux/resources/testedplatforms.html. The site contains more distributions as they become available.
- The name of a Linux system cannot exceed 230 characters.

Browser requirements for using the IBM zAware GUI

To take full advantage of the IBM zAware graphical user interface (GUI), you must use one of the following browsers. Edit your browser options to enable JavaScript, Cascading Style Sheets (CSS), and cookies. Disable software that blocks pop-up windows, especially if you are using keyboard controls rather than the mouse to use the GUI.

- Mozilla Firefox Extended Support Release (ESR) 45
- Microsoft Windows Internet Explorer (IE) 11, 10, or 9. Compatibility View must be turned off for all versions of IE.

Other browsers and browser release levels might work but are not tested; if you use them, some IBM zAware functions might not be available and page content might not display correctly.

Where to find hardware planning and corequisite software information

For the most recent hardware planning and corequisite software information, go to IBM Resource Link:

`http://www.ibm.com/servers/resourcelink`

- For hardware updates, click **Tools** on the navigation pane. Then, click **Machine information** under **Servers**, and enter your enterprise number, customer number, or machine serial number for the host system (CPC). You must register with IBM to search machine information.
- For software updates, click **Fixes** on the navigation pane. Then, click **Preventative Service Planning buckets (PSP)** under **Preventive actions**, and check the PSP bucket for the appropriate z Systems server:
  - For a z14, the 3906DEVICE PSP bucket
  - For a z13, the 2964DEVICE PSP bucket
  - For a z13s, the 2965DEVICE PSP bucket
  - For a zEC12, the 2827DEVICE PSP bucket
  - For a zBC12, the 2828DEVICE PSP bucket

Chapter 3. Project plan for configuring and using IBM zAware

System planners and installation managers collaborate with specialized IT personnel to plan, configure, and manage IBM zAware. The following checklists provide a task summary, identify the IT role or skill required for each task, and provide links to further details.

Phase 1: Planning

The planning phase includes identifying the IBM zAware host system and z/OS monitored clients, and determining datacenter resources required for IBM zAware server operation.

Table 1. Planning checklist for the IBM zAware environment

✓	Task summary:	IT role / skills:	Where to find instructions:
	Plan the configuration of the IBM zAware environment.	System planners and installation managers	Chapter 7, "Planning your IBM zAware environment," on page 39
	Plan the LPAR characteristics of the IBM zAware partition.	System planner	"Estimating processor and memory resources" on page 49
	Plan the network connections required for the IBM zAware partition and each z/OS or Linux monitored client.	Network administrator	"Planning network connections and capacity" on page 52
	Plan the physical storage capacity required to support the IBM zAware server and its monitored clients.	Storage administrator	"Planning persistent storage configuration and capacity" on page 59
	Plan the security requirements for the IBM zAware server, its monitored clients, and users of the IBM zAware graphical user interface (GUI).	Security administrator	Chapter 9, "Planning for security," on page 75
	Plan for using the IBM zAware GUI.	System planner	Chapter 10, "Planning to use the IBM zAware GUI," on page 81
	Plan to create initial IBM zAware models for monitored clients.	System programmer	Chapter 11, "Planning to create IBM zAware models," on page 87
	Plan to add IBM zAware to your local list of diagnostic resources, and to use IBM zAware data through an exiting system management product, such as IBM Operations Analytics for z Systems.	System planner	Chapter 27, "Enabling system management products to use IBM zAware data," on page 271
	Plan to add IBM zAware to your local list of diagnostic resources, and to use IBM zAware data through IBM Operations Analytics for z Systems.		

Phase 2: Configuring the IBM zAware partition

This configuration phase encompasses first-time setup tasks for the IBM zAware partition.

Table 2. Configuration checklist for the IBM zAware partition

✓	Task summary:	IT role / skills:	Where to find instructions:
	Verify that your installation meets the prerequisites for using IBM zAware.	System programmer	Chapter 2, "Prerequisites for configuring and using IBM zAware," on page 13

Table 2. Configuration checklist for the IBM zAware partition (continued)

✓	Task summary:	IT role / skills:	Where to find instructions:
	Configure network connections for the IBM zAware partition through the Hardware Configuration Definition (HCD) or the Input/Output Configuration Program (IOCP).	Network administrator	Step 1 on page 96 in Chapter 12, "Configuring network connections and storage for the IBM zAware partition," on page 95
	Configure persistent storage for the IBM zAware partition through the HCD or IOCP.	Storage administrator	Step 2 on page 97 in Chapter 12, "Configuring network connections and storage for the IBM zAware partition," on page 95
	Define the LPAR characteristics of the IBM zAware partition through the Hardware Management Console (HMC).	System programmer	"Configuring the IBM z Systems Secure Service Container for IBM zAware" on page 27
	Define network settings for the IBM zAware partition through the HMC.	Network administrator	"Configuring the IBM z Systems Secure Service Container for IBM zAware" on page 27.
	Activate the IBM zAware partition through the HMC.	System programmer	Follow the operating procedures at your company for activating a partition on a z13 or z14 CPC.

Phase 3: Configuring the IBM zAware server and its monitored clients

This configuration phase encompasses first-time setup tasks for the IBM zAware server and its z/OS monitored clients.

Table 3. Configuration checklist for the IBM zAware server and monitored clients

✓	Task summary:	IT role / skills:	Where to find instructions:
	Assign storage devices for the IBM zAware server through the IBM zAware GUI.	Storage administrator	Step 2 on page 100 in Chapter 13, "Configuring storage, security, and analytics for the IBM zAware server," on page 99
	Optional: Replace the default Secure Sockets Layer (SSL) certificate that is configured in the IBM zAware server.	Security administrator	Step 3 on page 102 in Chapter 13, "Configuring storage, security, and analytics for the IBM zAware server," on page 99
	Configure an LDAP repository or local file-based repository for authenticating users of the IBM zAware GUI.	Security administrator or LDAP administrator	• To configure an LDAP repository, see step 4 on page 103 in Chapter 13, "Configuring storage, security, and analytics for the IBM zAware server," on page 99. • To configure a file-based repository, see "Setting up a local repository to secure access to the IBM zAware GUI" on page 108.
	Authorize users or groups to access the IBM zAware GUI.	Security administrator	Step 5 on page 106 in Chapter 13, "Configuring storage, security, and analytics for the IBM zAware server," on page 99
	Modify the configuration values that control IBM zAware analytics operation.	System programmer	Step 8 on page 107 in Chapter 13, "Configuring storage, security, and analytics for the IBM zAware server," on page 99

Table 3. Configuration checklist for the IBM zAware server and monitored clients (continued)

✓	Task summary:	IT role / skills:	Where to find instructions:
	Configure a network connection for each z/OS or Linux monitored client. If necessary, update firewall settings.	Network administrator	• Step 1 on page 112 in "Configuring z/OS monitored clients to send data to the IBM zAware server" on page 111. • Step 1 on page 130 in "Configuring Linux on z Systems monitored clients to send data to the IBM zAware server" on page 129
	Prepare each z/OS system to be monitored for operation with IBM zAware.	z/OS system programmer	"Configuring z/OS monitored clients to send data to the IBM zAware server" on page 111
	Prepare each Linux system to be monitored for operation with IBM zAware.	Linux administrator	"Configuring Linux on z Systems monitored clients to send data to the IBM zAware server" on page 129
	Build a model of normal system behavior for IBM zAware to use for analysis.	z/OS system programmer or Linux administrator	• "Creating an IBM zAware model for new z/OS monitored clients" on page 118 • "Creating an IBM zAware model for new Linux on z Systems monitored clients" on page 132

Phase 4: Daily operations

In this phase, the primary activity is viewing IBM zAware analytical data to find and diagnose anomalies in the behavior of z/OS and Linux monitored clients. z/OS system programmers, Linux system programmers, and experienced application programmers are the most likely IT personnel to participate in this activity. Topics in the following parts describe this ongoing activity and other occasional management tasks.

- Part 5, "Managing and using the IBM zAware server," on page 137
- "Setting up the User Profile to send email alerts" on page 83
- Part 6, "Advanced topics for managing IBM zAware," on page 219

Part 2. Installing and upgrading to IBM zAware V3.1

Topics in this part explain the installation and migration considerations that system planners, installation managers and network, storage and security administrators need to know before they start planning to configure IBM zAware and its operating environment.

Topics covered in this part are:

Chapter 4. Overview of installing or upgrading to IBM zAware V3.1

The installation and upgrade depends on the server model on which you are hosting IBM zAware. Use the following topic to decide which instruction set to use for the installation or upgrade to IBM zAware V3.1 and above.

Use one or more of the following instructions to install or upgrade to IBM zAware V3.1 and above.

- To install or upgrade to IBM zAware V3.1 on z Systems server and above, use the following tasks:
  - Chapter 5, “Installing or upgrading IBM zAware V3.1 on z Systems servers,” on page 25
  - “Configuring the IBM z Systems Secure Service Container for IBM zAware” on page 27
- To install or upgrade to IBM zAware V3.1 on an zEnterprise server, follow the instructions for the following tasks:
  - Chapter 6, “Installing IBM zAware V3.1 on an IBM zEnterprise server,” on page 29
  - “Modifying the Bootstrap Configuration for IBM zAware” on page 33

Chapter 5. Installing or upgrading IBM zAware V3.1 on z Systems servers

The following information provides the procedures for installing or upgrading IBM z Advanced Workload Analysis Reporter (IBM zAware) V3.1 to a Secure Service Container on an IBM z Systems (z Systems) server.

Before you begin

If you are upgrading to IBM zAware V3.1, from the IBM zAware GUI, go to **Configuration** > **Utilities**, and save your configuration. For instructions, see Chapter 25, "Restoring IBM zAware configuration data," on page 265.

If you are installing IBM zAware V3.1 for the first time, review the following topics before you begin:

- Chapter 2, "Prerequisites for configuring and using IBM zAware," on page 13
- Chapter 3, "Project plan for configuring and using IBM zAware," on page 17
- Part 3, "Planning to configure IBM zAware," on page 37

Important: You must have 6 GB of free storage to install the new IBM zAware V3.1 image.

About this task

Use the following task to upgrade or install IBM zAware V3.1 on a Linux partition in an IBM z Systems server.

Procedure

1. Order *IBM Operations Analytics for z Systems* with the IBM zAware feature code (5698-ABH) from IBM Shopz at https://www.ibm.com/software/shopzseries/ShopzSeries_public.wss. If you are installing IBM zAware for the first time, complete the following steps:
   a. Review the information in *IBM Operations Analytics for z Systems IBM zAware Program Directory*, GI13-4170-00, which is included with your order.
   b. Review the complete information about network connections, processors, memory, and storage for the IBM zAware partition, see Chapter 8, "Estimating data center resource requirements," on page 49.

   If you are upgrading to IBM zAware V3.1, be sure to note which storage device numbers are in use by IBM zAware. You need to know the exact numbers when you log back in to the IBM zAware GUI.
2. Go to the HMC, and enter the activation profile for the IBM zAware partition.
   a. Go to the **General** tab, and ensure that the IBM z Systems Secure Service Container is selected. (The IBM z Systems Secure Service Container is the partition that previously ran in zACI-mode.)
   b. Go to the IBM z Systems Secure Service Container tab, and then select **Secure Container Installer**.
   c. Apply the changes to the activation profile, and then save and close. You are back to the HMC.
3. Deactivate the LPAR, and then reactivate it. To see the status of the activation, go to the HMC, right-click the **IBM zAware LPAR**, and then select **Daily** > **Operating system messages**. When the activation is complete, you can find the IP address for the software container installer in the messages. The following example shows how to locate the IP address.

```
Please connect to the IBM z Systems Secure Service Container web UI via your browser
The server is listening on: 192.0.2.0
Network Interface Summary:
Interface IP Address
=====================================================================
enccw0.0.1a80 [IPv4]  192.0.2.0
enccw0.0.1a80 [IPv6] 2001:0DB8:0:0:0:0:0:0/32
```

Figure 6. Locating the IP address

4. Open the installer for IBM z Systems Secure Service Container in your browser. For example, if **192.0.2.0** is the IP address you found in the messages, you want to enter **https://192.0.2.0** in your browser.
5. Click the plus sign in the upper right portion of the Installer, and then **Browse** to select the IBM zAware V3.1 image.
   a. Set the storage device to be formatted and store the IBM zAware software.

      Important:
      - The storage device **must** be different than any previously used IBM zAware devices.
      - You must have 6 GB of free storage to install the IBM zAware V3.1 software image.

      When you are done, the storage device is formatted and the IBM zAware software is installed.
   b. Click the **Reboot Automatically** check box. The restart process might take up to 15 minutes.
6. Enter the IP address that you previously detected to log in to the IBM zAware V3.1 GUI in a browser. For example, enter **https://192.0.2.0** in your browser (as described earlier).
   - If IBM zAware is new to your installation, you are ready to complete the following steps:
     – Part 4, “Configuring IBM zAware and its monitored clients,” on page 93
     – Part 5, “Managing and using the IBM zAware server,” on page 137.
   - If you are upgrading to IBM zAware V3.1, complete the following steps.
     a. Add the existing storage devices back by using the **Preserve** option as shown in Figure 7 on page 27. For more information, review Chapter 13, “Configuring storage, security, and analytics for the IBM zAware server,” on page 99.

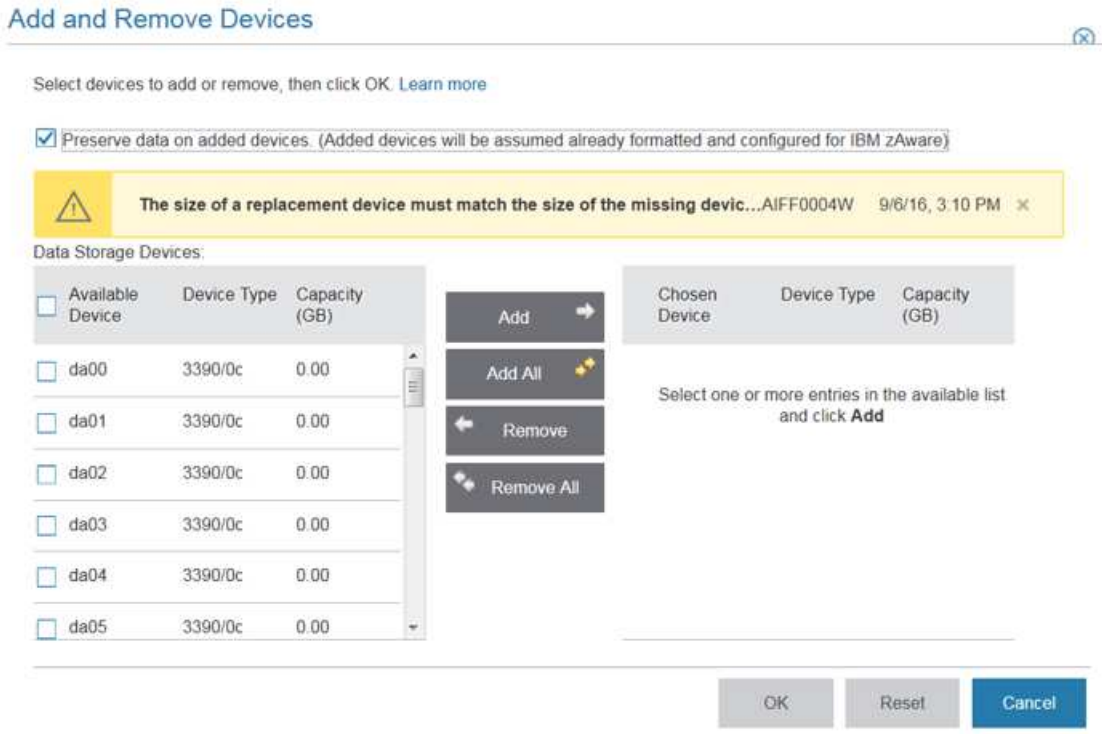

Figure 7. Preserving IBM zAware device data

b. Restore your existing configurations through the **Configuration** > **Utilities** tab. When the IBM zAware partition is reactivated at any time after its initial configuration, IBM zAware attempts to reconnect to the previously configured Lightweight Directory Access Protocol (LDAP) server. For more information, see "LDAP Settings tab" on page 185.

Configuring the IBM z Systems Secure Service Container for IBM zAware

Use the links in the following task to configure a logical partition (LPAR) for IBM zAware in an IBM z Systems Secure Service Container.

Before you begin

If you upgrade or install IBM zAware V3.1 or later on an IBM zEnterprise BC12 (zBC12) or IBM zEnterprise EC12 (zEC12), see Chapter 6, "Installing IBM zAware V3.1 on an IBM zEnterprise server," on page 29

About this task

Install the IBM z Systems Secure Service Container for IBM zAware. The service container is a partition that runs in a secure service container (previously called zACI-mode partition) for IBM z Systems servers. For complete information, see the following publication.

- The *IBM z Systems Secure Service Container User's Guide*, SC28-6978-02, in the IBM Support Portal at https://ibm.biz/Bd2Kqe.

Results

When you complete the steps, the secure service container for IBM zAware is configured. You are ready to install IBM zAware on a supported IBM z Systems (z Systems) server. For more information, see Chapter 5, “Installing or upgrading IBM zAware V3.1 on z Systems servers,” on page 25.

Chapter 6. Installing IBM zAware V3.1 on an IBM zEnterprise server

The following topic describes how to upgrade or install IBM zAware V3.1 on an IBM zEnterprise EC12 (zEC12) or IBM zEnterprise BC12 (zBC12).

Before you begin

The following instructions assume that a system administrator created a Linux partition and defined the network connections and storage devices for the partition.

If you are installing IBM zAware for the first time, review the following topics before you begin:

- Chapter 2, "Prerequisites for configuring and using IBM zAware," on page 13
- Chapter 3, "Project plan for configuring and using IBM zAware," on page 17
- Part 3, "Planning to configure IBM zAware," on page 37

About this task

Use the following procedure to install or upgrade to IBM zAware V3.1 in a Linux LPAR.

Procedure

1. Order *IBM Operations Analytics for z Systems* with the IBM zAware feature code (5698-ABH) from IBM Shopz at https://www.ibm.com/software/shopzseries/ShopzSeries_public.wss. If you are installing IBM zAware on IBM zEnterprise EC12 (zEC12) or IBM zEnterprise BC12 (zBC12) for the first time, use the following steps.
   a. Review the information in *IBM Operations Analytics for z Systems IBM zAware Program Directory*, GI13-4170-00, which is included with your order.
   b. Review the complete information about network connections, processors, memory, and storage for the IBM zAware partition, see Chapter 8, "Estimating data center resource requirements," on page 49.
   c. Note the exact storage device to which you are installing the Linux Software Container.

   If you are upgrading to IBM zAware V3.1, note all of the storage device numbers that are in use by IBM zAware. You must know the exact numbers when you log in to IBM zAware V3.1.
2. Copy the DVD installation image to an FTP server.
3. Run the bootstrap configuration script that is included with your order to configure the Login and Network settings. For instructions, see "Modifying the Bootstrap Configuration for IBM zAware" on page 33.
4. Add the completed `bootstrapSettings.json` file to the FTP server where you copied the installation image.
5. Point to the IBM zAware V3.1 Service Container that came with your order, named "`IBM Operations Analytics for z Systems Service Container installer`" to the FTP server. For example, in Figure 8 on page 30, you want to FTP the `IBM Operations Analytics for z Systems Service Container installer` file to the Linux LPAR that the system administrator created. In the following example, the LPAR is named "ZAWAREB".

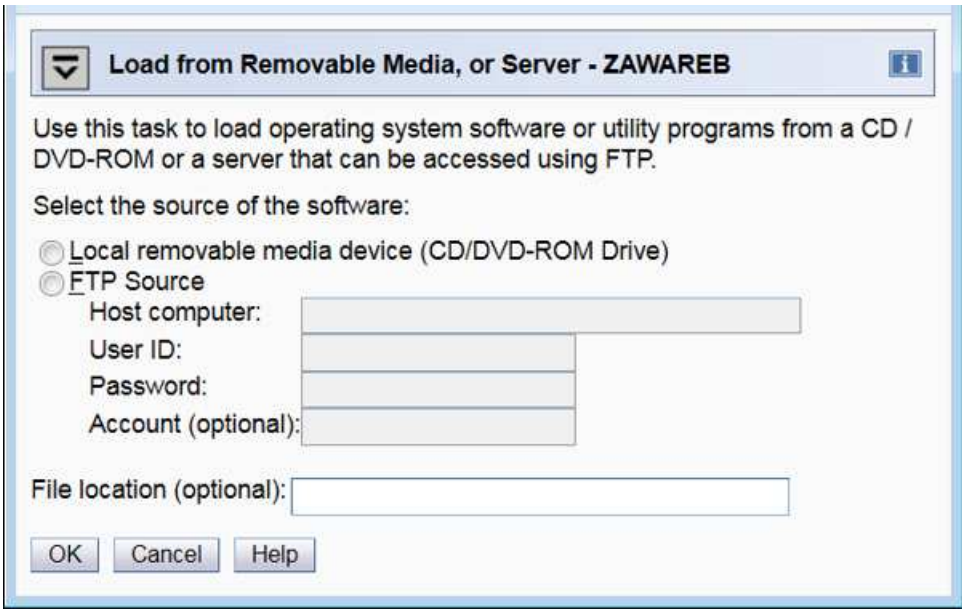

Figure 8. FTP option in Load from Removable Media or Server

6. Install the IBM zAware Software Appliance, to the exact storage device number that you noted earlier. Figure 9 represents an example only.

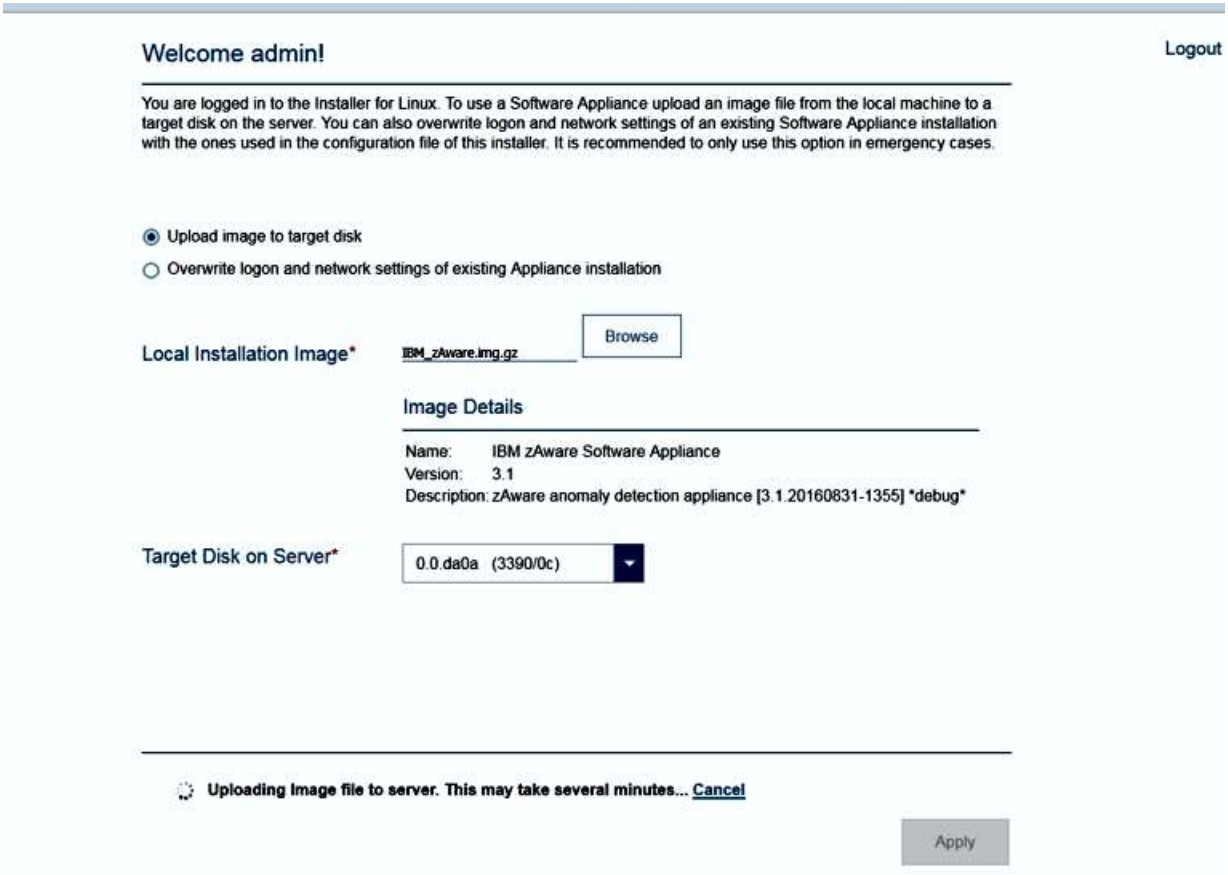

Figure 9. Load the IBM zAware Software Appliance

Note: In Figure 9 on page 30, the Login and Network settings are subject to change. You can disregard the "...to only use this option in emergency cases" statement.
You will see a confirmation window, as shown in Figure 10, which explains that a manual reboot is required after the IBM zAware Software Appliance is installed.

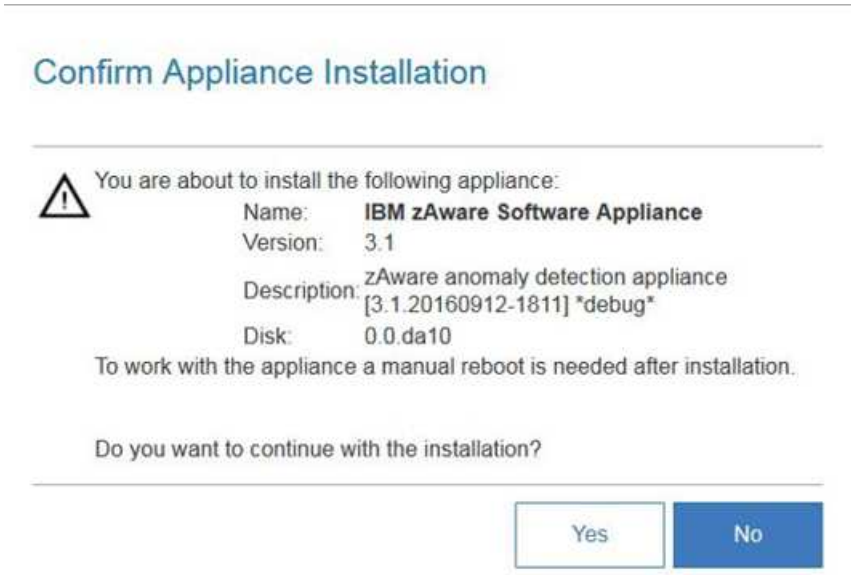

Figure 10. Confirm Appliance Installation window

7. Go to the HMC, and restart the LPAR on which you installed IBM zAware Software Appliance. When you are done, the **Confirmation Appliance Installation** window indicates success as shown in Figure 11.

Figure 11. Success in the Appliance Installation window

8. Go to the HMC to set the **Load Address** field for the IBM zAware Software Appliance storage device. Click **CPC Operational Customization** > **Customize/Delete Activation Profiles**, and then enter the device number in the **Load Address** field (this is the same storage device that you noted

earlier.) Apply the changes, and then save and close. When you are done, except for the **Load Address** field and LPAR name, the **Customize Image Profile** looks similar to the Figure 12 window.

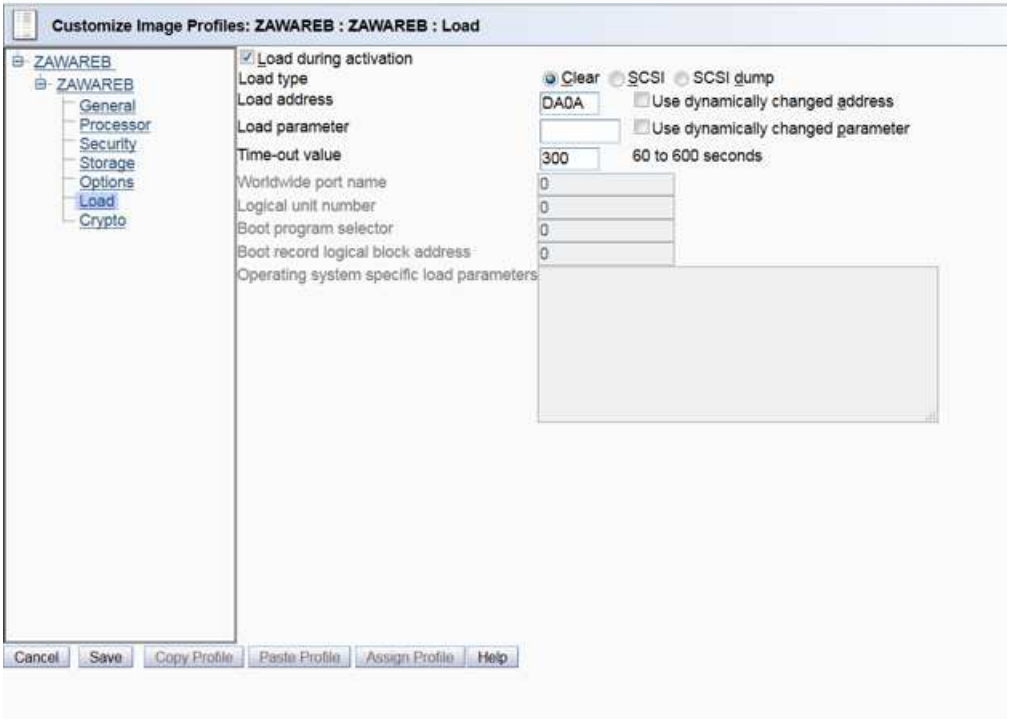

Figure 12. Customize Image Profile

9. Go back to the HMC, and reactivation the partition that you created.
10. Connect to the IP address that was configured for the network adapter of the Linux partition. If you do not know the IP address, open the HMC, and then click **Daily** > **Operating system messages**. Look for the IP address, which looks similar to the following example.

```
Please connect to the Container Installer web UI via your browser
The server is listening on: 192.0.2.0
Network Interface Summary:
Interface IP Address
======================================================================
enccw0.0.1a80 [IPv4]  192.0.2.0
enccw0.0.1a80 [IPv6] 2001:0DB8:0:0:0:0:0:0/32
```

11. Log in to IBM zAware V3.1 and configure the storage, security, and analytics.
    - If this is an IBM zAware V3.1 new installation, you are ready to complete the following steps:
      - Part 4, “Configuring IBM zAware and its monitored clients,” on page 93
      - Part 5, “Managing and using the IBM zAware server,” on page 137.
    - If you are upgrading to IBM zAware V3.1, complete the following steps.
      a. Add the existing direct access storage device (DASD) back by using the **Preserve** option as shown in Figure 13 on page 33. For more information, review Chapter 13, “Configuring storage, security, and analytics for the IBM zAware server,” on page 99.

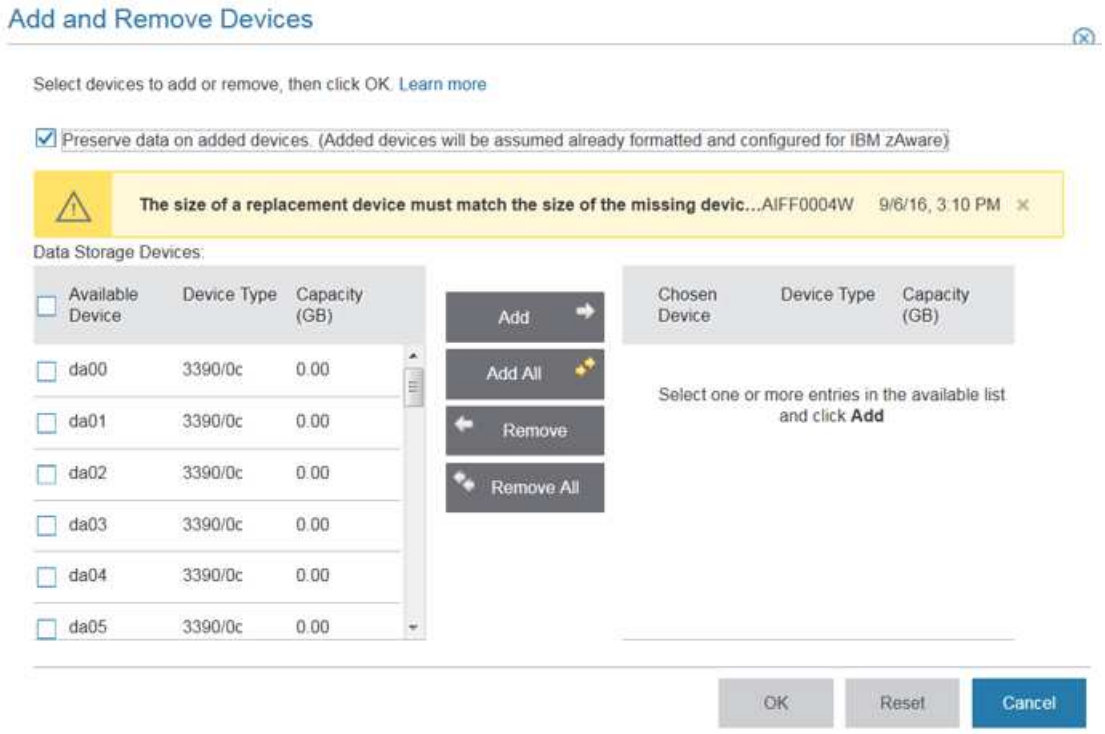

Figure 13. Example for preserving device data

b. Restore your existing configurations through the **Configuration** > **Utilities** tab. When the IBM zAware partition is reactivated at any time after its initial configuration, IBM zAware attempts to reconnect to the previously configured Lightweight Directory Access Protocol (LDAP) server. For more information, see "LDAP Settings tab" on page 185.

Activation profile information in IBM Knowledge Center

http://www.ibm.com/support/knowledgecenter/HW11R/com.ibm.hwmca.kc_hmc.doc/introductiontotheconsole/introductionaboutactivationprofiles.html

Modifying the Bootstrap Configuration for IBM zAware

Use the IBM z Advanced Workload Analysis Reporter (IBM zAware) **Bootstrap Configuration** when you install or upgrade to IBM zAware V3.1 on an IBM zEnterprise System (zEnterprise) server.

Before you begin

The Bootstrap Configuration utility is packaged with IBM zAware V3.1 to upgrade or install IBM zAware on an zEnterprise server. The package contains the following files and one folder:

- `createBoostrapConfig.bat` for Windows
- `createBoostrapConfig.sh` for Linux
- The "lib" folder that contains the JAR file.

The bootstrap utility requires Java™ 1.7 to run. You must set either the JAVA_HOME environment variable, or you must be able to locate "java" in the user's PATH.

About this task

The bootstrap configuration script is packaged with IBM zAware. Use it to configure the user and network settings.

Procedure

1. Run the correct script for your operating system configuration.
   a. For Windows users, run **createBootstrapConfig.bat** from a command prompt.
   b. For Linux users, run **./createBoostrapConfig.sh** from a command line.

   After the utility runs, it prompts you for a password. The password is encrypted, and then inserted into the `boostrapSettings.json` template file that gets created. You can also update the newly created or existing `boostrapSettings.json` file to give it a new password.
2. Open the `boostrapSettings.json` file, and then modify it to contain the following information.
   - Host name
   - User ID
   - CHPID
   - DNS
   - Gateway for the Service Container

   By running the utility with the information already filled out in the `bootstrapSettings.json` file, you can update the new password and keep the remaining fields in the file intact.
   The following example shows a completed `bootstrapSettings.json` file with a condensed password hash.

```
{
  "hostName": "zawareb",
  "masterUser": "admin",
  "masterPasswordHash": "$6$UtFPJ9J5LVCk/lrx$FfZxpi/iex..../WykqreOEjv.zu/",
  "masterEncrypted": "unused",
  "setAdminCredentials": false,
  "setNetworkSettings": false,
  "networkCards": [
     {
        "CHPID": "0x2",
        "staticIPv4": "192.10.2.107\/24",
        "vlan": "507"
     }
  ],
  "DNS": [
     "192.10.16.2",
     "192.0.130.50"
  ],
  "defaultGateway": "192.10.2.1"
}
```

Figure 14. Completed bootstrap file example

What to do next

Go back to Chapter 6, "Installing IBM zAware V3.1 on an IBM zEnterprise server," on page 29 to complete the installation or upgrade.

Part 3. Planning to configure IBM zAware

Topics in this part explain the planning considerations that system planners, installation managers and network, storage and security administrators need to know before they start configuring IBM zAware and its operating environment.

Topics covered in this part are:

Chapter 7. Planning your IBM zAware environment

The IBM zAware environment consists of an IBM zAware partition, the IBM zAware server that is running on that partition, and all of the monitored clients that are connected to the server. This topic provides illustrations of sample configurations for the IBM zAware environment, and introduces planning considerations that are covered in more detail in other planning chapters.

Supported servers that can host IBM zAware

Every supported configuration must include an IBM zAware host system, which is the central processor complex (CPC) in which the IBM zAware partition runs. Only the following servers can be host systems for IBM zAware:

- An IBM z14 (z14)
- An IBM z13 (z13) or IBM z13s (z13s)
- An IBM zEnterprise EC12 (zEC12) or IBM zEnterprise BC12 (zBC12)

The IBM zAware partition that runs on the host system requires the following resources:

- A shared or dedicated Open Systems Adapter (OSA) port, with an IP address that is either dedicated or assigned through Dynamic Host Connection Protocol (DHCP).
  - For OSA-Express4S or later generation features, IBM zAware can use only port 0.
  - With DHCP-type IP addresses, use of a domain name system (DNS) server is required.
- Shared or dedicated Integrated facilities for Linux (IFLs) or central processors (CPs).
- Storage and memory resources that are sufficient to support the IBM zAware server that runs on the partition and the clients that the server monitors.

Chapter 8, "Estimating data center resource requirements," on page 49 provides more detail about these resource requirements.

Supported types of IBM zAware monitored clients

IBM zAware supports z/OS and Linux systems as monitored clients. The number of monitored clients is limited only by the resources that are assigned to the IBM zAware partition.

A single instance of an IBM zAware server can monitor any combination of supported clients. For example, one IBM zAware server can monitor only z/OS systems, only Linux systems, or both z/OS and Linux systems.

z/OS systems

IBM zAware supports z/OS systems that run in z/OS partitions or as z/VM guests. z/OS monitored clients must meet the following requirements:

- The z/OS system must be running on a supported server:
- The z/OS system must be configured as a single-system sysplex (monoplex), a system in a multisystem sysplex, or a member of a Parallel Sysplex.
- The system must be running a supported release of the z/OS operating system.
- The z/OS system must be using the operations log (OPERLOG) as the hardcopy medium.
- The z/OS system name and sysplex name must uniquely identify the system to be monitored. IBM zAware identifies each monitored client by sysplex and system name, in the format *sysplex_name*.*system_name*; for example: SYSPLEX1.SYSA. IBM zAware cannot monitor more than one system with the same sysplex and system name combination.

Linux systems

To become monitored clients of the IBM zAware server, Linux systems must meet the following requirements.

- A monitored Linux system can run in its own logical partition, or as a z/VM guest, on a supported z Systems server. The z/VM operating system must be a version that is supported for the z Systems server on which it runs. Supported z/VM versions are listed in the Preventative Service Planning (PSP) bucket for the z Systems server.
- Supported servers are:
  - An IBM z14 (z14)
  - An IBM z13 (z13) or IBM z13s (z13s)
  - An IBM zEnterprise EC12 (zEC12) or IBM zEnterprise BC12 (zBC12)

  Although an IBM zAware partition cannot be defined or activated on a host system that has IBM Dynamic Partition Manager (DPM) mode enabled, IBM zAware can monitor Linux systems that run on Dynamic Partition Manager-enabled servers.
- The syslog daemon for the Linux monitored system must be configured to send messages over a plain TCP transport layer to port 2003. The messages must be formatted according to the Internet Engineering Task Force (IETF) syslog protocol RFC 5424, which includes 4-digit years and time zone information. Additionally, each individual message that is transmitted must be preceded by the length of the message; this convention is known as octet framing. IBM zAware supports either rsyslog or syslog-ng as the syslog daemon on the monitored system.
  - The Linux system must correctly, consistently, and uniquely identify itself in the host name portion of the syslog message. IBM zAware interprets different but equivalent host name specifications to be different systems.
  - Each Linux system must be configured to send its syslog directly to the IBM zAware server, without consolidation with other Linux syslogs.
  - When sending syslog messages, the Linux system must provide a correct time stamp, including the Coordinated Universal Time (UTC) offset.
  - For IBM zAware to produce valuable analysis results, the syslog daemon must be configured to send at least the default level of messages, or more. With more message data, IBM zAware can more quickly build a quality model and produce valuable analysis results; message filtering through the syslog daemon has the opposite effect.
- The Linux operating system must be a distribution that was tested for the z Systems server on which it runs. The distributions that support RFC 5424 include:
  - SUSE Linux Enterprise Server (SLES) 10 or later.
  - Red Hat Enterprise Linux (RHEL) 6 or later.
  - Ubuntu 16.04

  For the recommended Linux on z Systems distribution levels and z Systems servers, see the IBM tested operating systems at this URL: www.ibm.com/systems/z/os/linux/resources/testedplatforms.html. The site contains more distributions as they become available.
- The name of a Linux system cannot exceed 230 characters.

Figure 15 on page 41 illustrates all types of monitored clients that are supported by IBM zAware Version 2.0 or later.

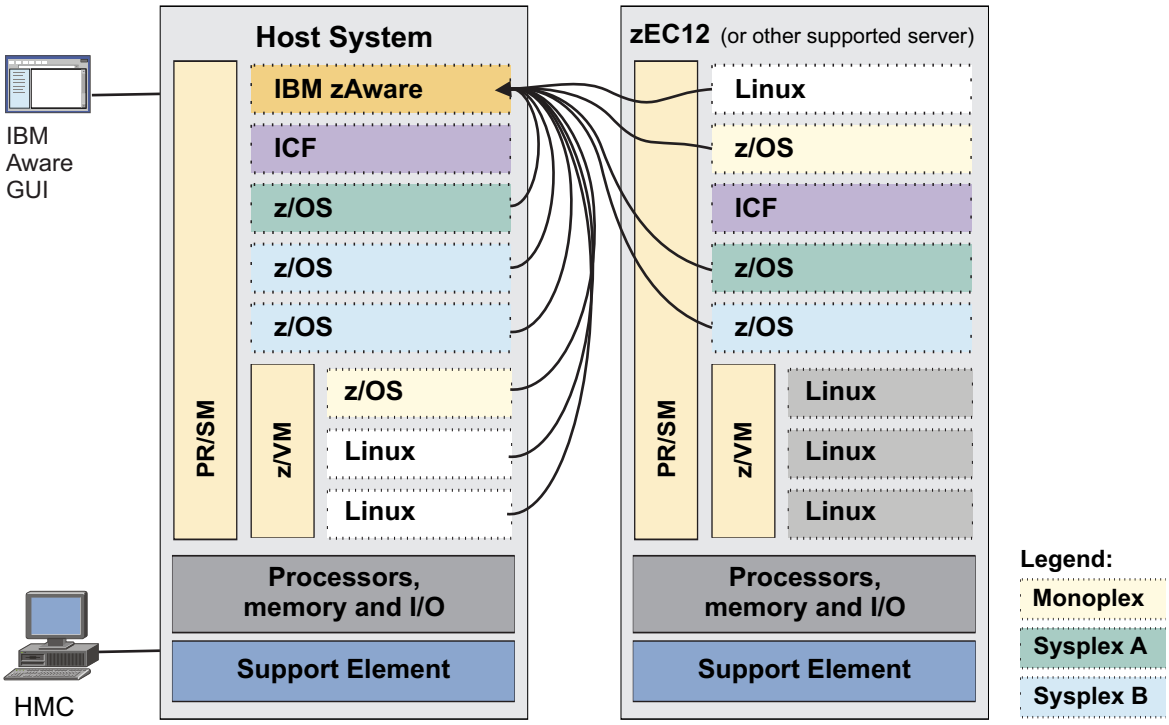

Figure 15. Types of monitored clients connected to a IBM zAware partition

The IBM zAware environment in Figure 15 contains one IBM zAware server that runs in a partition on the host system, and monitored clients that run on either the host system or another supported server, such as the zEC12. Monitored clients that run on a z Systems server other than the host system are called *remote clients*.

- The host system contains partitions in which the following monitored clients run:
  - One z/OS system in Sysplex A (highlighted in green)
  - Two systems in Sysplex B (highlighted in blue)
  - One z/OS system, which is configured as a single-system sysplex (monoplex), running as a z/VM guest
  - Two Linux systems that are running as z/VM guests
- The zEC12 on the right contains the remote clients:
  - One Linux system that is running in a PR/SM™ partition
  - One z/OS system, which is configured as a single-system sysplex (monoplex)
  - Another z/OS system in Sysplex A (highlighted in green)
  - Another z/OS system in Sysplex B (highlighted in blue)

  The three z/VM Linux guests also can become monitored clients, but are not currently configured for monitoring.

In addition to the previously listed system prerequisites, the workload type and message traffic determine which z/OS systems are candidates for monitoring through IBM zAware.

- z/OS systems that run production workloads are good choices because of their importance to your business, and your reliance on their availability. z/OS production systems also tend to have high-volume, consistent message traffic that IBM zAware can use to build a useful model of normal system behavior. Chapter 11, "Planning to create IBM zAware models," on page 87 describes the specific attributes of message traffic that are required to create an IBM zAware model.
- z/OS quality assurance (QA) systems are also well-suited for monitoring through IBM zAware.

- Test and development systems, however, might not warrant configuration as monitored clients. Because of their inherent unpredictability, you might find it difficult to generate a model of normal system behavior that yields useful analytical data.

Linux systems that run production workloads are also good candidates for monitoring through IBM zAware. In addition to the previously listed system prerequisites, each Linux system must belong to an administrator-defined model group for IBM zAware to analyze message traffic and display analysis results. A model group is a collection of one or more systems that handle the same type of workload, and thus can be expected to exhibit similar behavior. Model group definitions are based on Linux system naming conventions.

Through model groups, multiple systems contribute to the generation of a single model for the group. The use of a model group supports dynamic activation and deactivation of Linux images, because IBM zAware can use the group model to analyze a Linux system as soon as it connects to the IBM zAware server.

Matching monitored clients to a specific IBM zAware server

The physical distance between an IBM zAware server and its monitored clients can determine the configuration of your IBM zAware environment. IBM testing experiences indicate acceptable IBM zAware capability and performance up to a maximum distance of 3500 kilometers (approximately 2174 miles) between the host system and monitored clients that run on other z Systems servers.

Figure 16 on page 43 illustrates an IBM zAware configuration that requires two separate servers because of the distance between the two company sites. For this example, the company configures two IBM zAware partitions to monitor the production systems that are running on four z Systems servers:

- In the host system for the Boston site (shown on the upper left of the figure), the IBM zAware server monitors clients that are on the host system and on two other supported z Systems servers at the Boston site. One z/OS system is configured as a single-system sysplex (monoplex); the remaining clients are members of either Sysplex A (highlighted in green) or Sysplex B (highlighted in blue).
- In the host system for the Los Angeles site (shown on the lower left of the figure), another IBM zAware server monitors clients that are only on the host system at the Los Angeles site. The z/OS clients are all members of Sysplex C (highlighted in dark blue). The Linux clients all run as z/VM guests. Because this host system supports IBM zAware Version 2.0 or later, several Linux images are also IBM zAware monitored clients.

Notice that each site that is illustrated in Figure 16 on page 43 has its own firewall, and that each IBM zAware environment— the server and all of its monitored clients— are protected within that firewall. Because IBM zAware does not provide any encryption or authentication for communication from monitored clients, the recommended configuration is to protect the IBM zAware and its clients with the same firewall.

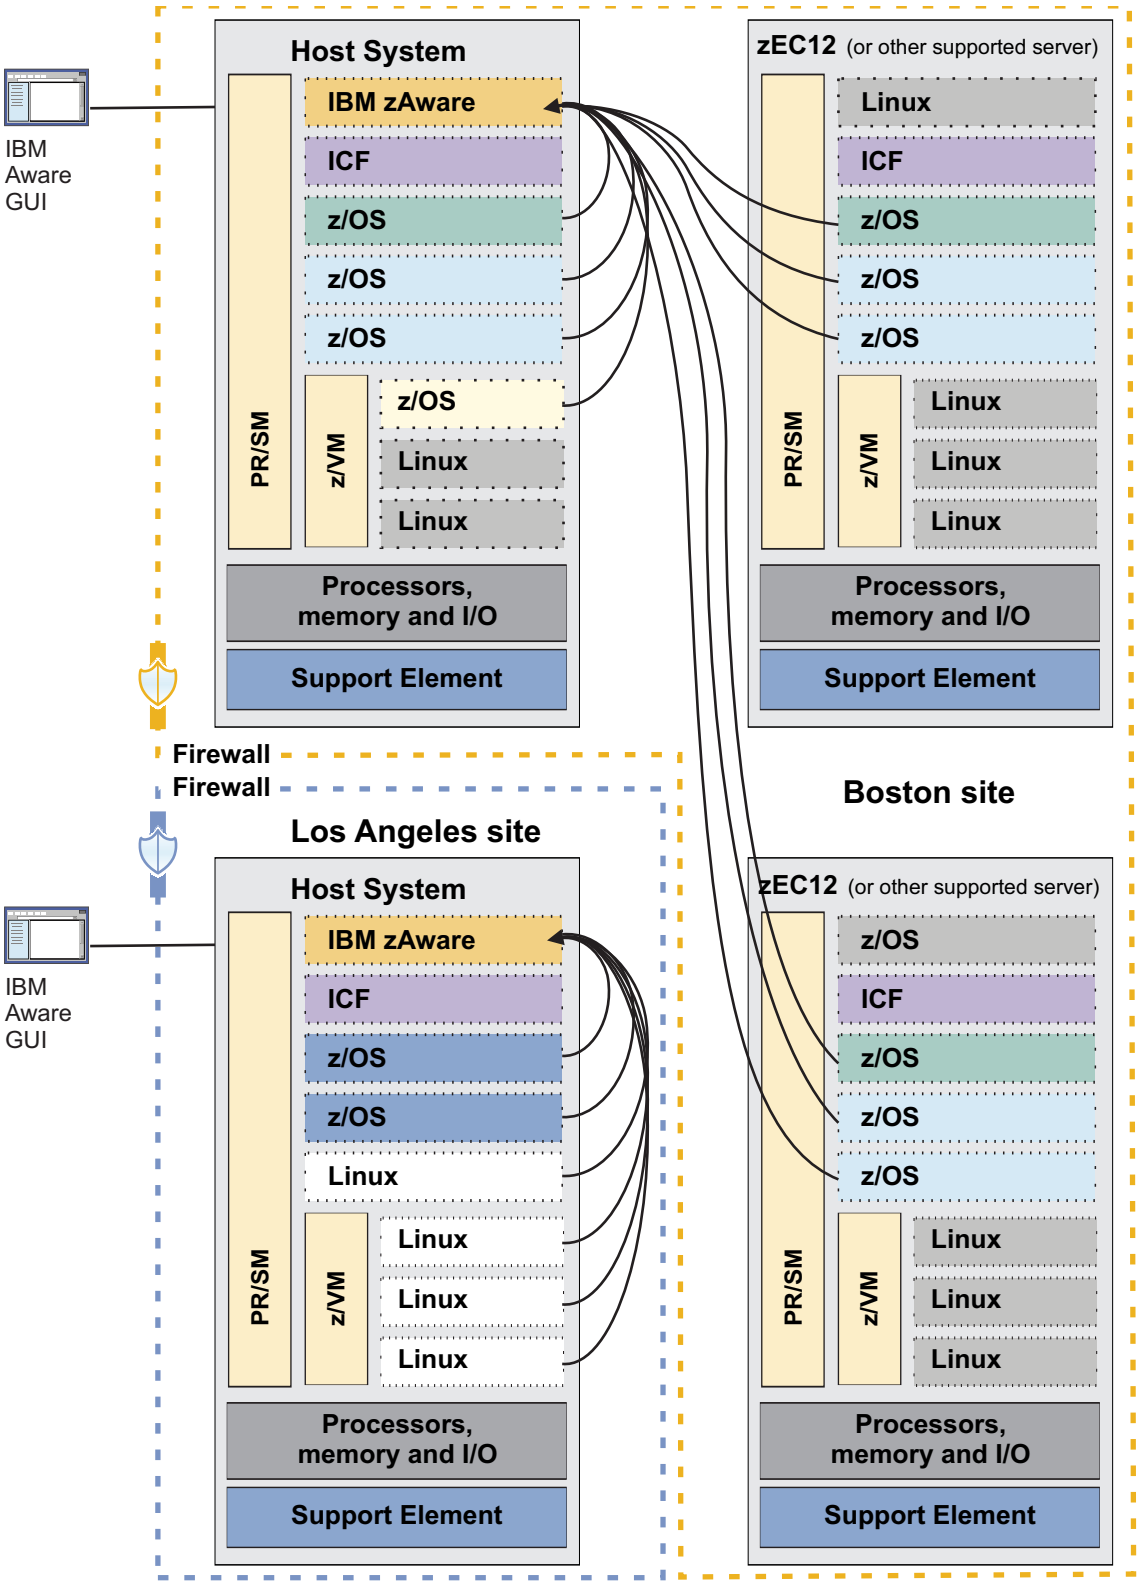

Figure 16. An IBM zAware configuration with two partitions, each supporting the clients in one of two physical sites

Figure 17 on page 44 illustrates another possible configuration with two IBM zAware partitions. Because the host system has multiple firewalls, the company configured each IBM zAware server and its clients to

be protected within the same firewall.

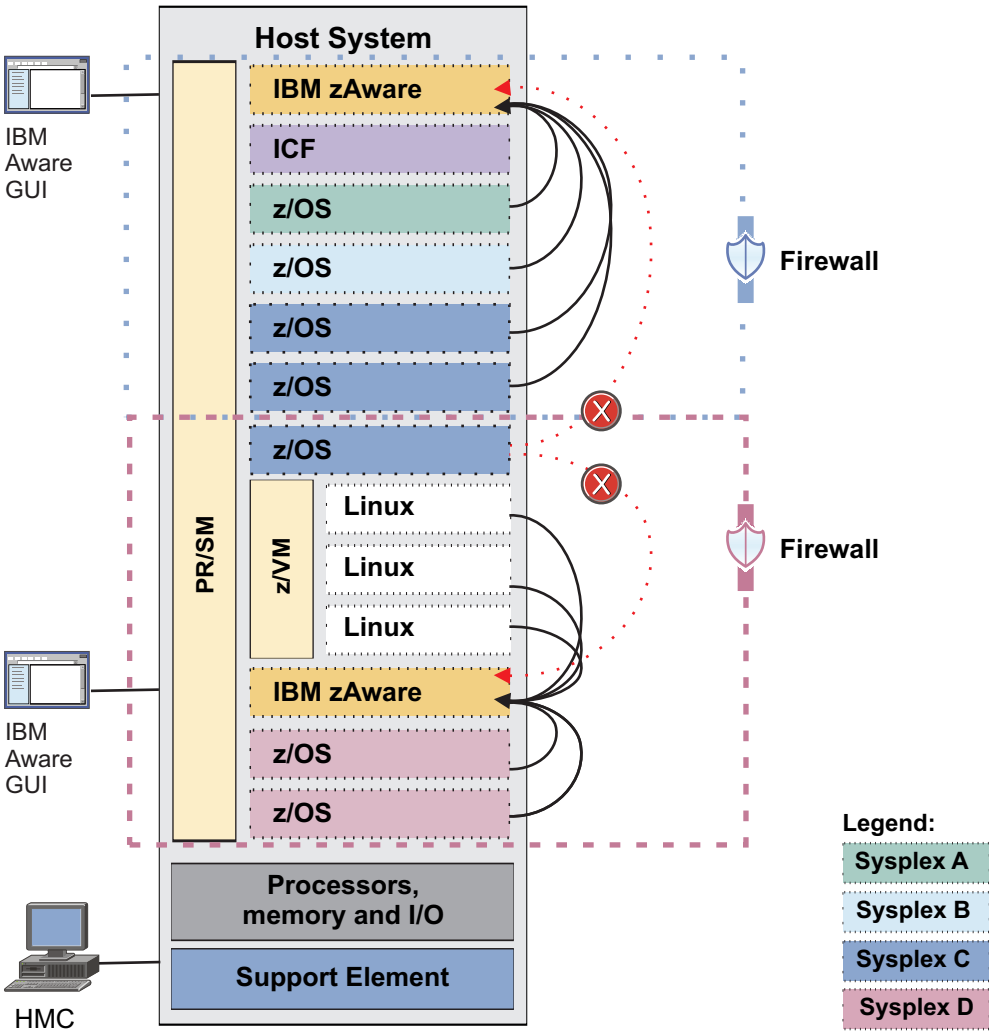

Figure 17. Two IBM zAware partitions in one host system with multiple firewalls

In Figure 17:

- All clients that are protected by the upper firewall (highlighted in blue) are connected to the IBM zAware server that is running in the partition that is protected by the upper firewall. The clients are members of Sysplex A, Sysplex B, or Sysplex C.
- All clients that are protected by the lower firewall (highlighted in pink) are connected to the IBM zAware server that is running in the partition that is protected by the lower firewall. The z/OS clients are members of Sysplex D; the Linux clients are run as z/VM guests.
- The dotted red lines from one Sysplex C client illustrate two possible configurations that are **not** recommended:
  - Configuring any monitored client to connect to an IBM zAware server that is protected by a different firewall.

    IBM zAware does not require clients to provide authentication credentials or to encrypt the data that they send. If your installation considers this data to be sensitive, you need to ensure that the communication between IBM zAware and its monitored clients occurs over secured networks that are configured with preexisting security mechanisms. One method of ensuring secure communication between the server and its clients is to configure them so that they are protected by the same firewall. More network and security considerations are covered in the following topics:

- "Planning network connections and capacity" on page 52
- "Securing communication between IBM zAware and its monitored clients" on page 75

– Using different IBM zAware servers to monitor clients that belong to the same sysplex.

  You can choose to configure only some members of the same sysplex as monitored clients; not all members of a sysplex must be connected to the IBM zAware server. In some cases, all sysplex members cannot be connected; for example, in a Geographically Dispersed Parallel Sysplex™ (GDPS®), the controlling system (K-system) cannot be a monitored client because you cannot use OPERLOG on a K-system. However, any members of the same multiple-system sysplex that you do configure as monitored clients must be clients of the same IBM zAware server.

Figure 18 on page 46 illustrates a more complex configuration that contains all types of supported monitored clients. The clients run in partitions on different types of z Systems servers that are protected by the same firewall, although they are physically separated in two different buildings at the company's Boston site.

- All monitored clients are connected to the IBM zAware server that is running in the partition on the host system.
- The z/OS system that is shown at the top of the z114 on the upper right is configured as a single-system sysplex (monoplex). This system is also a monitored client that is sending data to the IBM zAware server.
- Each member of Sysplex A (highlighted in green) runs on a different z Systems server; all members are connected and sending data to the IBM zAware server.
- Similarly, the members of Sysplex B (highlighted in blue) and Sysplex C (highlighted in dark blue) run on different z Systems servers, and all members are connected and sending data to the IBM zAware server.
- The members of Sysplex D (highlighted in pink) run as z/VM guests on the z196 and the zEC12; only those on the z196 are connected and sending data to the IBM zAware server.
- Because this host system supports IBM zAware Version 2.0 or later, several Linux images are also IBM zAware monitored clients. The Linux clients on the z196 run as z/VM guests; those on the zEC12 run in PR/SM partitions.

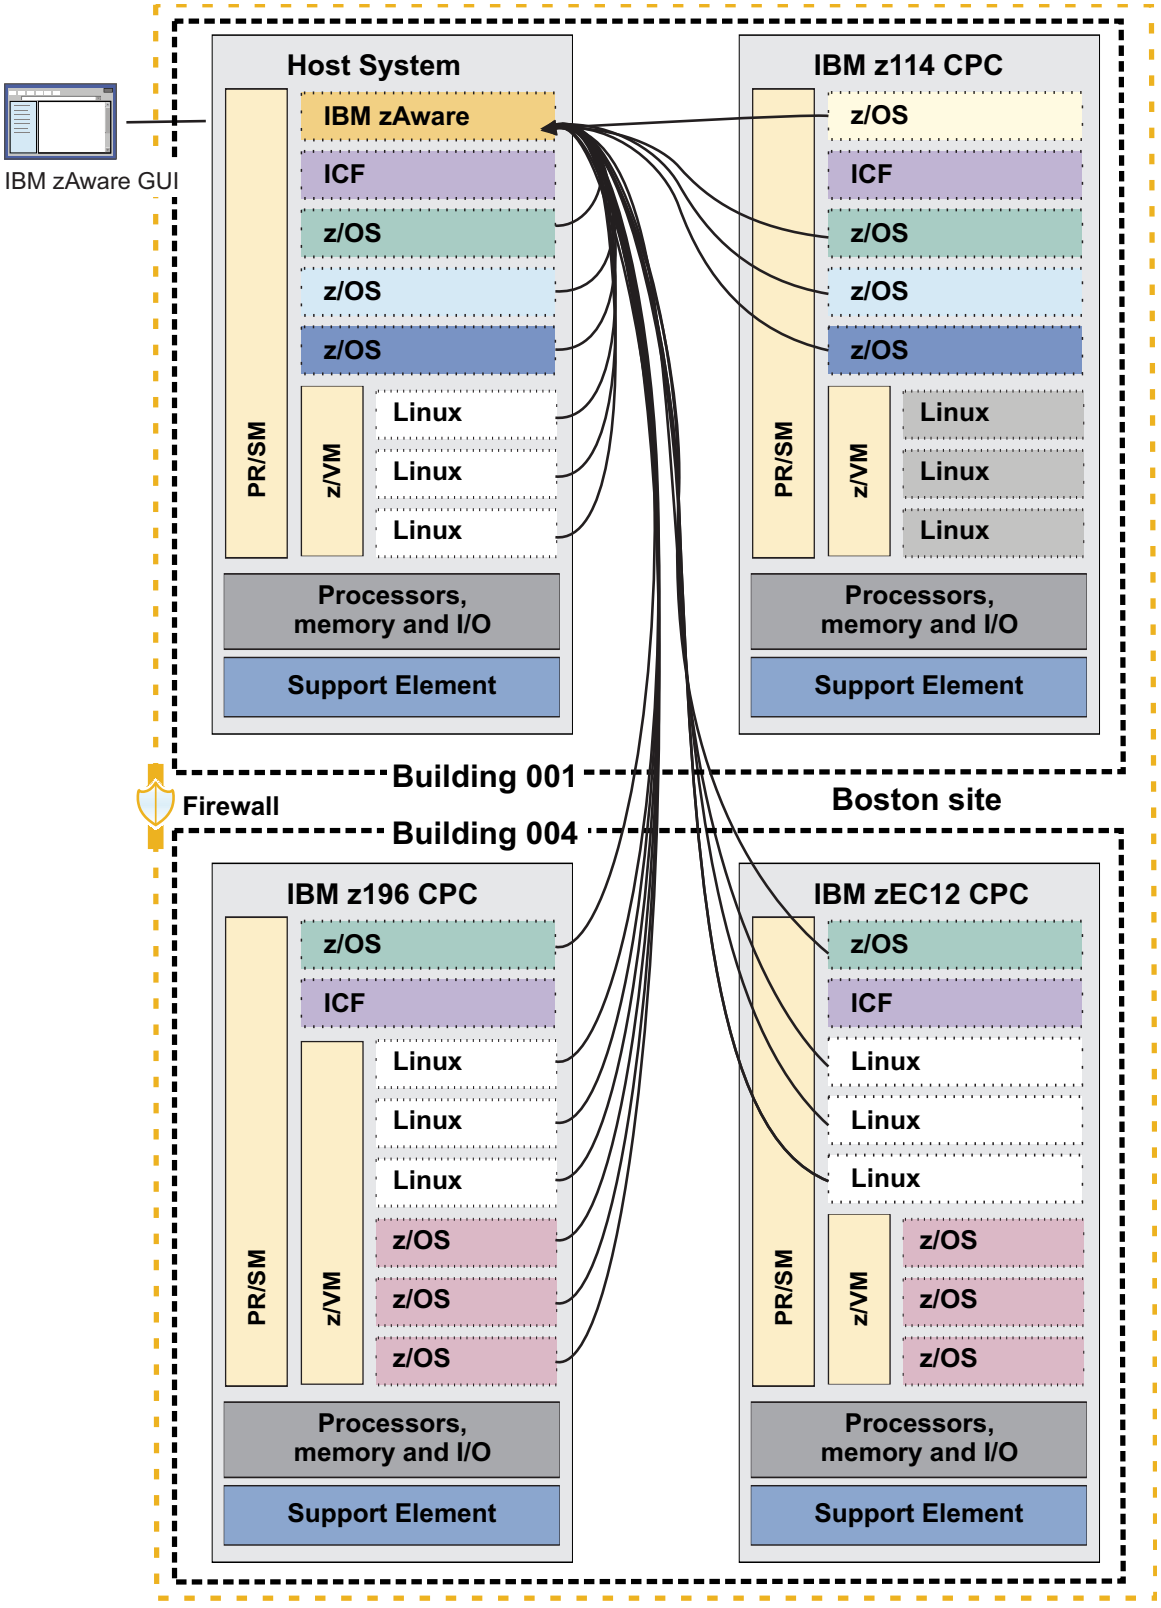

Figure 18. An IBM zAware environment that spans two physical buildings

Summary of planning considerations for the IBM zAware environment

- Chapter 8, "Estimating data center resource requirements," on page 49 provides more information about specific resource requirements for the IBM zAware partition.

- Monitored clients must meet specific configuration requirements for IBM zAware analysis, but can run on any supported z Systems server, including an IBM zAware host system. For more planning and configuration information, see the following topics.
  - Chapter 11, “Planning to create IBM zAware models,” on page 87
  - Chapter 14, “Configuring z/OS monitored clients for IBM zAware analysis,” on page 111
  - Chapter 15, “Configuring Linux on z Systems monitored clients for IBM zAware analysis,” on page 129
- The recommended configuration for an IBM zAware environment— the partition, the server, and its monitored clients— is to protect them within the same firewall.
- The physical distance between an IBM zAware server and its monitored clients can determine the configuration of your IBM zAware environment. IBM testing experiences indicate acceptable IBM zAware capability and performance up to a maximum distance of 3500 kilometers (approximately 2174 miles) between the host system and monitored clients that run on other z Systems servers.
- Members of the same z/OS sysplex must be connected to the same IBM zAware server.

Chapter 8. Estimating data center resource requirements

The processor, memory, network and storage resources that IBM zAware requires vary by installation. The following topics provide guidelines for determining these resource requirements.

Estimating processor and memory resources

According to IBM test results, the amount of required processing resource for the IBM zAware partition varies depending on the number of monitored clients and their combined total rate of message traffic, and also on the phase of IBM zAware operation. Use the guidelines in this topic to determine the amount of processor and memory resource that your installation requires.

Determining the rate of message traffic for a z/OS system

To determine the rate of message traffic for the purposes of estimating processor resource requirements, you can use one of the following options. With either option, the key metric is the total number of message lines, rather than the number of unique message identifiers (IDs).

- IBM message analysis program

  Using the IBM message analysis program is perhaps the easiest way to analyze the message traffic for a given z/OS monitored client. Through this program, you can analyze z/OS SYSLOG data sets to determine the message rate per second, as well as the number and frequency of unique message IDs. The message analysis program is available on the z/OS Tools and Toys web site at the following URL:

  `http://www-03.ibm.com/systems/z/os/zos/features/unix/bpxa1ty2.html`

  To find the message analysis program, search the table of download packages for the MSGLG610 package. The latest version of MSGLG610 produces a report listing the message IDs that adhere to the IBM zAware training criteria.
- Manual calculation

  As an alternative to using the message analysis program, you can calculate the maximum message rate by completing the following steps:
  1. Select a peak workload interval for the z/OS system and store the OPERLOG messages for that interval.
  2. Choose a precise 10-minute interval from the stored messages, and count the messages issued during that interval. For any multi-line messages that might have been issued during the interval, count each line.
  3. Divide the message count by 600 to obtain the message rate per second.

Calculating processor and memory resources

The following guidelines are based on IBM testing of the two operations phases:

- The initial priming and training phase for the IBM zAware server, during which your installation:
  - Uses the z/OS bulk load client for IBM zAware to transfer prior data for monitored clients.
  - Requests the server to build a model for each client from the transferred data.
- Normal operations, during which the primary IBM zAware activity is the analysis of current data from monitored clients.

If an IBM zAware partition does not have a sufficient amount of processor resources or memory assigned to it, users of the IBM zAware GUI will likely notice poor response times for page refreshes and other requests.

Processor resources

On a supported host system, your installation can assign only one of two processor types for the IBM zAware partition: Integrated facilities for Linux (IFLs) or central processors (CPs). The IFLs or CPs can be shared or dedicated. The available processor types vary by host system; for example, both IFLs and CPs are available on a z13.

Use the guidelines in the following tables to determine the amount of processor resource that your installation requires. Keep in mind that IFLs run at full capacity, but CPs can run at full capacity or various subcapacity settings, depending on the host system model. These guidelines are based on IFLs or CPs that are running at full capacity.

- For an IBM zAware configuration of z/OS monitored clients only, see Table 4.
- For an IBM zAware configuration of Linux on z Systems monitored clients only, see Table 5.
- For an IBM zAware configuration of both z/OS and Linux on z Systems monitored clients, see Table 6 on page 51.

Table 4. Processor capacity guidelines for an IBM zAware partition that monitors z/OS clients only

Client configuration and message rates	Guidelines for processor capacity on a z Systems host system
Up to 10 z/OS clients, for a total maximum rate of 500 messages per second	**For initial priming and training** Approximately 25% of one IFL or 25% of one full-capacity CP **For analysis** Approximately 20% of one IFL or 20% of one full-capacity CP
Up to 10 z/OS clients, for a total maximum rate of 1500 messages per second	**For initial priming and training** Approximately 80% of one IFL or 80% of one full-capacity CP **For analysis** Approximately 40% of one IFL or 40% of one full-capacity CP

If you connect more than 10 z/OS monitored clients during a 15-minute time period when the maximum message rate per second is approximately 1500, the capacity of the configured IFLs or CPs might be overrun during the initial priming and training phase. To avoid this potential condition, you can configure additional IFLs or CPs for IBM zAware to use, based on the guidelines in Table 4.

Table 5. Processor capacity guidelines for an IBM zAware partition that monitors Linux clients only

Client configuration and message rates	Guidelines for processor capacity on a z Systems host system
Up to 10 Linux clients, for a total maximum rate of 10 KB per day	**For initial priming and training** Approximately 20% of one IFL or 20% of one full-capacity CP **For analysis** Approximately 20% of one IFL or 20% of one full-capacity CP
Up to 10 Linux clients, for a total maximum rate of 1 MB per day	**For initial priming and training** Approximately 35% of one IFL or 35% of one full-capacity CP **For analysis** Approximately 35% of one IFL or 35% of one full-capacity CP

Table 5. Processor capacity guidelines for an IBM zAware partition that monitors Linux clients only (continued)

Client configuration and message rates	Guidelines for processor capacity on a z Systems host system
For each additional set of 30 Linux clients, with each set totaling a maximum rate of 10 KB up to 1 MB per day	**For initial priming and training** Approximately 80% of one IFL or 80% of one full-capacity CP **For analysis** Approximately 35% of one IFL or 35% of one full-capacity CP

If you connect more than 10 Linux monitored clients during a 15-minute time period when the maximum message rate per day is approximately 1 MB, the capacity of the configured IFLs or CPs might be overrun during the initial priming and training phase. To avoid this potential condition, you can configure additional IFLs or CPs for IBM zAware to use, based on the guidelines in Table 5 on page 50

Table 6. Processor capacity guidelines for an IBM zAware partition that monitors both z/OS and Linux clients

Client configuration and message rates	Guidelines for processor capacity on a z Systems host system
• Up to 10 z/OS clients, for a total z/OS maximum rate of 500 messages per second, **plus** • Up to 10 Linux clients, for a total Linux maximum rate of 10 KB per day	**For initial priming and training** Approximately 40% of one IFL or 40% of one full-capacity CP **For analysis** Approximately 30% of one IFL or 30% of one full-capacity CP
• Up to 10 z/OS clients, for a total z/OS maximum rate of 1500 messages per second, **plus** • Up to 10 Linux clients, for a total Linux maximum rate of 1 MB per day	**For initial priming and training** Approximately 80% of one IFL or 80% of one full-capacity CP **For analysis** Approximately 40% of one IFL or 40% of one full-capacity CP
For each additional set of 30 Linux clients, with each set totaling a maximum rate of 10 KB up to 1 MB per day	**For initial priming and training** Approximately 80% of one IFL or 80% of one full-capacity CP **For analysis** Approximately 35% of one IFL or 35% of one full-capacity CP

If you connect more than 10 monitored clients of each type during a 15-minute time period when the maximum z/OS message rate per second is approximately 1500, or the maximum rate per day is approximately 1 MB, the capacity of the configured IFLs or CPs might be overrun during the initial priming and training phase. To avoid this potential condition, you can configure additional IFLs or CPs for IBM zAware to use, based on the guidelines in Table 6.

Memory

Use the following guidelines to determine the amount of memory that an IBM zAware partition requires for your installation. These guidelines apply for an IBM zAware partition on a z13 and a z14.

- You must allocate a minimum of 6144 MB (6 GB) to activate the IBM zAware partition. This amount of memory is sufficient to support a relatively small number of monitored clients (six or fewer) with relatively light message traffic (500 messages per second).
- You need to assign an extra 256 MB of memory for each monitored client; the amount is the same, regardless of the type of monitored client. Use the following formula for determining the amount of memory to assign to the partition.

  ```
  6144MB + (256MB * number of z/OS clients) + (256MB * number of Linux clients)
  ```

- The suggested maximum amount is 512 GB of memory.

Note: Beginning with MCL N98812.037 or MCL P08444.002, IBM zAware creates a swap file to avoid potential memory constraints that might develop when many monitored clients are connected over a long period. For more information, see "Estimating required storage capacity and selecting devices" on page 61.

See the following topics for additional information:

- "Configuring the IBM z Systems Secure Service Container for IBM zAware" on page 27
- "Monitoring processor, memory, and storage resources" on page 216 contains information about monitoring and adjusting these resources.

Planning network connections and capacity

After you decide which z/OS systems to monitor with IBM zAware, you need to plan the network connections that are required to enable communication between the IBM zAware server and its clients. Your choices are dependent on the types of network connections that the IBM zAware server supports and on the location of the monitored clients.

The IBM zAware server requires a shared or dedicated Open Systems Adapter (OSA) port and an IP address for the partition in which the server runs.

- For OSA-Express4S or later generation features, IBM zAware can use only port 0.
- The IP address must be either dedicated or assigned through Dynamic Host Connection Protocol (DHCP). With DHCP-type IP addresses, use of a domain name system (DNS) server is required.

The server supports the Transmission Control Protocol (TCP) and the User Datagram Protocol (UDP).

Supported network connection types and capacity requirements

The IBM zAware server supports the following types of network options, which are illustrated in figures in the topic "Sample network configuration."

- A customer-provided data network that provides Ethernet connectivity through an OSA channel.
- A HiperSockets™ subnet within the host system.

  HiperSockets is a hardware feature that provides high performance internal communications between logical partitions within the same CPC, without the use of any additional or external hardware equipment.
- The intraensemble data network (IEDN) on the host system.

  The IEDN is a private high-speed network for application data communications within and between nodes of an ensemble. The IBM zAware server also supports the use of HiperSockets over the IEDN.

Regardless of the network connection types used in an IBM zAware environment, IBM test results indicate that the additional message traffic between the IBM zAware server and any monitored clients is relatively insignificant, even for monitored systems that produce a high volume of messages. For example, during testing of normal IBM zAware operations with message traffic up to a maximum of 1500 messages per second, the reported transfer rate on the network between the server and its monitored clients ranged from 10 to 20 kilobytes per second. So the use of IBM zAware does not require any special consideration in network capacity planning.

Sample network configuration

The simple configuration in Figure 19 on page 53 shows the IBM zAware server and all of its monitored clients residing on the same host system.

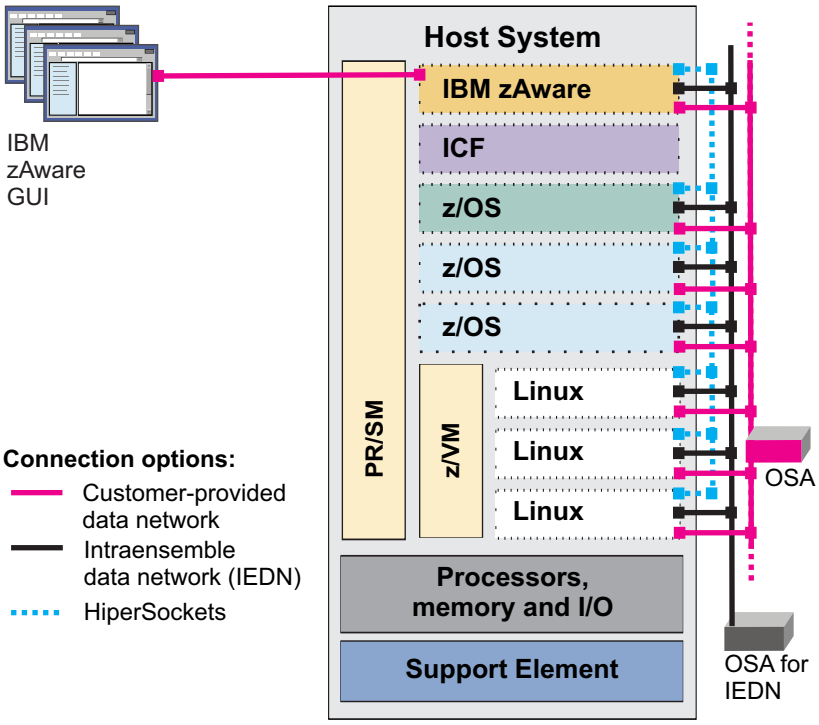

Figure 19. Supported network connection options for the IBM zAware environment

For this type of configuration, you need to configure the following network connections:

- An OSA channel path for browser access to the IBM zAware server. Of the supported network options, the OSA channel is the most logical choice for allowing browser connections to the server, so users can view the analytical data for the monitored clients through the IBM zAware graphical user interface (GUI).
- Any one of the supported network options for connecting an IBM zAware server to monitored clients: an OSA channel, HiperSockets subnet, or the IEDN. In this case, factors other than the location of monitored clients might influence your choice. For example, you might need to consider security mechanisms or the possible impact to current network traffic of additional message traffic from the monitored clients to the server.

If you expand the IBM zAware environment to include monitored systems that reside on other CPCs at your installation, your options are different. For example, suppose that you want to monitor z/OS clients that belong to two sysplexes, A and B, which have some system images on the IBM zAware host system and others on another supported CPC, as shown in Figure 20 on page 54. For the clients that do not reside on the host system, your only option is an OSA channel because the server on which they reside cannot use the IEDN or HiperSockets, both of which provide an internal network channel only within a single CPC. Similarly, an OSA channel is the only option for monitoring Linux systems that reside on a CPC other than the host system. Depending on the complexity of the IBM zAware environment that you plan to configure, you might define all three types of supported network connections.

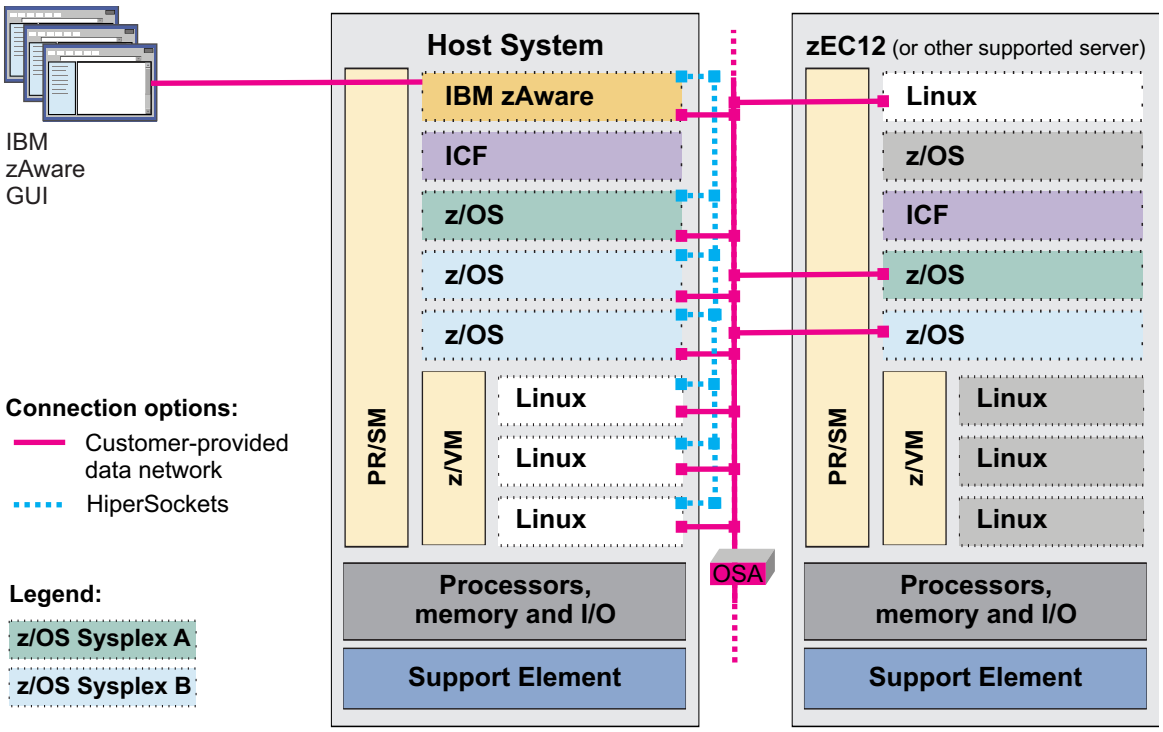

Figure 20. Network connections for an IBM zAware partition supporting clients in the host system and one zEC12 CPC

Overview of the network configuration procedure

When you configure the network connections for an IBM zAware partition, you select an IP address type for each connection: Dynamic Host Connection Protocol (DHCP), link local addressing, or static IPv4 or IPv6 address.

- With local link addressing and static IPv4 or IPv6 addresses, monitored clients and GUI users access the IBM zAware server through the dedicated IP address of the partition.
- With DHCP, the IP address of the partition can change. When you select a DHCP address type, you also must set up a domain name system (DNS) server to resolve a host name for the IBM zAware partition. With DHCP, both monitored clients and GUI users need to use a host name rather than an IP address to successfully connect to the IBM zAware server.

To define the network connections, you need to modify configuration files for the host system, the IBM zAware partition, and monitored clients. Figure 21 on page 55 shows the configuration files that you need to modify.

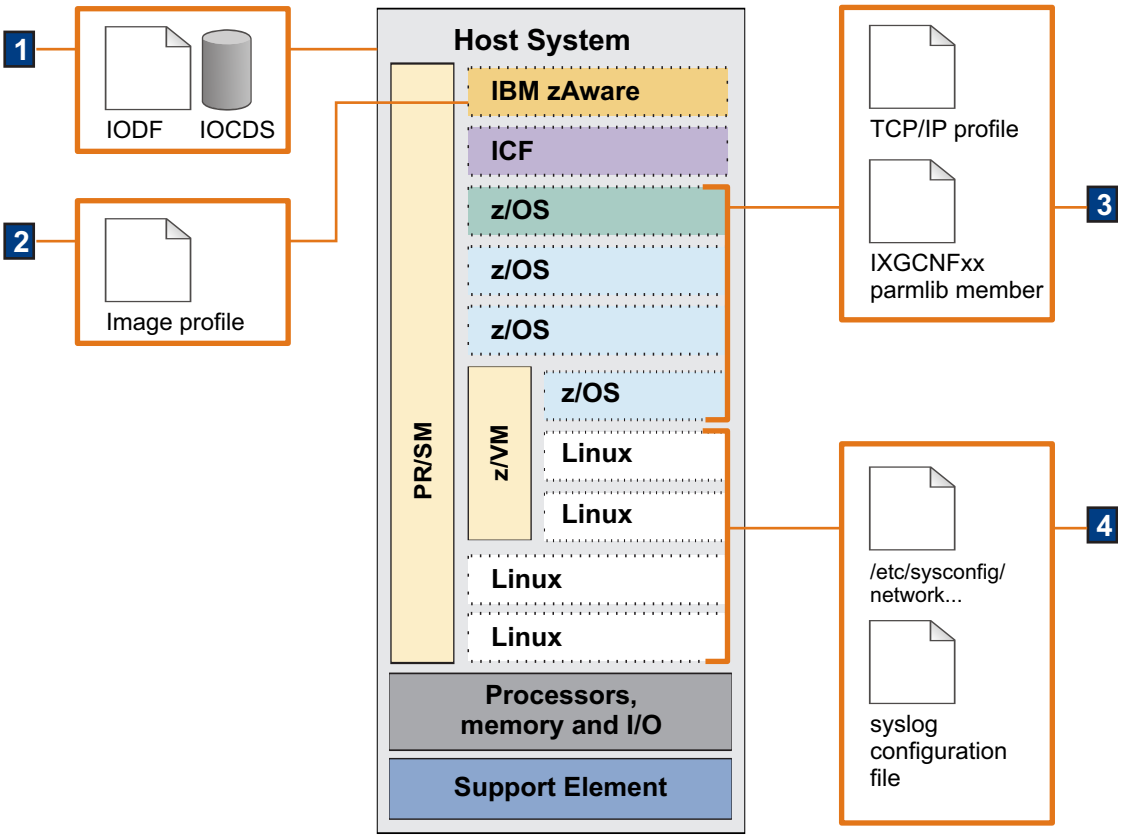

Figure 21. Configuration files for network connections

1. The I/O definition file (IODF) and input/output configuration data set (IOCDS) for the IBM zAware host system contain definitions that associate specific channel paths to the IBM zAware partition.

   The supported channel path identifier (CHPID) types are:
   - **OSD** for an OSA data network.
   - **OSX** for the IEDN.
   - **IQD** for HiperSockets or HiperSockets over the IEDN.

     Only the following functions are supported IQD channel parameters:
     - Basic HiperSockets
     - IEDN Access, only when the CPC is a member of an ensemble

     The External Bridge function is not a supported IQD channel parameter.

   Channel paths can be dedicated, reconfigurable, shared, or spanned.

   If you plan to configure monitored clients to use an existing HiperSockets subnet, for example, the IODF or IOCDS must include a CHPID type of **IQD** for the IBM zAware partition. The IODF or IOCDS also must include a channel path definition with an **OSD** type for GUI browser access to the IBM zAware server.

   For optimal performance and operations, configure the IBM zAware partition such that it has access to only those channel path identifiers (CHPIDs), control units, and I/O devices that are required for network connectivity and storage. You can use the following HCD mechanisms to limit access:
   - Image access and candidate lists in channel path definitions
   - The explicit device candidate list for I/O devices
2. The image profile for the IBM zAware partition contains the IP address and host name of the partition, as well as network adapter definitions (such as the IP address type) for each type of network connection. These definitions correspond to the channel path types in the IODF or IOCDS.

For example, to match the HiperSockets CHPID in the IODF or IOCDS, the image profile must include a network adapter definition with a CHPID type of **IQD**. That network adapter definition includes the IP address and other details about the HiperSockets subnet. Similarly, another network adapter definition with a CHPID type of **OSD** is required for GUI browser access to the IBM zAware server.

3. The TCP/IP profiles and system logger configuration settings for z/OS monitored clients also must contain specific network configuration settings.
   - The TCP/IP profile must contain a statement that corresponds to the channel path types in the IODF or IOCDS for the IBM zAware host system.

     For example, for connectivity through the HiperSockets subnet, the profile must contain a DEVICE and LINK MPCIPA statement with a device name of IUTIQD*xx*, where *xx* is the CHPID number that matches the hexadecimal value specified on the CHPID type IQD definition statement for the IBM zAware host system. Similarly, for GUI browser access to the IBM zAware server, the TCP/IP profile must contain an INTERFACE statement to define an IPAQENET or IPAQENET6 interface with CHPIDTYPE OSD.

     Note: z/OS clients that use HiperSockets must be configured to use layer 3 (which is the default layer) or the client cannot successfully connect to the IBM zAware server. For these HiperSockets, the use of address resolution protocol (ARP) must be disabled. These restrictions do not apply for HiperSockets over the IEDN.
   - The system logger configuration settings in the IXGCNFxx parmlib member must contain the following information:
     – The IP address or host name specified in the image profile for the IBM zAware partition.
     – The port number associated with the IBM zAware partition. The port number is 2001.
4. Various configuration files for Linux monitored clients must contain specific network configuration settings.
   - Various files in `/etc/sysconfig/network` must contain statements that correspond to the network connections in use for z/VM guests or, if the Linux system runs in a logical partition, the network connections in use for the PR/SM LPAR.
   - The syslog configuration file must contain a destination statement with the IP address and port number for the IBM zAware server. The destination statement syntax depends on which syslog daemon is in use for the Linux monitored client.

In summary, for each type of network connection that you want to use, you must have corresponding definitions in the configuration files. These corresponding definitions connect the IBM zAware partition and its monitored clients to the same network. The IBM zAware partition can be connected to more than one network, as necessary, to monitor systems at your installation.

- For additional information about the channel path types that the IBM zAware server supports, see "Summary of channel path types."
- For a network planning checklist, see "Task summary and configuration checklist for network administrators" on page 57.

Summary of channel path types

Table 7 on page 57 contains a summary of supported channel path types, usage, and sources of additional information.

Table 7. Supported channel path types for the IBM zAware partition

Channel path type	Description	IBM zAware usage	Additional information
OSD	Ethernet connectivity through an Open Systems Adapter (OSA) channel	• Use for monitored clients that reside on any supported z Systems CPC in the IBM zAware environment. • Use for GUI browser connections to the IBM zAware server.	*zEnterprise System, System z10®, System z9® and eServer™ zSeries Open Systems Adapter-Express® Customer's Guide and Reference*, SA22-7935 **Note:** For OSA-Express4S or later generation features, IBM zAware can use only port 0.
OSX	Connectivity through the intraensemble data network (IEDN)	Use only for monitored clients that reside on the IBM zAware host system or on other nodes in the same zEnterprise ensemble.	*zEnterprise System, System z10, System z9 and eServer zSeries Open Systems Adapter-Express Customer's Guide and Reference*, SA22-7935 **Note:** For OSA-Express4S or later generation features, IBM zAware can use only port 0.
IQD	Connectivity through HiperSockets	Use only for monitored clients that reside on the IBM zAware host system. Use the IQD channel path type for HiperSockets. Only the following functions are supported IQD channel parameters: • Basic HiperSockets • IEDN Access, only when the CPC is a member of an ensemble The External Bridge function is not a supported IQD channel parameter.	*z/OS Communications Server: IP Configuration Guide*, SC31-8775

Task summary and configuration checklist for network administrators

- Table 8 provides a summary of network administration tasks and links to additional information.
- Table 9 on page 58 and Table 10 on page 58 are checklists that a network administrator can use to complete the networking step in "Configuring the IBM z Systems Secure Service Container for IBM zAware" on page 27.

Table 8. Task summary for network administrators

✓	Task summary:	Where to find instructions:
	Collaborate with the security administrator to determine whether any additional network definitions are required to ensure secure communications between the IBM zAware server and its monitored clients.	"Securing communication between IBM zAware and its monitored clients" on page 75
	Collaborate with the security administrator to determine whether any additional network definitions are required to ensure communication between the IBM zAware server and the Lightweight Directory Access Protocol (LDAP) server.	Your installation has the option to configure user authentication through the use of an LDAP repository. If a firewall exists between the IBM zAware partition and the LDAP server, the IBM zAware partition must be able to use the port that is used by the LDAP server. Also, if your installation uses a domain name instead of an explicit IP address for the LDAP server, the IBM zAware partition must be able to reach a DNS server for hostname-to-IP resolution.

Table 8. Task summary for network administrators (continued)

✓	Task summary:	Where to find instructions:
	Define channel paths for networks in the I/O definition file (IODF) and input/output configuration data set (IOCDS) for the IBM zAware host system.	Step 1 on page 96 in Chapter 12, "Configuring network connections and storage for the IBM zAware partition," on page 95
	Define network settings in the image profile for the IBM zAware partition.	See "Configuring the IBM z Systems Secure Service Container for IBM zAware" on page 27
	Define network settings in the TCP/IP profile and system logger configuration file for each z/OS monitored client.	Steps 1 on page 112 and 3 on page 113 in "Configuring z/OS monitored clients to send data to the IBM zAware server" on page 111.
	Define network settings in various files in `/etc/sysconfig/network` and in the syslog configuration file for each Linux monitored client.	Steps 1 on page 130 and 3 on page 130 in "Configuring Linux on z Systems monitored clients to send data to the IBM zAware server" on page 129.

Table 9. Checklist for IBM zAware partition network settings

✓	Host name or IP address[1]	Port	Master ID	Default gateway	DNS server[2]	Secondary DNS server
Sample	IBMzAWLPAR	2001	ZAIADMIN	198.51.100.10	198.51.100.64	—

Footnotes for Table 9:
1. Specify a host name if you select DHCP as the IP address type for the network adapter (channel path).
2. Specify a primary (and secondary, if appropriate) DNS server if you select DHCP as the IP address type for the network adapter (channel path).

Table 10. Checklist for IBM zAware partition network adapters

✓	CHPID type[1]	CHPID	Interface or device name	IP address type[2]	VLAN	IP address	Mask / prefix
Sample	OSD	46	OSAQDIO26	DHCP	1211	198.51.100.11	32

Footnotes for Table 10:
1. One of the following: OSD (Ethernet connectivity over an OSA channel), OSX (intraensemble data network), or IQD (HiperSockets). The type determines which of the remaining columns you need to complete.
2. One of the following: DHCP, link local, static IPv4 or static IPv6 address. The type determines which of the remaining columns you need to complete.

Planning persistent storage configuration and capacity

To provide analytical data for monitored clients, IBM zAware requires continuous access to a set of Extended Count Key Data (ECKD™) direct-access storage devices (DASD). You must have one separate ECKD device with a minimum of 8 GB to 15GB storage, to store the IBM zAware code. If the DASD size is many times larger than this range, it is possible that the disk formatting takes too long time, and the installation times out before the image is uploaded. The additional number of storage devices that are in that set depends, in part, on the number of monitored clients in the IBM zAware environment and on the adjustable retention times that IBM zAware uses to manage stored data. Because of the way IBM zAware uses and depends on the availability of storage for its operation, careful planning is a critical phase for preparing to configure the IBM zAware environment.

IBM zAware stores the following analytical data on DASD:
- Current® data from each z/OS or Linux monitored client, as well as priming data, if any.
- IBM zAware models for each z/OS monitored client and for each model group, which is an administrator-defined group of Linux clients.
- Analysis results for each z/OS or Linux monitored client.

IBM zAware sets default retention times for each of these types of analytical data and, through an automated process, removes the data when the retention time has elapsed.

To plan for configuring and managing persistent storage use in an IBM zAware environment, review the following topics:

- "IBM zAware storage use and planning considerations" describes IBM zAware storage use and dependencies that necessitate careful planning to avoid or quickly recover from potential storage-related problems.
- "Estimating required storage capacity and selecting devices" on page 61 provides guidelines for estimating and providing the persistent storage resources for your installation.
- Two examples provide an overview of the storage configuration process, along with the planning considerations for two different IBM zAware environments:
  - "Example: Storage configuration for normal operations" on page 62 describes how your installation might configure in-use and backup storage for a single instance of IBM zAware.
  - "Example: Storage configuration for multiple IBM zAware partitions" on page 66 describes how your installation might configure persistent storage for an IBM zAware environment that consists of:
    - One instance of IBM zAware (the primary) for normal operations.
    - Another instance of IBM zAware (the alternate) to be used in switchover situations.

  An instance of IBM zAware is called an IBM zAware *server*.
- "Summary of planning considerations for persistent storage" on page 71 provides a summary of planning considerations, along with storage usage notes and best practices.
- "Task summary and configuration checklist for storage administrators" on page 73 provides a planning checklist for storage administrators.

IBM zAware storage use and planning considerations

To avoid or quickly recover from potential storage-related problems, storage administrators need to consider the information in Table 11 on page 60.

Table 11. Storage use, planning considerations, and best practices for IBM zAware storage

IBM zAware storage use	Why this fact is important for planning	Best practices
Through the I/O definition file (IODF) and input/output configuration data set (IOCDS) for the IBM zAware host system, the IBM zAware partition has access to all storage devices that are physically attached to the CPC and defined to the same network connection that is defined for the partition.	When the IBM zAware partition is activated, IBM zAware uses the IOCDS, and any existing IBM zAware storage configuration information, to populate the **Data Storage** page with a list of storage devices and their status (available, in use, and so on). Only when an administrator uses that list to select available devices to add to the storage configuration, IBM zAware formats and initializes the selected devices, overwriting any of the data that might be stored on those devices. **Unless IBM zAware access is restricted to only specific CPC storage devices, an administrator might inadvertently cause the loss of critical data by adding a storage device that is not intended for IBM zAware use.**	• When defining storage devices in the IODF or IOCDS for the IBM zAware host system, the storage administrator can use image access and candidate lists for channel path definitions to allow only the IBM zAware partition to access specific devices. Using the explicit device candidate list is an alternative method of restricting access to specific devices. • As a precaution, you can use channel path definition lists, the explicit device candidate list, or operating system mechanisms to prevent the applications that are running in other CPC partitions from using storage devices that are intended for IBM zAware use. The only exception to this practice is the CPC partition, if any, that your installation is using to back up the storage that the IBM zAware server is using.
After IBM zAware formats and initializes storage devices for its use, the added devices constitute the in-use set of storage devices. Although IBM zAware tracks the status and capacity of individual storage devices that it is using, it treats the set of in-use devices as a single logical volume.	**If one or more individual devices become unavailable through any method other than removal through the IBM zAware GUI, IBM zAware effectively loses access to all of its stored data and operations stop.** Only the following corrective actions can resolve this condition: 1. If possible, try to reattach the device and reactivate the IBM zAware partition. 2. If the device cannot be reattached, you must replace it with an equivalent device containing a backup copy of the data that was stored on the unavailable device. The original device and the backup device must be the same size. 3. If your installation does not have backup copies of IBM zAware data, you must deactivate the IBM zAware partition and reconfigure the IBM zAware environment.	• If you must remove in-use devices from the CPC storage configuration, first use the IBM zAware GUI to remove those devices from the IBM zAware configuration. Then you can use other removal methods (such as changing the IOCDS) without affecting IBM zAware operations. • Because IBM zAware does not automatically replicate any of its stored data, maintaining a backup copy of this data is highly recommended. For replication, your installation can consider using IBM FlashCopy® or one of several Data Facility Storage Management Subsystem (DFSMS) copy services. • When replicating IBM zAware data from in-use data sets to backup data sets, remember to manage the backup set as a single entity. • When you use the IBM zAware GUI to replace a missing in-use device with its equivalent backup device, use the **Preserve data** option when adding the backup device, which prevents IBM zAware from overwriting data on the backup device.

Figure 22 on page 61 shows a configuration in which DASD attached to the IBM zAware host system has been configured exclusively for use by IBM zAware. Another partition in this configuration also can access the DASD, but only for the purpose of creating and maintaining a backup copy of IBM zAware data.

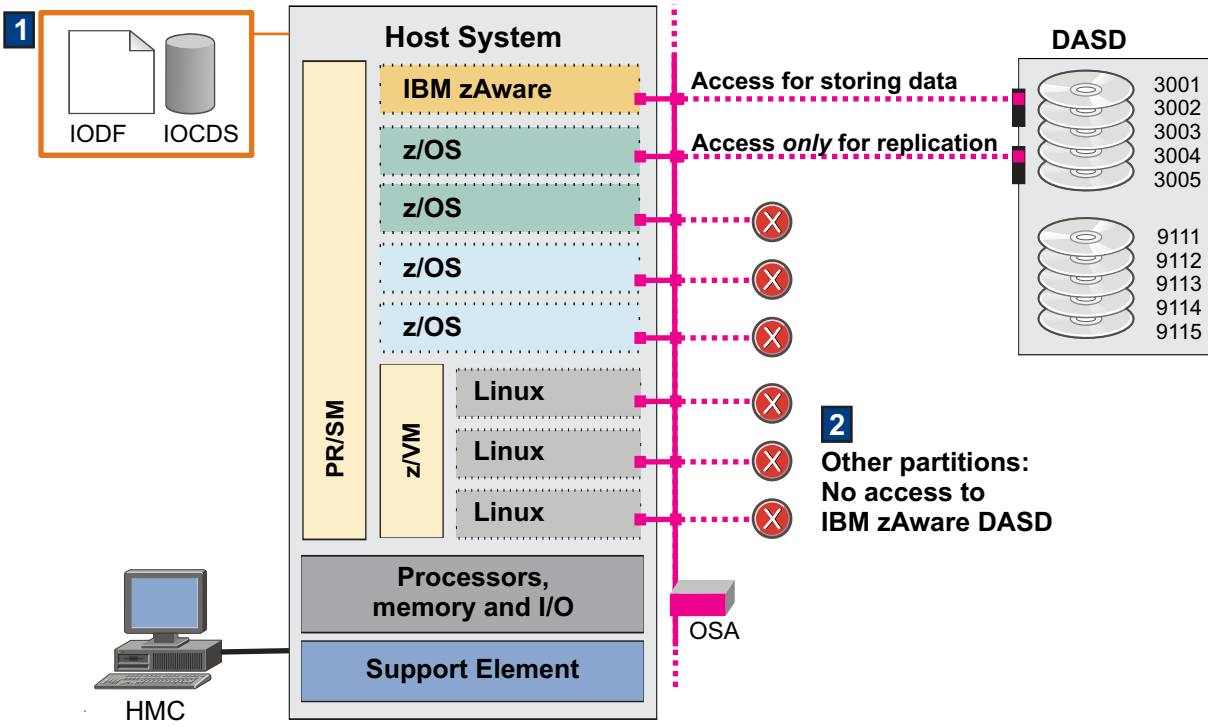

Figure 22. DASD configuration for normal operations, with access for one z/OS partition only to back up IBM zAware data

1. The I/O definition file (IODF) and input/output configuration data set (IOCDS) for the IBM zAware host system contain definitions for the physical storage devices that are attached to the CPC. Figure 22 shows only a subset of the DASD attached to the CPC.
   - Through the use of channel path definition lists, the storage administrator has configured this subset of DASD exclusively for use by the IBM zAware partition.
   - The storage administrator also has configured one additional partition to access the DASD, but only for the purpose of backing up IBM zAware data, which is a highly recommended practice. The system used for backing up data can reside on the IBM zAware host system or on another z Systems server in your installation.

     For replication, your installation can consider using IBM FlashCopy or one of several Data Facility Storage Management Subsystem (DFSMS) copy services, including Extended Remote Copy (XRC) and Peer-to-Peer Remote Copy (PPRC) solutions. Another possible alternative is using DFSMShsm to copy data. In contrast to real-time replication solutions, DFSMShsm requires deactivating the IBM zAware partition before copying data, then reactivating the partition after the copy operation completes. Non-IBM replication products also are available for use.
2. Other partitions on the IBM zAware host system do not have access to the DASD that IBM zAware uses. Partitions that reside on other z Systems server in your installation also must not have access to storage devices that are exclusively for IBM zAware use.

   You can use channel path definition lists, the explicit device candidate list, or operating system mechanisms to prevent the applications that are running in other CPC partitions from using storage devices that are intended for IBM zAware use.

Estimating required storage capacity and selecting devices

Storage requirements vary depending on the retention times for each type of analytical data and on the number of monitored systems that you plan to connect to IBM zAware. Start with 500 GB of storage for IBM zAware to use, plus 4 - 5 GB of storage for each monitored system.

If you increase the number of monitored clients, you need to configure an extra 4 - 5 GB of storage for each monitored system. If you increase the retention times of instrumentation data, training models, or analysis results, you also might need to increase the amount of persistent storage that IBM zAware can use. To determine whether you need to add storage devices, periodically use the **Administration** > **Configuration** > **Data Storage** tab to monitor the list of assigned storage devices, their status, and capacity.

Note: Beginning with MCL N98812.037 or MCL P08444.002, IBM zAware creates a swap file to avoid potential memory constraints that might develop when many monitored clients are connected over a long period. IBM zAware reserves a portion of in-use storage for the swap file, which it creates and manages automatically, using the file only when necessary. Administrator intervention is not required when IBM zAware has access to enough DASD for the swap file. If IBM zAware does not have enough in-use storage, it issues message AIFP0022E to indicate how much more storage is required.

After estimating the required capacity, the storage administrator selects specific physical storage devices to reserve for the IBM zAware partition.

- Your installation can use Extended Count Key Data (ECKD) direct-access storage devices (DASD) to store IBM zAware data.
- Your installation can select a combination of small volumes or volumes of different sizes to provide the storage capacity that your IBM zAware environment requires. These volumes cannot be SMS-managed volumes.
- If any devices are to be used for storing backup copies of IBM zAware data, your installation must define two physically separate but equivalent sets of storage devices:
  - One set for IBM zAware to use for normal operations.
  - Another set for storing backup copies of data.

  The number of storage devices in each set must match, and each backup device must be equivalent in size to the device from which the data is copied. These number and size requirements also apply for configurations that contain primary and alternate IBM zAware partitions.
- Storage devices that are configured for use by IBM zAware must be available for write operations. Specifically, the devices cannot be target devices in a FlashCopy relationship, secondary devices of a PPRC configuration, or otherwise write-inhibited.

After selecting specific physical storage devices, the storage administrator establishes and communicates conventions for persistent storage use to IBM zAware administrators. For example, suppose that your installation plans to set up an environment with one IBM zAware server, and plans to set up replication to copy its data. To implement this plan, the storage administrator might designate two sets of 3390 DASD for IBM zAware and communicate the following instructions and conventions:

- Use only devices 3001-3005 for normal operations.
- Devices 9111-9115 are reserved for replication. Naming conventions identify which backup device matches each in-use device; for example, 9111 contains the backup copy for data stored on 3001, 9112 contains the copy of data on 3002, and so on.
- If you change the set of in-use devices by adding or removing devices through the GUI, make sure that you adjust replication accordingly. To successfully replace an in-use device with its equivalent backup device, the set of in-use devices must match the set of backup devices in number of devices, size of devices, and content.

"Example: Storage configuration for normal operations" illustrates how IBM zAware administrators can use these instructions and follow conventions to configure and manage IBM zAware storage.

Example: Storage configuration for normal operations

Figure 23 on page 63 shows a sample IBM zAware environment that contains one IBM zAware partition.

Through the I/O definition file (IODF) and input/output configuration data set (IOCDS) for the IBM zAware host system, the partition has access to two physically separate sets of equivalent storage devices:

3001-3005 and 9111-9115. The storage administrator has designated devices 3001-3005 for IBM zAware data, and devices 9111-9115 to contain backup copies of that data. This type of configuration enables you to quickly restore IBM zAware operations after the loss of an in-use storage device, which can result from conditions such as a control unit failure on the device.

After the IBM zAware partition is activated, all of these devices are displayed as available for use through the Data Storage page in the IBM zAware graphical user interface (GUI). The numbered areas of Figure 23 illustrate the sequence of tasks that an IBM zAware administrator might follow to correctly configure these devices for normal operations and replication.

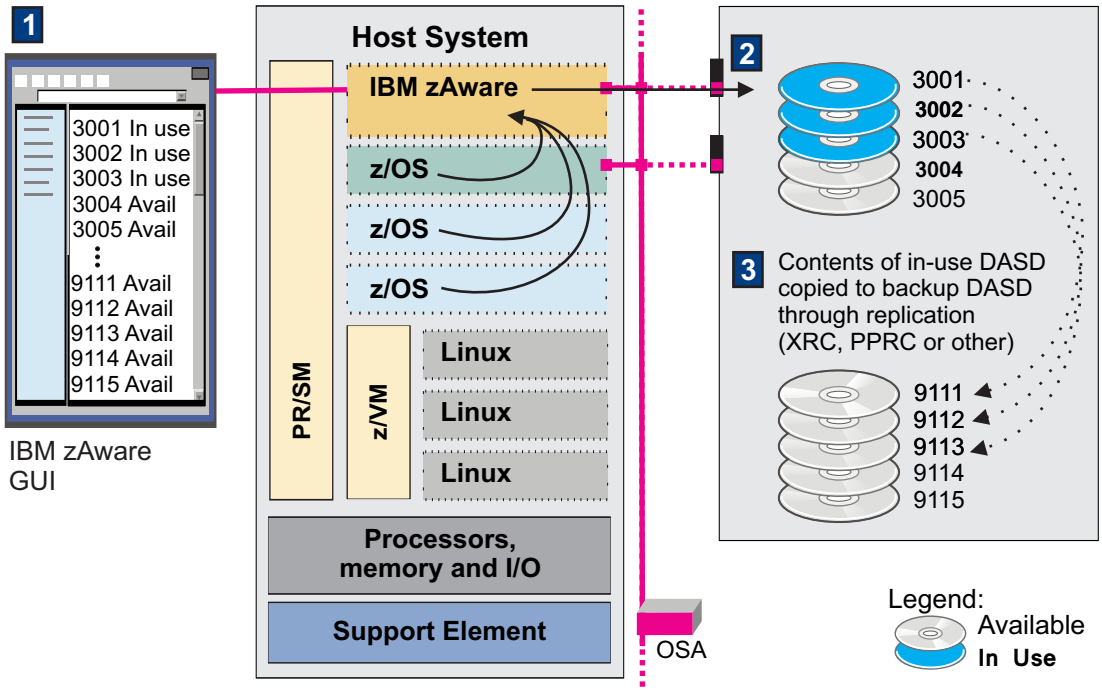

Figure 23. Storage configuration for the IBM zAware server for normal operations and data replication

1. Using the Data Storage page in the IBM zAware GUI, an IBM zAware administrator selects devices for IBM zAware to use for normal operations.

   Before IBM zAware can store data in any of the 300x devices, an administrator must use the **Add and Remove Devices** function to select and add those storage devices.

   - For this sample configuration, an administrator has initially added only storage devices 3001-3003, reserving the other 300x devices for any additional capacity that might be needed when more monitored clients are added to the environment. This action causes IBM zAware to format and initialize the added devices. The GUI window depicted in Figure 23 shows the results of this selection process: the displayed status for devices 3001-3003 is now "In use".
   - Note that the administrator has *not* used **Add and Remove Devices** to configure the 911x devices, which are displayed in the GUI as "Available" devices. In a configuration that includes only one IBM zAware server, you cannot add backup devices until they are needed to replace an in-use device that is no longer available.

2. With in-use devices available to store data, an administrator can finish configuring the IBM zAware environment by configuring and connecting monitored clients, sending priming data, and requesting IBM zAware to build models of client behavior.

   In Figure 23, monitored clients that send data to IBM zAware are shown in partitions running on the IBM zAware host system, but additional monitored clients can run in partitions on other z Systems servers. IBM zAware parses and stores the following data on in-use devices 3001-3003:
   - Current data from each z/OS or Linux monitored client, as well as priming data, if any.

- IBM zAware models for each z/OS monitored client and for each model group, which is an administrator-defined group of Linux clients.
- Analysis results for each z/OS or Linux monitored client.

3. After normal IBM zAware operations begin, an administrator can set up replication for the in-use devices. For this replication to be successfully completed, both the number and sizes of devices in the backup set must match the number and sizes of in-use storage devices.

   For replication, your installation can consider using IBM FlashCopy or one of several Data Facility Storage Management Subsystem (DFSMS) copy services, including Extended Remote Copy (XRC) and Peer-to-Peer Remote Copy (PPRC) solutions. Another possible alternative is using DFSMShsm to copy data. In contrast to real-time replication solutions, DFSMShsm requires deactivating the IBM zAware partition before copying data, then reactivating the partition after the copy operation completes. Non-IBM replication products also are available for use.

   Figure 23 on page 63 shows that one of the monitored clients, the z/OS partition highlighted in light green, has access to IBM zAware in-use storage devices 3001-3005, as well as their equivalent backup set of devices, 9111-9115. An administrator has set up this z/OS image with a replication solution that copies data from the set of in-use devices to the backup set:
   - Device 9111 contains a copy of the data from device 3001.
   - Device 9112 contains a copy of the data from device 3002.
   - Device 9113 contains a copy of the data from device 3003.

   If you change the set of in-use devices by adding or removing devices through the GUI, make sure that you adjust replication accordingly. To successfully replace an in-use device with its equivalent backup device, the set of in-use devices must match the set of backup devices in number of devices, size of devices, and content.

Figure 24 illustrates how this backup data can be used when an in-use storage device becomes unavailable. Starting from the right side of the figure, the numbered areas illustrate the sequence of events and recovery tasks that an IBM zAware administrator can follow to resolve a “missing in-use device” condition.

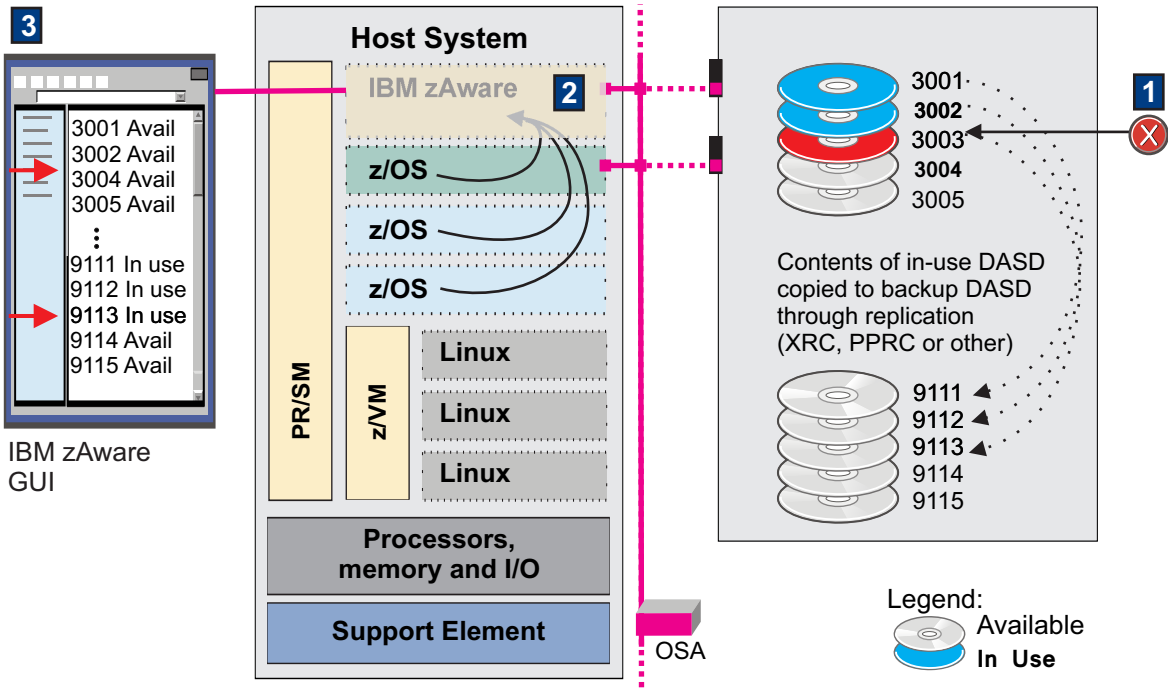

Figure 24. What happens when an in-use storage device becomes unavailable

1. In-use storage device 3003 is no longer attached to the IBM zAware partition. It might have been removed or disconnected through storage operations that are not provided through the **Data Storage** page of the IBM zAware GUI, such as:
   - Replacing the I/O definition file (IODF) for the host system with an IODF that does not contain the in-use storage devices for IBM zAware.
   - Using the Support Element (SE) to take offline one or more channel paths (CHPIDs) for storage devices or for the network through which those devices are connected to the IBM zAware partition.
2. When an in-use storage device becomes unavailable, IBM zAware operations stop. To a user of the IBM zAware GUI, the GUI appears to hang. On the SE for the IBM zAware host system, hardware messages indicate input/output (I/O) problems that are related to the loss of access to physical storage devices.

   At this point, an IBM zAware administrator can take only the following corrective actions to resolve this condition:
   a. If possible, try to reattach the device and reactivate the IBM zAware partition.
   b. If the device cannot be reattached, you must replace it with an equivalent device containing a backup copy of the data that was stored on the unavailable device. The original device and the backup device must be the same size.
   c. If your installation does not have backup copies of IBM zAware data, you must deactivate the IBM zAware partition and reconfigure the IBM zAware environment.
3. Assuming that the administrator could not reattach missing device 3003 but was able to reactivate the IBM zAware partition, the administrator must replace all in-use 300x devices with their equivalents in the 911x set of devices.
   a. When the administrator first logs in to the GUI, IBM zAware presents the **Data Storage** page with message AIFF0002W, which indicates that IBM zAware overwrites data on added storage devices.
   b. From the **Data Storage** page, the administrator clicks **Add and Remove Devices** to remove devices 3001 and 3002.
   c. When the removal operation is complete and no storage devices are listed as in use, the administrator uses **Add and Remove Devices** again, to select the replacement devices.

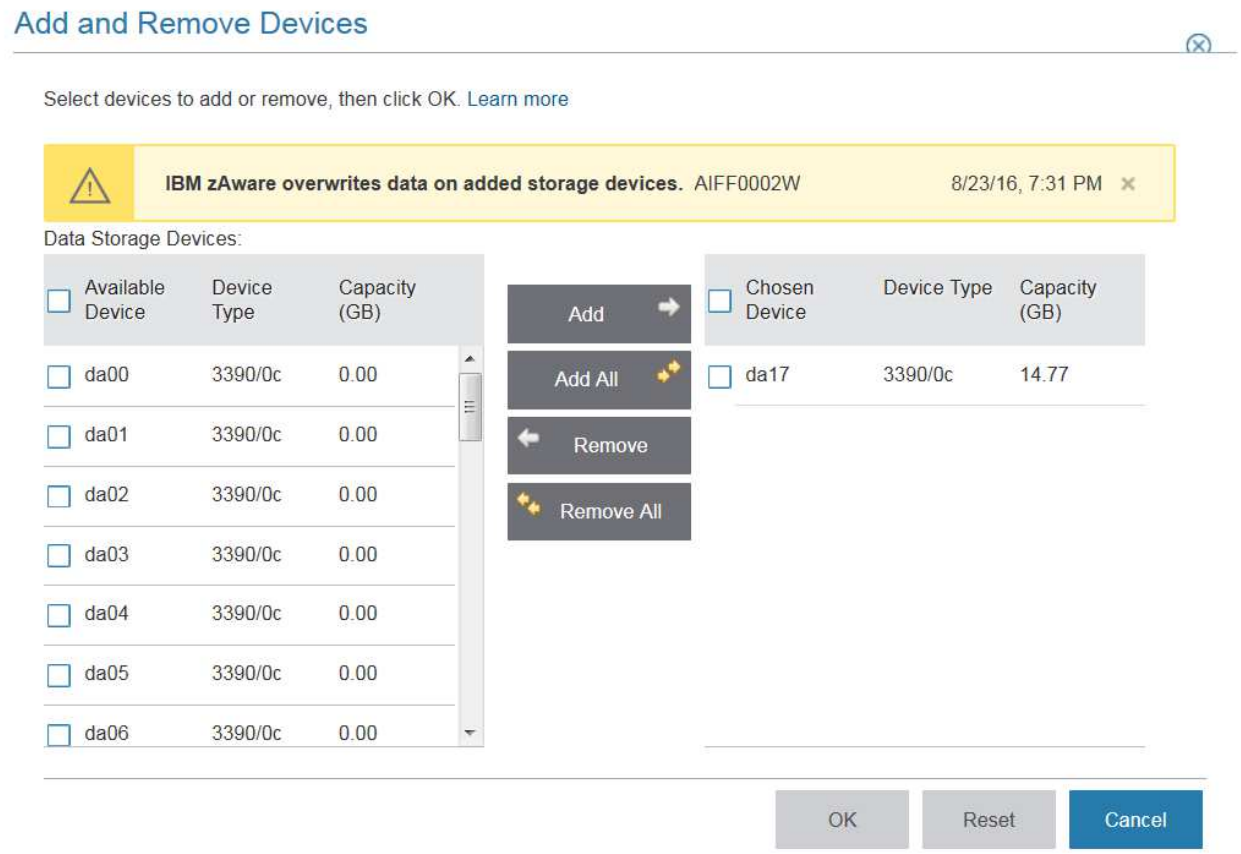

Figure 25. The Add and Remove Devices *window*

The GUI window depicted in Figure 24 on page 64 shows the results of the missing device condition and the administrator's corrective action:

- Devices 3001 and 3002 are listed as Available in the Data Storage Devices table.
- Device 3003 is no longer listed in the table.
- Devices 9111, 9112, and 9113 are listed as in-use devices.

If the replicated data on a backup device is back-level, IBM zAware cannot provide analytical data for the dates between the last day of replication and the date when the in-use storage device became unavailable. Also, during the time between the detection of the missing device condition and its successful resolution, any data that monitored clients attempt to send to IBM zAware is lost. If you consider this missing data to be critical for analyzing monitored clients, you can use the z/OS bulk load client for IBM zAware to resend any client OPERLOG data that might be missing.

Example: Storage configuration for multiple IBM zAware partitions

Your installation can configure more than one IBM zAware partition, with one for normal operations and another reserved for switchover situations. This type of configuration enables you to quickly restore IBM zAware operations after a failure.

- To quickly recover from a control unit (CU) failure on a storage device, set up an alternate IBM zAware on either the same or a different host system.

- To quickly recover from a central processor complex (CPC) failure, set up an alternate IBM zAware on a different host system. The alternate host system might have the IBM zAware disaster recovery (DR) feature installed, but this feature is not required.

When your installation configures a primary IBM zAware partition and an alternate IBM zAware partition for switchover situations, only one IBM zAware server can be active at a given time but both servers must have access to the same data.

- To correctly configure the partition in which the alternate server runs, use the same IP address as you defined for the primary partition. Doing so guarantees that you cannot have multiple IBM zAware servers running simultaneously, and also eliminates the need to reconfigure the TCP/IP settings of monitored clients if you have to switch from using the primary server to the alternate server.
- To correctly configure persistent storage for primary and alternate IBM zAware partitions, your installation must define physically separate but equivalent sets of storage devices for each partition, and also set up replication to copy the content of the primary storage devices to the alternate storage devices. For data replication to be successful, the number of storage devices in the primary set must match the number of devices in the alternate set. Additionally, each alternate device must be equivalent in size to the primary device.

The series of figures in this topic illustrate how to configure storage for a primary IBM zAware server and an alternate server that reside in separate host systems. For this example, assume that the storage administrator has designated devices 3001-3005 for IBM zAware data, and devices 9111-9115 to contain backup copies of that data.

Figure 26 shows the configuration for the primary IBM zAware server. Through the I/O definition file (IODF) and input/output configuration data set (IOCDS) for the IBM zAware host system, the partition has access to *only one* of the two physically separate sets of equivalent storage devices: 3001-3005. After the IBM zAware partition is activated, only the 300x devices are displayed as available for use through the Data Storage page in the IBM zAware graphical user interface (GUI). The numbered areas of Figure 26 illustrate the sequence of tasks that an IBM zAware administrator might follow to correctly configure these devices for normal operations and replication.

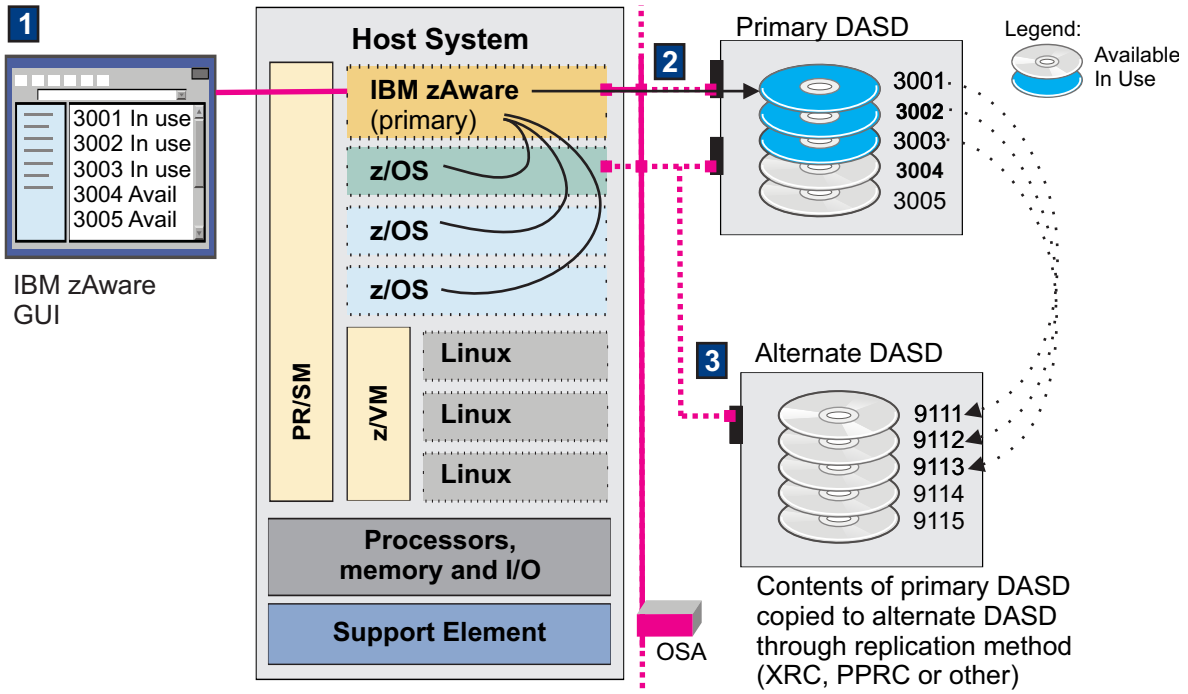

Figure 26. DASD configuration for the primary IBM zAware server on one host system

1. Using the Data Storage page in the IBM zAware GUI, an IBM zAware administrator selects devices for the primary IBM zAware server to use for normal operations. Note that only the 300x devices are listed in the GUI as available for use by the primary server.

   Before IBM zAware can store data in any of the 300x devices, an administrator must use the **Add and Remove Devices** function to select and add those storage devices. This action causes IBM zAware to format and initialize the added devices.

   For this sample configuration, an administrator has initially added only storage devices 3001-3003, reserving the other 300x devices for any additional capacity that might be needed when more monitored clients are added to the environment. The GUI window depicted in Figure 26 on page 67 shows the results of this selection process: the displayed status for devices 3001-3003 is now "In use".
2. With in-use devices available to store data, an administrator can finish configuring the primary IBM zAware environment by configuring and connecting monitored clients, sending priming data, and requesting IBM zAware to build models of client behavior.

   In Figure 26 on page 67, monitored clients that send data to the primary IBM zAware server are shown in partitions running on the same host system, but additional monitored clients can run in partitions on other z Systems servers. IBM zAware parses and stores the following data on in-use devices 3001-3003:
   - Current data from each z/OS or Linux monitored client, as well as priming data, if any.
   - IBM zAware models for each z/OS monitored client and for each model group, which is an administrator-defined group of Linux clients.
   - Analysis results for each z/OS or Linux monitored client.
3. After normal IBM zAware operations begin, an administrator can set up replication for the in-use devices. For this replication to be successfully completed, both the number and sizes of devices in the backup set match the number and sizes of in-use storage devices.

   For replication, your installation can consider using IBM FlashCopy or one of several Data Facility Storage Management Subsystem (DFSMS) copy services, including Extended Remote Copy (XRC) and Peer-to-Peer Remote Copy (PPRC) solutions. Another possible alternative is using DFSMShsm to copy data. In contrast to real-time replication solutions, DFSMShsm requires deactivating the IBM zAware partition before copying data, then reactivating the partition after the copy operation completes. Non-IBM replication products also are available for use.

   Figure 26 on page 67 shows that one of the monitored clients, the z/OS partition highlighted in light green, has access to IBM zAware in-use storage devices 3001-3005, as well as their equivalent backup set of devices, 9111-9115.

   An administrator has set up this z/OS image with a replication solution that copies data from the set of in-use devices to the backup set:
   - Device 9111 contains a copy of the data from device 3001.
   - Device 9112 contains a copy of the data from device 3002.
   - Device 9113 contains a copy of the data from device 3003.

With backup devices and data available for the primary IBM zAware server, an administrator can set up the alternate server to prepare for potential switchover situations. The numbered areas of Figure 27 on page 69 illustrate the sequence of tasks that an administrator follows to correctly configure the alternate IBM zAware server.

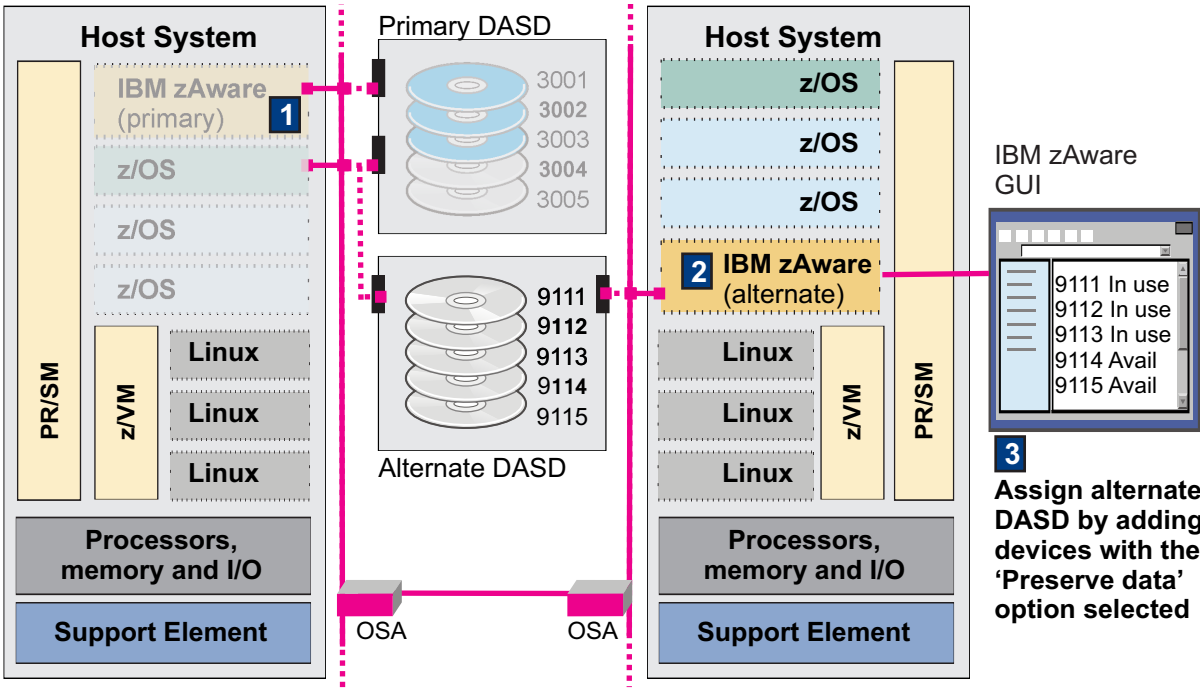

Figure 27. Configuration process for the alternate IBM zAware server on a different host system

1. The administrator disconnects any monitored clients and deactivates the partition in which the primary IBM zAware server is running.
2. To configure the alternate server, the administrator follows a configuration process that is identical to the configuration steps required for the primary server, with the exception of the available storage devices:
   a. The administrator must check the IODF and IOCDS for the host system on which the alternate partition resides. Through the HCD channel path definitions, the administrator limits access to storage devices 9111-9115 to only the alternate IBM zAware partition.
   b. Through the Hardware Management Console (HMC) for the alternate host system, the administrator customizes an image profile for the alternate IBM zAware partition, using the same IP address as the address defined for the primary partition.
   c. Through the HMC for the alternate host system, the administrator activates the alternate IBM zAware partition.
3. The administrator logs in to the IBM zAware GUI to configure storage and security for the alternate IBM zAware server. Because only devices 3001-3003 are being used by the primary IBM zAware server, the administrator plans to assign only the corresponding alternate devices: 9111-9113.
   a. When the administrator first logs in to the GUI, IBM zAware presents the Data Storage page because no persistent storage has been assigned yet.
   b. From the **Data Storage** page, the administrator clicks **Add and Remove Devices** to select the alternate devices. The Add and Remove Devices window contains the **Preserve data** option, which is intended for use *only* when adding a storage device that contains replicated data from an in-use device, and is displayed only when no other devices are in use.
   c. The administrator selects both the **Preserve data** option and devices 9111-9113, then clicks **Add** to add these devices to the list of in-use storage devices. Because the **Preserve data** option was selected, IBM zAware does *not* format the devices to be added. The replicated data on those devices is preserved and usable.

      The GUI window depicted in Figure 27 shows the results of this device assignment: Devices 9111-9113 are listed as in-use devices for the alternate IBM zAware server.

d. Except for the storage devices to be added, the configuration of the alternate server can exactly match that of the primary server. For example, if your installation is using an existing Lightweight Directory Access Protocol (LDAP) server for user authentication to the primary server, the administrator can configure the same LDAP server for use with the alternate server.

After the persistent storage and security configuration for the alternate server is complete, the administrator completes the following steps to reactivate the primary IBM zAware environment:

1. Deactivate the alternate IBM zAware partition.
2. Reactivate the primary IBM zAware partition.
3. Reconnect the monitored clients.

During normal operations, an administrator might need to change the set of in-use devices for the primary server by adding or removing devices through the GUI. Because the in-use set and backup set of devices must be equivalent for a switchover to be successful, the administrator also must adjust replication and the set of storage devices for the alternate server so both of the primary and alternate sets match in number of devices, size of devices, and content.

After an administrator completes the initial configuration of both the primary and alternate servers, either IBM zAware server can detect and report mismatches only when its partition is activated. Remember that only one IBM zAware server can be active at a time. The following example illustrates how, on activation, the primary IBM zAware reports mismatches between its set of in-use devices and the set of devices for the alternate IBM zAware server. Similarly, on activation, the alternate IBM zAware can detect and report mismatches between its set of in-use devices and the primary set.

- If a mismatch occurs because the administrator removed a device from the alternate set, the activated primary IBM zAware server issues message AIFP0012I to indicate that a device was removed from the alternate set, and takes corrective action by removing the equivalent device from its in-use set of devices. In this case, no administrator intervention is required.
- If a mismatch occurs because the administrator added a device to the alternate set, the activated primary IBM zAware issues message AIFP0013E to indicate that a device is missing from its in-use set. In this case, IBM zAware operations stop until an administrator successfully adds an equivalent device to the storage configuration of the primary IBM zAware server.

Figure 28 on page 71 illustrates how an administrator can switch the IBM zAware environment from the primary to the alternate server when a switchover situation occurs.

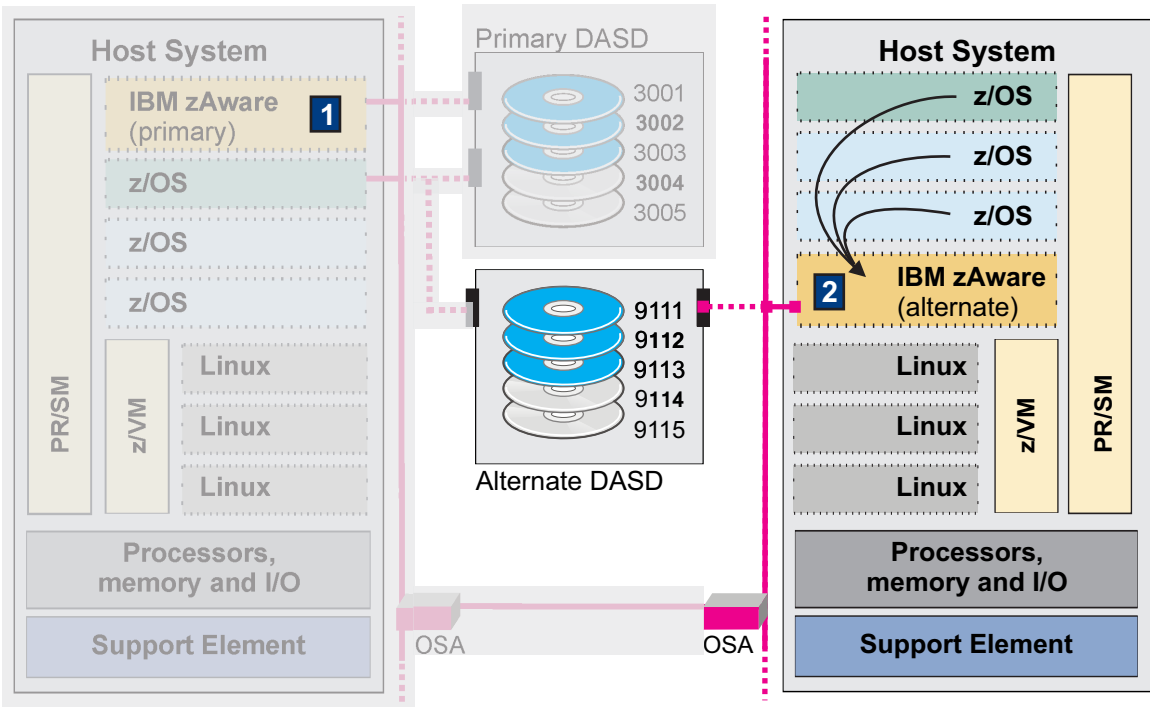

Figure 28. What happens when a switchover situation occurs

1. The host system for the primary IBM zAware experiences a CPC failure. IBM zAware operations stop. To a user of the IBM zAware GUI, the GUI appears to hang.
2. The administrator activates the alternate IBM zAware on its host system. Because its configuration is the same as that of the primary, with the exception of storage devices, the administrator can reIPL the monitored clients and reconnect them to the alternate server.

   Depending on the timing of the CPC failure and the replication schedule for backing up IBM zAware data, the data on devices 9111-9113 might be back-level. In this case, the alternate IBM zAware cannot provide analytical data for the dates between the last day of replication and the date and time when the administrator switched the IBM zAware environment to the alternate CPC.

Summary of planning considerations for persistent storage

Table 12 on page 72 provides a summary of planning considerations and best practices for IBM zAware storage configuration.

Table 12. Planning considerations and best practices for IBM zAware storage configuration

Planning consideration	Usage notes and best practices
Storage device requirements	• Your installation can use Extended Count Key Data (ECKD) direct-access storage devices (DASD) to store IBM zAware data. • Your installation can select a combination of small volumes or volumes of different sizes to provide the storage capacity that your IBM zAware environment requires. These volumes cannot be SMS-managed volumes. • If any devices are to be used for storing backup copies of IBM zAware data, your installation must define two physically separate but equivalent sets of storage devices: – One set for IBM zAware to use for normal operations. – Another set for storing backup copies of data. The number of storage devices in each set must match, and each backup device must be equivalent in size to the device from which the data is copied. These number and size requirements also apply for configurations that contain primary and alternate IBM zAware partitions. • Storage devices that are configured for use by IBM zAware must be available for write operations. Specifically, the devices cannot be target devices in a FlashCopy relationship, secondary devices of a PPRC configuration, or otherwise write-inhibited.
Partition access to storage devices	• When defining storage devices in the IODF or IOCDS for the IBM zAware host system, the storage administrator can use image access and candidate lists for channel path definitions to allow only the IBM zAware partition to access specific devices. Using the explicit device candidate list is an alternative method of restricting access to specific devices. • As a precaution, you can use channel path definition lists, the explicit device candidate list, or operating system mechanisms to prevent the applications that are running in other CPC partitions from using storage devices that are intended for IBM zAware use. The only exception to this practice is the CPC partition, if any, that your installation is using to back up the storage that the IBM zAware server is using.
Minimum storage amount and the number of monitored clients	Storage requirements vary depending on the retention times for each type of analytical data and on the number of monitored systems that you plan to connect to IBM zAware. Start with 500 GB of storage for IBM zAware to use, plus 4 - 5 GB of storage for each monitored system. If you increase the number of monitored clients, you need to configure an extra 4 - 5 GB of storage for each monitored system.
Adjustable retention times for stored data	IBM zAware stores the following analytical data on DASD: • Current data from each z/OS or Linux monitored client, as well as priming data, if any. • IBM zAware models for each z/OS monitored client and for each model group, which is an administrator-defined group of Linux clients. • Analysis results for each z/OS or Linux monitored client. IBM zAware sets default retention times for each of these types of analytical data and, through an automated process, removes the data when the retention time has elapsed. Reducing these default retention times might reduce the amount of storage you need to configure for IBM zAware; if necessary, you can reduce the default retention times through the **Administration** > **Configuration** > **Analytics** page in the IBM zAware GUI.

Table 12. Planning considerations and best practices for IBM zAware storage configuration (continued)

Planning consideration	Usage notes and best practices
Backup copies of IBM zAware data	Creating and maintaining a backup copy of IBM zAware data is recommended for the following reasons: • IBM zAware does not automatically replicate any of its data. • If a storage device is damaged, disconnected, or removed from the CPC, IBM zAware effectively loses access to its data and cannot continue to analyze data from monitored clients. Only the following corrective actions can resolve this condition: 1. If possible, try to reattach the device and reactivate the IBM zAware partition. 2. If the device cannot be reattached, you must replace it with an equivalent device containing a backup copy of the data that was stored on the unavailable device. The original device and the backup device must be the same size. 3. If your installation does not have backup copies of IBM zAware data, you must deactivate the IBM zAware partition and reconfigure the IBM zAware environment. Guidelines for configuring and managing backup devices and data: • If any devices are to be used for storing backup copies of IBM zAware data, your installation can define physically separate but equivalent sets of storage devices, one set for IBM zAware to use for normal operations (the in-use set), and another set for storing backup copies of data. The number of storage devices in the in-use set must match the number of devices in the backup set. Additionally, each backup device must be equivalent in size to the in-use device. • The same number and size requirements apply for configurations containing primary and alternate IBM zAware servers. Your installation must define physically separate but equivalent sets of storage devices for each server, with the same number of storage devices in the primary set and the alternate set. Additionally, each alternate device must be equivalent in size to the primary device. • For replication, your installation can consider using IBM FlashCopy or one of several Data Facility Storage Management Subsystem (DFSMS) copy services, including Extended Remote Copy (XRC) and Peer-to-Peer Remote Copy (PPRC) solutions. Another possible alternative is using DFSMShsm to copy data. In contrast to real-time replication solutions, DFSMShsm requires deactivating the IBM zAware partition before copying data, then reactivating the partition after the copy operation completes. Non-IBM replication products also are available for use. • If you change the set of in-use devices by adding or removing devices through the GUI, make sure that you adjust replication accordingly. To successfully replace an in-use device with its equivalent backup device, the set of in-use devices must match the set of backup devices in number of devices, size of devices, and content. This requirement also applies when switching between primary and alternate IBM zAware servers. • If the replicated data on a backup device is back-level, IBM zAware cannot provide analytical data for the dates between the last day of replication and the date and time when the administrator replaced an in-use storage device with its equivalent backup device.

Task summary and configuration checklist for storage administrators

- Table 13 on page 74 provides a summary of storage administration tasks and links to additional information.
- Table 14 on page 74 is a checklist that a storage administrator can use to complete step 2 on page 97 in Chapter 12, “Configuring network connections and storage for the IBM zAware partition,” on page 95.

Table 13. Task summary for storage administrators

✓	Task summary:	Where to find instructions:
	Configure persistent storage for the IBM zAware partition through the Hardware Configuration Definition (HCD).	Step 2 on page 97 in Chapter 12, "Configuring network connections and storage for the IBM zAware partition," on page 95
	Assign storage devices for the IBM zAware server through the **Administration** > **Configuration** > **Data Storage** page in the IBM zAware graphical user interface (GUI).	Step 2 on page 100 in Chapter 13, "Configuring storage, security, and analytics for the IBM zAware server," on page 99
	Manage storage devices that the IBM zAware server is using through the GUI **Administration** > **Configuration** > **Data Storage** page.	"Adding and removing storage devices" on page 195
	(Optional but highly recommended) Use the replication method of your choice to back up the storage devices that the IBM zAware server is using	The product documentation for the replication method that you choose. If you are using a DFSMS copy service or DFSMShsm, see the relevant z/OS DFSMS topics in the z/OS Information Center at this URL:http://publib.boulder.ibm.com/infocenter/zos/v1r13/
	(Optional) Configure an alternate IBM zAware partition for use in switchover situations	Chapter 26, "Setting up multiple IBM zAware partitions for switchover situations," on page 267

Table 14. Checklist for Extended Count Key Data (ECKD) storage devices

✓	Name	Size[1]	Type	Location	Device number	Volume serial
Sample	PoolEckd01	10017	3390-9	SYSTEM1	9051	AC3231
Backup storage devices[2]						

Footnotes for Table 14:

1. The size in cylinders for ECKD devices. Only numbers 0-9 are allowed. The maximum value is 2^{63} -1 (9,223,372,036,854,775,807).
2. To store backup copies of IBM zAware data through replication, provide one backup device for each in-use device; the backup device must match the size of the in-use device.

Chapter 9. Planning for security

IBM zAware does not require clients to provide authentication credentials or to encrypt the data that they send. If your installation considers this data to be sensitive, you need to ensure that the communication between IBM zAware and its monitored clients occurs over secured networks that are configured with preexisting security mechanisms. IBM zAware does provide security mechanisms through which you can limit access to the IBM zAware graphical user interface (GUI), through which users can view the analytical data for clients and modify the operational controls for the IBM zAware server.

Securing communication between IBM zAware and its monitored clients

As stated and illustrated in Chapter 7, "Planning your IBM zAware environment," on page 39, the recommended configuration for an IBM zAware environment is contained within a single security zone protected by a firewall.

For a z/OS monitored client to be correctly configured, the z/OS system logger must have Security Authorization Facility (SAF) authorization to send data to the IBM zAware server. Additional security options, which you can use for authentication, data privacy, and data integrity, depend on the type of network connection that the server and clients use. As noted in "Planning network connections and capacity" on page 52, the supported network connections are:

- Ethernet connectivity through an Open Systems Adapter (OSA) channel. If you are using an OSA channel for the IBM zAware environment, you can restrict access to the server and client IP addresses or to specific networks by using SERVATH and NETACCESS profiles.

  Using NETACCESS statements, z/OS Communications Server can map networks, subnetworks and IP addresses to SAF resource names. Users that are not permitted access to a particular SAF resource are not allowed to communicate with the corresponding network, subnetwork, or IP address.

  If you are using the OSX channel type for the IBM zAware environment, you can implement a virtual local area network (VLAN) to enforce isolation between networks.
- Connectivity through the intraensemble data network (IEDN). If you are using the IEDN for the IBM zAware environment, you can implement VLANs to enforce isolation between networks.
- Connectivity through HiperSockets. If you are using HiperSockets, you can implement VLANs to enforce isolation between networks. Because a HiperSocket is an in-memory socket, it inherits the normal z Systems memory protections.

Firewall consideration: If communication from remote monitored clients must pass through a firewall to the IBM zAware server, you might need to configure the firewall to allow incoming connections to the server on port 2001. For additional information about securing network connections, see *z/OS Communications Server IP Configuration Guide*, SC31-8775.

Securing communication between IBM zAware and GUI users

IBM zAware provides the following security mechanisms that your installation can configure to limit access to the IBM zAware GUI.

Master user ID and password

Your installation defines the master user ID in the image profile of the IBM zAware partition. The password does not expire and IBM zAware does not provide any mechanism that requires you to change the password on a regular basis. If you forget the password, you can restore access by setting a new default user ID and password through in the image profile.

After configuring the IBM zAware environment and activating the IBM zAware partition, you must use the default master user ID to initially log in to the IBM zAware GUI. This master user ID has authority to perform any task that is available through the IBM zAware GUI.

See "Configuring the IBM z Systems Secure Service Container for IBM zAware" on page 27 for instructions for using the Hardware Management Console (HMC) to define the master user ID in the image profile of the partition.

Server SSL certificate

A Secure Sockets Layer (SSL) certificate is automatically generated for IBM zAware when your installation initially activates the IBM zAware partition. The certificate is not signed by a certificate authority (CA). Therefore, the first time you log in to the IBM zAware graphical user interface (GUI), the browser displays a warning message because it does not recognize the default SSL certificate. You can resolve this problem by replacing the default SSL certificate with a certificate that is signed by a certificate authority of your choice. Doing so provides secure communication between the IBM zAware server and the browsers of all authorized users.

If you do not replace the automatically generated certificate, users can bypass the browser error message by adding a security exception, but cannot verify that they are connected to a legitimate IBM zAware partition. The automatically generated certificate is valid for one year from the initial activation of the IBM zAware partition. It is not automatically renewed. You must manually replace it with a new one after one year.

If you decide to replace the automatically generated certificate, which is the recommended practice for improved security, you can use any third-party certificate authority of your choice, or your installation can provide an internal certificate authority for certificate signing tasks. IBM zAware does not renew these replacement certificates; in this case, managing replacement certificates becomes the responsibility of the security administrator.

The required format for replacement certificates is Base64 encoded X509 certificate blocks.

Chapter 13, "Configuring storage, security, and analytics for the IBM zAware server," on page 99 provides instructions for replacing the default self-signed SSL certificate, and provides sample certificate blocks to illustrate the content that you might receive from a certificate authority. When you receive a reply from a certificate authority, the reply might contain a chain of certificates, starting with the signed server certificate. Then, it is possibly followed by certificates from one or more intermediate CAs and finally, the self-signed certificate of the CA. When you supply this information in the GUI, provide the entire certificate chain, starting with the signed server certificate through the self-signed certificate of the CA.

In some cases, the CA reply that you receive is delivered in a public-key cryptography standards (PKCS) #7 file. You must extract the certificates from the file before pasting them into the GUI. One method of extracting certificates from a PKCS #7 file is to use the **openssl pkcs7** command; for more information, see the OpenSSL Project website at the following URL.

www.openssl.org/

User authentication for the IBM zAware GUI

You can authorize users to access the IBM zAware server through the use of an existing Lightweight Directory Access Protocol (LDAP) repository or, alternatively, through the use of a local file-based repository.

For simplicity, using only an LDAP repository is the preferred option. However, you might want to define one or two user IDs in a local repository so you can access the IBM zAware GUI when the LDAP server is unavailable. If you configure an LDAP repository and also define users or groups in a local repository, both sets of users or groups are available through the IBM zAware GUI. Do not define the same user ID in more than one repository; results are not predictable.

- The IBM zAware GUI provides pages through which you can configure an LDAP repository. Chapter 13, "Configuring storage, security, and analytics for the IBM zAware server," on page 99 provides instructions for configuring an LDAP repository through the IBM zAware GUI.
- To use a local file-based repository instead of an LDAP repository, use the instructions in Chapter 24, "Setting up a local repository to secure access to the IBM zAware GUI," on page 261. Those instructions explain how to use the zAware GUI to add users or groups to, or remove them from, the repository.

LDAP data uses a tree structure of user information, as shown in Figure 29. When IBM zAware is configured to use LDAP, the IBM zAware administrator can retrieve user entries from the LDAP server and assign an IBM zAware security role, Administrator or User, to each of those users so they can successfully log in to the IBM zAware GUI.

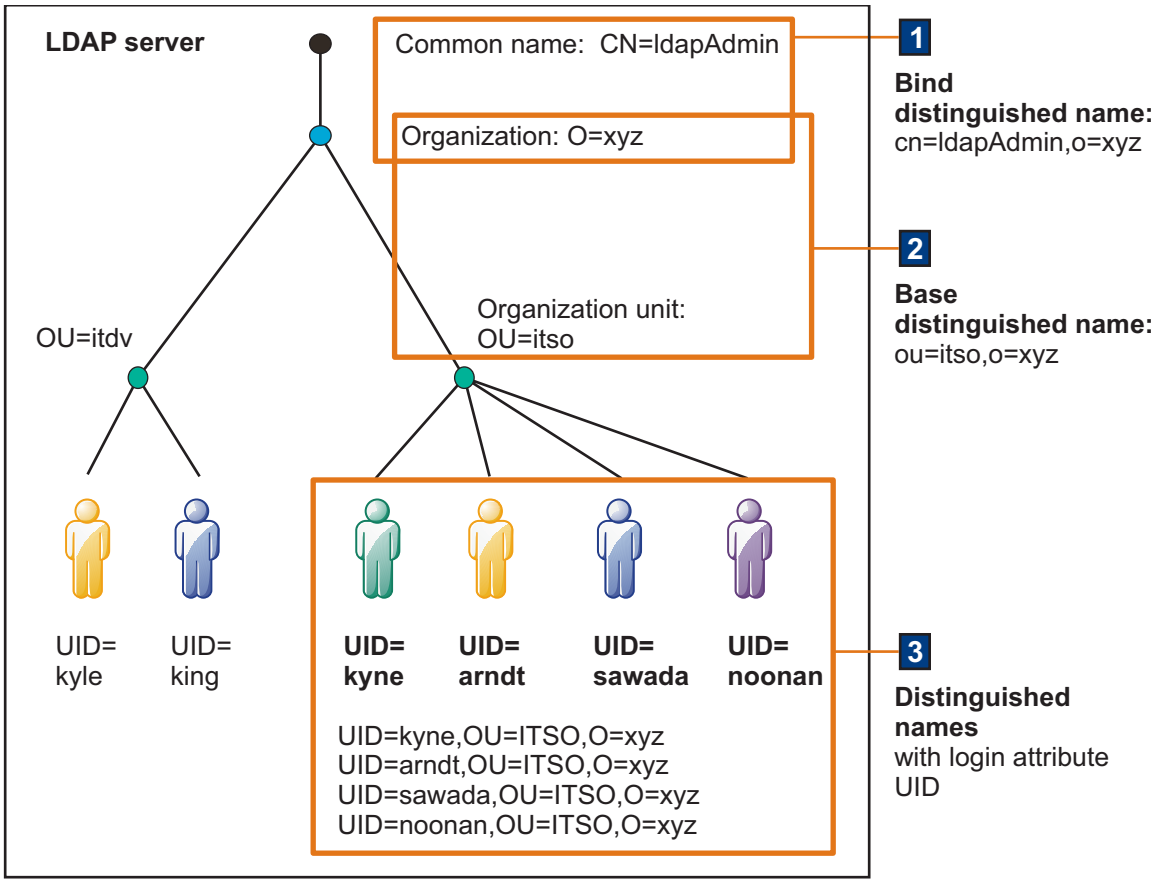

Figure 29. LDAP directory information tree

To configure IBM zAware to use LDAP for user authentication, an IBM zAware administrator supplies details about the LDAP server and its directory through the **Configuration** > **Security** > **LDAP Settings** tab in the GUI. The settings include the host name and port number of an LDAP server to which IBM zAware will communicate, or bind, to authenticate users.

1. A valid bind distinguished name and bind password are optional settings. The bind distinguished name specifies the LDAP entry that IBM zAware uses to bind to LDAP for searches and retrieval of user entries; if this information is not specified, the IBM zAware server binds anonymously. The bind distinguished name consists of the common name (CN) and the organization name (O); for example, `cn=ldapAdmin,o=xyz`
2. The base distinguished name is required. The base distinguished name specifies the portion of the LDAP directory tree that IBM zAware searches for entries. The base distinguished name consists of the organization unit (OU) and the organization name (O); for example, `ou=itso,o=xyz`
3. In addition to the base distinguished name, IBM zAware requires a login attribute (for example, UID), and at least one user object class value to access the list of distinguished names of users to be authenticated. The login attribute specifies what LDAP attribute is used as the login attribute on the IBM zAware GUI login prompt, and the user object classes specify the type of LDAP objects that are searched and retrieved by IBM zAware.

Figure 30 illustrates sample values entered in the General settings fields on the **LDAP Settings** tab.

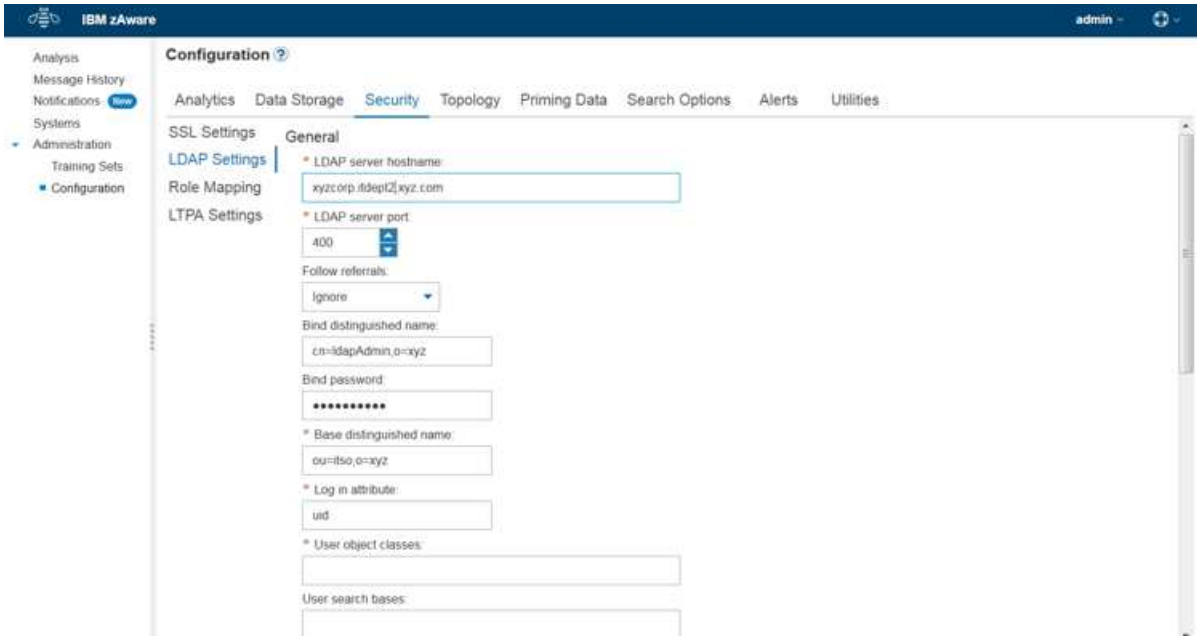

Figure 30. Sample General LDAP settings

If you are using LDAP groups, IBM zAware also can retrieve group objects from LDAP. To configure IBM zAware to authenticate users in an LDAP group, an IBM zAware administrator enters values in the Group settings fields on the **LDAP Settings** tab. Required Group settings include group object classes and group member attributes.

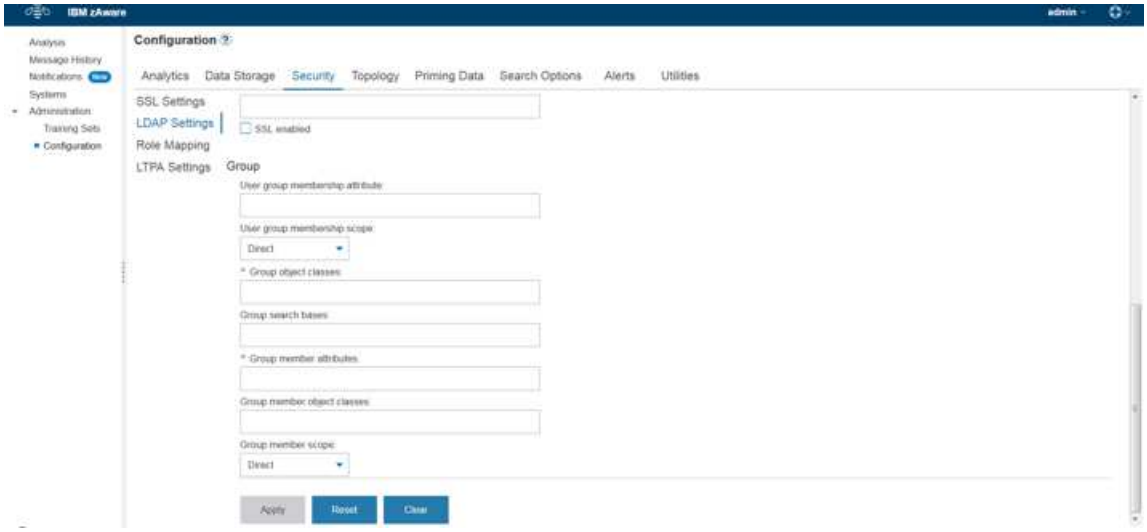

Figure 31. Sample LDAP Group settings

After an IBM zAware administrator applies the LDAP settings, IBM zAware can retrieve user IDs from the LDAP server, and assign users and groups to IBM zAware specific roles. These roles determine which GUI pages and functions are available to each user or group. After an LDAP user ID has been assigned a role, the person using that ID can then authenticate to the IBM zAware GUI. IBM zAware issues a bind to the LDAP server using the data entered (user ID and

password) in the IBM zAware GUI login prompt. If the bind is successful, IBM zAware allows the user access. The user ID that is entered in the IBM zAware login panel is determined by the login attribute field of the **LDAP Settings** tab. For example, if the login attribute field is set to UID and an LDAP entry has the distinguished name of `kyne,ou=itso,o=xyz`, then the person must enter the user ID `kyne` as the login name on the IBM zAware login prompt.

Role-based access to IBM zAware GUI functions

To grant access to the GUI, you need to map authorized users and groups to specific roles: either **Administrator** or **User**. User IDs or groups that are assigned to the Administrator role can access and use any task in the GUI, including those listed under **Administration** in the GUI navigation tree. Administrators have the authority to customize the IBM zAware environment and operations; for example:

- Through the **Analytics** page, an administrator can change data retention and training values that affect analytical operations.
- Through the **Data Storage** page, an administrator can add to or remove storage devices from the IBM zAware configuration.
- Through the **Security** page, an administrator can configure user authentication controls.

User IDs or groups that are assigned to the User role cannot view any of the pages under **Administration** in the GUI navigation tree. User IDs or groups assigned to the User role can view only the following pages and use only the actions as noted:

- On all views of the **Analysis** page, all controls and actions are permitted.
- On the **Interval** page, all controls and actions are permitted except for modifying the non-IBM rules status for a specific message ID. Only administrators can view and change a rules status value. IBM rules cannot be changed.
- On the **Notifications** page, all actions are disabled.
- On the **Systems** > **System Status** tab, all actions are disabled.
- On the **Systems** > **Model Groups** tab, all actions except for **Search Systems** are disabled.

Chapter 13, “Configuring storage, security, and analytics for the IBM zAware server,” on page 99 provides instructions for mapping user IDs or groups to specific roles.

Browser session timeout setting

By default, browser sessions time out after 12 hours (720 minutes). Your installation can change this setting through the **LTPA Settings** tab on the **Administration** > **Configuration** > **Security** page in the IBM zAware GUI.

Task summary for security administrators

Table 15 provides a summary of security administration tasks and links to additional information.

Table 15. Task summary for security administrators

✓	Task summary:	Where to find instructions:
	Collaborate with the network administrator if you have determined that additional network definitions are required to ensure secure communications between the IBM zAware server and its monitored clients.	• Additional security options, which you can use for authentication, data privacy, and data integrity, depend on the type of network connection that the server and clients use. • If communication from remote monitored clients must pass through a firewall to the IBM zAware server, you might need to configure the firewall to allow incoming connections to the server on port 2001. For additional information about securing network connections, see *z/OS Communications Server IP Configuration Guide*, SC31-8775.

Table 15. Task summary for security administrators (continued)

✓	**Task summary:**	**Where to find instructions:**
	Collaborate with the network administrator to determine whether any additional network definitions are required to ensure communication between the IBM zAware server and the Lightweight Directory Access Protocol (LDAP) server.	Your installation has the option to configure user authentication through the use of an LDAP repository. If a firewall exists between the IBM zAware partition and the LDAP server, the IBM zAware partition must be able to use the port that is used by the LDAP server.
	(Optional) Secure the communication between the IBM zAware server and browsers by requesting and importing a digital certificate.	Step 3 on page 102 in Chapter 13, "Configuring storage, security, and analytics for the IBM zAware server," on page 99.
	Configure the LDAP repository or a local file-based repository for storing user access information.	Step 4 on page 103 in Chapter 13, "Configuring storage, security, and analytics for the IBM zAware server," on page 99.
	Authorize users or groups to access the IBM zAware GUI.	Step 5 on page 106 in Chapter 13, "Configuring storage, security, and analytics for the IBM zAware server," on page 99.
	Allow users to browse the operations log (OPERLOG).	Step 6 on page 115 in "Configuring z/OS monitored clients to send data to the IBM zAware server" on page 111.
	Authorize the z/OS system logger to communicate with the IBM zAware server.	Step 3 on page 113 in "Configuring z/OS monitored clients to send data to the IBM zAware server" on page 111.

Chapter 10. Planning to use the IBM zAware GUI

Through the IBM zAware graphical user interface (GUI), you can view analytical data that indicates which system is experiencing deviations in behavior, when the anomalies occurred, and details about unusual messages and message patterns. Using this information, you can take corrective action for these anomalies before they develop into more visible problems.

The GUI provides pages through which you can accomplish many tasks. Some include the following tasks:

- Checking the status of monitored clients that are connected to the IBM zAware server.
- Modifying the IBM zAware models for monitored clients.
- Configuring analytics settings, storage, security, analytics settings, topology, priming data, and more.

The following sections cover the basics to getting started

Browser prerequisites

To take full advantage of the IBM zAware GUI, you must use one of the following browsers. Edit your browser options to enable JavaScript, Cascading Style Sheets (CSS), and cookies. Disable software that blocks pop-up windows, especially if you are using keyboard controls rather than the mouse to use the GUI.

- Mozilla Firefox Extended Support Release (ESR) 45
- Microsoft Windows Internet Explorer (IE) 11, 10, or 9. Compatibility View must be turned off for all versions of IE.

Other browsers and browser release levels might work but are not tested; if you use them, some IBM zAware functions might not be available and page content might not display correctly.

IBM zAware title bar

The title bar of the IBM zAware V3.1 GUI contains the user functions in a menu that displays as the name of user who is logged in to IBM zAware V3.1 The menu contains the **User Profile** and **Logout** function. To open, click the down arrow next to the name. For more information, see “Setting up the User Profile to send email alerts” on page 83. When you are ready to **Logout** of IBM zAware, click **Logout**.

Table 16. Help menu

Menu Action	Description
Help	The Help  menu contains the options for embedded help, Container Settings, Environment Checker, and About. To see information about using the actions, from any IBM zAware page, click  , and then click **Help**.
Container Settings	Click **Container Settings** to view the IBM z Systems Secure Service Container.

Table 16. Help menu (continued)

Menu Action	Description
Environment Checker	To take full advantage of the IBM zAware GUI, you must use one of the following browsers. Edit your browser options to enable JavaScript, Cascading Style Sheets (CSS), and cookies. Disable software that blocks pop-up windows, especially if you are using keyboard controls rather than the mouse to use the GUI. • Mozilla Firefox Extended Support Release (ESR) 45 • Microsoft Windows Internet Explorer (IE) 11, 10, or 9. Compatibility View must be turned off for all versions of IE. Other browsers and browser release levels might work but are not tested; if you use them, some IBM zAware functions might not be available and page content might not display correctly. IBM zAware provides an environment checker that evaluates your browser level and settings. To use it, click the down arrow next to the Help icon ( ) on the IBM zAware header, and select **Environment Checker**. The Environment Checker window opens in a new browser tab, and presents the evaluation results, as illustrated in Figure 32. If the current setting for a particular option does not meet requirements, the display includes a warning icon for that setting. **IBM® z Advanced Workload Analysis Reporter (IBM zAware) - Environment Checker** The environment checker tool has inspected your workstation for compliance with IBM zAware. **Environment Option** / **Settings as of Tue Aug 23 2016 15:58:08 GMT-0400 (Eastern Standard Time)** / **Requirements** JavaScript / JavaScript enabled / JavaScript must be enabled. Cookies / Cookies enabled / Cookies must be enabled. DOM Storage / DOM Storage enabled / DOM Storage is required only for convenience features, such as saved selections and bookmarks. Pop-up Windows / Pop-up windows enabled / Pop-up windows must be enabled to view the IBM zAware online help. Screen Resolution / 1536 by 864 / Minimum screen resolution of 1024 by 768. Browser Content Dimensions / 1536 by 715 / Minimum browser content dimensions of 800 by 600. Browser Name and Version Browser User-Agent value / Supported Browser Mozilla/5.0 (Windows NT 6.1; WOW64; rv:45.0) Gecko/20100101 Firefox/45.0 / Supported browsers: Mozilla Firefox 45 Internet Explorer 11 Other browsers and release levels might work, but have not been tested; if you use them, some IBM zAware functions might not be available and page content might not display correctly. IBM zAware Version / Version: 3.1.0 Build ID: 20160822_1358 (Mon Aug 22 18:19:57 2016 UTC) / — *Figure 32. Sample environment checker results*
About	About IBM z Advanced Workload Analysis Reporter (IBM zAware) Version 3.1.0 Build *build identifier* *date and time* Licensed Materials - Property of IBM Corp. © Copyright IBM Corporation 2012, 2019. IBM, the IBM logo, and ibm.com® are 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. Linux is a registered trademark of Linux Torvalds in the United States, other countries, or both. Microsoft and Windows are trademarks of Microsoft Corporation in the United States, other countries, or both. Java and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

Network connections

As noted in "Planning network connections and capacity" on page 52, the most logical network option for browser access to the IBM zAware server is an Open Systems Adapter (OSA) channel for a customer-provided data network. This OSA channel path is defined to the IBM zAware partition on which the server runs. For users to access the IBM zAware GUI through this network channel path, port 80 (HTTP) and port 443 (HTTPS) must be open for inbound communication from users and outbound communication from the IBM zAware server.

Security summary

The following list provides a summary of security considerations that are related to the configuration and use of the IBM zAware GUI. Chapter 9, "Planning for security," on page 75 provides more details.

- Browser certificate authority (CA) certificate
- Master user ID and password
- User authentication for the IBM zAware GUI
- Role-based access to IBM zAware GUI functions
- Browser session timeout setting

Integration with system management products

IBM zAware provides an application programming interface (API) through which existing alerting products can be enhanced by including IBM zAware data into their alerting framework. For example, if IBM Tivoli OMEGAMON XE for z/OS detects a service level agreement (SLA) violation, it can use IBM zAware anomaly information to confirm that the SLA violation needs immediate attention. Through the IBM zAware API, system management products can request and receive IBM zAware analytical data in XML format. This data is equivalent to the information that is available through the Analysis views and Interval page in the IBM zAware GUI.

- IBM Operations Analytics for z Systems includes IBM zAware. For more information, see "Linking to IBM Operations Analytics for z Systems for message analysis" on page 168.
- Starting with Tivoli OMEGAMON on z/OS V5.1.1, IBM zAware data is consolidated with performance and other information to support diagnoses of problems and to include in OMEGAMON XE on z/OS situations. OMEGAMON XE on z/OS provides a workspace through which users can display, manage, and customize IBM zAware data.
- Your installation can configure other system management products, such as IBM Tivoli NetView® for z/OS, to use IBM zAware data.
- Your installation can configure the z/OS Management Facility (z/OSMF) so that users can start the IBM zAware GUI from the z/OSMF **Links** page.

For more information, see Chapter 27, "Enabling system management products to use IBM zAware data," on page 271.

Setting up the User Profile to send email alerts

Set up the **User Profile** to send email alerts that notify you when an anomaly occurs.

Before you begin

Before you set up the **User Profile** to receive email alerts, define the Simple Mail Transfer Protocol (SMTP) that your installation uses for email. For more information, see "Defining an SMTP email server" on page 85.

About this task

The anomaly scores, message alerts, or both are used to generate the email alerts. Before you complete the fields in the **User Profile** menu, ensure that you understand the minimum and maximum anomaly scores for the IBM z Advanced Workload Analysis Reporter (IBM zAware).

Procedure

1. Click **Name** > **User Profile**. In the **Personal Info** pane, enter the email address that receives the alerts, and then click **Save**. If the SMTP server is properly defined, you can send a test email. For more information, see "Defining an SMTP email server" on page 85.
2. Set up the type of alert you want to receive in the **My Alerts** pane. To receive alerts, you must select at least one option in the **Interval Anomaly Score** or the **Message Alerts**.
   - **Interval Anomaly Score** - Select the minimum score, the maximum score, or both.
   - **Messages not in the current model** - Check the box for notification about new messages that were not seen in the current model.
   - **Messages never encountered before** - Check the box for notification about messages that were never seen in the previous models.
3. Click **Save** to complete the **User Profile** configuration.

Results

You successfully set up IBM zAware to alert you by email when an anomaly in the range that you defined occurs.

Table 17. Required fields for email alerts

Required Fields	Description
Email Address	The email address that receives the alerts when an interval anomaly occurs.
My Alerts	To receive an email alert, you must specify at least one option from the **My Alerts** pane. For example, either one **Interval Anomaly Score** entry or one **Message Alerts**. You can also configure the profile to notify you about all events.

Example

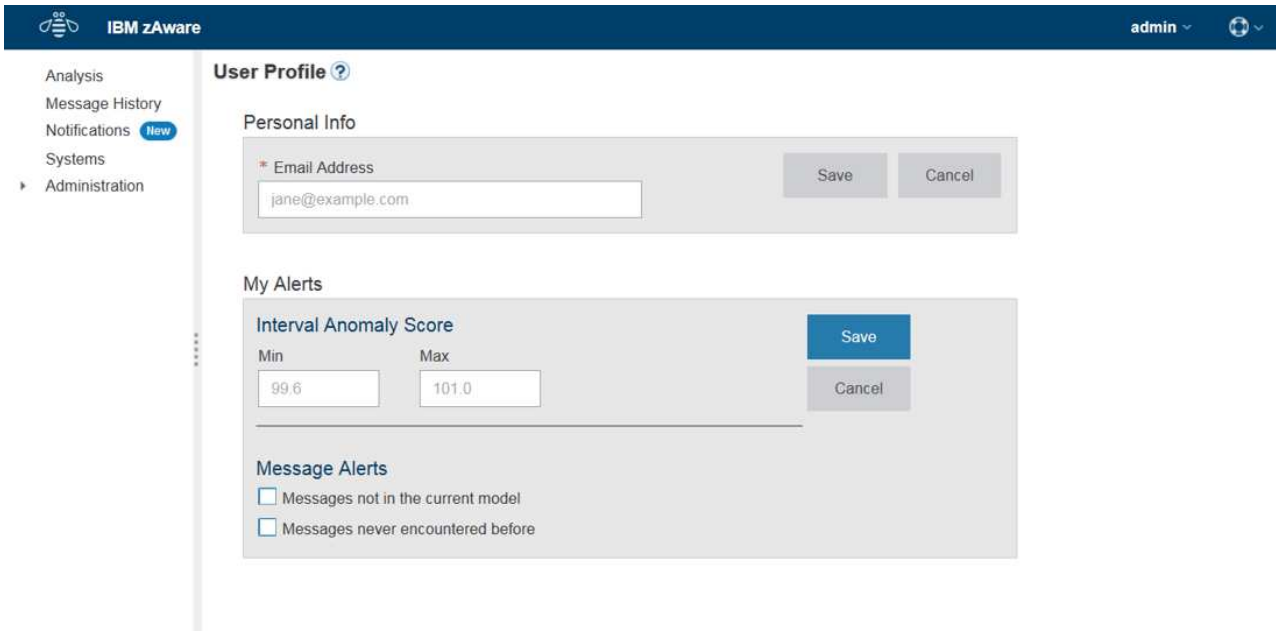

Figure 33. Example of a generic User Profile window

Related information:

"Understanding how IBM zAware calculates and displays anomaly scores " on page 4

Defining an SMTP email server

An administrator can use the following procedure to define the Simple Mail Transfer Protocol (SMTP) email server to send and receive emails from IBM z Advanced Workload Analysis Reporter (IBM zAware).

About this task

Before you can start receiving email alerts (through the **User's Name** > **User Profile** menu), you must define an SMTP server.

Procedure

1. Click **Configuration** > **Alerts**. You are on the **Email Server** page.
   a. Enter the fully qualified IP address for the SMTP in the **Server Address** field.
   b. Enter the TCP/IP port that your installation uses for SMTP connections in the **Port** field.
2. Click **Save** to save your entry or **Cancel** to exit.

Results

Table 18. SMTP configuration

Required Fields	Description
Server Address	Enter the fully qualified SMTP server address.
Port	Enter the TCP/IP port number that your installation uses for SMTP communications. Check with your network administrator to ensure that the correct port is defined to receive the alerts.

Related tasks:

"Setting up the User Profile to send email alerts" on page 83
Set up the **User Profile** to send email alerts that notify you when an anomaly occurs.

Chapter 11. Planning to create IBM zAware models

To provide analytical results for any monitored client, the IBM zAware server requires a model of normal system behavior to use for comparison to data that each monitored client is currently sending. The type of monitored client determines the data that IBM zAware uses to build a model, and the options that administrators have for managing the build process.

- IBM zAware builds and uses one model for each z/OS system. z/OS system models can be built using both OPERLOG and system log (SYSLOG) data, which is available only when an administrator uses priming data to build a z/OS model. For more information about z/OS models, see "Planning to create IBM zAware models for z/OS monitored clients."
- IBM zAware builds one model for a group of Linux systems with similar workloads, and uses that model to compare to current syslog data from each system in the group. IBM zAware administrators who manage Linux support determine which Linux systems belong to a particular group, and define the group and its members through the IBM zAware GUI. For more information about models for Linux groups, see "Planning to create IBM zAware models for Linux on z Systems monitored clients" on page 90.

Planning to create IBM zAware models for z/OS monitored clients

To provide analytical data for a z/OS monitored client, the IBM zAware server requires a model of normal system behavior to use for comparison. IBM zAware builds a model for each z/OS monitored client with message data from that client. You have two options for building a model: waiting for the server to build a model from data collected over a specific time period, or priming the server with prior data. This priming option is recommended because analysis can start shortly after the model is built. Regardless of the option that you choose, you must provide sufficient data for the IBM zAware server to successfully build a z/OS model.

Message data requirements for building a z/OS system model

IBM zAware requires sufficient message data to determine what constitutes normal behavior for a given monitored client. When processing message data from a system, the IBM zAware server determines which messages are issued during routine system events, such as starting a batch job or a particular subsystem. For such system events, the server identifies and recognizes the pattern of messages that are associated with each event. The message patterns are called *clusters* and define the normal context for the messages in the cluster. If the IBM zAware server does not receive a sufficient amount of data for a given system, it is not able to detect and recognize these message clusters.

For a z/OS production system, which is typically a very stable system that produces high-volume, consistent message traffic, IBM zAware can build a model that uses the system's standard behavior as the desired behavior. For less stable systems, such as test systems or development sandbox environments, modelling standard behavior can be more difficult. By default, 90 consecutive calendar days of message data is required for IBM zAware to learn and model the behavior of a given system. This default value, which is known as the *training period* for IBM zAware analytics, covers multiple unusual but predictable events, such as end-of-the-month processing.

Your installation can modify the training period, based on your knowledge of the workloads running on z/OS monitored clients. If you change this value, make sure that the new training period is long enough to include several occurrences of normal events for all systems that you plan to connect to the IBM zAware server. This training period applies for all monitored clients; you cannot define a different training period for each client.

Within the training period, the message traffic for the system must contain a minimum of 250 unique message IDs. Each of those unique IDs must be issued at least once in three different 10-minute intervals during the training period. IBM zAware recognizes messages IDs that conform to the z/OS standard, which consists of a component identifier, a message number, and an action code, in that order. IBM zAware also is capable of recognizing message IDs that do not completely conform to this z/OS standard. Because of this capability, the total count of unique message IDs for a monitored client can consist of messages issued by IBM products, non-IBM products, and possibly your own application programs.

Using the IBM message analysis program is perhaps the easiest way to analyze the message traffic for a given z/OS monitored client. Through this program, you can analyze z/OS SYSLOG data sets to determine the message rate per second, as well as the number and frequency of unique message IDs. The message analysis program is available on the z/OS Tools and Toys web site at the following URL:

`http://www-03.ibm.com/systems/z/os/zos/features/unix/bpxa1ty2.html`

To find the message analysis program, search the table of download packages for the MSGLG610 package. The latest version of MSGLG610 produces a report listing the message IDs that adhere to the IBM zAware training criteria.

The IBM Redbooks® publication *Extending z/OS System Management Functions with IBM zAware*, SG24-8070, describes how to use the prior version of the message analysis program. This Redbooks publication is available at the following URL:

`http://www.redbooks.ibm.com/`

After you determine the training period that you need for your z/OS systems, you can choose which option to use for training IBM zAware:

- "Option 1: Waiting for the server to build a z/OS model," which uses current OPERLOG data for the model.
- "Option 2: Transferring z/OS priming data to build a model" on page 89, which uses prior SYSLOG data for the model.

If you recently converted your z/OS system to use OPERLOG, which is a requirement for IBM zAware monitored clients, and that z/OS system was formerly configured or still uses the JES3 DLOG for its hardcopy log, you must use option 1 because option 2 requires SYSLOG data in a specific format. You can use the SYSLOG data on JES2 systems to prime the server but you cannot use SYSLOG data in the JES3 DLOG format for priming.

Option 1: Waiting for the server to build a z/OS model

When you wait for the server to build a model, the training period determines when analytical data will be available for a monitored client. The *training period* is the number of consecutive calendar days that the IBM zAware server uses to identify the instrumentation data to include in training models. By default, the training period is 90 days.

When you wait for the server to build a model, you have to connect the monitored client and make sure that it sends message data to the IBM zAware server for the duration of the training period. The monitored client must meet operating system and configuration requirements to send its OPERLOG data to the server.

When the training period ends, IBM zAware server automatically attempts to create the initial model for the monitored client. If the automatic build of the initial model is successful, IBM zAware can begin to analyze current data that it receives from the client. From this point on, IBM zAware uses the training interval value to determine when to automatically rebuild the system model. The *training interval* is the number of consecutive calendar days between automatic builds of system behavior models. By default, the training interval for z/OS monitored systems is 30 days.

Until a model is successfully built, the server is not able to provide analytical data. In this case, when you view the **Analysis** page in the IBM zAware graphical user interface (GUI), the monitored client is listed in the page display but no anomaly scores are provided for this client. To verify that the client is connected and sending data, use the **System Status** page in the GUI.

Option 2: Transferring z/OS priming data to build a model

Instead of waiting for the IBM zAware server to collect data over the course of the training period, you can prime the server by transferring prior data from the hardcopy or system logs of monitored clients, and requesting the server to build a model for each client from the transferred data.

The priming process consists of several phases:

1. The transfer of prior data to the IBM zAware server. To transfer this priming data, you configure a log stream and run the z/OS bulk load client for IBM zAware through a REXX exec on a z/OS system.

   The process of transferring priming data could require several hours or more, depending on a number of factors that include:
   - The priority of the job that runs the REXX exec for the z/OS bulk load client
   - The amount of priming data to be sent, and whether any of that priming data resides on migrated data sets
   - The network configuration and traffic at your installation

   For example, if you run the REXX exec at a very high priority to send 46000 tracks of priming data that is archived, the transfer might take approximately 10 to 15 minutes. The process can take longer if the z/OS bulk load client runs at a lower priority, or if network or system conditions are not favorable when the REXX exec runs.
2. The assignment of priming data to the correct sysplex.

   In contrast to data that the IBM zAware server receives from the z/OS system logger running on a monitored client, the priming data from the z/OS bulk load client does not include the name of the sysplex to which the monitored client belongs. Without the sysplex name, the IBM zAware server cannot associate the priming data with the appropriate sysplex. You use the **Administration** > **Configuration** > **Priming Data** page in the IBM zAware GUI to assign the received priming data to the appropriate sysplex. After the priming data is associated with the appropriate sysplex, you can request the IBM zAware server to build the model.
3. The training of the IBM zAware server, which results in a model of normal system behavior.

   To build the model for a specific monitored client, you have two options:
   - You can use the **Request Training** action on the **Administration** > **Training Sets** page. Any data that the z/OS system logger is currently sending does not become part of the model for the client until you request training again or the IBM zAware server automatically rebuilds the model. This priming option is recommended because analysis can start shortly after the model is built.
   - You can wait for the next scheduled training, during which the IBM zAware server automatically uses the priming data to build the model. In this case, any data that the z/OS system logger is currently sending becomes part of the model for the client.

Note that analytical data is not available for the dates for which you supplied priming data, unless the monitored client was connected and sending data to the IBM zAware server on those dates. The server does not analyze priming data; it uses priming data only for creating the model of system behavior.

Through the z/OS bulk load client, you can transfer data for one or more monitored clients by identifying the sequential data sets that contain the priming data. The sequential data sets can contain only SYSLOG data that is stored in hardcopy log 2-digit year (HCL) or 4-digit year (HCR) format. If any data set has been archived, the z/OS bulk load client can recall the data set, transfer its contents, and migrate the data set.

The recommended approach to priming is to complete the following tasks in sequence.

1. Configure a Quality Assurance (QA) system as a IBM zAware monitored client. Use the instructions in "Configuring z/OS monitored clients to send data to the IBM zAware server" on page 111.
2. From the QA system, run the z/OS bulk load client to send priming data and build models for all the systems that you want to configure as monitored clients.
   a. Transfer only a portion of the priming data for a monoplex or sysplex to validate the configuration of the IBM zAware environment.
   b. When you verify that the transfer was successful, transfer the remaining data for one monoplex or sysplex at a time. If you transfer priming data for multiple sysplexes through one invocation of the REXX exec, priming data for some systems can be overlaid.

   To send priming data and build models, use the instructions in "Creating an IBM zAware model for new z/OS monitored clients" on page 118.
3. Configure the systems for which you transferred priming data as IBM zAware clients.

Although the primary use of the z/OS bulk load client is to quickly build the initial model for a monitored client, you can use it after initial setup as well. For example, if a monitored client is disconnected from the IBM zAware server for an extended period of time, and you believe that the message traffic generated on that system during that time is important to include in the system model, you can use the z/OS bulk load client to transfer that generated message data. Remember, however, the server does not analyze priming data, so analytical data is not available for the time period during which the monitored client was disconnected. Also, the process of assigning priming data results in automatic recycling of the analytics engine and the disconnection of all monitored clients, so you need to determine whether the missing message data is worth this disruption to your IBM zAware environment.

Planning to create IBM zAware models for Linux on z Systems monitored clients

Because the message traffic on Linux systems often can be relatively light, and because Linux images are typically configured in pools of dynamically activated images, IBM zAware is designed to provide analysis results for Linux systems through the use of model groups. Through model groups, multiple systems contribute to the generation of a single model for the group; the more systems in the group, the more data IBM zAware can use to build the model.

Defining model groups and their member systems

A model group is a collection of one or more systems that handle the same type of workload, and thus can be expected to exhibit similar behavior. IBM zAware administrators who manage Linux support determine which Linux systems belong to a particular group, and define the group and its members through the IBM zAware GUI. A model group definition consists of a name for the group, an optional description, a membership rule that is based on Linux system naming conventions, and a membership evaluation order.

IBM zAware provides one predefined model group, named UNASSIGNED, for Linux systems with names that do not match any administrator-defined membership rules. IBM zAware does not provide analysis results for Linux monitored systems that belong to the UNASSIGNED model group, nor does it build a model for that group.

When considering Linux systems to group together in a single model group, use the following guidelines:

- Group together Linux systems that support very similar workloads. For example, group a set of Linux web servers in one model group, and a set of Linux database servers in another model group.
- Use a consistent naming convention for Linux systems to belong to the same model, and follow a system naming convention that reflects the function of or workload supported by different Linux systems. The name of a Linux system cannot exceed 230 characters.

Building a model for a model group

IBM zAware builds one model for a group of Linux systems with similar workloads, and uses that model to compare to current syslog data from each system in the group. To build a robust model of Linux system behavior, IBM zAware generally needs a minimum of 120 days of message data. Analysis can begin, however, as soon as the system data that is available for training meets the criteria for building a valid model.

To build the initial model for a newly defined Linux model group as quickly as possible, IBM zAware uses an early training schedule that it calculates when at least one of the Linux systems in the model group is first connected to the IBM zAware server. IBM zAware automatically schedules early training every seven days, starting from the first day for which IBM zAware has data available. For example, assume that an IBM zAware administrator defines a new model group and connects the member systems on day 1 of the current month. Assuming that IBM zAware starts collecting data from the member systems on that day, Figure 34 illustrates the early training schedule for the new model group.

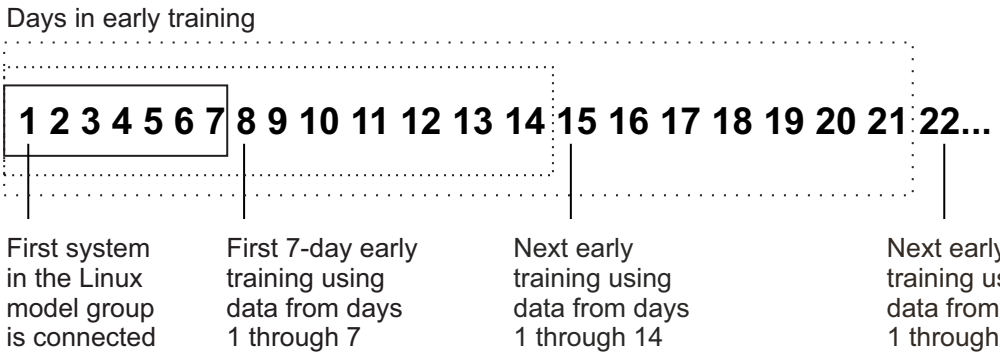

Figure 34. Early training schedule for Linux

Depending on the quality of the available system data that IBM zAware uses for the first early training, the initial model might or might not be successfully built.

- If a model is successfully built, IBM zAware begins analyzing current data from the monitored systems in the group, and calculates the date for the next automatic training. The seven-day early training schedule continues until at least one of the group members is connected to IBM zAware for the configured training period; at that point, IBM zAware uses the configured training interval to schedule automatic training. By default, the configured training interval for Linux systems is 30 days.
- If an automatic training attempt fails and a model is not available, IBM zAware automatically retries the training attempt the next day and, if necessary, every following day until a model is successfully built. Analysis cannot begin until a model is successfully built.

The first model that is successfully built might be a limited model. A limited model is the result when the system data provided for training enables IBM zAware to successfully build a model but, because the system data lacks sufficient variety, the resulting model does not provide enough information for IBM zAware to distinguish between unusual and normal behavior. In this case, IBM zAware displays the Limited Model icon ( ) in Analysis views and on the Interval page to indicate analysis results that were produced using a limited model. To avoid producing an initial model that is limited, follow these guidelines:

- Define only a small number of model groups for the Linux systems to be monitored.
- Avoid modifying or deleting that small number of model group definitions until the first configured training period has elapsed. By default, the configured training interval for Linux systems is 120 days.
- If they are available, use archived Linux system logs as priming data for the model. In Figure 34, the early training schedule does not start until day 8 because priming data is not available. If an administrator provides priming data for a model group, IBM zAware can start the early training before

day 8 because more days of data are available for training. These extra days also can contribute to the quality of system data that IBM zAware uses for the first early training.

IBM zAware administrators can manually request training at any time. If an administrator manually requests training and a model is successfully built, IBM zAware recalculates the scheduled date for the next automatic build, using the date on which the model was created and either the early seven-day schedule or the configured training interval, whichever is in effect.

For related information, see the following topics:

- For step-by-step instructions for building the model for a Linux model group, see "Creating an IBM zAware model for new Linux on z Systems monitored clients" on page 132.
- For more information about limited models, see "Understanding how IBM zAware calculates and displays anomaly scores " on page 142.
- For a detailed example of early training for Linux, see "Linux example: Allowing IBM zAware to collect the data for the initial group model" on page 223.

Part 4. Configuring IBM zAware and its monitored clients

Topics in this part provide step-by-step instructions for configuring the IBM zAware environment, which includes the IBM zAware partition, the IBM zAware server, and its monitored clients. Systems programmers and administrators use these configuration tasks primarily for first-time setup.

Topics covered in this part are:

Chapter 12. Configuring network connections and storage for the IBM zAware partition

Use this procedure to learn how to configure network connections and persistent storage for the IBM zAware partition. This procedure is intended for experienced system programmers, network administrators, or storage administrators who are responsible for configuring z Systems servers and their peripheral hardware devices.

Before you begin

For optimal performance and operations, configure the IBM zAware partition such that it has access to only those channel path identifiers (CHPIDs), control units, and I/O devices that are required for network connectivity and storage.

- Your installation must have correctly configured a supported host system, or central processor complex (CPC), on which you can configure the IBM zAware partition. For a list of supported host systems, engineering change (EC) numbers, and IBM zAware feature codes, see Chapter 2, "Prerequisites for configuring and using IBM zAware," on page 13.

  Planning information and instructions for configuring and activating a z13s, z13, or z14 are available in the following:

  – *IBM z14 Technical Guide*, SG24-8451
  – *IBM z13s Technical Guide*, SG24-8294
  – *IBM z13 Technical Guide*, SG24-8251
  – The HMC/SE topics in IBM Knowledge Center, at http://www.ibm.com/support/knowledgecenter/
- To complete most of the configuration tasks in this procedure, you use either the Hardware Configuration Definition (HCD) or the Input/Output Configuration Program (IOCP). Depending on the tool you are using, you might need to see one of the following books:

  – *z/OS Hardware Configuration Definition User's Guide*, SC33-7988, and *z/OS Hardware Configuration Definition Scenarios*, SC33-7987
  – *System z Input/Output Configuration Program User's Guide for ICP IOCP*, SB10-7037
- Make sure that you have reviewed the network planning considerations in "Planning network connections and capacity" on page 52, and have acquired the appropriate resources. Peripheral devices for network connections must be installed and attached to the host system before you begin this procedure.
- Make sure that you have reviewed the storage planning considerations in "Planning persistent storage configuration and capacity" on page 59, and have acquired the appropriate resources. Peripheral devices for storage must be installed and attached to the host system before you begin this procedure.

Attention: To avoid the potential loss of critical system and application data on storage devices that are connected to the CPC, the storage administrator must use either channel path definition lists or the explicit candidate device list in HCD to ensure that only the IBM zAware partition can access the devices that are intended for IBM zAware use. The only exception to this practice is the CPC partition, if any, that your installation is using to back up the storage that the IBM zAware server is using.

About this task

Through either HCD or the IOCP, you can define network connections and storage devices for the IBM zAware partition in the input/output configuration data set (IOCDS) for the CPC. HCD supplies an interactive dialog to generate the I/O definition file (IODF) and subsequently the IOCDS, whereas IOCP

generates the IOCDS without the use of an IODF. Using HCD is preferred for generating the I/O configuration because HCD performs validation checking as you enter data, which helps minimize the risk of errors. The information in this procedure is tailored for HCD users.

The steps in this procedure provide the information that you need to update the IOCDS for the IBM zAware partition. Additional network and storage configuration tasks are required to complete the initial setup for the IBM zAware environment; these tasks are documented in the following topics:

- "Configuring the IBM z Systems Secure Service Container for IBM zAware" on page 27
- Chapter 13, "Configuring storage, security, and analytics for the IBM zAware server," on page 99
- "Configuring z/OS monitored clients to send data to the IBM zAware server" on page 111

Procedure

1. Through HCD, add the IBM zAware partition to the I/O configuration for the host system (CPC).
   a. Add the IBM zAware partition to the list of partitions for the host system.
      - The name that you provide for the partition must exactly match the name that you use for the LPAR image profile that you create, using the configuration procedure in "Configuring the IBM z Systems Secure Service Container for IBM zAware" on page 27.
      - The partition usage field marks a partition to be used for coupling facility support or for operating system usage. For partition usage, enter `OS` for operating system.
   b. Define or update channel path definitions for the IBM zAware host system.

      The recommended practice is to configure the IBM zAware partition such that it has access to only those channel path identifiers (IDs), control units, and I/O devices that are required for network connectivity and storage. When network connections are defined in the IODF or IOCDS for the IBM zAware host system, the network administrator can use the following HCD constructs to limit IBM zAware access to specific resources.
      - Image access and candidate lists in channel path definitions
      - The explicit device candidate list for I/O devices

      1) Assign channel paths to the IBM zAware partition for network connections. Table 19 lists the types of channel paths that you can assign.

Table 19. Supported channel path types for the IBM zAware partition

Channel path type	Description	IBM zAware usage	Additional information
OSD	Ethernet connectivity through an Open Systems Adapter (OSA) channel	• Use for monitored clients that reside on any supported z Systems CPC in the IBM zAware environment. • Use for GUI browser connections to the IBM zAware server.	*zEnterprise System, System z10, System z9 and eServer zSeries Open Systems Adapter-Express Customer's Guide and Reference*, SA22-7935 **Note:** For OSA-Express4S or later generation features, IBM zAware can use only port 0.
OSX	Connectivity through the intraensemble data network (IEDN)	Use only for monitored clients that reside on the IBM zAware host system or on other nodes in the same zEnterprise ensemble.	*zEnterprise System, System z10, System z9 and eServer zSeries Open Systems Adapter-Express Customer's Guide and Reference*, SA22-7935 **Note:** For OSA-Express4S or later generation features, IBM zAware can use only port 0.

Table 19. Supported channel path types for the IBM zAware partition (continued)

Channel path type	Description	IBM zAware usage	Additional information
IQD	Connectivity through HiperSockets	Use only for monitored clients that reside on the IBM zAware host system. Use the IQD channel path type for HiperSockets. Only the following functions are supported IQD channel parameters: • Basic HiperSockets • IEDN Access, only when the CPC is a member of an ensemble The External Bridge function is not a supported IQD channel parameter.	*z/OS Communications Server: IP Configuration Guide*, SC31-8775

2) If necessary, change the channel path mode. Channel paths can be dedicated, reconfigurable, shared, or spanned.

   DED Dedicated; if you want only one logical partition to access a channel path, specify that channel path as dedicated. You cannot reconfigure a dedicated channel path. This is the default mode.

   REC Reconfigurable; if you want only one logical partition at a time to access a channel path and you want to be able to reconfigure the channel path from one partition to another, specify that channel path as reconfigurable.

   SHR Shared; if you want more than one logical partition to access a channel path simultaneously, specify that channel path as shared.

   SPAN
   Spanned; if in XMP processors for certain channel types, you want to have a shared channel accessed by partitions from multiple logical channel subsystems, specify that channel path as spanned.

c. Add the IBM zAware partition to the access list for each channel path that you have selected.

d. Verify your configuration changes. One way to verify your changes is to build an IOCP input data set and view it to check the partition list for the channel path that you defined.

   As an alternative, you can view your changes graphically if your system meets the appropriate HCD prerequisites. On the HCD Channel Path List panel:
   1) Select the channel paths and press **Enter**.
   2) Select **View graphically** and press **Enter**.

2. Configure persistent storage for the IBM zAware partition.

   • When defining storage devices in the IODF or IOCDS for the IBM zAware host system, the storage administrator can use image access and candidate lists for channel path definitions to allow only the IBM zAware partition to access specific devices. Using the explicit device candidate list is an alternative method of restricting access to specific devices.

   • As a precaution, you can use channel path definition lists, the explicit device candidate list, or operating system mechanisms to prevent the applications that are running in other CPC partitions from using storage devices that are intended for IBM zAware use.

     The only exception to this practice is the CPC partition, if any, that your installation is using to back up the storage that the IBM zAware server is using.

   The IBM zAware partition uses direct access storage devices (DASD) that are attached to the host system. For DASD storage only, complete the following general steps.

   a. On the storage controller, define volumes.

   b. Through HCD, complete the following steps:

1) Define the channel path for the storage device. Add the IBM zAware partition to the access list for the channel paths for the storage devices.
2) Define the control unit (controller).
3) Define the storage device, using the explicit candidate list as appropriate for your planned IBM zAware environment.

3. Verify your configuration changes. One way to verify your changes is to build an IOCP input data set and view it to check the partition list for the channel path that you defined.

   As an alternative, you can view your changes graphically if your system meets the appropriate HCD prerequisites. On the HCD Channel Path List panel:
   a. Select the channel paths and press **Enter**.
   b. Select **View graphically** and press **Enter**.

Results

Network connections and persistent storage are configured for the IBM zAware partition.

What to do next

Activate your configuration according to your company operating procedures. Depending on your environment, you might be able to dynamically update the IODF; another option is to complete the following steps.

1. Through HCD:
   a. Build a production IODF.
   b. Build an IOCDS.
2. Through the HMC:
   a. Perform power on reset (POR) for the IBM zAware host system, with the new IOCDS.
   b. Define the new partition as instructed in "Configuring the IBM z Systems Secure Service Container for IBM zAware" on page 27.

Chapter 13. Configuring storage, security, and analytics for the IBM zAware server

Use this procedure to configure persistent storage, set up security, and configure settings for the analytics engine in the IBM zAware server. Depending on the roles and responsibilities for your IT organization, systems programmers, storage administrators, and security administrators might be required to collaborate to complete all of the steps in this procedure.

Before you begin

- The IBM zAware partition must be defined according to the instructions in "Configuring the IBM z Systems Secure Service Container for IBM zAware" on page 27. The IBM zAware partition also must be activated.
- You need to know the default master user ID and password and the URL to access and log in to the IBM zAware graphical user interface (GUI). This information is derived from theIBM z Systems Secure Service Container page of the image profile that you created in "Configuring the IBM z Systems Secure Service Container for IBM zAware" on page 27.

  The URL includes the IP address or host name that is assigned to the IBM zAware partition:

  `https://`*`ip_address`*`/zAware/` **or** `https://`*`host_name`*`/zAware/`

  The "zAware" portion of the URL is case-sensitive.
- Make sure that the browser you plan to use meets the requirements and recommendations that are listed in Chapter 10, "Planning to use the IBM zAware GUI," on page 81.
- Make sure that you review the planning considerations in "Planning persistent storage configuration and capacity" on page 59 and have a list of the storage devices that are intended for IBM zAware use at your installation.

Attention: The recommended practice is to configure the IBM zAware partition such that it has access to only those channel path identifiers (IDs), control units, and I/O devices that are required for network connectivity and storage. If this practice is not followed, an IBM zAware administrator might inadvertently assign storage devices that are in use by other partitions. To avoid the potential loss of critical system and application data on storage devices that are connected to the IBM zAware host system, make sure that you use the GUI to assign only those storage devices that are intended for IBM zAware use. Check with your storage administrator to make sure that you select the appropriate storage devices to assign to the IBM zAware server.

- Make sure that you review the planning considerations in Chapter 9, "Planning for security," on page 75.

  If you plan to replace the default self-signed SSL certificate with a certificate that is signed by a certificate authority (CA) of your choice, you might need to process the CA reply before you can paste it into the appropriate field in the IBM zAware GUI, as instructed in step 3 on page 102.
  - The required format for replacement certificates is Base64 encoded X509 certificate blocks.
  - When you receive a reply from a certificate authority, the reply might contain a chain of certificates, starting with the signed server certificate. Then, it is possibly followed by certificates from one or more intermediate CAs and finally, the self-signed certificate of the CA. When you supply this information in the GUI, provide the entire certificate chain, starting with the signed server certificate through the self-signed certificate of the CA.
  - In some cases, the CA reply that you receive is delivered in a public-key cryptography standards (PKCS) #7 file. You must extract the certificates from the file before pasting them into the GUI. One method of extracting certificates from a PKCS #7 file is to use the **openssl pkcs7** command; for more information, see the OpenSSL Project website at the following URL.

    www.openssl.org/

- If you plan to use an existing Lightweight Directory Access Protocol (LDAP) server to authorize access to the IBM zAware GUI, the network administrator must configure the network connections to ensure that the IBM zAware server can access the LDAP server. IBM zAware cannot save LDAP settings unless it can communicate with the LDAP server when you apply the new or changed settings.

  If a firewall exists between the IBM zAware partition and the LDAP server, the IBM zAware partition must be able to use the port that is used by the LDAP server.

About this task

Before the IBM zAware server can receive and analyze data from monitored clients, you need to configure persistent storage, set up security, and configure settings for the analytics engine. You can accomplish most of these tasks through the IBM zAware GUI, by using the default master user ID.

When you log in to the IBM zAware GUI for the first time, the browser displays a warning message because the default Secure Sockets Layer (SSL) certificate that is used by the IBM zAware server is not signed by a trusted certificate authority (CA). As part of the security configuration steps in this procedure, you can resolve this error by replacing the default SSL certificate with a certificate that is signed by a certificate authority of your choice. Doing so provides secure communication between the IBM zAware server and the browsers of all authorized users.

If you decide to replace the automatically generated certificate, which is the recommended practice for improved security, you can use any third-party certificate authority of your choice, or your installation can provide an internal certificate authority for certificate signing tasks. IBM zAware does not renew these replacement certificates; in this case, managing replacement certificates becomes the responsibility of the security administrator.

Also, as part of the security configuration steps in this procedure, you can authorize users to access the IBM zAware GUI by using an LDAP repository. You can configure user authentication through an LDAP repository or alternatively by using a local file-based repository. For simplicity, using only an LDAP repository is the preferred option. However, you might want to define one or two user IDs in a local repository so you can access the IBM zAware GUI when the LDAP server is unavailable. If you configure an LDAP repository and also define users or groups in a local repository, both sets of users or groups are available through the IBM zAware GUI.

Procedure

1. From a browser, enter the URL for the IBM zAware GUI to display the IBM zAware welcome page.

   The URL includes the IP address or host name that is assigned to the IBM zAware partition:

   `https://`*`ip_address`*`/zAware/` **or** `https://`*`host_name`*`/zAware/`

   The "zAware" portion of the URL is case-sensitive.

   a. If the browser presents a warning message because it does not recognize the default SSL certificate for the IBM zAware server, bypass the warning message by adding a security exception. Step 3 on page 102 explains how to replace the default SSL certificate with another certificate to permanently prevent this error.
   b. Click **Log in** to open the IBM zAware User Login window.
   c. Enter the default master user ID and password and click **Log in**.
   d. If the **Administration** category in the navigation pane is not expanded, click the link to display the administration tasks. Click **Configuration** to display the **Configuration** page.
2. Assign storage devices for the IBM zAware server to use for storing analysis results, system behavior models, and data from monitored clients.

Attention: The recommended practice is to configure the IBM zAware partition such that it has access to only those channel path identifiers (IDs), control units, and I/O devices that are required for network connectivity and storage. If this practice is not followed, an IBM zAware administrator might inadvertently assign storage devices that are in use by other partitions. To avoid the potential loss of critical system and application data on storage devices that are connected to the IBM zAware host system, make sure that you use the GUI to assign only those storage devices that are intended for IBM zAware use. Check with your storage administrator to make sure that you select the appropriate storage devices to assign to the IBM zAware server.

a. Click the **Data Storage** tab on the **Configuration** page as shown in Figure 35 on page 102. The GUI populates the Data Storage Devices table with a list of the devices that are available and connected to the IBM zAware partition. You can sort the list by clicking any one of the column headings in the Data Storage Devices table. Until you assign these devices to the IBM zAware server, their status is "Available".
b. On the **Data Storage** tab, click **Add and Remove Devices**. The GUI opens the Add and Remove Devices window.
c. To add a storage device for the IBM zAware server to use, select one or more devices in the **Devices Available** list and click either **Add >** or **Add All >>** to move the devices to the **Devices in Use** list. You do not have to assign all devices in the list unless the server requires the total capacity.

Attention: Do not use **Add All** if any of the available storage devices are shared. If a device is shared and in use by another application, data will be lost or overwritten if the IBM zAware server formats the device.

   Although the GUI provides a **Preserve data** option for adding storage devices, which prevents IBM zAware from overwriting data on the device to be added, use this option only when assigning a storage device that contains a backup copy of IBM zAware data.
d. When you complete moving devices to the **Devices in Use** list, click **OK** to assign those devices.

The IBM zAware server formats the devices that you moved to the **Devices in Use** list. While the formatting process is in progress, the device status is "Being Added"; when the formatting process is complete, the device status is "In Use". As part of the formatting process, the volume serial (VOLSER) for the device is renamed.

Depending on the number of devices that you assign, this formatting process might take some time. For example, IBM test experiences indicate that the IBM zAware server requires approximately 10 minutes to format and initialize a 3390 model 9 device. Periodically click **Refresh** to update the information in the Data Storage Devices table. Next, sort by clicking the **Status** column heading twice to display the devices with status other than "Available" at the top of the list.

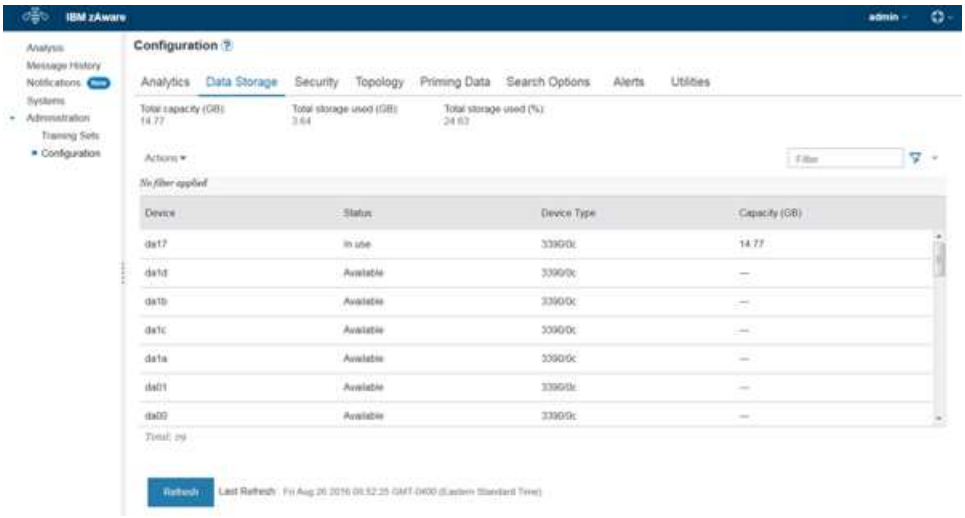

Figure 35. Data Storage Configuration

3. Optional: Secure the communication between the IBM zAware server and browsers by generating a request for an SSL certificate and importing the reply from the certificate authority.

   This process might take several days to complete, depending on the time that the certificate authority requires to receive and process your request, and to send the reply to you. You can complete other tasks in the IBM zAware GUI while you wait for a reply from the certificate authority.

   a. Click **Security** > **SSL Settings** tab on the **Configuration** page. The **SSL Settings** tab displays information about the default SSL certificate that is configured in the IBM zAware server.

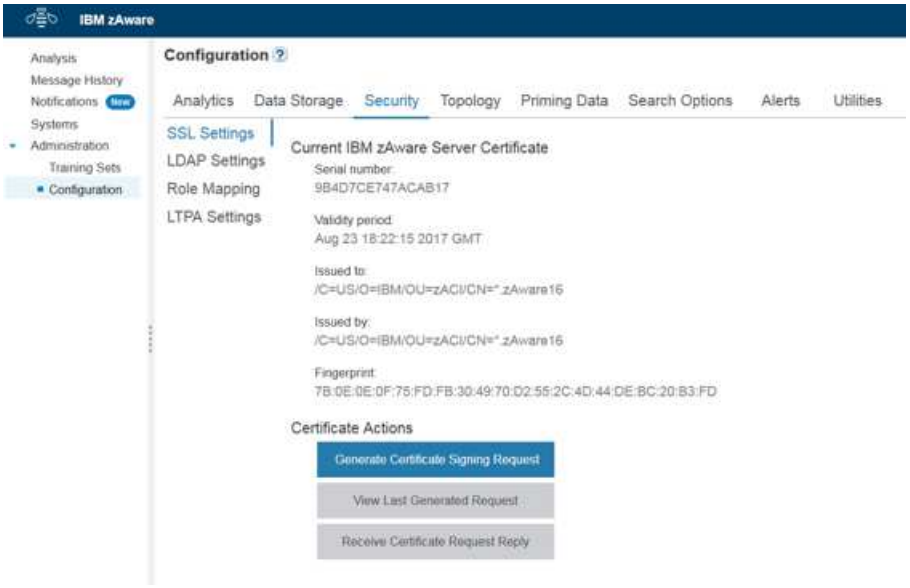

Figure 36. SSL Settings Configuration

   b. Under Certificate Actions, click **Generate Certificate Signing Request** to create a certificate signing request (CSR) to send to the certificate authority of your choice.

      Provide the appropriate information for the fields in Table 20 on page 103.

Table 20. Fields displayed on the page before the CSR is generated

Field	Description
Common name	Verify the host name or IP address of the IBM zAware partition. IBM zAware loads this field with a value that matches the host name or IP address that is specified in the image profile of the IBM zAware partition. The common name is required.
Organization	Enter the name of your company. The value that you supply for this field must be a string of length 1-64. The organization name is optional.
Organizational unit	Enter the company organization or department name. The value that you supply for this field must be a string of length 1-64. The organizational unit is optional.
Locality	Enter the city in which your company is located. The value that you supply for this field must be a string of length 1-128. The city is optional.
State or province	Enter the state or province in which your company is located. The value that you supply for this field must be a string of length 1-128. The state or province is optional.
Postal code	Enter the postal code for your company address. The value that you supply for this field must be a string of length 1-16. The postal code is optional.
Country code	Enter the two-character abbreviation for the country in which your company is located. The value that you supply for this field must be a string of length 1-2. The country code is optional.

c. Click **Generate** to generate the certificate request.

d. Click ► **Generated Request Input** to display and verify the generated request input.

e. From the Generated Request text area, copy the generated certificate signing request and submit it to the certificate authority. Follow the procedures that are specified by the certificate authority for submitting requests.

f. Click **Close** to return to the main **SSL Settings** tab.

g. When you receive a reply from the certificate authority, extract the certificates, if necessary, and return to the main **SSL Settings** tab.

   When you receive a reply from a certificate authority, the reply might contain a chain of certificates, starting with the signed server certificate. Then, it is possibly followed by certificates from one or more intermediate CAs and finally, the self-signed certificate of the CA. When you supply this information in the GUI, provide the entire certificate chain, starting with the signed server certificate through the self-signed certificate of the CA.

   See Appendix B, “Sample certificate authority (CA) reply,” on page 289 for a sample CA reply that contains a certificate chain, and for an illustration of the required format for pasting the reply into the GUI.

h. Click **Receive Certificate Request Reply**.

i. Paste the reply into the Certificate Authority reply text area and click **Receive** to import the CA reply into the IBM zAware server.

   Make sure that you do not insert any lines or spaces between the end of one certificate and the beginning of the next certificate. When you paste certificate replies in the GUI, make sure that you include all of the content, including the header `-----BEGIN CERTIFICATE-----` through and including `-----END CERTIFICATE-----`

The main **SSL Settings** tab now displays information from the received CA reply.

4. Configure the LDAP repository for storing user access information.

   As an alternative, you can use a local file-based repository instead of an existing LDAP repository. For instructions, see “Setting up a local repository to secure access to the IBM zAware GUI” on page 108. Do not define the same user ID in more than one repository; results are not predictable.

To configure user authentication by using an existing LDAP repository, click the **LDAP Settings** tab on the **Security** tab and supply appropriate values for the following fields. The LDAP administrator for your installation can either complete this step or provide the information that you need to do so.

When you or the LDAP administrator enter values for all required fields and any optional fields that you want to specify, click **Apply** to store these LDAP configuration values. If necessary, click **Restore** to restore the LDAP configuration values that were in effect before you clicked **Apply**.

When you click **Apply**, the GUI displays a confirmation message to indicate that the web server must be restarted for your changes to be applied. This process might take a considerable amount of time to complete. Click **OK** to confirm that you want to apply your changes, or click **Cancel**.

Table 21. General LDAP settings

Setting	Description
LDAP server hostname	Enter the resolvable host name or IP address of the LDAP server to which you want to connect. The host name is required.
LDAP server port	Enter the port on which the LDAP server listens for TCP/IP connections. The value can range from 0 - 65535. The port is required.
Follow referrals	A referral is an entity that is used to redirect a client request to another LDAP server. A referral contains the names and locations of other objects. It is sent by the server to indicate that the information the client requested can be found at another location, possibly at another server or several servers. Select one of the following options: **Follow** Indicates that referrals to other LDAP servers will be followed. **Ignore** Indicates that referrals to other LDAP servers will be ignored. This option is selected by default. A selection is required.
Bind distinguished name	Enter the distinguished name used to bind to the LDAP repository. The name must be a string of length 0-512. If no name is specified, the server binds anonymously. The name is optional.
Bind password	Enter the password used to bind to the LDAP directory. The password must be a string of length 0-512. The password is optional.
Base distinguished name	Enter the distinguished name of a base entry in the repository. The name must be a string of length 1-512. The name is required.
Login attribute	Enter the LDAP attribute of a user entity used to login to the IBM zAware GUI. The attribute must uniquely identify a user in the directory. Typical login attributes include *uid*, *mail*, *primaryuserid*, and so on. The value must be a string of length 1-512. The default value is *uid*. The login attribute is required.
User object classes	Enter the object class or classes that are associated with user entities in the LDAP repository. Delimit multiple object classes with semicolons (;). Typical object classes include *Person*, *ePerson*, *inetOrgPerson*, and so on. The value must be a string of length 1-512. At least one user object class is required.
User search bases	Specify the base object of the directory (or level of the directory) from which to start a search for user entities in the LDAP repository. The search bases must be subtrees of the base distinguished name. Delimit multiple search bases with semicolons (;). The value must be a string of length 1-512. The user search base is optional. If unspecified, the base distinguished name is used.
SSL enabled	Select this option to enable secure socket communication to the LDAP server. If selected, you must supply the SSL certificate in the **LDAP server certificate** field. By default, SSL is disabled.

Table 21. General LDAP settings (continued)

Setting	Description
LDAP server certificate	Enter the Base64 encoded certificate that is required to validate the certificate of the LDAP server. This certificate should be the signer of the server certificate for the LDAP repository. A certificate is required only when SSL is enabled.

Table 22. Group LDAP settings

Setting	Description
User group membership attribute	Enter the LDAP attribute of a user entity that indicates the groups to which an entry belongs. If your LDAP server does not support the group membership attribute, do not specify this attribute. The value must be a string of length 1-512. The user group membership attribute is optional.
User group membership scope	Select the scope of the user group membership attribute. You can select one of the following options: **Direct** Indicates that the attribute contains only immediate members of the group without members of subgroups. This option is selected by default. **Nested** Indicates that the attribute contains direct members and members nested within subgroups of this group. **All** Indicates that the attribute contains all direct, nested, and dynamic members. A selection is required if a value is specified in the **User group membership attribute** field.
Group object classes	Enter the object class or classes that are associated with group entities in the LDAP repository. Delimit multiple object classes with semicolons (;). Typical object classes include *groupOfNames*, *groupOfUniqueNames*, and so on. The value must be a string of length 1-512. At least one group object class is required.
Group search bases	Specify the base object of the directory (or level of the directory) from which to start a search for group entities in the LDAP repository. The search bases must be subtrees of the base distinguished name. Delimit multiple search bases with semicolons (;). The value must be a string of length 1-512. The group search base is optional. If unspecified, the base distinguished name is used.
Group member attributes	For each group object class specified in the **Group object classes** field, indicate the LDAP attribute of a group entity in the object class that contains the members of the group. Delimit multiple group member attributes with semicolons (;). A value is required, and must be a string of length 1-512. For example, if the group object classes specification is *groupOfNames;groupOfUniqueNames*, the group member attributes specification might be *member;uniqueMember*.
Group member object classes	For each attribute specified in the **Group member attributes** field, specify the object classes of the group that uses the member attribute. The value must be a string of length 1-512. The group member object classes are optional. If unspecified, the member attributes apply to all group object classes.

Table 22. Group LDAP settings (continued)

Setting	Description
Group member scope	Select the scope of the group member attribute. You can select one of the following options: **Direct** Indicates that the member attribute contains only direct members. This option is selected by default. **Nested** Indicates that the member attribute contains both direct and nested members. **All** Indicates that the member attribute contains direct, nested, and dynamic members. A selection is required.

5. Authorize users or groups to access the IBM zAware GUI. You can map authorized users and groups to specific roles: either **Administrator** or **User**. Users or groups with Administrator authority can use any task in the GUI. Users or groups with User authority can view only the following pages and use only the actions as noted:
   - On all views of the **Analysis** page, all controls and actions are permitted.
   - On the **Interval** page, all controls and actions are permitted except for modifying the non-IBM rules status for a specific message ID. Only administrators can view and change a rules status value. IBM rules cannot be changed.
   - On the **Notifications** page, all actions are disabled.
   - On the **Systems** > **System Status** tab, all actions are disabled.
   - On the **Systems** > **Model Groups** tab, all actions except for **Search Systems** are disabled.

   a. Map users or groups in the LDAP repository to specific roles by using the IBM zAware GUI. Click **Security** > **Role Mapping** and complete the following steps.

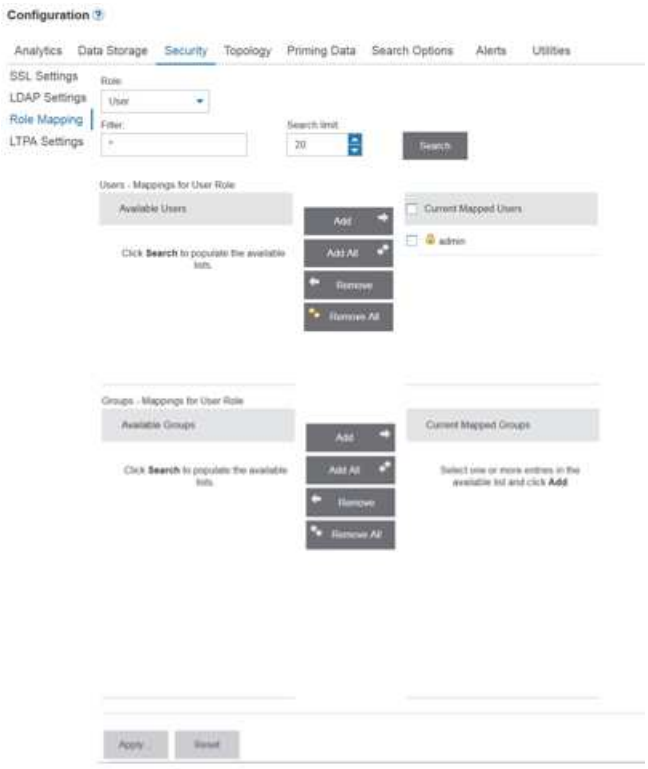

Figure 37. Role Mapping Configuration

1) Select either **Administrator** or **User** as the role to which you want to map particular users or groups. The IBM zAware server populates the **Current mapped users** and **Current mapped groups** lists with all users or groups that are currently mapped to the selected role. Initially, only the default master user ID appears in the **Current mapped users** list for both the Administrator role and the User role.
2) To add users or groups to the selected role, provide a filter value to populate the **Available users** and **Available groups** lists with matching user and group entries in the LDAP repository. You can specify an asterisk (*) as a wildcard value at any position in the filter value. An asterisk (*) is the default filter value.
3) Enter a search limit value to limit the number of matching entries that display in the **Available groups and users** list. This limit applies to both groups and users so a search limit of 20 might return 40 entries: 20 groups and 20 users. The default value for this limit is 20. You can replace the default value with any value from 1 through 200.
4) Click **Search** to apply the filter and search limit values. The IBM zAware server populates the **Available users** and **Available groups** lists with entries that match the filter value, up to the search limit value.
5) To add a user or group to the selected role, select one or more entries in the **Available users** and **Available groups** lists and click either **Add >** or **Add All >>** to copy the entries to the **Current mapped users** and **Current mapped groups** lists.
6) When you finish adding entries to the **Current mapped users** and **Current mapped groups** lists, click **Apply** to store your changes in the LDAP repository. You might need to scroll down the page to find **Apply**.
7) The IBM zAware GUI displays the Apply Role Mappings window, through which you can check the role assignments that you selected. If the changes are correct, click **Apply**; if you need to make further changes, click **Cancel** to return to the **Role Mapping** tab.

When you click **Apply**, the GUI displays a confirmation message to indicate that the web server must be restarted for your changes to be applied. This process might take a considerable amount of time to complete. Click **OK** to confirm that you want to apply your changes, or click **Cancel**.

6. Optional: Change the configuration value that determines the duration of a browser session.

   By default, browser sessions time out after 12 hours (720 minutes). To change this setting, complete the following steps.

   a. Click the **LTPA Settings** tab on the **Security** tab.
   b. In the **LTPA timeout** field, click the arrows to select a value in minutes. The allowable range of values is 10 - 525600 minutes (365 days).
   c. Click **Apply** to save the new value.

   When you click **Apply**, the GUI displays a confirmation message to indicate that the web server must be restarted for your changes to be applied. This process might take a considerable amount of time to complete. Click **OK** to confirm that you want to apply your changes, or click **Cancel**.

7. If necessary, update firewall settings to ensure secure communications between the IBM zAware server and its monitored clients.

   Although the configurations in Chapter 7, “Planning your IBM zAware environment,” on page 39 show both the IBM zAware server and its monitored clients within the boundary of a firewall, you can set up a configuration in which communication crosses firewall boundaries. In this case, you need to determine whether unsecured communication is an acceptable risk. If it is not an acceptable risk, you must provide your own method of securing this communication. For more information, see “Securing communication between IBM zAware and its monitored clients” on page 75.

8. Optional: Check the configuration values that control IBM zAware analytics operation and adjust them, if necessary. Click the **Analytics** tab on the **Configuration** page. Depending on the type of monitored systems that you plan to connect to the IBM zAware server, click one or more of the following tabs to view the analytics settings: **z/OS** or **Linux**.

All fields on each of the **Analytics** tabs contain default values that represent reasonable estimates for IBM zAware analytics. These estimates might not be appropriate for monitored clients at your installation, so you might need to change the default values according to your knowledge of client workloads.

The values on each tab are global settings that apply for all monitored clients of a specific type; some default values vary by type. You cannot specify different date ranges for individual monitored clients but you can manage the training dates used for each monitored client through the **Administration** > **Training Sets** page.

If you alter any of the default values on any of the **Analytics** page tabs, click **Apply** to save them. For setting descriptions and the default values for each type of monitored system, see "Specifying settings for the analytics engine" on page 199

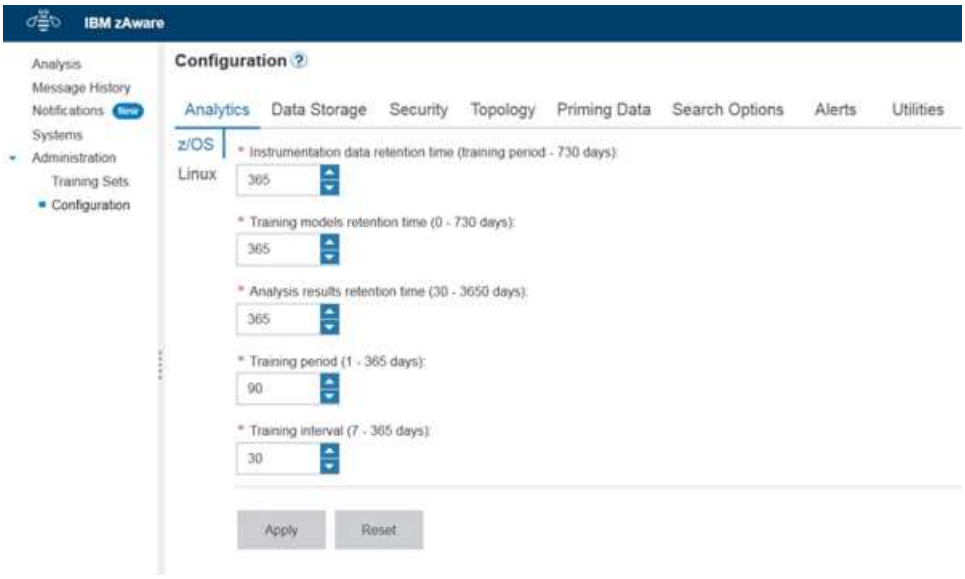

Figure 38. Analytics Configuration

After you click **Apply** to store new configuration values for analytics operation, the IBM zAware GUI displays a message that indicates whether it successfully stored your changes.

Results

The IBM zAware server is fully configured and ready to receive data from monitored clients.

What to do next

Depending on the type of monitored systems that you plan to connect to the IBM zAware server, use the instructions in one or more of the following topics.

- "Configuring z/OS monitored clients to send data to the IBM zAware server" on page 111
- "Configuring Linux on z Systems monitored clients to send data to the IBM zAware server" on page 129

Setting up a local repository to secure access to the IBM zAware GUI

Your installation has the option to provide user authentication to the IBM zAware graphical user interface (GUI) through either an existing Lightweight Directory Access Protocol (LDAP) repository or a local file-based repository. For simplicity, using only an LDAP repository is the preferred option. However, you might want to define one or two user IDs in a local repository so you can access the IBM zAware GUI when the LDAP server is unavailable. If you configure an LDAP repository and also define users or

groups in a local repository, both sets of users or groups are available through the IBM zAware GUI. Use this procedure to add users or groups to a local file-based repository.

Before you begin

- Defining users or groups to the local repository is done via the IBM zAware GUI. You can also delete users and groups, delete group members, and change a user's password via the GUI.
- You also need to log in to the GUI with a user ID that has the appropriate authority to add or define users or groups. This user ID can be the default master user ID and password that was defined in the image profile for the IBM zAware partition on the host system, or another user ID that is assigned to the IBM zAware Administrator role.
- Make sure that you have reviewed the planning considerations in Chapter 9, "Planning for security," on page 75.
- Prepare a list of user IDs or groups to define in the local repository. Do not define the same user ID in more than one repository; results are not predictable.

About this task

To define users or groups to the local repository, you need use the IBM zAware GUI.

To configure an existing LDAP repository for user authentication, see Chapter 13, "Configuring storage, security, and analytics for the IBM zAware server," on page 99 for instructions.

Procedure

1. Log in to the Integrated Solutions Console, providing a user ID and password for a user assigned to the IBM zAware Administrator role.
2. In the navigation tree, select **Users and Groups**.
3. Define one or more new users in the local repository.
   a. Click **Manage Users**.
   b. Click **Create** and supply the required information for the new user on the Create a User page.
   c. Optional: Click **Group Membership** to assign the user to a group.
      1) Optionally, supply a group name filter.
      2) Click **Search** to search for defined groups that match the filter.
      3) Select one or more groups to which you want to assign the new user, and click **Add**.
      4) Click **Close** to return to the Create a User page.
   d. Click **Create**.
   e. Repeat these steps, as necessary, to create additional users.
4. Optional: Define one or more new groups to the local repository:
   a. Click **Manage Groups**.
   b. Click **Create** and supply the required information for the new group.
   c. Click **Create**.
   d. Repeat these steps, as necessary, to create additional groups.

Results

The user IDs and passwords that you added are defined to the local repository. If you deleted any users or groups, those users or groups have been removed from the local repository.

What to do next

- Use the instructions in "Assigning users or groups to a role" on page 187 to assign each user ID or group in the local repository to a specific IBM zAware role.

- If you need to delete a user or group from the local repository, complete the following steps.
  1. In the navigation tree of the Integrated Solutions Console, select **Users and Groups**.
  2. Click **Manage Users** or **Manage Groups**, depending on whether you need to delete a user or group.
  3. Optionally, supply a filter value and click **Search** to search for defined users or groups that match the filter.
  4. Select the user or group that you want to delete and click **Delete**.
  5. Click **Delete** again to confirm the deletion.
  6. Repeat these steps, as necessary, to delete more users or groups.

Chapter 14. Configuring z/OS monitored clients for IBM zAware analysis

To configure z/OS monitored clients for IBM zAware analysis, you need to modify z/OS system constructs to send operations log (OPERLOG) data from that system to the IBM zAware server and, optionally, to prime the server with message data to build a model of system behavior.

For additional details, see the following topics:

1. "Configuring z/OS monitored clients to send data to the IBM zAware server"
2. "Creating an IBM zAware model for new z/OS monitored clients" on page 118

Configuring z/OS monitored clients to send data to the IBM zAware server

Use this procedure as an overview for configuring z/OS monitored clients to send data to the IBM zAware server for analysis. This procedure is intended primarily for skilled z/OS system programmers who have experience with configuring and managing systems in a Parallel Sysplex. Depending on the roles and responsibilities defined for your IT organization, you might need the assistance of network or security administrators to correctly configure secure connectivity within the IBM zAware environment.

Before you begin

- Make sure that your installation has completed the following steps for defining network connections for the hardware in the IBM zAware environment:
  - Step 1 on page 96 in Chapter 12, "Configuring network connections and storage for the IBM zAware partition," on page 95.
  - "Configuring the IBM z Systems Secure Service Container for IBM zAware" on page 27.

  Consider using the checklist in "Task summary and configuration checklist for network administrators" on page 57 as an aid for step 1 on page 112 in this procedure.
- Make sure that your installation has completed the procedure in Chapter 13, "Configuring storage, security, and analytics for the IBM zAware server," on page 99.
- List the z/OS systems to become monitored clients of the IBM zAware server. For information about the types of z/OS systems to monitor, see Chapter 7, "Planning your IBM zAware environment," on page 39. IBM zAware supports z/OS systems that run in z/OS partitions or as z/VM guests. z/OS monitored clients must meet the following requirements:
  - The z/OS system must be configured as a single-system sysplex (monoplex), a system in a multisystem sysplex, or a member of a Parallel Sysplex.
  - The system must be running a supported release of the z/OS operating system.
  - The z/OS system must be using the operations log (OPERLOG) as the hardcopy medium.
  - The z/OS system name and sysplex name must uniquely identify the system to be monitored. IBM zAware identifies each monitored client by sysplex and system name, in the format *sysplex_name*.*system_name*; for example: SYSPLEX1.SYSA. IBM zAware cannot monitor more than one system with the same sysplex and system name combination.
- Log in to the z/OS system with a user ID that has the authority to complete the following tasks.
  - Access and modify members of the SYS1.PARMLIB data set.
  - Access and modify the LOGR couple data set and log stream usage.
  - Access the z/OS USS segment.

To determine the authority required to issue specific z/OS commands to accomplish these tasks, see the list of MVS™ commands, z/OS Security Server (RACF®) access authorities and resource names in *z/OS MVS System Commands*.

About this task

This procedure provides an overview of the steps required to configure each z/OS monitored client to send operations log (OPERLOG) data to the IBM zAware server. The details for each step are documented in various books in the z/OS product library; in this procedure, the appropriate books are listed for your reference. The z/OS product library is available in IBM Knowledge Center at http://www.ibm.com/support/knowledgecenter/. From the IBM Knowledge Center welcome page:

1. Use the Table of Contents to go to **IBM Operating Systems**, expand the appropriate platform.
2. Click **z/OS** and select the appropriate z/OS version and release that you are using.
3. Then navigate to one of the following book collections for each z/OS element:
   - z/OS Communications Server
   - z/OS MVS
   - z/OS Security Server RACF
   - z/OS UNIX System Services

Table 68 on page 286 contains a summary of related updates in the z/OS product library.

Procedure

1. Configure a network connection from each z/OS monitored client to the IBM zAware server. To determine whether this network configuration step is required, you can ping the IBM zAware server from the z/OS system. From the TSO command panel or the READY prompt on the z/OS system, enter the **PING** or **TRACERTE** command with either the IP address or host name of the IBM zAware server. If the command is successful, skip to step 2 on page 113.
   a. Update the TCP/IP profile for the z/OS monitored client, as necessary, for the channel paths that are assigned to the IBM zAware server.
      - For Ethernet connectivity, use an INTERFACE statement to define an IPAQENET or IPAQENET6 interface with CHPIDTYPE OSD.
      - For connectivity through the intraensemble data network (IEDN), use an INTERFACE statement to define an IPAQENET or IPAQENET6 interface with CHPIDTYPE OSX. Specify the CHPID parameter with the 2-character hexadecimal value that matches the value specified on the CHPID type OSX definition statement for the IBM zAware host system.
      - For connectivity through HiperSockets, use a DEVICE and LINK MPCIPA statement to define the HiperSockets device. Specify a device name of IUTIQD*xx*, where *xx* is the CHPID number that matches the hexadecimal value specified on the CHPID type IQD definition statement for the IBM zAware host system.

        Make sure that the z/OS client is configured to use layer 3 to connect to the IBM zAware server.

      If you make any changes to the domain name system (DNS) or local host file, you need to refresh the resolver.
   b. Start the modified TCP/IP stacks. For example, use the **VARY TCPIP** command to start the desired device, where *tcpipproc* is the name of the TCP/IP profile and *devicename* is the name of the device:

      ```
      VARY TCPIP,tcpipproc,START,devicename
      ```
   c. Verify the TCP/IP connection between the z/OS monitored client and the IBM zAware server. To verify the status of devices and links defined to the TCP/IP stack, use the **DISPLAY TCPIP** command to request NETSTAT information. For example:

      ```
      D TCPIP,procname,NETSTAT,DEVLINKS
      ```

      In the resulting display, check for the following:

- The device name and type match the TCP/IP profile definitions and the device is in a ready status.
- The link name and type match the TCP/IP profile definitions and the link is in a ready status.

d. If necessary, update firewall settings to ensure secure communications between the IBM zAware server and its monitored clients.

   Although the configurations in Chapter 7, "Planning your IBM zAware environment," on page 39 show both the IBM zAware server and its monitored clients within the boundary of a firewall, you can set up a configuration in which communication crosses firewall boundaries. In this case, you need to determine whether unsecured communication is an acceptable risk. If it is not an acceptable risk, you must provide your own method of securing this communication. For more information, see "Securing communication between IBM zAware and its monitored clients" on page 75.

For additional details, see the following books:
- *z/OS Communications Server: IP Configuration Guide*
- *z/OS Communications Server: IP Configuration Reference*
- *z/OS Communications Server: IP System Administrator's Commands*

2. Configure the z/OS monitored client as a single-system sysplex (monoplex), a system in a multisystem sysplex, or a member of a Parallel Sysplex. To determine whether the z/OS monitored client is configured correctly already, issue the **DISPLAY XCF,SYSPLEX,ALL** command and check the resulting message display for the system mode.

   If message IXC337I indicates that the system is running in XCF-LOCAL mode, modify the PLEXCFG parameter in the IEASYSxx parmlib member for the z/OS client. Specify one of the following parameter values:
   - PLEXCFG=MONOPLEX for a single-system sysplex
   - PLEXCFG=MULTISYSTEM for a system in a multisystem sysplex or a member of a Parallel Sysplex

   You can specify PLEXCFG=ANY, in which case the system mode is determined by settings in the COUPLExx and CLOCKxx parmlib members. In this case, make sure that those parmlib settings *do not* result in the z/OS monitored client running in XCF-LOCAL mode.

   For additional details, see the following topics:
   - "Planning parmlib members for a sysplex" in *z/OS MVS Setting Up a Sysplex*, SA22-7625.
   - IEASYSxx, COUPLExx, and CLOCKxx parmlib member descriptions in *z/OS MVS Initialization and Tuning Reference*, SA22-7592.

3. Configure the z/OS system logger to send data to the IBM zAware server. The following steps provide an overview of the required updates for the system logger. The primary references for details are:
   - Planning topics in *z/OS MVS Setting Up a Sysplex*, including "Preparing for z/OS IBM zAware log stream client usage."
   - Parmlib member descriptions in *z/OS MVS Initialization and Tuning Reference*.

   To define or update a log stream that contains data to be sent to the IBM zAware server, you must use a user ID with Security Authorization Facility (SAF) update access to the IXGZAWARE_CLIENT resource in the FACILITY class.

   a. Make sure the LOGR CDS format level is at least HBB7705.

      To determine what format level is in use for a sysplex, enter the following command and check the resulting message display.

      ```
      D XCF,COUPLE,TYPE=LOGR
      ```

      If the LOGR CDS format level is not HBB7705, your installation needs to run the format CDS utility IXCL1DSU with the DATA TYPE(LOGR) and ITEM NAME(SMDUPLEX) NUMBER(1) options. For more information, see the topic about LOGR parameters for the format utility in *z/OS MVS Setting Up a Sysplex*.

   b. Set up the authority that the z/OS system logger requires to send data to the IBM zAware server.

- For TCP/IP connectivity to the IBM zAware server, the IXGLOGR address space must have a z/OS UNIX System Services segment.
- Give the IXGLOGR address space superuser authority so the z/OS system logger can establish a TCP/IP connection to the IBM zAware server when z/OS UNIX System Services is not completely initialized. For example:

  ```
  ADDUSER IXGLOGR OMVS(UID(0) HOME('/')).
  ```

c. Configure SYS1.PARMLIB members that define or control system logger operations.

1) Create an IXGCNFxx parmlib member to define communication and log buffering details. Specify the `ZAI` statement with the following parameters:

SERVER(*host_name*|*IP_address***)**
Specifies the host name (as defined by the DNS server) or the IPv4 or IPv6 address that identifies where the IBM zAware server is running. If the z/OS client is connecting to the IBM zAware server over a network that uses the Dynamic Host Connection Protocol (DHCP), you must specify the host name for the IBM zAware partition.

PORT(*number***)**
Identifies the port number associated with the IBM zAware server. The port number is 2001.

LOGBUFMAX(*value***)**
Identifies the maximum amount of storage buffers (in gigabytes) to be used by system logger for managing z/OS monitored client data that is being sent to the IBM zAware server.

LOGBUFWARN(*nn***)**
Identifies the amount (as a percentage) of used buffer space for which the z/OS system logger starts issuing the error message IXG375E. Consider setting up automation for message IXG375E to avoid losing any data to be sent to the IBM zAware server.

LOGBUFFULL(MSG|QUIESCE)
Specifies the action that the z/OS system logger is to take when the log stream buffers become full.

MSG
z/OS system logger continues to send log data to the IBM zAware server, and keeps a count of the number of log blocks that could not be buffered and were lost. When buffers become available again, message IXG383I is issued to indicate that log data was not sent.

QUIESCE
z/OS system logger disconnects the TCP/IP connection to the IBM zAware server, releases the buffers, and issues message IXG382I for each log stream that is quiesced. A SETLOGR FORCE,ZAICONNECT or SET IXGCNF command is required to reconnect the z/OS system logger to the IBM zAware server.

The following example shows the **ZAI** statement with parameters and sample values:

```
ZAI
 SERVER(zserver.loc.com)
 PORT(2001)
 LOGBUFMAX(1)
 LOGBUFWARN(80)
 LOGBUFFULL(MSG)
```

2) Add the IXGCNFxx system parameter to the IEASYSxx parmlib member. This system parameter specifies the IXGCNFxx parmlib member to be used when the z/OS system logger starts or is restarted. In the IEASYSxx parmlib member, use the syntax format `IXGCNF=xx`

d. Issue the **SET IXGCNF=***xx* command to apply the updated system logger configuration.

4. Use the **DISPLAY LOGGER** command to verify the configuration updates. The following **DISPLAY** command requests the z/OS system logger to use current configuration options to communicate with the IBM zAware server.

   ```
   DISPLAY LOGGER,STATUS,ZAI,VERIFY
   ```

   The resulting display contains general system logger status, the state of the z/OS monitored client, and ZAI statement parameter options, along with an indication of whether the "verify communication" request succeeded or failed. If the display contains "ZAI VERIFY INITIATED", check for messages in the range IXG37x-IXG38x with the text "DISPLAY ZAI,VERIFY" included for the verification results. The following system logger message indicates successful communications between the z/OS monitored client and the IBM zAware server:

   ```
   IXG380I ZAI LOGSTREAM CLIENT ESTABLISHED FOR DISPLAY ZAI,VERIFY
   ```

   For additional information about the **DISPLAY LOGGER** command, see *z/OS MVS System Commands*.

5. Determine whether the z/OS monitored client is using OPERLOG as the hardcopy medium. Issue the **DISPLAY CONSOLES** command and check the resulting message display for information about the hardcopy medium.
   - If message CNZ4100I indicates that the system is not configured to use OPERLOG, continue to step 6.
   - If message CNZ4100I indicates that the system is using OPERLOG already, complete the following steps.

   a. Update the existing OPERLOG log stream with the parameters required for IBM zAware.

      ZAI(YES)
      Specifies that the log stream data is to be sent to the IBM zAware server.

      ZAIDATA('*value*')
      Provides an optional value to be passed to the IBM zAware server. For example, you can specify 'OPERLOG' as the value for the ZAIDATA keyword.

      To update the OPERLOG log stream with the administrative data utility IXCMIAPU, use a SYSIN DD statement similar to the following sample.

      ```
      //SYSIN DD *
      DATA TYPE(LOGR) REPORT(YES)
      UPDATE LOGSTREAM NAME(SYSPLEX.OPERLOG)
      ZAI(YES)
      ZAIDATA('OPERLOG')
      ```

   b. Create the required security definitions for the z/OS Security Server (RACF), or an equivalent security product, to allow users to browse the operations log. In the following example, the SYSPLEX.OPERLOG of the LOGSTRM resource CLASS is given READ permission, which allows all users to browse the operations log. *userid1* has UPDATE access level, which allows *userid1* to delete records from the log stream.

      ```
      RDEFINE LOGSTRM SYSPLEX.OPERLOG UACC(READ)
      PERMIT SYSPLEX.OPERLOG CLASS(LOGSTRM) ID(userid1)
      ACCESS(UPDATE) SETROPTS CLASSACT(LOGSTRM)
      ```

   c. Continue to step 7 on page 116.

6. Configure the z/OS monitored client to use OPERLOG as the hardcopy medium. Use the following steps as a model for configuring OPERLOG at your installation.

   a. Using the administrative data utility IXCMIAPU, define the corresponding coupling facility structure in the coupling facility resource management (CFRM) policy. For example:

      ```
      STRUCTURE NAME(OPERLOG)
      SIZE(40448)
      INITSIZE(40448)
      PREFLIST(FACIL01,FACIL02)
      ```

   b. Activate the CFRM policy. You can activate the policy through the COUPLExx parmlib member or through the SETXCF command. For example:

```
SETXCF START,POLICY,TYPE=CFRM,POLNAME=policy_name
```

c. Create the OPERLOG log stream, using the following parameters that are required for IBM zAware.

ZAI(YES)
Specifies that the log stream data is to be sent to the IBM zAware server.

ZAIDATA('*value*')
Provides an optional value to be passed to the IBM zAware server. For example, you can specify 'OPERLOG' as the value for the ZAIDATA keyword.

The following sample illustrates JCL and control statements for using the administrative data utility to define an OPERLOG log stream:

```
//OPERLOG JOB CLASS=A,MSGCLASS=A
//POLICY EXEC PGM=IXCMIAPU
//SYSPRINT DD SYSOUT=A
//SYSIN DD *
DATA TYPE(LOGR)
DEFINE STRUCTURE NAME(OPERLOG)
LOGSNUM(1)
MAXBUFSIZE(4092)
AVGBUFSIZE(512)
DEFINE LOGSTREAM NAME(SYSPLEX.OPERLOG)
STRUCTNAME(OPERLOG)
LS_DATACLAS(LOGR24K)
HLQ(IXGLOGR)
LS_SIZE(2560)
LOWOFFLOAD(0)
HIGHOFFLOAD(80)
STG_DUPLEX(NO)
RETPD(0)
AUTODELETE(NO)
ZAI(YES)
ZAIDATA('OPERLOG')
```

d. Create the required security definitions for the z/OS Security Server (RACF), or an equivalent security product, to allow users to browse the operations log. In the following example, the SYSPLEX.OPERLOG of the LOGSTRM resource CLASS is given READ permission, which allows all users to browse the operations log. *userid1* has UPDATE access level, which allows *userid1* to delete records from the log stream.

```
RDEFINE LOGSTRM SYSPLEX.OPERLOG UACC(READ)
PERMIT SYSPLEX.OPERLOG CLASS(LOGSTRM) ID(userid1)
ACCESS(UPDATE) SETROPTS CLASSACT(LOGSTRM)
```

e. Define the hardcopy device as OPERLOG in the HARDCOPY statement of the CONSOLxx parmlib member. You can change this setting using the **VARY** command:

```
V OPERLOG,HARDCPY
```

f. If your system is configured as a monoplex, create a DASD-only log stream for OPERLOG. To create a DASD-only log stream, you need to define or update the system logger couple data set (LOGR CDS) with a large enough log stream records (LSR) value to allow sufficient space for managing the DASD-only log stream for this z/OS monitored client. Review the planning considerations in *z/OS MVS Setting Up a Sysplex*, including the topics about planning DASD space for system logger and managing log data.

For more information about setting up OPERLOG, see the topic on preparing to use system logger applications in *z/OS MVS Setting Up a Sysplex*.

7. Start sending log stream data to the IBM zAware server. Issue the **SETLOGR** command:

```
SETLOGR FORCE,ZAICONNECT,LSNAME=SYSPLEX.OPERLOG
```

The **ZAICONNECT** parameter directs system logger to attempt a socket connection from the z/OS monitored client to the IBM zAware server, as defined in the current IXGCNFxx parmlib member.

The system responds with the following messages:

```
IXG651I SETLOGR FORCE ZAICONNECT COMMAND ACCEPTED FOR LOGSTREAM=SYSPLEX.OPERLOG
IXG386I ZAI LOGSTREAM CLIENT CONNECT ATTEMPT IN PROGRESS FOR LOGSTREAM SYSPLEX.OPERLOG
STATUS:  ATTEMPTING SOCKET CREATE
IXG386I ZAI LOGSTREAM CLIENT CONNECT ATTEMPT IN PROGRESS FOR LOGSTREAM SYSPLEX.OPERLOG
STATUS:  SOCKET CREATE SUCCESSFUL
IXG386I ZAI LOGSTREAM CLIENT CONNECT ATTEMPT IN PROGRESS FOR LOGSTREAM SYSPLEX.OPERLOG
STATUS:  ATTEMPTING SOCKET CONNECT
IXG386I ZAI LOGSTREAM CLIENT CONNECT ATTEMPT IN PROGRESS FOR LOGSTREAM SYSPLEX.OPERLOG
STATUS:  SOCKET CONNECT SUCCESSFUL
IXG386I ZAI LOGSTREAM CLIENT CONNECT ATTEMPT IN PROGRESS FOR LOGSTREAM SYSPLEX.OPERLOG
STATUS:  INITIATING SOCKET VALIDATION
IXG386I ZAI LOGSTREAM CLIENT CONNECT ATTEMPT IN PROGRESS FOR LOGSTREAM SYSPLEX.OPERLOG
STATUS:  SOCKET VALIDATION SUCCESSFUL
IXG380I ZAI LOGSTREAM CLIENT ESTABLISHED FOR LOGSTREAM SYSPLEX.OPERLOG
```

For detailed information about log stream status, you can issue the following **DISPLAY LOGGER** command:

```
D LOGGER,C,LSN=SYSPLEX.OPERLOG,D
```

The system responds with message IXG601I, as shown in the following sample:

```
IXG601I   15.40.21  LOGGER DISPLAY   FRAME  1     F      E   SYS=SY2
CONNECTION INFORMATION BY LOGSTREAM FOR SYSTEM SY2
LOGSTREAM                 STRUCTURE        #CONN  STATUS
---------                 ---------        ------ ------
SYSPLEX.OPERLOG           IXGLOGR_STR1     000001 OFFLOAD IN PROGRESS
  DUPLEXING: STAGING DATA SET
    STGDSN: LOGGER.SYSPLEX.OPERLOG.SY2
            VOLUME=SMSVL4  SIZE=000060 (IN 4K)   % IN-USE=050
  GROUP: PRODUCTION   ZAI CLIENT: YES - CONNECTED
  ZAIDATA: NO_ZAIDATA
    LOG BLOCKS SENT TO SERVER OK: 0000000008, FAILED: 0000000000
  JOBNAME: CONSOLE   ASID: 0009
    R/W CONN: 000000 / 000001
    RES MGR./CONNECTED: *NONE*   / NO
LOGSTREAM                 STRUCTURE        #CONN  STATUS
    IMPORT CONNECT: NO

NUMBER OF LOGSTREAMS:  000001
```

For additional information about the **SETLOGR** and **DISPLAY LOGGER** commands, see *z/OS MVS System Commands*.

Results

The z/OS system is established as an IBM zAware monitored client.

When z/OS monitored clients send current data to the IBM zAware server, the server uses both the sysplex name and system name passed in this data to uniquely identify data for a specific monitored client. To view the status for monitored clients that are sending current data, navigate to the **System Status** tab on the **Systems** page in the IBM zAware GUI.

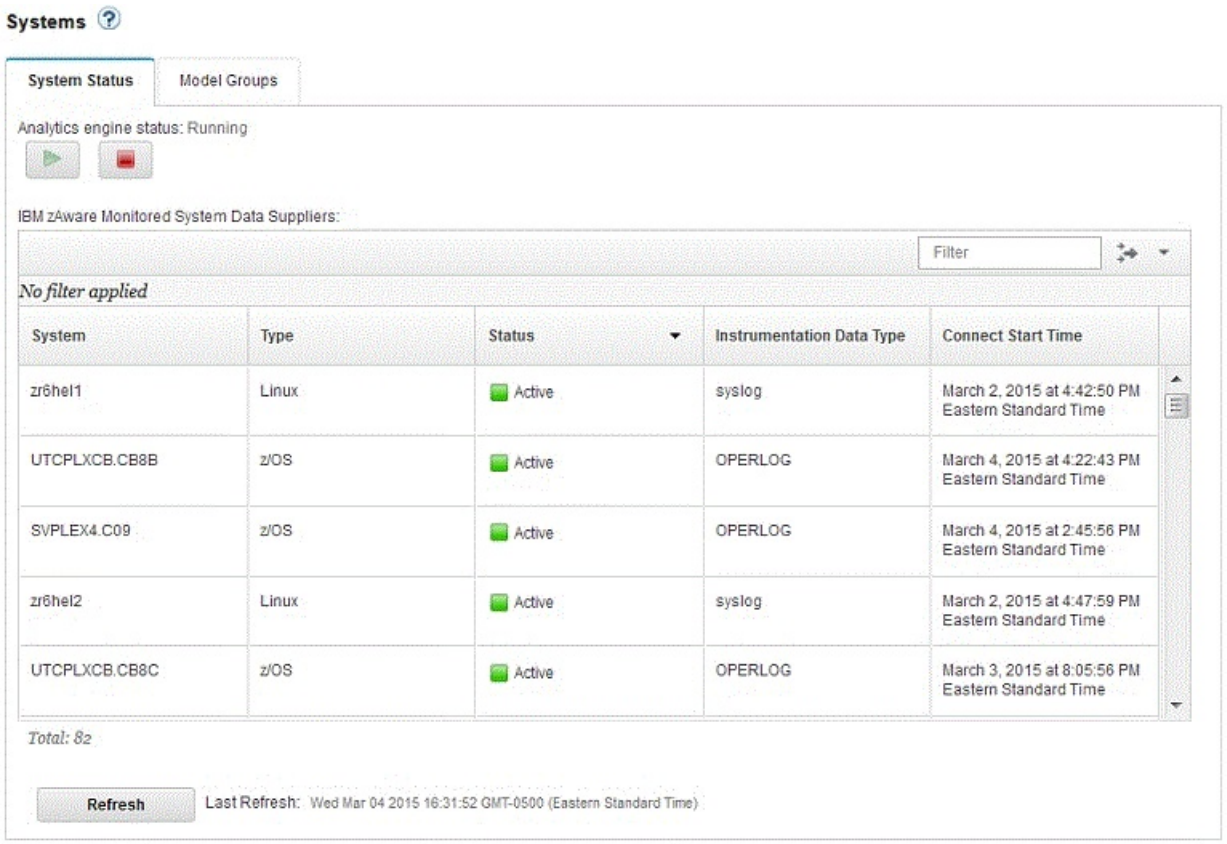

Figure 39. Systems Status page

What to do next

- To prime the IBM zAware server to begin data analysis as quickly as possible, follow the instructions in "Creating an IBM zAware model for new z/OS monitored clients."
- Issue the **SETLOGR** command as necessary during normal operations for the monitored client.
  - If an unscheduled IPL or system logger restart occurs after the z/OS system is established as a monitored client, you do not have to issue the **SETLOGR** command to reconnect the client and IBM zAware server. Following an IPL or system logger restart, communication is attempted automatically after the OPERLOG log stream is connected on the z/OS system.
  - To prepare for a scheduled IPL, consider disconnecting the monitored client to avoid reconnection attempts and the messages associated with those attempts. For example, you can complete the following shutdown procedure:
    1. Issue `SETLOGR FORCE,ZAIQUIESCE` with additional parameters, as necessary, to disconnect a specific or all IBM zAware log stream clients that are running on the z/OS system.
    2. Follow the procedure that your installation uses to shut down TCP/IP and z/OS UNIX System Services.
- If necessary, see *z/OS MVS Diagnosis: Reference* for information about resolving system logger errors related to log stream processing for z/OS monitored clients.

Creating an IBM zAware model for new z/OS monitored clients

Use this procedure to prime the IBM zAware server with data to create a model of normal system behavior for z/OS monitored clients. The estimated amount of data for building the most accurate models is 90 days of data for each client. Your installation can modify this data requirement, which is known as the *training period* for IBM zAware analytics, based on your knowledge of the workloads running on z/OS monitored clients. Instead of waiting for the IBM zAware server to collect data over the

course of the training period, however, you can prime the server by transferring prior data for monitored clients, and request the server to build a model for each client from the transferred data. This procedure provides instructions for priming the IBM zAware server and building a model from the priming data. This procedure is intended for skilled system programmers who have experience with the z/OS systems that are monitored clients.

Before you begin

- Make sure that your installation has completed the procedures in:
  - Chapter 13, "Configuring storage, security, and analytics for the IBM zAware server," on page 99.
  - "Configuring z/OS monitored clients to send data to the IBM zAware server" on page 111.
- To ensure that you have enough priming data to successfully build a model, review the information in "Planning to create IBM zAware models for z/OS monitored clients" on page 87.
- Identify the z/OS system from which you plan to send priming data to the IBM zAware server, and make sure that you configure this priming system according to the instructions in "Configuring z/OS monitored clients to send data to the IBM zAware server" on page 111.
- List the systems for which you plan to send priming data; the location, names, and sizes of the sequential data sets that contain their priming data; and the order in which you want to send priming data.
  - These systems must be configured as described in "Configuring z/OS monitored clients to send data to the IBM zAware server" on page 111 before you send their priming data.
  - The z/OS bulk load client can process sequential data sets that contain only SYSLOG data that is stored in hardcopy log 2-digit year (HCL) or 4-digit year (HCR) format.

  The recommended approach is to transfer only a portion of the priming data for a monoplex or sysplex to verify the configuration of the IBM zAware environment. When the transfer is successful, you can transfer the remaining data for one monoplex or sysplex at a time. If you transfer priming data for multiple sysplexes through one invocation of the REXX exec, priming data for some systems can be overlaid.
- To configure and run the z/OS bulk load client for IBM zAware on the priming z/OS system, the user ID under which the z/OS bulk load client runs must have authority to use log streams and to read the SYSLOG or OPERLOG archives that constitute the priming data. To run the z/OS bulk load client, and to define the model log stream that it uses, you must use a user ID with Security Authorization Facility (SAF) update access to the IXGZAWARE_CLIENT resource in the FACILITY class.

  To use the z/OS bulk load client, you also need to check and modify storage attributes for the log stream that it uses. See *z/OS DFSMSdfp Storage Administration* if you need additional details about creating SMS classes and other related DASD storage attributes for the log stream.
- To log in to the IBM zAware graphical user interface (GUI), you need to know the URL.

  The URL includes the IP address or host name that is assigned to the IBM zAware partition:

  `https://ip_address/zAware/` **or** `https://host_name/zAware/`

  The "zAware" portion of the URL is case-sensitive.

  To assign priming data and use it to create the model of system behavior, you use specific administration functions in the IBM zAware GUI. To use these functions, you must log in to the GUI with a user ID that is assigned to the Administrator role.

About this task

After your installation completes the procedure in "Configuring z/OS monitored clients to send data to the IBM zAware server" on page 111, the IBM zAware server is receiving current data from the z/OS system logger running on z/OS monitored clients. However, the server cannot use this data for analysis until a model of normal system behavior exists. The estimated amount of data for building the most accurate models is 90 days of data for each client. Your installation can modify the number of days required for this training period, based on your knowledge of the workloads running on z/OS monitored clients. This training period applies for all monitored clients; you cannot define a different training period

for each client. Chapter 13, “Configuring storage, security, and analytics for the IBM zAware server,” on page 99 contains information about the training period for analytics.

Instead of waiting for the IBM zAware server to collect data over the course of the training period, you can prime the server by transferring prior data from the hardcopy or system logs of monitored clients, and request the server to build a model for each client from the transferred data. To transfer this priming data, you configure a log stream and run the z/OS bulk load client through a REXX exec on a z/OS system. You can run the REXX exec for the z/OS bulk load client in the TSO/E foreground or in the background as a batch job. This procedure explains how to run the exec as a batch job after editing sample JCL in SYS1.SAMPLIB.

Through the z/OS bulk load client, you can transfer data for one or more monitored clients by identifying the sequential data sets that contain the priming data. If any data set has been archived, the z/OS bulk load client can recall the data set, transfer its contents, and migrate the data set. The REXX exec is designed for use with direct-access storage device (DASD) data sets, not tape data sets.

The process of transferring priming data could require several hours or more, depending on a number of factors that include:

- The priority of the job that runs the REXX exec for the z/OS bulk load client
- The amount of priming data to be sent, and whether any of that priming data resides on migrated data sets
- The network configuration and traffic at your installation

For example, if you run the REXX exec at a very high priority to send 46000 tracks of priming data that is archived, the transfer might take approximately 10 to 15 minutes. The process can take longer if the z/OS bulk load client runs at a lower priority, or if network or system conditions are not favorable when the REXX exec runs.

In contrast to data that the IBM zAware server receives from the z/OS system logger running on a monitored client, the priming data from the z/OS bulk load client does not include the name of the sysplex to which the monitored client belongs. Without the sysplex name, the IBM zAware server cannot associate the priming data with the appropriate sysplex. You use the **Administration** > **Configuration** > **Priming Data** page in the IBM zAware GUI to assign the received priming data to the appropriate sysplex. After the priming data is associated with the appropriate sysplex, you can request the IBM zAware server to build the model.

To build the model for a specific monitored client, you have two options:

- You can use the **Request Training** action on the **Administration** > **Training Sets** page. Any data that the z/OS system logger is currently sending does not become part of the model for the client until you request training again or the IBM zAware server automatically rebuilds the model. This priming option is recommended because analysis can start shortly after the model is built.
- You can wait for the next scheduled training, during which the IBM zAware server automatically uses the priming data to build the model. In this case, any data that the z/OS system logger is currently sending becomes part of the model for the client.

The time required for an automatic or manual training request to complete depends on the amount of priming data that the bulk loader sent; typically, the training process takes only several minutes to complete but might take longer for large amounts of priming data. After the training status shown on the **Training Sets** page changes from “In progress” to “Complete”, a model containing the priming data is available for the IBM zAware server to use for analysis, when the client is connected and sending current data.

Procedure

1. On the z/OS system from which you plan to send priming data, configure the z/OS bulk load client. The z/OS bulk load client load modules reside in SYS1.MIGLIB and sample files reside in SYS1.SAMPLIB.

a. Decide where you want the z/OS bulk load client load modules to reside. You have the option to copy the AIZBLKR and AIZBLKM load modules from SYS1.MIGLIB into an authorized or unauthorized library. If you decide to leave the load modules in SYS1.MIGLIB, however, make sure that you either add SYS1.MIGLIB to the LINKLIST concatenation or to the JCL JOBLIB or STEPLIB parameter for the job that runs the REXX exec.

b. Copy the z/OS bulk load client sample files from SYS1.SAMPLIB into your JCL library.

   1) Copy AIZBLK to a JCL data set. AIZBLK contains a sample job that completes the following steps:
      a) Defines a model log stream to establish the attributes of the target log stream that the z/OS bulk load client creates and uses for priming data. The z/OS bulk load client writes priming data into this target log stream.
      b) Defines and populates a control data set that lists the sequential data sets containing priming data to be transferred.
      c) Invokes the z/OS bulk load client to transfer the priming data.
      d) Deletes the model log stream and the target log stream that was used to send the priming data.
   2) Copy AIZBLKE from SYS1.SAMPLIB to a data set in your SYSEXEC concatentation. AIZBLKE contains the sample REXX exec to run the AIZBLKR load module.

      If the target data set is VB format, make sure the numbers in columns 72-80 are deleted before you copy AIZBLKE into the data set.

c. For the sequential data sets that contain priming data, check the following values.
   - If the priming data contains ANSI carriage control characters, check that the RECFM attribute of the data set indicates ANSI control characters, such as VBA or FBA.
   - If the priming data contains machine control characters, check that the RECFM attribute of the data set indicates machine control characters, such as VBM or FBM.

d. Check the `g.MaxImportBytes` value in the AIZBLKE REXX exec. This value controls the maximum amount of data that can be imported per job. You might need to change this value to accommodate the amount of priming data that you plan to send to the IBM zAware server.

e. Modify the sample JCL according to the customization instructions in the AIZBLK file. To run the REXX exec as a batch job, you need to modify the sample JCL. Make sure that you make the following changes:
   - Check the value set for the REGION parameter; depending on the amount of data you are sending, you might need to specify REGION=0M on the step for running the REXX exec.
   - Update the model log stream name to match naming standards at your installation.
   - Check the logger keywords and parameters as noted in the following list. For more information, see the topic about LOGR keywords and parameters for the administrative data utility in *z/OS MVS Setting Up a Sysplex*, SA22-7625.
     – Check the parameters that control the use of offload data sets and, if necessary, modify the values to match the following code sample:

       ```
       AUTODELETE(YES)
       HIGHOFFLOAD(60)
       LOWOFFLOAD(30)
       ```

     – Check that the DFSMS data classes for staging data sets and offload data sets are set to the recommended control interval (CI) sizes. In the following statements, the LOGR4K and LOGR24K values are sample data class names; substitute the appropriate data class names that are defined in the DFSMS configuration for your installation.

       ```
       STG_DATACLAS(LOGR4K)
       LS_DATACLAS(LOGR24K)
       ```

     – Check the LS_SIZE and STG_SIZE parameters to ensure that they are set to reasonable values, based on the amount of priming data you are sending. If the size of the log stream data sets are established through DFSMS data class definitions, you can omit the LS_SIZE and STG_SIZE parameters.

LS_SIZE(*size*)
Specifies the size, in 4K blocks, of the log stream offload DASD data sets for the log stream being defined.

STG_SIZE(*size*)
Specifies the size, in 4K blocks, of the DASD staging data set for the log stream being defined.

Keep in mind that the recommended approach is to complete the following tasks in sequence; these tasks determine which data sets you select for the first invocation of the z/OS bulk load client. You need to load the data sets in chronological order, from oldest to newest.

1) Transfer only a portion of the priming data for a monoplex or sysplex to verify the configuration of the IBM zAware environment. In the IBM zAware GUI, you can verify the configuration by navigating to the **Administration** > **Configuration** > **Priming Data** tab. Check the "Priming data by systems" list for the system name of the z/OS monitored client associated with the priming data, and check the "Sysplex Topology" list for the name of the associated sysplex.
2) When you verify that the transfer was successful, transfer the remaining data for one monoplex or sysplex at a time. The following sample code shows how to specify multiple logs for a system named CB8E.

```
AIZBLKE -------.ZAI.CONTROL ADDSYSLOGDSN + LOGWRTR.LOGCB8E.G0784V00
AIZBLKE -------.ZAI.CONTROL ADDSYSLOGDSN + LOGWRTR.LOGCB8E.G0785V00
```

If you transfer priming data for multiple sysplexes through one invocation of the REXX exec, priming data for some systems can be overlaid.

f. Update security definitions to allow the z/OS bulk load client to access the log stream and the data sets that contain priming data. The z/OS bulk load client requires the following authorization through the z/OS Security Server (RACF), or an equivalent security product:
- UPDATE access for the log stream
- READ access for the priming data sets

g. Check the connection between the priming z/OS system and the IBM zAware server. Issue the following MVS command:

```
DISPLAY LOGGER,STATUS,ZAI,VERIFY
```

The system responds with message IXG601I; the following sample illustrates the message display.

```
SYSTEM LOGGER STATUS
SYSTEM  SYSTEM LOGGER STATUS
------  --------------------
SY1     ACTIVE

ZAI LOGSTREAM CLIENTS: AVAILABLE
BUFFERS IN USE: 00 GB 0000 MB
ZAI VERIFY INITIATED, CHECK FOR MESSAGES IXG37X, IXG38X

LOGGER PARAMETER OPTIONS
KEYWORD            SOURCE   VALUE
-----------------  -------- ------------------------------------------
ZAI
 SERVER            IPL (NN) HOST.ZAWARE.SERVER.LOCATION
 PORT              DEFAULT  2001
 LOGBUFMAX         DEFAULT  02
 LOGBUFWARN        DEFAULT  75
 LOGBUFFULL        DEFAULT  MSG
IXG386I ZAI LOGSTREAM CLIENT CONNECT ATTEMPT IN PROGRESS FOR DISPLAY ZAI,VERIFY
STATUS:  ATTEMPTING SOCKET CREATE
IXG386I ZAI LOGSTREAM CLIENT CONNECT ATTEMPT IN PROGRESS FOR DISPLAY ZAI,VERIFY
STATUS:  SOCKET CREATE SUCCESSFUL
IXG386I ZAI LOGSTREAM CLIENT CONNECT ATTEMPT IN PROGRESS FOR DISPLAY ZAI,VERIFY
STATUS:  ATTEMPTING SOCKET CONNECT
IXG386I ZAI LOGSTREAM CLIENT CONNECT ATTEMPT IN PROGRESS FOR DISPLAY ZAI,VERIFY
STATUS:  SOCKET CONNECT SUCCESSFUL
```

```
IXG386I ZAI LOGSTREAM CLIENT CONNECT ATTEMPT IN PROGRESS FOR DISPLAY ZAI,VERIFY
STATUS:  INITIATING SOCKET VALIDATION
IXG386I ZAI LOGSTREAM CLIENT CONNECT ATTEMPT IN PROGRESS FOR DISPLAY ZAI,VERIFY
STATUS:  SOCKET VALIDATION SUCCESSFUL

IXG380I ZAI LOGSTREAM CLIENT ESTABLISHED
FOR DISPLAY ZAI,VERIFY
```

2. On the priming z/OS system, run the z/OS bulk load client to send priming data to the IBM zAware server.
   a. Submit the edited JCL sample for processing. The time required for the z/OS bulk load client to transfer the priming data depends on how much data is in the data sets listed in the control file. If a data set has been archived, the transfer process requires more time for the z/OS bulk load client to recall the data set, transfer its contents, and migrate the data set.

      When the job completes, check the job log for the return code from the z/OS bulk load client. Possible return code values are:

      0 Successful
      4 Request not valid
      8 Request failed
      12 Request stopped; the maximum number of bytes has been exceeded

      A return code of 0 indicates that the priming data is queued for transmission to the IBM zAware server.
   b. Check for an indication that the z/OS system logger sent the priming data. Look in the system log (SYSLOG) for IXG38x or IXG37x messages that report problems with the transfer of data between the z/OS system logger and the IBM zAware server.
   c. Verify that the IBM zAware server has received the priming data. Through the IBM zAware GUI, check the "Priming data by systems" list on the **Administration** > **Configuration** > **Priming Data** tab for the system name of the z/OS monitored client associated with the priming data.

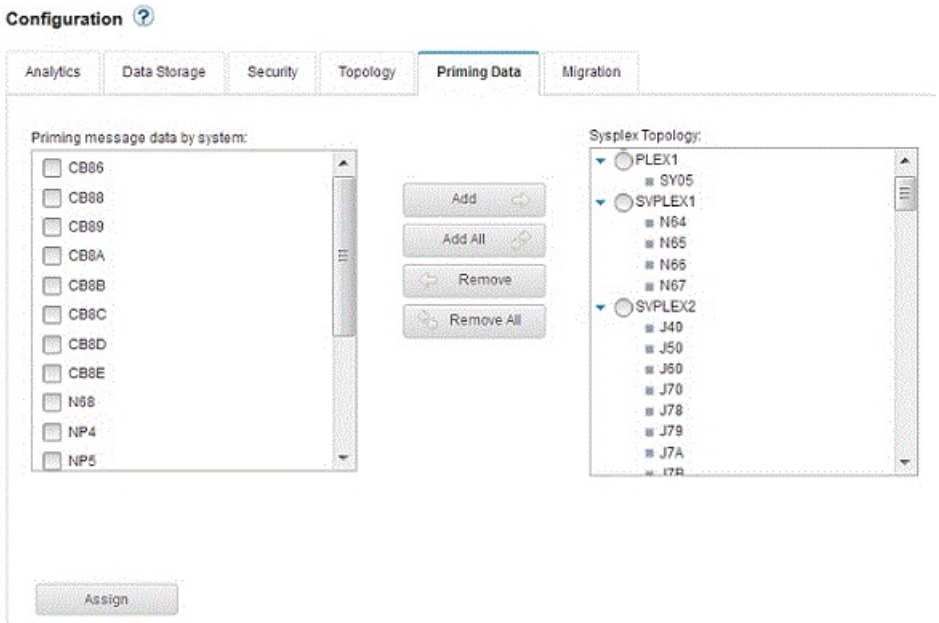

Figure 40. Priming Data: z/OS monitored client

   d. Edit and resubmit the JCL, as necessary, to finish sending the priming data for the z/OS monitored clients.
3. Through the IBM zAware GUI, assign the received priming data to the appropriate sysplex. After the data transfer from the priming z/OS system is complete, the IBM zAware server indicates the data

received by adding the system name in the "Priming data by systems" list on the **Administration** > **Configuration** > **Priming Data** tab. The system name is the name of the z/OS monitored client associated with the priming data. Because the z/OS bulk load client can send data for more than one monitored client at a time, several systems might be listed.

The following steps explain how to assign priming data from monitored clients (systems) by moving those systems from the "Priming data by systems" list to the Sysplex Topology list on the **Priming Data** page. You do not have to assign all systems in the list until you are ready to do so. Unassigned systems remain in the "Priming data by systems" list until you add them to the sysplex topology.

a. In the Sysplex Topology list, select the sysplex to which you want to assign systems.
b. In the "Priming message data by systems" list, select the systems that you want to move to the selected sysplex and click **Add >**. If you want to assign all of the systems in the "Priming message data by systems" list to the same sysplex, you do not have to select them first; instead, you can use **Add All >>** to move all systems in the list to the selected sysplex. After you click **Add >** or **Add All >>**, the systems are moved from the "Priming message data by systems" list to the Sysplex Topology list. They are displayed under the selected sysplex, with the parenthetical phrase "data to be assigned" displayed after the system name. If necessary, expand the sysplex topology to see the list of systems for the selected sysplex.
c. As necessary, repeat steps a and b to move each system in the "Priming message data by systems" list to the appropriate sysplex in the Sysplex Topology list.
d. When you have finished moving systems from the "Priming message data by systems" list to the appropriate sysplex, click **Assign** to apply your changes.
e. Review and confirm your changes by clicking **OK** on the **Assign Priming Data** window.

IBM zAware assigns the priming data to the appropriate sysplex. During this process, IBM zAware recycles its analytics engine. When the analytics engine is stopped or recycled, IBM zAware disconnects all monitored systems from the server. When z/OS monitored clients are disconnected from the server, they automatically attempt to reconnect and continue to buffer data for approximately 10 minutes. If the analytics engine is restarted within that time, the z/OS system reconnects and sends the buffered data to the IBM zAware server. If the z/OS system times out before the analytics engine is restarted, any buffered data is lost and you must manually reconnect the system. To reconnect a z/OS system, issue the SETLOGR command.

```
SETLOGR FORCE,ZAICONNECT,LSN=SYSPLEX.OPERLOG
```

4. Verify that the transferred data is available for the IBM zAware server to use.
   a. From the **Administration** > **Training Sets** > **z/OS Systems** tab, select the system name of the z/OS monitored client.

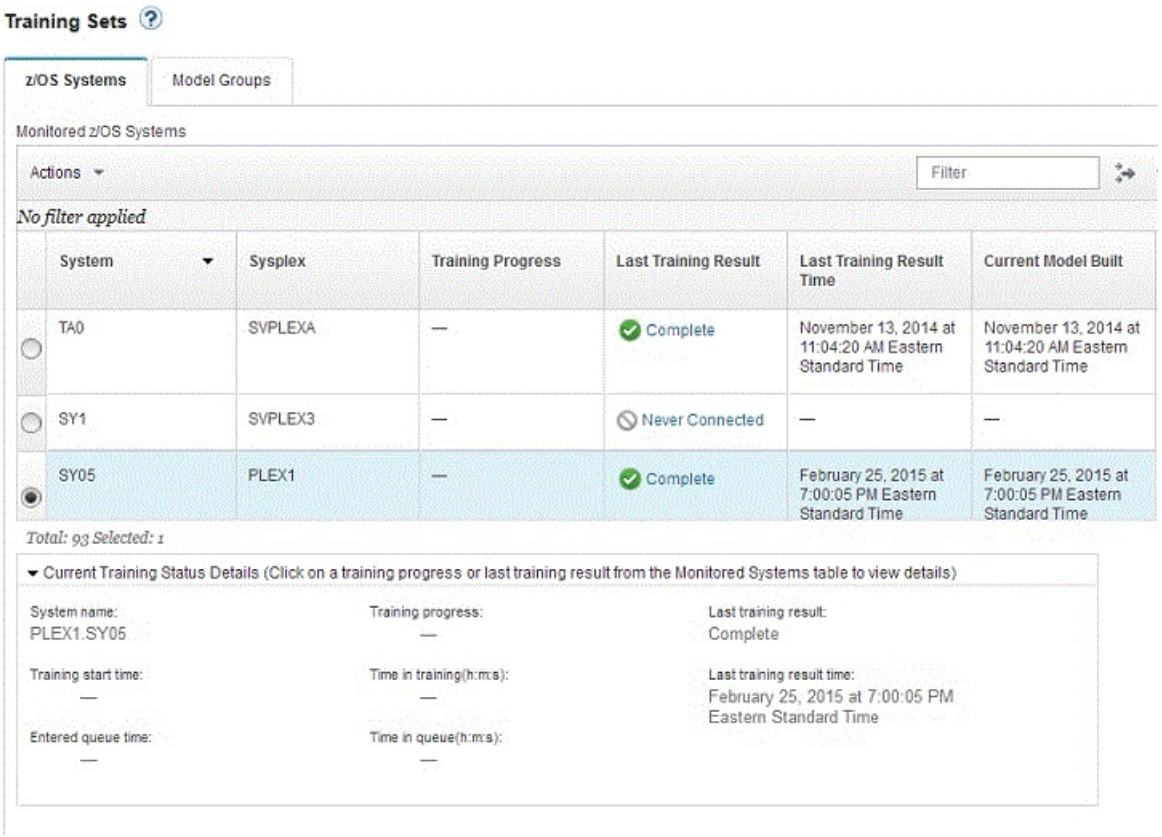

Figure 41. Training Sets: z/OS monitored client

b. From the **Actions** list, select **Manage Model Dates**.

c. From the **Manage Model Dates** page, select **Next Training Period Model Dates** from the **Model dates** list.

   Use the calendar view to determine days for which transferred data is available. Calendar days that are not marked as “Excluded” or “Unavailable” identify the dates for which the IBM zAware server has data to use.

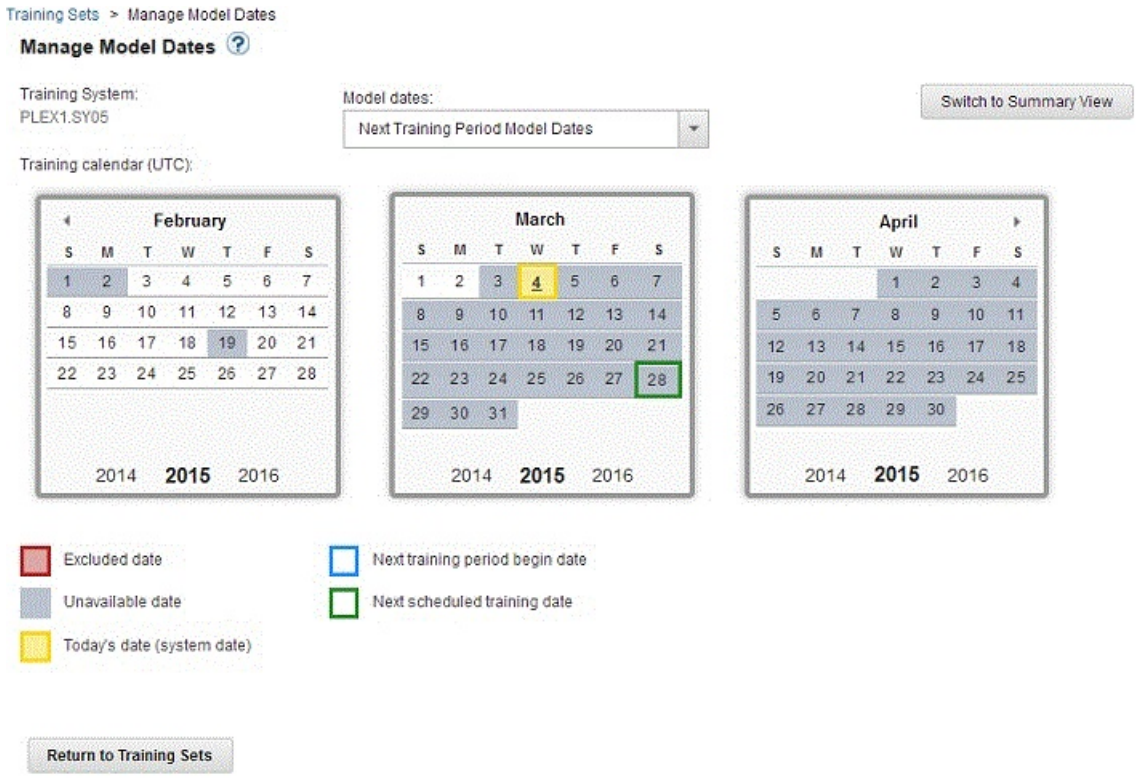

d. Click **Return to Training Sets** to return to the previous page.

e. Repeat these steps, as necessary, for each system for which data was transferred.

5. Optional: Request the IBM zAware server to build models from the priming data. Otherwise, you can wait for the next scheduled training and the IBM zAware server automatically uses the priming data when it builds the model.

   a. Through the **Administration** > **Training Sets** > **z/OS Systems** tab, select one monitored client (system).

   b. From the **Actions** list, select **Request Training** to build a model for the selected client.

   c. Confirm your request by clicking **OK** on the **Request Training** window. On the **Administration** > **Training Sets** > **z/OS Systems** tab, the Current Training Status column for the monitored client (system) that you selected contains either "In Progress" or "In Queue", with the queue position indicated.

   d. Repeat these steps, as necessary, for each system for which priming data was transferred.

   The time required for the training request to complete depends on the amount of priming data that the bulk loader sent; the training process might take several hours to complete. To track progress, use any of the following techniques.

   - Click **Refresh** to update the **z/OS Systems** tab content.
   - Click ► to expand the **Current Training Status Details** section, then click on any value in the Current Training Status column to view details.

   For an explanation of the **z/OS Systems** tab content, see "Training sets for z/OS systems" on page 232.

Results

After the training process completes for a given client (system), a model containing the priming data is available for the IBM zAware server to use for analysis, when the client is connected and sending current data.

What to do next

View current analysis data for monitored clients as described in Chapter 16, “Viewing and using analytical data to monitor and diagnose system behavior,” on page 139.

Chapter 15. Configuring Linux on z Systems monitored clients for IBM zAware analysis

To configure Linux on z Systems monitored clients for IBM zAware analysis, you need to configure the Linux syslog daemon to send systems log (syslog) data from that Linux system to the IBM zAware server. Also, an administrator must define model groups to match Linux system host names so IBM zAware can build models and produce analysis results.

For additional details, see the following topics:

1. "Configuring Linux on z Systems monitored clients to send data to the IBM zAware server"
2. "Creating an IBM zAware model for new Linux on z Systems monitored clients" on page 132

Configuring Linux on z Systems monitored clients to send data to the IBM zAware server

Use this procedure as an overview for configuring Linux on z Systems monitored clients to send data to the IBM zAware server for analysis. This procedure is intended primarily for skilled Linux administrators. Depending on the roles and responsibilities defined for your IT organization, you might need the assistance of network or security administrators to correctly configure secure connectivity within the IBM zAware environment.

Before you begin

- Make sure that your installation has completed the following steps for defining network connections for the hardware in the IBM zAware environment:
  - Chapter 12, "Configuring network connections and storage for the IBM zAware partition," on page 95.
  - "Configuring the IBM z Systems Secure Service Container for IBM zAware" on page 27.

  Consider using the checklist in "Task summary and configuration checklist for network administrators" on page 57 as an aid for step 1 on page 130 in this procedure.
- Make sure that your installation has completed the procedure in Chapter 13, "Configuring storage, security, and analytics for the IBM zAware server," on page 99.
- List the Linux systems to become monitored clients of the IBM zAware server. For additional information about supported configurations of monitored systems, see Chapter 7, "Planning your IBM zAware environment," on page 39.

  To become monitored clients of the IBM zAware server, Linux systems must meet the following requirements.
  - A monitored Linux system can run in its own logical partition, or as a z/VM guest, on a supported z Systems server. The z/VM operating system must be a version that is supported for the z Systems server on which it runs. Supported z/VM versions are listed in the Preventative Service Planning (PSP) bucket for the z Systems server.
  - Supported servers are:
    - An IBM z14 (z14)
    - An IBM z13 (z13) or IBM z13s (z13s)
    - An IBM zEnterprise EC12 (zEC12) or IBM zEnterprise BC12 (zBC12)

    Although an IBM zAware partition cannot be defined or activated on a host system that has IBM Dynamic Partition Manager (DPM) mode enabled, IBM zAware can monitor Linux systems that run on Dynamic Partition Manager-enabled servers.

– The syslog daemon for the Linux monitored system must be configured to send messages over a plain TCP transport layer to port 2003. The messages must be formatted according to the Internet Engineering Task Force (IETF) syslog protocol RFC 5424, which includes 4-digit years and time zone information. Additionally, each individual message that is transmitted must be preceded by the length of the message; this convention is known as octet framing. IBM zAware supports either rsyslog or syslog-ng as the syslog daemon on the monitored system.
  - The Linux system must correctly, consistently, and uniquely identify itself in the host name portion of the syslog message. IBM zAware interprets different but equivalent host name specifications to be different systems.
  - Each Linux system must be configured to send its syslog directly to the IBM zAware server, without consolidation with other Linux syslogs.
  - When sending syslog messages, the Linux system must provide a correct time stamp, including the Coordinated Universal Time (UTC) offset.
  - For IBM zAware to produce valuable analysis results, the syslog daemon must be configured to send at least the default level of messages, or more. With more message data, IBM zAware can more quickly build a quality model and produce valuable analysis results; message filtering through the syslog daemon has the opposite effect.
– The Linux operating system must be a distribution that was tested for the z Systems server on which it runs. The distributions that support RFC 5424 include:
  - SUSE Linux Enterprise Server (SLES) 10 or later.
  - Red Hat Enterprise Linux (RHEL) 6 or later.
  - Ubuntu 16.04

  For the recommended Linux on z Systems distribution levels and z Systems servers, see the IBM tested operating systems at this URL: www.ibm.com/systems/z/os/linux/resources/testedplatforms.html. The site contains more distributions as they become available.
– The name of a Linux system cannot exceed 230 characters.

• Make sure that you use a user ID with the appropriate authority to configure the network connections of the syslog daemon on your Linux system.

About this task

This procedure provides an overview of the steps required to configure Linux monitored clients to send system log (syslog) data to IBM zAware.

Procedure

1. Configure a network connection from each Linux monitored client to the IBM zAware server. To determine whether this network configuration step is required, you can ping the IBM zAware server from the Linux system. From a Linux shell, enter the **PING** or **TRACEROUTE** command with either the IP address or host name of the IBM zAware server.
2. If necessary, update firewall settings to ensure secure communications between the IBM zAware server and its monitored clients.

   Although the configurations in Chapter 7, “Planning your IBM zAware environment,” on page 39 show both the IBM zAware server and its monitored clients within the boundary of a firewall, you can set up a configuration in which communication crosses firewall boundaries. In this case, you need to determine whether unsecured communication is an acceptable risk. If it is not an acceptable risk, you must provide your own method of securing this communication. For more information, see “Securing communication between IBM zAware and its monitored clients” on page 75.
3. Configure the Linux syslog daemon to send syslog messages to the IBM zAware server.
   a. In the syslog configuration file, add a destination statement that contains the IP address and port for the IBM zAware server. Use the following sample statements as models. For complete configuration instructions, see the product documentation for the type of syslog daemon that you are using.

Sample configuration statement for rsyslog

```
*.* @@(o)xxx.xxx.xxx.xxx:2003;RSYSLOG_SyslogProtocol23Format
```

In the sample, replace `xxx.xxx.xxx.xxx` with the IP address of the IBM zAware partition.

- `*.*` is the rsyslog.conf selector field and specifies the facility and priority of messages that are to be sent to IBM zAware.
- `@@` indicates that messages are to be forwarded using plain TCP.
- `(o)` enables octet framing.
- The value `2003` specifies the port on which IBM zAware listens for incoming Linux syslog messages.
- `RSYSLOG_SyslogProtocol23Format` is the name of the rsyslog template for generating RFC5424 formatted syslog entries.

Sample configuration statement for syslog-ng versions that support the syslog() driver

For syslog-ng, you need to add a source and a log statement, as well as a destination statement. Note that you can reuse an existing source definition from your syslog-ng configuration.

```
source src { ... };
destination zaware { syslog("xxx.xxx.xxx.xxx" transport("tcp") port(2003)); };
log {  source(src); destination(zaware); }
```

In the sample, replace *src* with the name of a messaging source, and replace `xxx.xxx.xxx.xxx` with the IP address of the IBM zAware partition.

- `source` is the syslog-ng keyword for declaring a source of messages, and the variable *src* represents the name of a messaging source.
- `destination` is the syslog-ng keyword for declaring a logging destination.
  - `zaware` is the name of the logging destination.
  - `syslog()` specifies the syslog-ng driver used to send RFC5424 formatted syslog messages to a remote destination. This driver automatically enables octet framing.
  - `transport("tcp")` indicates that messages should be forwarded using plain TCP.
  - The value `2003` specifies the port on which IBM zAware listens for incoming Linux syslog messages.
- `log` is the syslog-ng keyword for defining a log path that maps message sources to destinations.
  - `source()` specifies the previously defined message source.
  - `destination(zaware)` specifies the IBM zAware log destination.

Sample configuration statement for syslog-ng versions prior to 3.0 that do not support the syslog() driver

```
destination zaware { tcp(xxx.xxx.xxx.xxx port(2003) template("0 <$PRI>1 $ISODATE
$HOST $PROGRAM $PID - - $MSGONLY\n")); };
```

In the sample, replace `xxx.xxx.xxx.xxx` with the IP address of the IBM zAware partition.

- `destination` is the syslog-ng keyword for declaring a logging destination.
  - `zaware` is the name of the logging destination.
  - `tcp()` specifies the syslog-ng driver used to send messages to a remote destination using plain TCP.
  - The value `2003` specifies the port on which IBM zAware listens for incoming Linux syslog messages.
- `template` defines a message format that approximates RFC5424 formatting. Octet framing is not supported by the tcp() driver; the template inserts a zero (0) for the framing value. IBM zAware correctly processes messages that are formatted using this template.

b. To achieve the best analysis results, make sure that the syslog configuration file does not contain any filter conditions that might suppress messages.
c. Save the configuration updates.

4. Activate the new syslog configuration by restarting the syslog daemon, using the appropriate command for the type of syslog daemon and for the type of initialization (init) process that is in use on the Linux system.

Results

The Linux system is established as an IBM zAware monitored client. To view the status for monitored clients that are sending current data, navigate to the **System Status** tab on the **Systems** page in the IBM zAware GUI.

What to do next

Determine whether the Linux system belongs to an administrator-defined model group, or to the UNASSIGNED model group. Before IBM zAware can provide analysis results for a Linux system, that system must belong to a model group other than the UNASSIGNED model group, and IBM zAware must have successfully built a model of system behavior for that model group.

- If an IBM zAware administrator has not defined any model groups through the IBM zAware GUI, the new monitored system becomes part of the UNASSIGNED model group. For information about defining model groups, see "Creating an IBM zAware model for new Linux on z Systems monitored clients."
- If an administrator has defined one or more model groups through the GUI, IBM zAware compares the Linux system name to the membership rule for each model group, according to the specified membership evaluation order. When IBM zAware finds the first matching membership rule, it assigns the Linux system to the model group associated with that rule. If the name does not match any membership rule, the new monitored system becomes part of the UNASSIGNED model group.

To find the model group to which the Linux system belongs, complete these steps:

1. Go to the **Model Groups** tab on the **Systems** page.
2. In the header of the Model Groups table, click **Actions** to open the Actions list.
3. Select **Search Systems** to open the Search Systems window.
4. Enter the name of the monitored system in the Name field and click **OK**. IBM zAware returns you to the **Model Groups** tab and, in the Model Groups table, identifies the group by displaying a checkmark next to the model group name.
5. To verify that the system belongs to the indicated model group, click the link for that model group in the Name column of the Model Groups table. In the Model Group Details pane, check the "Known matching member systems" table for the system name.

Creating an IBM zAware model for new Linux on z Systems monitored clients

Use this procedure to define a model group for a collection of Linux monitored clients, and to build a model of system behavior for the group. Before IBM zAware can provide analysis results for a Linux system, that system must belong to a model group other than the UNASSIGNED model group, and IBM zAware must have successfully built a model of system behavior for that model group. This procedure requires the insight of skilled Linux administrators who have experience with the Linux systems that are to become monitored clients. IBM zAware administrators need their insight to appropriately define rules that group the Linux systems together for analysis.

Before you begin

- Make sure that your installation has completed the following procedures:
  - Chapter 13, "Configuring storage, security, and analytics for the IBM zAware server," on page 99.

– "Configuring Linux on z Systems monitored clients to send data to the IBM zAware server" on page 129.

- Prepare a list the monitored Linux systems to be grouped together, and the names of the model groups that you want to create. For guidelines about which systems to group together, and an explanation of early training for model groups, see "Planning to create IBM zAware models for Linux on z Systems monitored clients" on page 90.
- Steps in this procedure require you to complete tasks using the IBM zAware graphical user interface (GUI). To log in to the GUI, you need to know the URL.

  The URL includes the IP address or host name that is assigned to the IBM zAware partition:

  `https://ip_address/zAware/` **or** `https://host_name/zAware/`

  The "zAware" portion of the URL is case-sensitive.

Procedure

1. Log in to the IBM zAware GUI with a user ID that is assigned to the Administrator role.
2. For each model group to be defined, verify that the Linux systems to belong to the model group have been connected to the IBM zAware server at least once. Each system in the group does not have to be currently connected to the IBM zAware server but, to achieve the best training results, keep as many Linux systems connected and sending as much data as possible, and avoid moving systems from one model group to another.
   a. Navigate to the **System Status** tab on the **Systems** page, and check the system names listed in the IBM zAware Monitored System Data Suppliers table.
   b. To ensure that you are viewing the most recent status information, click **Refresh**.
   c. If necessary, click the Type column header to sort the table contents to list the monitored clients by type.
   d. Use the Status column and the Connect Start Time column to determine which systems are connected to (Active) or disconnected from (Inactive) the IBM zAware server.
   e. If any Linux systems are not listed in the IBM zAware Monitored System Data Suppliers table, those missing systems have not successfully connected to the IBM zAware server. In this case, complete the following steps.
      1) For each of the missing systems, check the syslog configuration and restart the syslog daemon, using the appropriate command for the type of syslog daemon and for the type of initialization (init) process that is in use on the Linux system.
      2) Click **Refresh** to update the contents of the IBM zAware Monitored System Data Suppliers table, and verify that the systems are listed.
3. If possible, send prior log data from each Linux system for IBM zAware to use as priming data for building models.
   a. On each Linux system, check the message log to make sure that the first message contains the correct host name for the Linux system. The message cannot contain "localhost".
   b. Using the following sample command syntax, upload the message logs in chronological order, sending the oldest log first.

      ```
      nc xxx.xxx.xxx.xxx 2003 < /var/log/messages
      ```

      In the sample:
      - Replace `xxx.xxx.xxx.xxx` with the IP address of the IBM zAware partition. The value `2003` specifies the port on which IBM zAware listens for incoming Linux syslog messages.
      - Note that you can replace `/var/log/messages` with the name of a raw, uncompressed file that contains archived syslog messages.
4. For each model group to be defined, create a model group definition through the **Systems** > **Model Groups** tab. Repeat these steps, as necessary, to create more than one model group.
   a. In the header of the Model Groups table, click **Actions** to open the Actions list.

b. Select **New Group** to open the Model Group Details pane. If that pane was already open, the field values are cleared and NEWGROUP is displayed in the Name field. For field descriptions, see "Fields in the Model Group Details pane" on page 214.

c. Type a new value in the Name field. The name can contain alphanumeric characters (A through Z, a through z, and 0 through 9), underscores (_), and blanks.

d. Optional: Provide a description in the Description field.

e. Required: In the "Membership rule" field, specify the text string for IBM zAware to use when assigning Linux systems to this model group. The text string is a full or partial Linux system name, which can be a fully qualified domain name, a hostname, or an IP address. The text string can contain alphanumeric characters (A through Z, a through z, and 0 through 9), periods (.), colons (:), dashes (–), and forward slashes (/). To specify a partial name or IP address, use an asterisk (*) or question mark (?) as a wildcard for any one character in the text string (for example, `LNXVM5*`).

f. Required: Specify a value for the "Membership evaluation order" field. When you specify an evaluation order, make sure that more specific membership rules are evaluated before more generic rules; otherwise, a Linux system might be assigned to the wrong group. For example, suppose that you have several systems with names that range from LNXVM50 to LNXVM59. If you define a group for them with a rule of LNXVM5*, that rule must be moved higher in the evaluation order than a more general rule, such as LNXVM*.

g. Click **Evaluate Membership** and check the results presented in the "Known matching member systems" table. If the membership rule for a different model group definition contains an error, IBM zAware prompts you to correct the error before you can successfully evaluate the membership for the model group that you are creating or modifying.

h. When you are satisfied with the membership results, click **Save** to save your changes.

IBM zAware assigns each connected Linux system to a model group. To do so, it compares the Linux system name to the membership rule for each model group, according to the specified membership evaluation order. When IBM zAware finds the first matching membership rule, it assigns the Linux system to the model group associated with that rule.

Results

For each Linux model group that you created in this procedure, IBM zAware detects how many days of data are available for building a model for the group, including in its calculation any priming data that you provided in step 3 on page 133. IBM zAware automatically schedules early training every seven days, starting from the first day for which IBM zAware has data available.

Depending on the quality of the available system data that IBM zAware uses for the first early training, the initial model might or might not be successfully built.

- If a model is successfully built, IBM zAware begins analyzing current data from the monitored systems in the group, and calculates the date for the next automatic training.
- If an automatic training attempt fails and a model is not available, IBM zAware automatically retries the training attempt the next day and, if necessary, every following day until a model is successfully built. Analysis cannot begin until a model is successfully built.

What to do next

- To track the progress of early training, go to the **Administration** > **Training Sets** > **Model Groups** tab and use any of the following techniques.
  - Click **Refresh** to update the **Model Groups** tab content.
  - Click ► to expand the **Current Training Status Details** section, then click on any value in the Current Training Status column to view details.

  When a training request completes successfully, you can view current analysis data for monitored clients, as described in Chapter 16, "Viewing and using analytical data to monitor and diagnose system behavior," on page 139.

- At any time, you can alter the training schedule by manually requesting training. From the **Administration** > **Training Sets** > **Model Groups** tab:
  1. Select the model group for which you want to request training.
  2. From the **Actions** list, select **Request Training**.
  3. Click **OK** to confirm that you want to build the model.

  If the model is successfully built, IBM zAware recalculates the scheduled date for the next automatic early training, using the date on which the model was created and either the early seven-day schedule or the configured training interval, whichever is in effect.
- If you add more Linux systems to the IBM zAware environment, use the instructions in "Managing groups of Linux monitored clients" on page 210 to make sure the new systems are assigned to the appropriate model group.

Part 5. Managing and using the IBM zAware server

Topics in this part describe the IBM zAware GUI functions that systems programmers, systems administrators, and experienced application programmers use for daily operations, which include viewing and analyzing data from monitored clients. Additional topics include management tasks for modifying the IBM zAware resources or operations.

Topics covered in this part are:

Chapter 16. Viewing and using analytical data to monitor and diagnose system behavior

Through the IBM zAware graphical user interface (GUI), you can view analytical data that indicates which system is experiencing deviations in behavior, when these anomalies occurred, and details about unusual messages and unusual message patterns. Using this information, you can take corrective action for these anomalies before they develop into more visible problems.

In the IBM zAware GUI, information in the Analysis page and Interval page can help you find and diagnose anomalies in at least three situations:
- When a problem occurs
- After a change has been made
- When a random problem occurs intermittently

The following topics provide descriptions of the Analysis page and Interval page and explain how to use the information in these pages to diagnose different types of anomalies:
- "Using the Analysis page to monitor and diagnose system behavior"
- "Using the Interval page to pinpoint the causes of system anomalies" on page 161

The topic, "Verifying planned system changes with IBM zAware" on page 170, describes how to use IBM zAware for purposes other than problem diagnosis.

Using the Analysis page to monitor and diagnose system behavior

Use the Analysis page to help you determine which monitored system is behaving abnormally, the time when the abnormal behavior occurred, and how many unique messages were issued at that time. The default presentation of the Analysis page is the Analysis Heat Map Table view, which is one of several available display formats for the analysis results that IBM zAware produces. Each Analysis view contains an Interval Anomaly Scores table in a different format; the content in the table varies, depending on the Analysis Source setting, through which you can filter the system groups or individual systems for which you want to display analysis results. To display the Analysis page, click **Analysis** in the navigation pane of the IBM zAware GUI.

To learn more details about the Analysis page views and the analytical results that they display, see the following topics:
- "The Analysis page views and when to use them"
- "Understanding how IBM zAware calculates and displays anomaly scores " on page 142
- "Interpreting anomaly scores and the system behavior that they reflect" on page 145

The Analysis page views and when to use them

All Analysis page views have similar controls through which you can modify the content that is displayed in the Interval Anomaly Scores table. The format of the display and the table toolbar controls vary according to the type of Analysis page view. The default presentation of the Analysis page is the Analysis Heat Map Table view, through which you can quickly identify the system group that contains a monitored system that is exhibiting anomalous behavior.

Analysis Heat Map Table

The Analysis Heat Map Table displays the peak anomaly scores per day and per hour for the groups or systems in the IBM zAware topology. In the Interval Anomaly Scores table that is shown in this view, each table cell represents 1 hour; the cell is colored to indicate the highest anomaly score calculated for a monitored system during that hour. The table cell contains the

anomaly score itself, which is a link through which you can change the view from group to system, or display more detailed information about a specific system. Table cells of interest are the darkest blue, gold, or orange colors. Use this display to quickly find which system groups or monitored systems have the highest anomaly scores within a day or hour.

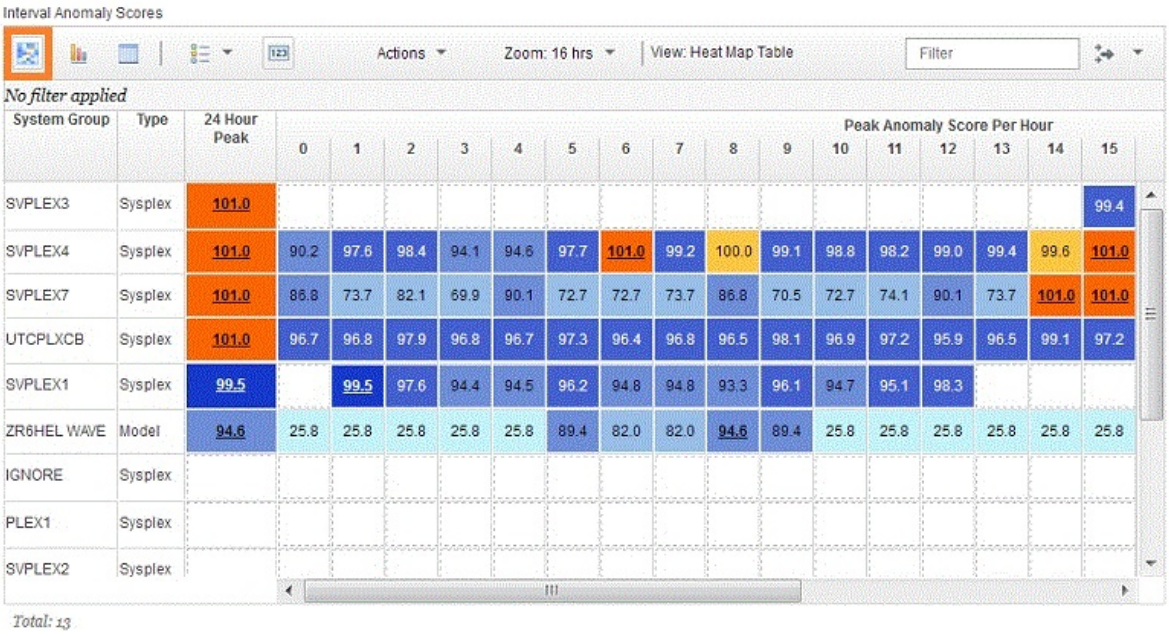
Interval Anomaly Scores

Actions | Zoom: 16 hrs | View: Heat Map Table | Filter

No filter applied

System Group	Type	24 Hour Peak	Peak Anomaly Score Per Hour														
			0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
SVPLEX3	Sysplex	101.0																99.4
SVPLEX4	Sysplex	101.0	90.2	97.6	98.4	94.1	94.6	97.7	101.0	99.2	100.0	99.1	98.8	98.2	99.0	99.4	99.6	101.0
SVPLEX7	Sysplex	101.0	86.8	73.7	82.1	69.9	90.1	72.7	72.7	73.7	86.8	70.5	72.7	74.1	90.1	73.7	101.0	101.0
UTCPLXCB	Sysplex	101.0	96.7	96.8	97.9	96.8	96.7	97.3	96.4	96.8	96.5	98.1	96.9	97.2	95.9	96.5	99.1	97.2
SVPLEX1	Sysplex	99.5		99.5	97.6	94.4	94.5	96.2	94.8	94.8	93.3	96.1	94.7	95.1	98.3			
ZR6HEL WAVE	Model	94.6	25.8	25.8	25.8	25.8	25.8	89.4	82.0	82.0	94.6	89.4	25.8	25.8	25.8	25.8	25.8	25.8
IGNORE	Sysplex																	
PLEX1	Sysplex																	
SVPLEX2	Sysplex																	

Total: 13

For more details about the controls and content displayed in the Analysis Heat Map Table, see "Analysis Heat Map Table" on page 146.

Analysis Graph

The Analysis Graph displays a bar graph for each group or system in the IBM zAware topology. In the Interval Anomaly Scores table that is shown in this view, each rectangle in a bar graph represents the anomaly score that IBM zAware recorded for a monitored system. IBM zAware records an anomaly score every 10 minutes. Rectangles of interest are the darkest blue, gold, or orange colors. Use this display to quickly find the time at which a monitored system exhibited high anomaly scores.

For more details about the controls and content displayed in the Analysis Graph, see "Analysis Graph view" on page 150.

Analysis Table

The Analysis Table view provides an accessible view of the analysis results for one or more systems in the IBM zAware topology. In the Interval Anomaly Scores table that is shown in this view, each system row contains the anomaly score that IBM zAware recorded in a given 10-minute period. Use this display as an accessible alternative to the Analysis Graph view, or to compare the interval scores of several systems in a tabular format.

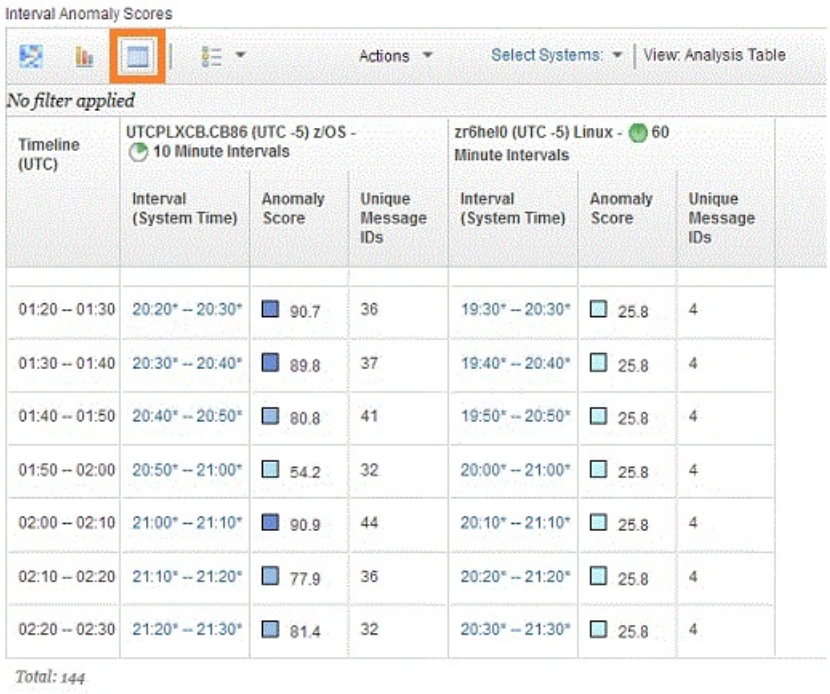

For more details about the controls and content displayed in the Analysis Table, see "Analysis Table" on page 155.

Regardless of the Analysis view that is in effect, the Interval Anomaly Scores table is blank until the following prerequisites are satisfied:

- At least one storage device has been added to the **Administration** > **Configuration** > **Data Storage** tab for IBM zAware to use for storing analysis results, system behavior models, and data from monitored clients (systems).
- At least one monitored system is connected to the IBM zAware server and is transferring current data to the server.
- IBM zAware has created a model of normal behavior for the individual z/OS system or for the Linux model group.

Chapter 11, "Planning to create IBM zAware models," on page 87 provides information about building models to accurately reflect normal behavior for an individual monitored system or for a model group. Ideally, the model represents a predictable, stable workload that generates the same artifacts when the monitored systems, subsystems, hardware, and applications are working as your installation expects them to function. Through its pattern recognition techniques and the process of building the model, which is called *training*, the IBM zAware server learns about the typical behavior of monitored systems or model groups, and their workloads. When a model exists for a monitored client or model group, and the client or group member is connected and sending current data to the IBM zAware server, the server can compare current data to the model to determine interval anomaly scores.

Understanding how IBM zAware calculates and displays anomaly scores

IBM zAware continuously analyzes the current data that connected monitored systems send to it. To produce meaningful analysis results for a monitored system, IBM zAware analyzes the most recent minutes of current data to compare to the model. These minutes are called the *analysis interval*, the length of which varies depending on the type of monitored system.

- For z/OS systems, which typically produce high-volume, consistent message traffic, IBM zAware requires 10 minutes of current data to produce an anomaly score.
- For Linux systems, which tend to produce lower volume, less consistent message traffic, IBM zAware requires 60 minutes of current data to produce an anomaly score.

An *interval anomaly score* indicates the relative difference in behavior of the monitored system, as compared to the system or group model. The IBM zAware server uses unsupervised machine learning and IBM rules to determine anomaly scores for all monitored clients.

- Through *unsupervised machine learning*, the IBM zAware server extracts and organizes message data to build a model of behavior for each z/OS monitored client or each Linux model group. This training process is repeated over time, with the frequency determined by the training interval, which enables the server to update the model with more recent system behavior.
  - Through the training process, the IBM zAware server determines which messages are issued during routine system events, such as starting a batch job or a particular subsystem. For such system events, the server identifies and recognizes groups of messages that are associated with each event. The message groups are called *clusters* and define the normal context for the messages in the cluster.

    When the server detects a specific message that is issued outside of its expected context (that is, without the other messages in the cluster), the server assigns a higher message anomaly score, which is combined with the other message anomaly scores in the interval to assign the interval anomaly score.

    IBM zAware also detects messages that are issued periodically; for example, a message that is issued every 11 minutes. This attribute affects the anomaly score when a periodic message is not issued as expected.

  - Also through the training process, IBM zAware determines the distribution of each unique message ID within a collection of intervals in the message data that is used for training. This distribution influences the interval anomaly score that the IBM zAware server displays for an interval of current data from the client.

  In summary, through the training process, the server learns about expected message patterns, and stores this information as part of the model for a specific client or group. IBM zAware uses this model data to determine interval scores when it analyzes current data that it receives from the client.
- Based on decades of experience, z/OS experts at IBM know which message IDs are likely to indicate potential problems. Message IXC101I, for example, indicates that a system is being removed from a sysplex. For a test system, this removal process could be reflected in the IBM zAware model for this system as a normal, expected behavior pattern. In the analytical data for this test system, you might expect the server to assign a low anomaly score and light blue color to any intervals that contain such a system removal, when message IXC101I is issued in context.

  However, message IXC101I might indicate a potential problem, whether or not it is issued in context. Because the removal of a system from a sysplex warrants further investigation, the IBM zAware server is programmed to assign the highest interval anomaly score to intervals in which message IXC101I is issued. IBM rules for other known messages can alter anomaly scores to a lesser degree.

  IBM zAware assigns rules only to messages that z/OS monitored systems issue.

In summary, comparison to the model, context, and IBM rules are key factors that contribute to the interval anomaly scores for systems in the Analysis page display.

The quality and quantity of the *training set*, which is the system data provided for training, has a direct effect on IBM zAware training and analysis results. For example, if the training set does not contain the minimum number of unique message IDs, IBM zAware cannot successfully create a model, and analysis results cannot be produced. In this case, IBM zAware issues a notification message that describes the training failure and suggests corrective action, such as increasing the amount of system data used for training.

Another example is the case in which the training set enables IBM zAware to successfully build a model but, because the system data lacks sufficient variety, the resulting model does not provide enough information for IBM zAware to distinguish between unusual and normal behavior. These limited models are more likely to be produced for systems that issue relatively few unique message IDs, and that do not issue any messages at all for a large number of time periods. When an interval is compared to a limited model, relatively insignificant changes in current system behavior can result in significant changes to anomaly scores. In this case, IBM zAware displays the Limited Model icon ( ) in Analysis views and on the Interval page to indicate analysis results that were produced using a limited model.

When a limited model is produced from a relatively large training set, the limited model condition is likely to be permanent. However, when a limited model is produced from a relatively small training set, more data and subsequent retraining might produce an improved model. If the model is improved, IBM zAware removes the icon for analysis results that were derived from the improved model. Any analysis results that were derived from the limited model continue to be marked with the Limited Model icon. Limited model indicators are not displayed for analysis results that were produced by IBM zAware Version 1.

When you are using high anomaly scores in an Analysis page display to find systems that are experiencing potential problems, first look at those systems for which the Limited Model icon is not displayed. When a limited model condition is indicated, carefully examine the specific interval results to determine whether the system is exhibiting unusual behavior that requires corrective action.

To display and record the results of analysis intervals, IBM zAware produces an *analysis snapshot* every 10 minutes for each monitored system. Each analysis snapshot is a point-in-time record of the anomaly score for an analysis interval. For example, the following snapshots:

- For a z/OS system, the snapshot that is recorded at 09:00 UTC represents the analysis score and number of unique messages that are issued by that system from 08:50 to 9:00 UTC. The next snapshot is taken at 09:10 UTC for the analysis interval from 09:00 to 09:10, and so on.
- For a Linux system, the snapshot that is recorded at 09:00 UTC represents the analysis score and number of unique messages that are issued by that system from 08:00 to 9:00 UTC. The next snapshot is taken at 09:10 UTC for the analysis interval from 08:10 to 09:10, and so on. Because of the 60-minute analysis interval, every snapshot for a Linux system overlaps with previous snapshots.

The analysis snapshot display varies depending on the Analysis page view.

- In the Analysis Graph view, the analysis snapshot is shown as rectangle in the Interval Anomaly Scores bar graph.

  In Figure 42, each colored rectangle in the Analysis Graph view represents an analysis snapshot. For a visual reminder that the snapshot for a Linux system represents analysis results over the prior 60 minutes, you can hover your cursor over any rectangle for a Linux system, and IBM zAware highlights the analysis interval with a transparent rectangle, as shown in the figure.

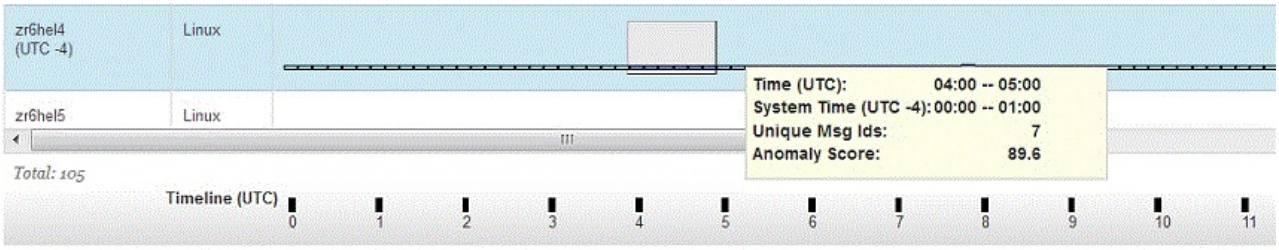

Figure 42. A sample Analysis Graph view illustrating analysis snapshots

- In the Analysis Table view, the analysis snapshot is one row in the Interval Anomaly Scores table.

  In Figure 43 on page 145, each row in the Analysis Table view represents an analysis snapshot. Note that, because the snapshot for a Linux system represents analysis results over the prior 60 minutes, the time ranges in the Interval (System Time) column represent 60-minute intervals.

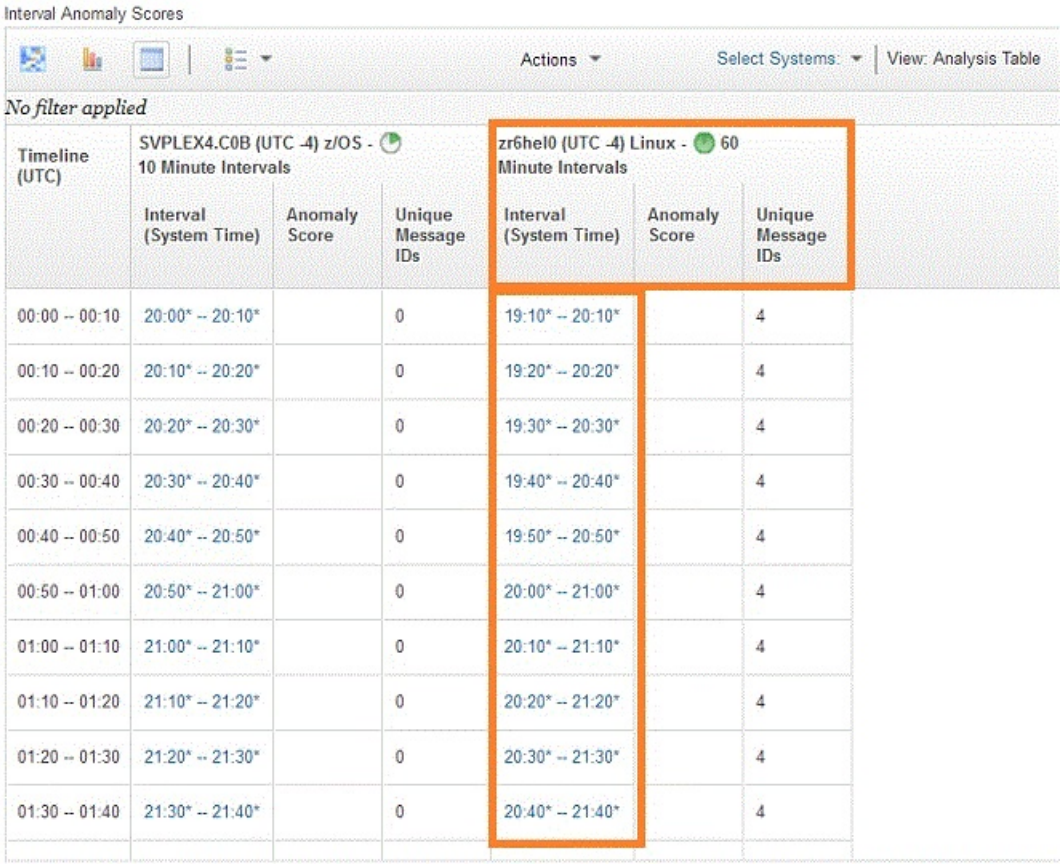
Interval Anomaly Scores

Actions | Select Systems: | View: Analysis Table

No filter applied

Timeline (UTC)	SVPLEX4.C0B (UTC -4) z/OS - 10 Minute Intervals			zr6hel0 (UTC -4) Linux - 60 Minute Intervals		
	Interval (System Time)	Anomaly Score	Unique Message IDs	Interval (System Time)	Anomaly Score	Unique Message IDs
00:00 -- 00:10	20:00* -- 20:10*		0	19:10* -- 20:10*		4
00:10 -- 00:20	20:10* -- 20:20*		0	19:20* -- 20:20*		4
00:20 -- 00:30	20:20* -- 20:30*		0	19:30* -- 20:30*		4
00:30 -- 00:40	20:30* -- 20:40*		0	19:40* -- 20:40*		4
00:40 -- 00:50	20:40* -- 20:50*		0	19:50* -- 20:50*		4
00:50 -- 01:00	20:50* -- 21:00*		0	20:00* -- 21:00*		4
01:00 -- 01:10	21:00* -- 21:10*		0	20:10* -- 21:10*		4
01:10 -- 01:20	21:10* -- 21:20*		0	20:20* -- 21:20*		4
01:20 -- 01:30	21:20* -- 21:30*		0	20:30* -- 21:30*		4
01:30 -- 01:40	21:30* -- 21:40*		0	20:40* -- 21:40*		4

Figure 43. A sample Analysis Table view highlighting 60-minute Linux analysis intervals

During the current 10-minute period, IBM zAware refreshes the analysis snapshot display with updated information every 2 minutes. At the end of the 10-minute period, the analysis snapshot becomes a persistent record of the results for the analysis interval.

Interpreting anomaly scores and the system behavior that they reflect

For each analysis interval, the IBM zAware server calculates an interval anomaly score and unique message ID count for each connected client that is sending current data. For each client, the server compares message patterns in the current data to the system or group model associated with that client.

The interval anomaly score indicates the difference in current behavior compared to the expected behavior that is reflected in the model. If the analysis interval contains messages that are relatively normal, common messages for that client, the server assigns a low score to the analysis interval and light blue color to the analysis snapshot. For example, suppose that you have configured a relatively stable test system as a IBM zAware monitored client. On this test system, various subsystems, such as Customer Information Control System (CICS®) and WebSphere® MQ, are recycled on a regular basis. This behavioral pattern is reflected in the client model that the server uses for analysis. When a current subsystem recycle completes normally, the intervals for subsystem recycling receive a low interval anomaly score and light blue color, because the pattern of messages issued during a successful recycle match an expected behavior in the model. However, if any unexpected messages are issued during a current subsystem recycle, the IBM zAware server assigns a higher interval anomaly score to those analysis intervals that contain the unexpected or unique messages. The analysis snapshots for those intervals are shown with a darker blue, gold, or orange color.

The possible interval anomaly scores are:

0 through 99.4
The analysis interval contains messages and message clusters that match or exhibit relatively insignificant differences in expected behavior, as defined in the IBM zAware model. A score of 0 is possible because the server eliminates all expected, in-context messages from its scoring calculation. A score of 0 indicates intervals that exhibit no difference in behavior when compared to the system or group model. The analysis snapshots for these analysis intervals are colored with the lightest blue shade.

Analysis intervals with scores that are greater than 0 but less than 99.5 contain some messages that are unexpected or issued out of context. Scores in this range indicate intervals that do not vary significantly from the system model. The analysis snapshots for these analysis intervals are colored with varying shades of blue.

99.5 Analysis intervals with this score contain rarely seen, unexpected, or out-of-context messages. Generally speaking, this score indicates analysis intervals with some differences from the system or group model but do not contain messages of much diagnostic value. The analysis snapshots for these analysis intervals are colored with the darkest blue shade.

99.6 - 100
Analysis intervals with this score contain rarely seen messages (these messages appear in the model only once or twice), or many messages that are unexpected or issued out of context. This score indicates analysis intervals with more differences from the system or group model; these intervals can contain messages that might help you diagnose anomalous system behavior. The analysis snapshots for these analysis intervals are the color gold.

101 Analysis intervals with this score exhibit the most significant differences from the system or group model; these intervals contain messages that merit investigation. The analysis snapshots for these analysis intervals are the color orange. IBM zAware assigns this score to analysis intervals that contain:
- Unusual or unexpected messages.
- Messages that IBM rules define as critical.
- A much higher volume of messages than expected.

Analysis Heat Map Table

The Analysis Heat Map Table displays peak anomaly scores per hour for the groups or systems that are indicated by the Analysis Source setting. Use this display to quickly find which system groups or monitored systems have the highest anomaly scores within a day or hour.

Depending on the Analysis Source setting, the Interval Anomaly Scores table shows analysis results in a "group view" of all or selected groups of monitored systems, or in a "system view" of all or selected individual monitored systems. By default, the Analysis Heat Map Table view displays analysis results for all monitored groups. The score in a table cell indicates the top anomaly score from an individual system within the group.

If IBM zAware used a limited model to calculate the anomaly score, the Limited Model icon ( ) is displayed in the System Group column. For more information about limited models, see "Understanding how IBM zAware calculates and displays anomaly scores " on page 142.

- To view detailed analysis results for systems within a group, click any one of the table cells in the row for that group. The Analysis Heat Map Table display changes to the system view, which shows one row for each system in the selected group. The top row contains the system with the highest anomaly score; the remaining system rows are sorted in descending order by anomaly score. If IBM zAware used a limited model to calculate the anomaly score, the Limited Model icon is displayed in the System column.
- To view more detailed analysis results for a particular system, click the anomaly score in the table cell for a particular hour. The Details pane opens to display a bar graph view of analysis results for the selected system, with a transparent rectangle that highlights the selected analysis snapshot.
- To return to the group view, click **Previous Group Selection** or **Change Source** to modify the display.

Analytical data might not be available for all systems for the date and time that you select for the Analysis page display. Data is not available under the following circumstances:

- The monitored system was added to the topology after the date you select for the Analysis page display.
- The monitored system is not connected to the IBM zAware server.
- The monitored system and the applications that run on it did not issue any messages.

If a monitored system is not connected during a time period, the table cell for that time period contains a dashed gray lining ( ). If a connected system did not send any message data during an analysis interval, the anomaly score for that interval is zero.

Figure 44 provides a sample Analysis Heat Map Table that displays a system view along with an expanded Details pane for a specific monitored system.

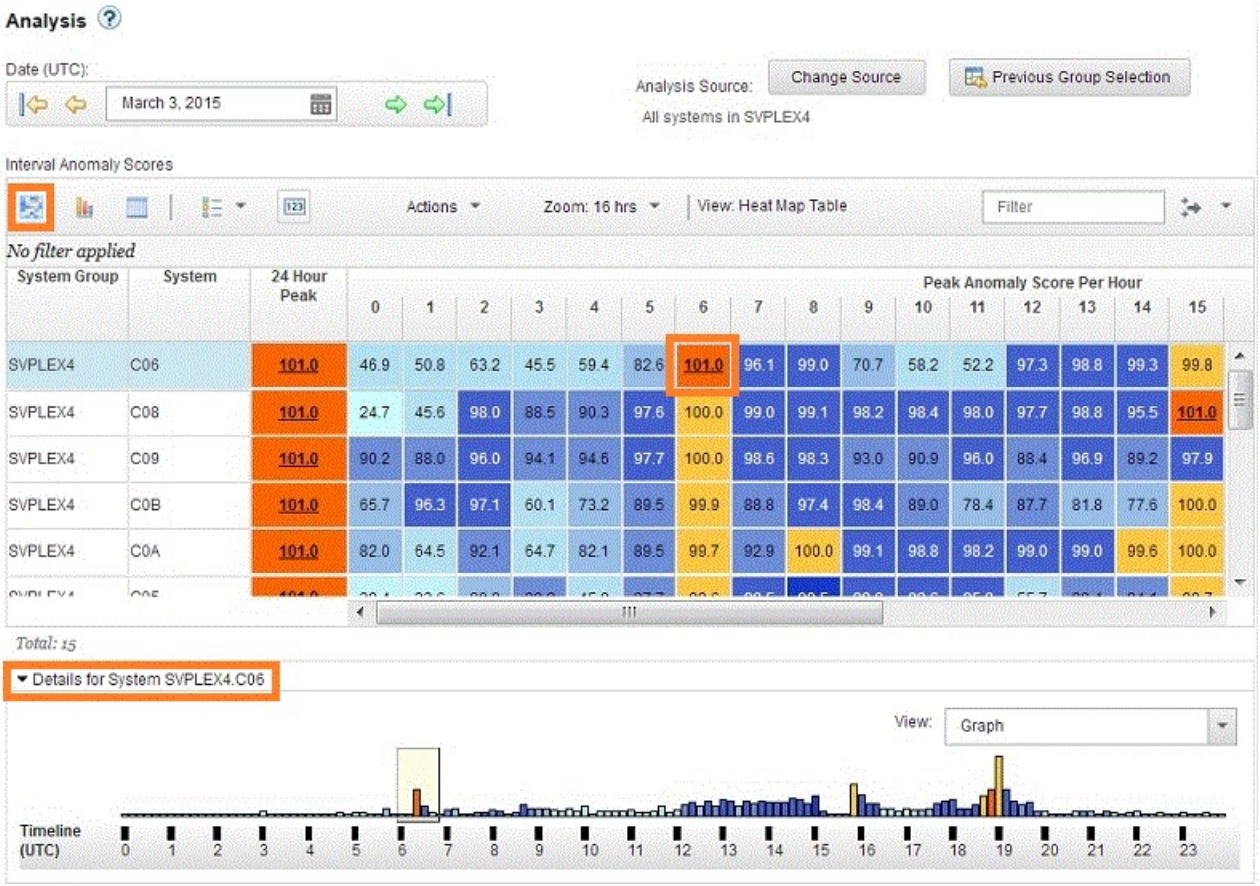

Figure 44. A sample Analysis Heat Map Table display of the system view with the Details pane for a specific system

The interface elements in the Analysis Heat Map Table view and the content of the Interval Anomaly Scores table are described in the following topics:

- "Fields and controls in the Analysis Heat Map Table"
- "Content of the Interval Anomaly Scores table in the Analysis Heat Map Table" on page 148
- "Tips for using the Analysis Heat Map Table" on page 150

Fields and controls in the Analysis Heat Map Table

All Analysis page views have similar controls through which you can modify the content that is displayed in the Interval Anomaly Scores table. The format of the display and the table toolbar controls vary according to the type of Analysis page view.

- The **Date** field displays the currently selected date. You can change the current day to any prior date for which IBM zAware has analytic data. You can type the date, select it from the calendar widget, or use the forward and back arrows to change the date.
- "Analysis source" identifies whether the display on the Analysis page shows analysis results in a "group view" of all or selected groups of monitored systems, or in a "system view" of all or selected individual monitored systems.
  - By default, the Analysis Heat Map Table view displays analysis results for all monitored groups. The score in a table cell indicates the top anomaly score from an individual system within the group.
  - Through **Change Source**, you can change the systems or groups that are listed in the display on the Analysis page.
- The Interval Anomaly Scores table contains the analytical data display. The table footer provides a total count of table entries, and the number of selected table entries. Depending on the total count, you might need to use the vertical scroll bar to view all of the table entries.

  The format of the display and the table toolbar controls vary according to the type of Analysis page view. The table toolbar indicates which view is in effect, and also contains controls through which you can change the display.
  - Click the view icons to toggle from one view to another: Analysis Heat Map Table ( ), Analysis Graph ( ), Analysis Table ( ).
  - Click the score key icon ( ) to display the interval anomaly score key; to close this display, either use the escape (Esc) key or click another area of the Analysis view. The score key correlates the color of an interval in the Analysis page display to its relative anomaly score. The interval anomaly score indicates unusual patterns of message IDs within an interval, as compared to the model of normal system behavior for a specific system.
  - Click the analysis numeric values icon ( ) to hide or show the anomaly score in each table cell. By default, anomaly scores are displayed in the Heat Map Table view.
  - The **Actions** list provides alternatives to using the toolbar icons to change the view and to display the interval anomaly score key. On the Heat Map Table only, the Actions list also provides an alternative to the analysis numeric values icon. Administrators only can select the **Manage System Groups** action to go to the **Systems** > **Model Groups** tab and edit Linux model groups.
  - Through **Zoom**, you can control how many hours of the day are displayed in the Analysis page.
  - The View field indicates which view is in effect for the current Analysis page display.

Content of the Interval Anomaly Scores table in the Analysis Heat Map Table

Table 23 provides a description of the items that are displayed in the Analysis Heat Map Table view of the Analysis page.

Table 23. Elements of the Interval Anomaly Scores table in the Analysis Heat Map Table

Table element	Description
System Group	Indicates the name of a z/OS sysplex or Linux model group in the IBM zAware topology.
Type	Indicates the type of system group. This column is displayed only in the group view of the Analysis Heat Map Table. Valid values are: **Sysplex** Indicates a collection of z/OS systems. **Model Group** Indicates an administrator-defined collection of Linux systems.
System	Provides the name of the monitored system. This column is displayed only in the system view of the Analysis Heat Map Table.

Table 23. Elements of the Interval Anomaly Scores table in the Analysis Heat Map Table (continued)

Table element	Description
24 Hour Peak	Indicates the highest anomaly score calculated for an individual system within the 24 hours that constitute the selected date. The table cell contains the anomaly score itself, which is a link through which you can change the view from group to system, or display more detailed information about a specific system. For a list of anomaly score ranges and their meaning, see "Interpreting anomaly scores and the system behavior that they reflect" on page 145.
Peak Anomaly Score Per Hour	Indicates the highest anomaly score calculated for a monitored system during each hour of the selected date. Each subcolumn indicates the start of each hour of the day according to the 24-hour clock, in Coordinated Universal Time (UTC). The table cell contains the anomaly score itself, which is a link through which you can change the view from group to system, or display more detailed information about a specific system. In either the group view or system view, the anomaly score displayed in the table cell is the highest anomaly score calculated for only one monitored system. • For a z/OS monitored client, the anomaly score in the table cell that is displayed for any hour is the highest anomaly score in a single 10-minute analysis interval within that hour. • For a Linux monitored client, the anomaly score displayed for any hour is the highest anomaly score that an individual system in the member group that is received during that hour and the prior 50 minutes. For example, the table cell for the sixth hour of the day indicates the highest score that is recorded from 5:10 AM to 7:00 AM. This time overlap occurs because the analysis interval for Linux systems is 60 minutes, and because IBM zAware records an analysis snapshot every 10 minutes. So the first 10-minute snapshot in the sixth hour represents the anomaly score calculated for the analysis interval from 5:10 AM to 6:10 AM. For an illustration, see Figure 45 on page 150. For a list of anomaly score ranges and their meaning, see "Interpreting anomaly scores and the system behavior that they reflect" on page 145.

Figure 45 on page 150 shows a sample Heat Map and Analysis Graph to that illustrates the overlapping analysis intervals for Linux systems.

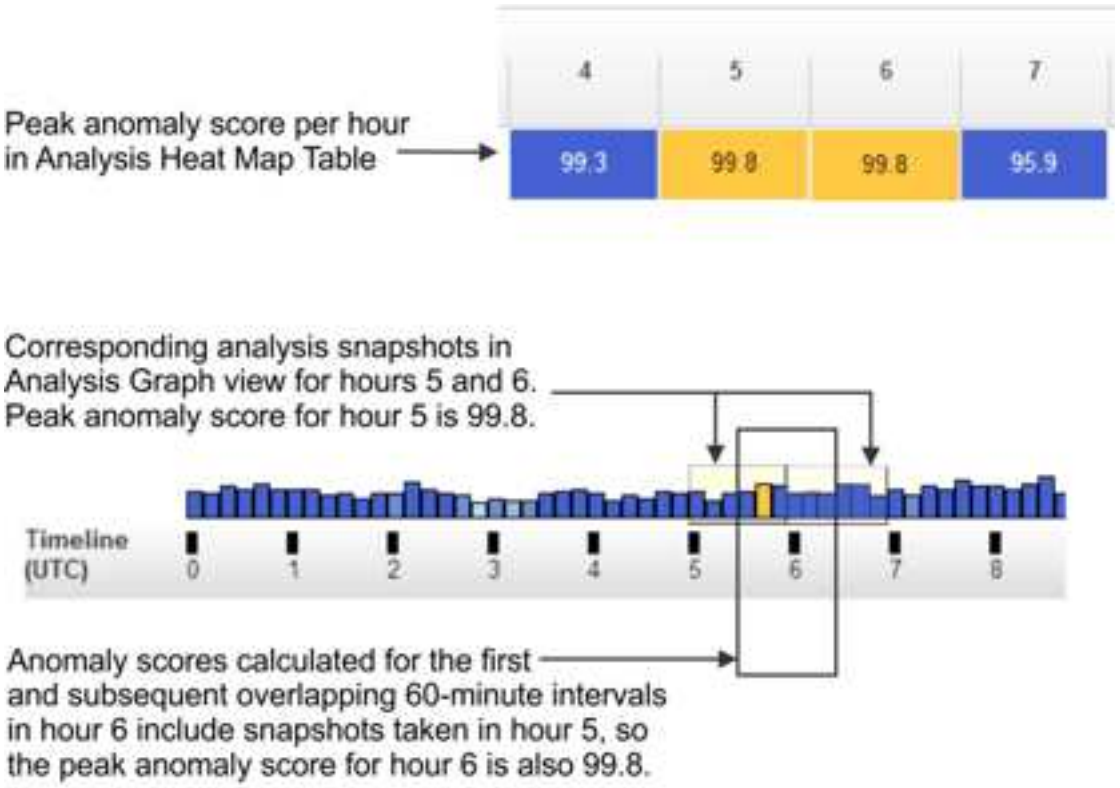

Figure 45. Peak hourly scores and their corresponding analysis intervals and snapshots for Linux systems

Tips for using the Analysis Heat Map Table

- When you use the Date field controls to scroll through days of analysis results, the displayed order of groups or systems likely changes because the sorting default behavior is based on the peak anomaly score. To track a particular group or system by date and preserve its table row position, first click the System Group or System column heading to change the sort order, and then use the Date controls to scroll back or forward by date.
- The "Analysis source" setting controls the display on all Analysis page views, and is preserved within a browser session, even when you go to other pages in the GUI. For example, when you drill down from a group view to a system view in the Analysis Heat Map Table, and switch to the Analysis Graph view, the system view is preserved until you use **Change Source** to select other systems or groups.
- When an interval is compared to a limited model, relatively insignificant changes in current system behavior can result in significant changes to anomaly scores. When you are using high anomaly scores in an Analysis page display to find systems that are experiencing potential problems, first look at those systems for which the Limited Model icon is not displayed. When a limited model condition is indicated, carefully examine the specific interval results to determine whether the system is exhibiting unusual behavior that requires corrective action.

Analysis Graph view

The Analysis Graph displays analysis results for the groups or systems that are indicated by the Analysis Source setting. Each bar graph in this view provides a clear visual indication of the time when abnormal behavior occurred and how many unique messages were issued then. Use this display to quickly find the time at which a monitored system exhibited high anomaly scores.

By default, the Analysis Graph view displays analysis results for all monitored groups, with one bar graph for each system. Each rectangle in a bar graph represents an analysis snapshot, which is a point-in-time record of the anomaly score for an analysis interval. The rectangle color indicates the

anomaly score, and its height is an approximate illustration of the number of unique messages that are issued during the analysis interval. Taller rectangles represent analysis intervals in which a larger number of unique messages were issued.

- Use **Change Source** to modify the display.
- The bar graph in each table row contains the analysis results for a specific monitored system. To view more detailed analysis results, click the rectangle to open the Interval page.
- If IBM zAware used a limited model to calculate the anomaly score, the Limited Model icon ( ) is displayed in the System column. For more information about limited models, see "Understanding how IBM zAware calculates and displays anomaly scores " on page 142.

Analytical data might not be available for all systems for the date and time that you select for the Analysis page display. Data is not available under the following circumstances:

- The monitored system was added to the topology after the date you select for the Analysis page display.
- The monitored system is not connected to the IBM zAware server.
- The monitored system and the applications that run on it did not issue any messages.

If a monitored system is not connected during a time period, the rectangle for that time period is outlined with a dashed gray line.

Figure 46 provides a sample illustration of the Analysis Graph view.

Figure 46. A sample Analysis Graph display of the system view

The interface elements in the Analysis Graph view and the content of the Interval Anomaly Scores table are described in the following topics:

- "Fields and controls in the Analysis Graph" on page 152
- "Content of the Interval Anomaly Scores table in the Analysis Graph" on page 152
- "Tips for using the Analysis Graph" on page 154

Fields and controls in the Analysis Graph

All Analysis page views have similar controls through which you can modify the content that is displayed in the Interval Anomaly Scores table. The format of the display and the table toolbar controls vary according to the type of Analysis page view.

- The **Date** field displays the currently selected date. You can change the current day to any prior date for which IBM zAware has analytic data. You can type the date, select it from the calendar widget, or use the forward and back arrows to change the date.
- "Analysis source" identifies whether the display on the Analysis page shows analysis results in a "group view" of all or selected groups of monitored systems, or in a "system view" of all or selected individual monitored systems.
  - The default analysis source for the Analysis Graph view is all monitored groups.
  - Through **Change Source**, you can change the systems or groups that are listed in the display on the Analysis page.
- The Interval Anomaly Scores table contains the analytical data display. The format of the display and the table toolbar controls vary according to the type of Analysis page view. The table toolbar indicates which view is in effect, and also contains controls through which you can change the display.
  - Click the view icons to toggle from one view to another: Analysis Heat Map Table ( ), Analysis Graph ( ), Analysis Table ( ).
  - Click the score key icon ( ) to display the interval anomaly score key; to close this display, either use the escape (Esc) key or click another area of the Analysis view. The score key correlates the color of an interval in the Analysis page display to its relative anomaly score. The interval anomaly score indicates unusual patterns of message IDs within an interval, as compared to the model of normal system behavior for a specific system.
  - The **Actions** list provides alternatives to using the toolbar icons to change the view and to display the interval anomaly score key. Administrators only can select the **Manage System Groups** action to go to the **Systems** > **Model Groups** tab and edit Linux model groups.
  - Through **Zoom**, you can control how many hours of the day are displayed in the Analysis page.
  - The View field indicates which view is in effect for the current Analysis page display.

Content of the Interval Anomaly Scores table in the Analysis Graph

Table 24 provides a description of the items that are displayed in the Analysis Graph view.

Table 24. Elements of the Interval Anomaly Scores table in the Analysis Graph

Item	Description
System Group	Provides the name of a z/OS sysplex or Linux model group in the IBM zAware topology.
Type	Indicates the type of system group. This column is displayed only in the group view of the Analysis Graph. Valid values are: **Sysplex** Indicates a collection of z/OS systems. **Model Group** Indicates an administrator-defined collection of Linux systems.
System	Provides the name of the monitored system. This column is displayed only in the system view of the Analysis Graph.

Table 24. Elements of the Interval Anomaly Scores table in the Analysis Graph (continued)

Item	Description
Anomaly Scores	Contains the bar graph display of analysis results for each monitored system that is identified in the "Analysis source" field. Each rectangle in a bar graph represents an analysis snapshot, which is a point-in-time record of the anomaly score for an analysis interval. The rectangle color indicates the anomaly score, and its height is an approximate illustration of the number of unique messages that are issued during the analysis interval. Taller rectangles represent analysis intervals in which a larger number of unique messages were issued. • For a quick view of analysis interval results, hover your cursor over any rectangle to display information about the analysis snapshot that the rectangle represents. This information is described in "Details for each analysis snapshot." • To view more detailed analysis results, click the rectangle to open the Interval page. For details about the Interval page, see "Using the Interval page to pinpoint the causes of system anomalies" on page 161.
Timeline (UTC)	Marks the start of each hour of the day according to the 24-hour clock, in Coordinated Universal Time (UTC).

Details for each analysis snapshot

For a quick view of analysis interval results, hover your cursor over any rectangle to display the following information about the analysis snapshot that the rectangle represents.

Time (UTC)
Indicates the start and end times for the analysis interval that the selected analysis snapshot represents, in Coordinated Universal Time (UTC), using the 24-hour clock.

System Time (UTC with offset)
Indicates the start and end times for the analysis interval that the selected analysis snapshot represents, in Coordinated Universal Time (UTC), with the offset, if any, for the local system time. Analysis interval times are marked with an asterisk (*) under two conditions:
- When an analysis interval extends into the day before the selected date.
- When the local time for the monitored system has a UTC offset that causes the Analysis display to contain analysis intervals that occurred on the day before or after the selected date.

Interval size (minutes)
The length of the analysis interval, which varies depending on the type of monitored system.
- For z/OS systems, the analysis interval is 10 minutes.
- For Linux systems, the analysis interval is 60 minutes.

Unique Message IDs
Provides the number of unique message identifiers (IDs) that were issued during the analysis interval.

Anomaly Score
Provides the anomaly score for the analysis interval that the selected analysis snapshot represents. For a description of each range of interval anomaly scores, see "Interpreting anomaly scores and the system behavior that they reflect" on page 145.

Limited Model
Indicates whether IBM zAware used a limited model to calculate the anomaly score for the interval. Valid values are Yes, No, or Unknown, which indicates temporary conditions under which IBM zAware cannot determine whether the model is limited. Limited model indicators are not displayed for analysis results that were produced by IBM zAware Version 1.

Tips for using the Analysis Graph

The bar graphs in the Interval Anomaly Scores table provide a clear visual indication of anomalous behavior through the height or color of interval rectangles. To find intervals that merit investigation, use the following tips.

- Color is the primary indicator of anomalous behavior, so look for dark blue, gold, or orange rectangles. The height of the rectangle is an additional indicator of unusual activity, with taller rectangles representing analysis intervals in which a larger number of unique messages were issued. For Linux systems, which typically have light message traffic, IBM zAware adjusts the rectangle height to improve visibility in the display. For any type of monitored system, hover your cursor over the rectangle to view the exact number of unique message IDs that were issued during the analysis interval.

  Depending on the initial display in the Analysis Graph view, you might need to use various controls to find these rectangles.
  - Use the scroll bar, if one is displayed, to view all systems in the list. The table footer provides a total count of table entries, and the number of selected table entries. Depending on the total count, you might need to use the vertical scroll bar to view all of the table entries.
  - Use **Change Source** to modify which systems are presented in the display. The "Analysis source" setting controls the display on all Analysis page views, and is preserved within a browser session, even when you go to other pages in the GUI. For example, when you drill down from a group view to a system view in the Analysis Heat Map Table, and switch to the Analysis Graph view, the system view is preserved until you use **Change Source** to select other systems or groups.

    For a description of the Change Analysis Source window, see "Change Analysis Source window" on page 159.
  - Use the date and time controls to focus on a specific time period. Note that analysis interval times are marked with an asterisk (*) under two conditions:
    - When an analysis interval extends into the day before the selected date.
    - When the local time for the monitored system has a UTC offset that causes the Analysis display to contain analysis intervals that occurred on the day before or after the selected date.

  Depending on the number of connected systems at your installation, the IBM zAware GUI might require some time to render the initial Analysis page display. For better performance, use **Change Source** to filter the display by sysplex or a selected set of systems. This filter setting remains in effect during the browser session or until you use **Change Source** to modify it.
- When an interval is compared to a limited model, relatively insignificant changes in current system behavior can result in significant changes to anomaly scores. When you are using high anomaly scores in an Analysis page display to find systems that are experiencing potential problems, first look at those systems for which the Limited Model icon is not displayed. When a limited model condition is indicated, carefully examine the specific interval results to determine whether the system is exhibiting unusual behavior that requires corrective action.
- When you find intervals of diagnostic interest, position your cursor over each rectangle to display the exact time of the interval, the number of unique message IDs issued during that interval, and the exact interval anomaly score. Investigate intervals with the following characteristics:
  - An interval anomaly score of 101. These intervals are represented by orange rectangles.
  - Multiple intervals with an interval anomaly score of 101.
  - A significant change in interval anomaly score from one interval to the next.
  - One of the following height and color combinations:
    - Short and dark-colored rectangles are interesting.
    - Tall and light rectangles can be interesting.
  - Intervals with no data are interesting.

  Also check prior intervals, which might yield additional information. Depending on the type of workload your installation runs on the monitored system, checking the system behavior on prior days (for example, the day before or the same day last week) might be helpful as well.

- Although the Analysis Graph is useful for finding unexpected problems or reported incidents, you also can use it to verify behavior after a planned change. For additional information, see "Verifying planned system changes with IBM zAware" on page 170.

Analysis Table

The Analysis Table provides an accessible view of the analysis results for one or more systems. Use this display as an accessible alternative to the Analysis Graph view, or to compare the interval scores of several systems in a tabular format.

By default, the Interval Anomaly Scores table in this view contains the analysis results for the first system that is displayed in the Analysis Graph view. Each row in the table display contains an analysis snapshot that IBM zAware records for the monitored system at a given time. For each analysis snapshot, the table row contains the anomaly score and the number of unique message IDs for the analysis interval.

- If IBM zAware used a limited model to calculate the anomaly score, the Limited Model icon ( ) is displayed in the monitored system header in the Interval Anomaly Scores table, which also contains a Limited Model column. For more information about limited models, see "Understanding how IBM zAware calculates and displays anomaly scores " on page 142.
- To view more detailed analysis results for the system at a given time, click an interval link in the Interval (System Time) column to open the Interval page.

Analytical data might not be available for all systems for the date and time that you select for the Analysis page display. Data is not available under the following circumstances:

- The monitored system was added to the topology after the date you select for the Analysis page display.
- The monitored system is not connected to the IBM zAware server.
- The monitored system and the applications that run on it did not issue any messages.

If a monitored system is not connected during a time period, the table row for that time period contains zero values in the Anomaly Score and Unique Message IDs columns. The Anomaly Score column also contains a gray square with a dashed outline.

Figure 47 on page 156 provides a sample illustration of the Analysis Table view.

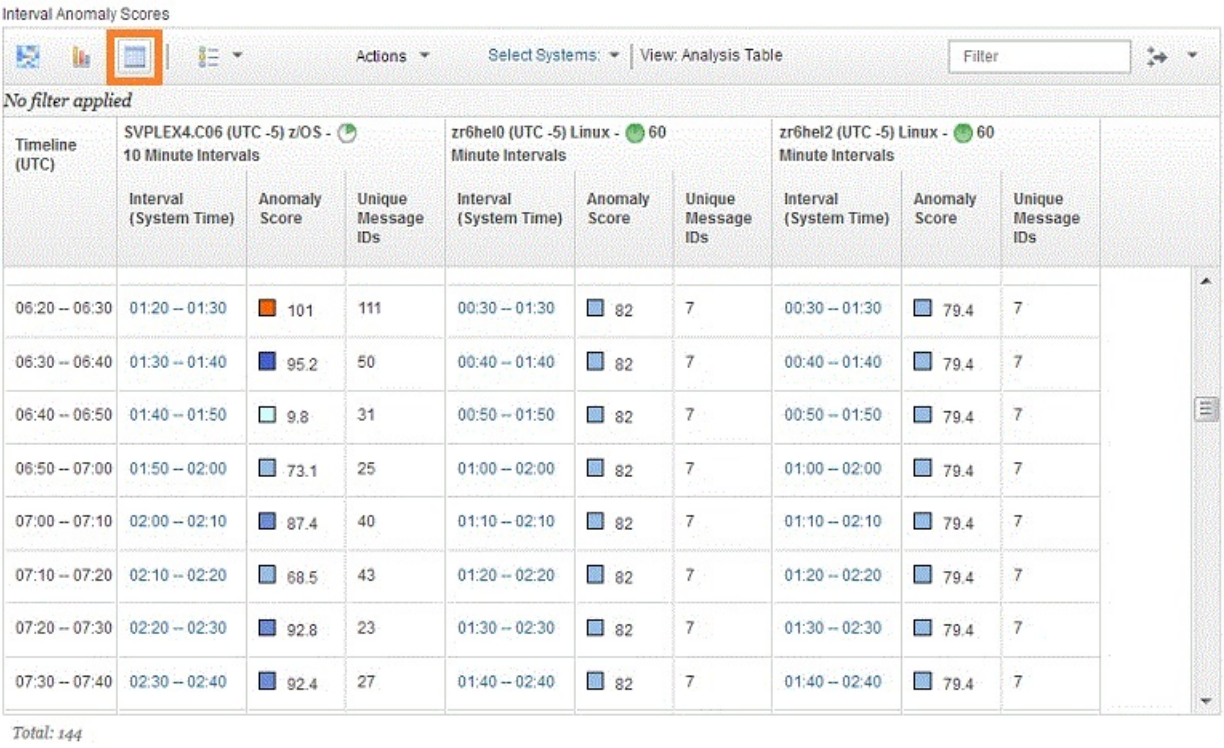

Figure 47. A sample Analysis Table view

The interface elements in the Analysis Table view and the content of the Interval Anomaly Scores table are described in the following topics:

- "Fields and controls in the Analysis Table"
- "Content of the Interval Anomaly Scores table in the Analysis Table" on page 157

Fields and controls in the Analysis Table

All Analysis page views have similar controls through which you can modify the content that is displayed in the Interval Anomaly Scores table. The format of the display and the table toolbar controls vary according to the type of Analysis page view.

- The **Date** field displays the currently selected date. You can change the current day to any prior date for which IBM zAware has analytic data. You can type the date, select it from the calendar widget, or use the forward and back arrows to change the date.
- "Analysis source" identifies whether the display on the Analysis page shows analysis results in a "group view" of all or selected groups of monitored systems, or in a "system view" of all or selected individual monitored systems. In the Analysis Table view, you can display analysis data only for one or more individual systems.
  - The default analysis source for the Analysis Table view is the first system that is listed on the Analysis Graph view.
  - Through **Change Source**, you can change the systems or groups that are listed in the display on the Analysis page.
- The Interval Anomaly Scores table contains the analytical data display.

  The format of the display and the table toolbar controls vary according to the type of Analysis page view. The table toolbar indicates which view is in effect, and also contains controls through which you can change the display.
  - Click the view icons to toggle from one view to another: Analysis Heat Map Table ( ), Analysis Graph ( ), Analysis Table ( )

– Click the score key icon ( ) to display the interval anomaly score key; to close this display, either use the escape (Esc) key or click another area of the Analysis view. The score key correlates the color of an interval in the Analysis page display to its relative anomaly score. The interval anomaly score indicates unusual patterns of message IDs within an interval, as compared to the model of normal system behavior for a specific system.
– The **Actions** list provides alternatives to using the toolbar icons to change the view and to display the interval anomaly score key. Administrators only can select the **Manage System Groups** action to go to the **Systems** > **Model Groups** tab and edit Linux model groups.
– The View field indicates which view is in effect for the current Analysis page display.
– Through **Select Systems**, you can add more systems to the default display.
  - z/OS system names are presented in the format *group_name*.*system_name*; for example: SYSPLEX1.SYSA
  - Linux system names are host names only; for example: LNXL3778

  The Analysis Source value determines which system names are in the **Select Systems** list. Use **Change Source** to modify the list.

Content of the Interval Anomaly Scores table in the Analysis Table

Table 25 provides a description of the items that are displayed in the Analysis Table view.

Table 25. Elements of the Interval Anomaly Scores table in the Analysis Table

Table element	Description
Time Line (UTC)	The Time Line (UTC) column contains the start and end of each 10-minute period for the selected date in the Date field, in Coordinated Universal Time (UTC), using the 24-hour clock.
Monitored system header	For each monitored system displayed in the table, a header spans three columns of analysis results for the system. This monitored systems header contains the following information: • The system name, which is presented in the format *group_name*.*system_name*; for example: SYSPLEX1.SYSA or LNXVM1.LNXL3778 • The UTC offset for the local time in effect for the central processor complex (CPC) on which the monitored system runs. • The system type, which is one of the following values: z/OS or Linux. • An icon and text that indicates the analysis interval length. An analysis interval is the length of time for which IBM zAware must collect current data from a monitored system to produce meaningful analysis results. – For z/OS systems, which typically produce high-volume, consistent message traffic, IBM zAware requires 10 minutes of current data to produce an anomaly score. – For Linux systems, which tend to produce lower volume, less consistent message traffic, IBM zAware requires 60 minutes of current data to produce an anomaly score.

Table 25. Elements of the Interval Anomaly Scores table in the Analysis Table (continued)

Table element	Description
Interval (System Time)	The Interval (System Time) column contains the start and end times of each analysis interval for the monitored system. Each end time corresponds to the end time of a 10-minute period listed in the Time Line (UTC) column. The system time is the local time in effect for the central processor complex (CPC) on which the monitored system runs, using the 24-hour clock. Analysis interval times are marked with an asterisk (*) under two conditions: • When an analysis interval extends into the day before the selected date. • When the local time for the monitored system has a UTC offset that causes the Analysis display to contain analysis intervals that occurred on the day before or after the selected date. • For z/OS systems, the analysis interval is 10 minutes. • For Linux systems, the analysis interval is 60 minutes. To view more detailed analysis results for the system at a given time, click an interval link in the Interval (System Time) column to open the Interval page.
Anomaly Score	The Anomaly Score column contains one colored square for each 10-minute scoring interval, along with the anomaly score for the analysis interval. Interval anomaly scores range 0 - 101. A score of 0 means that the IBM zAware server detected no difference from the system model during the analysis interval; the analysis interval contains expected messages and message clusters that match the model. A score of 101 means that the server detected significant differences from the system model; the analysis interval contains one or more new messages or many messages that are issued out of context. The server also assigns a score of 101 for the following conditions: • The analysis interval contains one or more messages that IBM rules define as a critical message. • The analysis interval receives an anomaly score that is greater than twice the largest analysis interval score in the model. Intervals with lower scores are colored with varying shades of blue. Intervals with higher scores are colored with the darkest shade of blue, or gold, or orange. For a description of each range of interval anomaly scores, see “Interpreting anomaly scores and the system behavior that they reflect” on page 145.
Unique Message IDs	Contains number of unique message identifiers (IDs) that were issued during the analysis interval.
Limited Model	Indicates whether IBM zAware used a limited model to calculate the anomaly score for the interval. Valid values are Yes, No, or Unknown, which indicates temporary conditions under which IBM zAware cannot determine whether the model is limited. Limited model indicators are not displayed for analysis results that were produced by IBM zAware Version 1.

Tips for using the Analysis Table

- The table footer provides a total count of table entries, and the number of selected table entries. Depending on the total count, you might need to use the vertical scroll bar to view all of the table entries.
- Note that analysis interval times are marked with an asterisk (*) under two conditions:
  - When an analysis interval extends into the day before the selected date.
  - When the local time for the monitored system has a UTC offset that causes the Analysis display to contain analysis intervals that occurred on the day before or after the selected date.
- To customize which systems are shown in the table, use the **Select Systems** list. Selected systems are presented in alphabetical order from left to right. You might have to use the scroll bar to view more than two systems at a time.

- The Analysis Source value determines which system names are in the **Select Systems** list. Use **Change Source** to modify the list. For a description of the Change Analysis Source window, see "Change Analysis Source window."
- The "Analysis source" setting controls the display on all Analysis page views, and is preserved within a browser session, even when you go to other pages in the GUI. For example, when you drill down from a group view to a system view in the Analysis Heat Map Table, and switch to the Analysis Graph view, the system view is preserved until you use **Change Source** to select other systems or groups.
- When an interval is compared to a limited model, relatively insignificant changes in current system behavior can result in significant changes to anomaly scores. When you are using high anomaly scores in an Analysis page display to find systems that are experiencing potential problems, first look at those systems for which the Limited Model icon is not displayed. When a limited model condition is indicated, carefully examine the specific interval results to determine whether the system is exhibiting unusual behavior that requires corrective action.

Change Analysis Source window

You can use the Change Analysis Source window provided in IBM zAware to change the system groups or individual systems listed in the display on the Analysis page. The currently displayed system groups or systems are listed under the "Analysis source" field.

"Analysis source" identifies whether the display on the Analysis page shows analysis results in a "group view" of all or selected groups of monitored systems, or in a "system view" of all or selected individual monitored systems.

This field can have one of the following values:

All Monitored Systems
: Indicates that analytic data is displayed for all the systems that IBM zAware is monitoring. This system view is the default display, and consists of z/OS systems, Linux systems, or both, depending on the types of systems that IBM zAware is monitoring.

All Monitored Groups
: Indicates that analytic data is displayed for all groups of monitored systems. This group view consists of z/OS sysplexes, Linux model groups, or both, depending on the types of systems that IBM zAware is monitoring.

List of system names
: Indicates that analytic data is displayed for only the systems in the list.

List of system group names
: Indicates that analytic data is displayed for only the groups in the list.

To select different systems or groups to include in the display, click **Change Source**.

To modify the current analysis source, click **Change Source** on the Analysis page to open the Change Analysis Source window, as illustrated in Figure 48 on page 160.

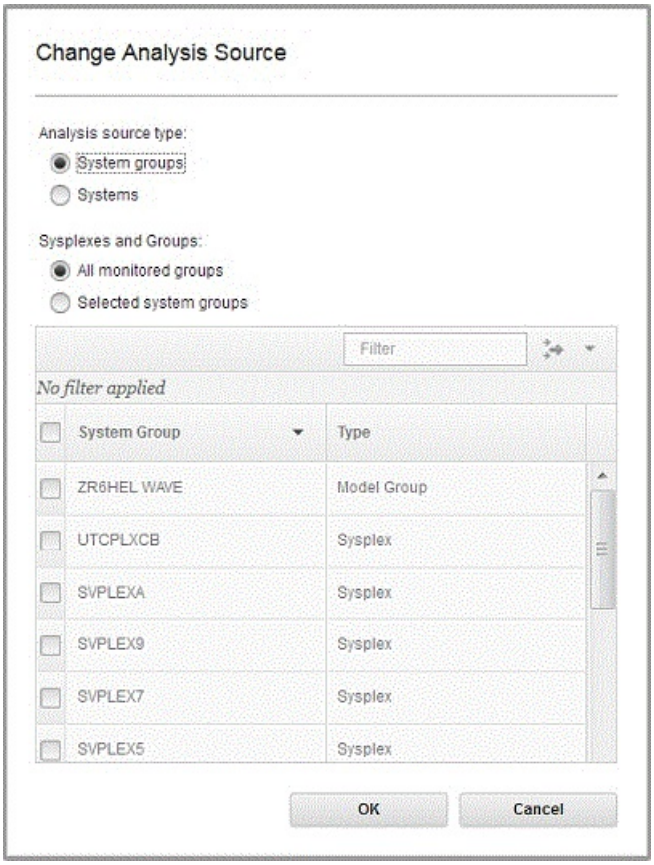

Figure 48. A sample Change Analysis Source window

For a description of the fields that are displayed in the Change Analysis Source window, see Table 26.

Fields in the Change Analysis Source window

Table 26 describes the fields that are displayed in the Change Analysis Source window.

Table 26. Fields in the Change Analysis Source window

Field	Description
Analysis source type	Controls the content that is displayed in the Change Analysis Source window. Select one of the following options: **System groups** Displays the **Sysplexes and Groups** list, and populates the list with the names of the system groups that currently exist in the topology. This option is selected by default; use it to select a group view in the Analysis page display. **Systems** Displays the **Systems** list, and populates the list with the names of all the monitored systems that currently exist in the topology. Use this option to select a system view in the Analysis page display.

Table 26. Fields in the Change Analysis Source window (continued)

Field	Description
Sysplexes and Groups	Lists the name of all system groups that currently exist in the topology. Select one of the following options: **All monitored system groups** Select this option to view analysis results for all system groups. **Selected system groups** Select one or more system groups in the list. The type of system group is indicated in the list: • Sysplex for a group of z/OS systems. • Model Group for a group of Linux systems. If necessary, use the Filter field to limit the list of group names.
Systems	Lists the name of all individual systems that currently exist in the topology. Select one of the following options: **All monitored systems** Select this option to view analysis results for all systems. **Selected systems** Select one or more systems in the list. The type of system and the system group to which each system belongs are also indicated in the list. If necessary, use the Filter field to limit the list of system names.
Previous Group Selection	Populates the **Sysplexes and Groups** list with the system groups in the previous Analysis Source setting for the Analysis Heat Map Table only. This control is displayed only when you initially used the Analysis Heat Map Table view to display system groups but drilled down to a system view. **Previous Group Selection** provides a quick method to return to the group view.

Using the Interval page to pinpoint the causes of system anomalies

You can use the Interval page to diagnose the problem that is causing a particular IBM zAware monitored system to behave abnormally during a specific analysis interval. While the Analysis page provides a clear visual indication of systems that are experiencing anomalous behavior, the Interval page helps you pinpoint and diagnose the causes of this behavior.

Through the diagnostic details in the Interval page display, you can answer these questions:
- What messages are unusual?
- How often did the unusual message get issued?
- Are messages issued in context within an expected pattern?
- Is a specific component or application issuing unusual messages?
- When did the message ID first appear?
- Did the message appear when expected?
- Did the message occur at an expected predictable time (for example, every 81 seconds)?

You can access the Interval page only by clicking on an analysis snapshot in one of the Analysis page views. An analysis snapshot is displayed in different formats, depending on the view that is in effect.
- In the system view of the Analysis Heat Map Table, click the anomaly score in the table cell for a particular hour. The Details pane opens to display a bar graph view of analysis results for the selected system, with a transparent rectangle that highlights the selected analysis snapshot. Click the highlighted bar graph rectangle to open the Interval page.
- In the system view of the Analysis Graph, click any rectangle in any bar graph to open the Interval page.
- In the Analysis Table, click on an interval link in the Interval (System Time) column to open the Interval page.

The Interval page contains a header section that provides details about the analysis interval itself, and the Messages table, which provides diagnostic details about the messages issued during the analysis interval. By default, only the most frequently used columns in the Messages table are shown in the display. To view all columns in the Messages table, click the Details column icon ( ) in the Messages table toolbar. The display changes to show more diagnostic columns on the right side of the Messages table. These additional columns are not displayed for analysis results that were generated before IBM zAware Version 2.0.

If a monitored system is not connected or has no message data to send during an analysis interval, the Messages table on the Interval page is blank except for a message that indicates why analytical data is not available.

Figure 49 shows a sample Interval page for a z/OS monitored system, with the initial subset of columns in the Messages table.

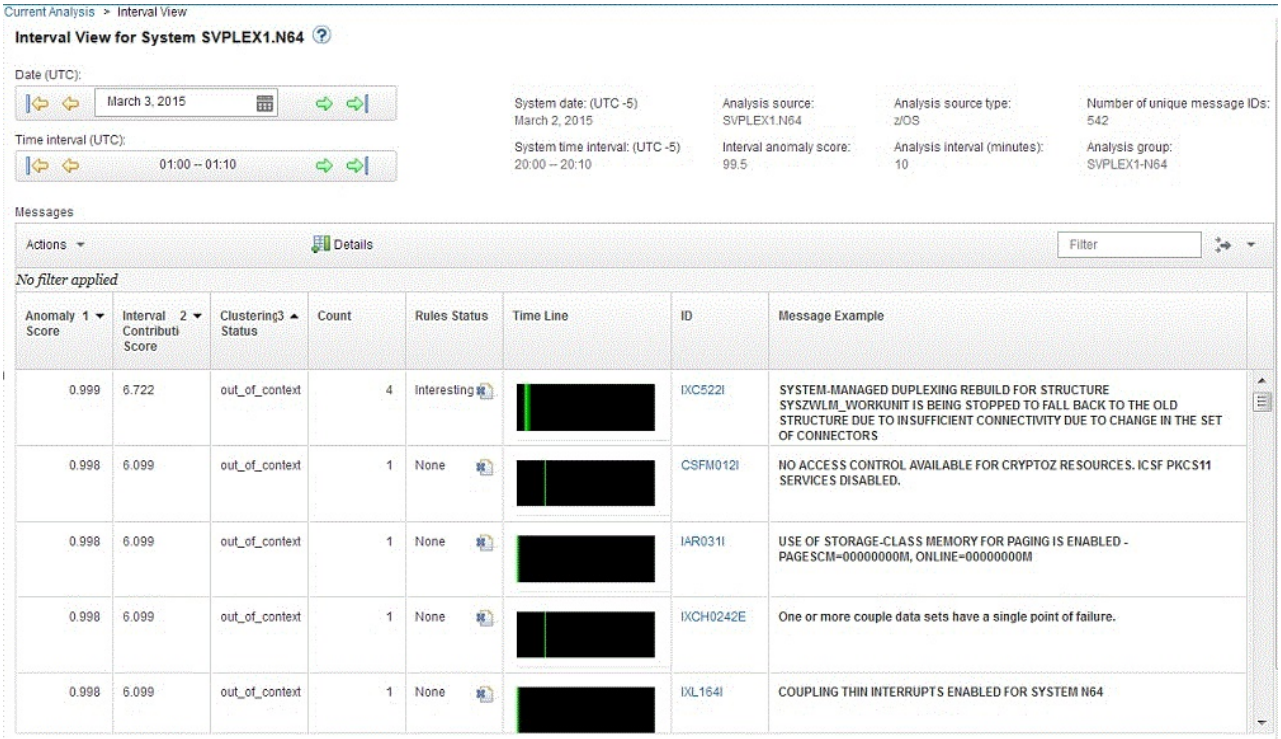

Figure 49. A sample Interval page display

The controls and content displayed in the Interval page are described in the following sections:

- "Fields and controls on the Interval page"
- "Content of the Messages table on the Interval page" on page 163
- "Tips for using the Interval page" on page 167

To return to the Analysis page, click **Return to Analysis**.

Fields and controls on the Interval page

The heading of the Interval page contains the name of the selected system in the format *group_name*`.`*system_name*; for example, `SYSPLEX1.SYSA` or `LNXVM1.LNXL3778`. The remainder of the Interval page display contains the following controls and fields.

- The **Date (UTC)** field displays the Coordinated Universal Time (UTC) date for the selected analysis interval. You can type the date, select it from the calendar widget, or use the forward and back arrows to change the date.
- The "System date" field displays the corresponding local date. The UTC offset for the geographic location of the monitored system is listed in the field label.
- The "Analysis source" field displays the name of the selected monitored system, in the format *group_name*.*system_name*; for example: SYSPLEX1.SYSA or LNXVM1.LNXL3778.
- The "Analysis source type" field identifies the type of monitored system, which is either z/OS or Linux.
- The "Number of unique message IDs" field indicates the number of unique message identifiers (IDs) that were issued during the analysis interval.
- The **Time interval (UTC)** field displays the UTC start and end times of the selected analysis interval. You can scroll through analysis intervals using the forward and back arrows.
- The "System time interval" field displays the corresponding local start and end times of the selected analysis interval. The UTC offset for the geographic location of the monitored system is listed in the field label.
- The "Interval anomaly score" field indicates the anomaly score for the selected analysis interval. For a description of each range of interval anomaly scores, see "Interpreting anomaly scores and the system behavior that they reflect" on page 145.

  If IBM zAware used a limited model to calculate the anomaly score, the Limited Model icon ( ) is displayed in this field. For more information about limited models, see "Understanding how IBM zAware calculates and displays anomaly scores " on page 142.
- The "Analysis interval (minutes)" field indicates the length of the selected analysis interval in minutes. The length varies depending on the analysis source type:
  - Analysis intervals for z/OS systems are 10 minutes in length.
  - Analysis intervals for Linux systems are 60 minutes in length.
- The "Analysis group" field displays the name of the system group to which the selected system belongs.
- The Messages table contains analysis results for every unique message issued during the analysis interval. The table toolbar contains the **Details** icon, the **Actions** list, a filter field, and clickable column headings through which you can modify the default sorting of message entries in the table. For more information about toolbar controls, see "Tips for using the Interval page" on page 167.
- **Return to Analysis** returns the display to the Analysis page view that was in effect before you clicked on an analysis snapshot.

Content of the Messages table on the Interval page

The Messages table contains details about every unique message issued during the analysis interval. If the same message ID was issued more than once during the selected interval, the Messages table contains only one entry for that unique message ID. Table 27 describes the details that are displayed for each message.

Table 27. Columns in the Messages table

Column	Description
Columns shown in the default display	
Anomaly Score	Indicates the difference in expected behavior for this specific message ID within the analysis interval. The message anomaly score is a combination of the interval contribution score for this message and the rule, if any, that is in effect for this message. Higher scores indicate greater anomaly so messages with high anomaly scores are more likely to indicate a problem. The message anomaly score ranges from 0 through 1.0.

Table 27. Columns in the Messages table (continued)

Column	Description
Interval Contribution Score	Indicates the relative contribution of this message to the anomaly score for the analysis interval. This interval score is a function of the following analysis results that are reported in the Messages table: Rarity Score, Clustering Status, Appearance Count, and Periodicity Status. Higher scores indicate greater contribution to the interval anomaly score.
Clustering Status	Indicates whether this message is part of a cluster, which is an expected pattern or group of messages associated with a routine system event (for example, starting a subsystem or workload). IBM zAware identifies and recognizes these patterns or groups, and the specific messages that constitute a specific cluster. When you analyze data from a monitored client, the server determines whether a specific message is expected to be issued within a specific cluster. A message that is issued out of context (without the other messages in the same cluster) might indicate a problem. Values for Clustering Status are: **New** IBM zAware did not previously detect this message in the model or detected one or more messages for the first time. **Unclustered** This message is not part of a defined cluster. **In context** IBM zAware expects this message to be issued within a specific cluster, and the message was issued as expected in the analysis interval. **Out of context** IBM zAware expects this message to be issued within a specific cluster, but the message was issued in a different context during the analysis interval.

Table 27. Columns in the Messages table (continued)

Column	Description
Rules Status	Indicates which rule, if any, applied to this message and affected its anomaly score for the analysis interval. This capability is available only for messages that z/OS monitored systems issue, so the Rules Status column is not displayed in the Messages table for a Linux system. The rule can be one of the following types: • Predefined by IBM. • Assigned by IBM zAware as a result of the analysis of training data. • Assigned by IBM zAware when an administrator has identified the message as one that IBM zAware is to ignore during analysis, either until the next model is built or until an administrator manually resets the ignore status. Possible values for this field are: **Critical** An IBM rule identifies this message as critical for diagnosing a potential system problem. For example, message IXC101I, which indicates that a system is being removed from a sysplex, is classified as critical. **Important** An IBM rule identifies this message as likely to indicate a problem. For example, message IEA911E, which indicates that an SVC dump was taken, is classified as important. **Interesting** An IBM rule identifies this message as indicative of a diagnostically useful event, such as a health check exception. **None** No rule is applied for this message. **Non-Interesting** A predefined IBM rule or an IBM zAware-assigned rule identifies this message as one with little or no diagnostic value. **User Ignored** An administrator identified the message as one that IBM zAware is to ignore during analysis, either until the next model is built or until an administrator manually resets the ignore status. For only those users with an ID mapped to the Administrator role, an icon ( ) is displayed in the column to indicate whether the rules status can be modified. Click the icon to open the Ignore Message Status window to view the current ignore status for this message ID. For a description of the Ignore Message Status window, see "Ignore Message Status window" on page 169.
Count	Specifies the number of times that this message was issued within the analysis interval.
Time Line	Provides an illustration of when this message was issued within the analysis interval. Each line represents a time period during the analysis interval in which the message was issued at least once. The length of the time period varies by the type of monitored system. • For z/OS systems, each line represents a 5-second time period. • For Linux systems, each line represents a 30-second time period. The browser zoom function affects the timeline: at 100% or lower, some lines are removed from the display. For the most accurate information, use the text-only format for this column content. To display the text-only format, position your cursor over the graphic display in the Time Line column on the Interval page, and click to open the **Time Line Summary** window. For a description of this text-only view, see "Time Line Summary window" on page 168.

Table 27. Columns in the Messages table (continued)

Column	Description
ID	Provides the message identifier. • For z/OS system messages, the message identifier is a direct link to the message query in IBM Knowledge Center. In the action menu, next to the message identifier, you can also click **Knowledge Center** or **View Message History** (for example, 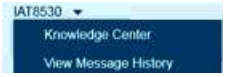). To display the search results for all occurrences of the z/OS message identifier across all clients that IBM zAware monitors, click **View Message History**, which opens the **Message History** page in a new browser tab. • For Linux messages, the message identifier itself is a link that you can click to open the **Message History** page in a new browser tab. The message identifier might be a known Linux system message identifier, or a message identifier that is generated by IBM zAware for its own use. You can also open a browser window and search for the Linux message description by using an internet search engine or Linux repository.
Example or Summary	Provides either the full message text for the first occurrence of this message within the analysis interval, or a summary of the common message text that was issued for each occurrence of the same message within the analysis interval. For summaries, only common text is displayed, with asterisks that replace any text that differs between occurrences of this message. By default, this column displays the full message text with the column heading **Example**. To change the display to the summary view, click **Actions** > **View Message Summary**. To reset the display to the default example view, click **Actions** > **View Message Full Text**.
Additional columns shown only after clicking the Details icon ( ). These additional columns are not displayed for analysis results that were generated before IBM zAware Version 2.0.	
Periodicity Status	Indicates whether this message has a tendency to recur at specific times, and whether the message recurred as expected within the analysis interval. Values for Periodicity Status are: **NEW** IBM zAware did not previously detect this message. **IN_SYNC** IBM zAware expects this message to be issued in a periodic pattern, and the message was issued as expected during the analysis interval. **NOT_IN_SYNC** IBM zAware expects this message to be issued in a periodic pattern, but the message was not issued as expected during the analysis interval. **NOT_PERIODIC** IBM zAware does not expect this message to be issued in a periodic pattern.
Periodicity Score	Indicates how the periodicity status of this message might contribute to the message anomaly score for the analysis interval. Higher scores generally indicate greater contribution to the message anomaly score.
Last Issued (UTC)	Indicates the UTC date and time when this message was last issued on the monitored system before the start of the current analysis interval. The time is displayed in 24-hour clock format.
Daily Frequency	Indicates the average number of analysis intervals in which the message is expected to be issued each day, according to analysis of the message data that IBM zAware uses for training.

Table 27. Columns in the Messages table (continued)

Column	Description
Rarity Score	Indicates how often this message was issued within the collection of analysis intervals that are used to build the model. Values range from 1 to 101: • A value of 1 indicates that the message is issued in almost all analysis intervals in the model. • A value of 100 indicates that the message is issued in almost none of the analysis intervals in the model. • A value of 101 indicates that this message ID was not issued in any analysis interval in the model.
Cluster ID	Provides the identifier of the cluster to which this message belongs. When the message is not part of a recognized cluster, the cluster ID is -1.

Tips for using the Interval page

To find those messages on the Interval page that have the most value for diagnosis, use the following tips:

- The message anomaly score is the primary indicator of diagnostic value, so the default organization of the Message table in the Interval page arranges message entries sorted first by **Anomaly Score** in descending order, from highest value to lowest, then by **Interval Contribution Score** in descending order, and then by **Clustering Status**, with new messages listed first.

  Sorting by **Interval Contribution Score**, in descending order, moves messages of little or no diagnostic value to the bottom of the Messages table. Given that the Messages table contains one entry for each unique message issued during the selected interval, this sorting helps you avoid scrolling through a considerable number of insignificant message entries.
- When an interval is compared to a limited model, relatively insignificant changes in current system behavior can result in significant changes to anomaly scores. When a limited model condition is indicated, carefully examine the specific interval results to determine whether the system is exhibiting unusual behavior that requires corrective action.
- To see all messages issued by a specific z/OS component, type the corresponding 3- or 4-character message prefix into the Filter field. For example, to filter the entries in the Messages table to contain only the messages that the z/OS cross-system coupling facility (XCF) component issued, type the message prefix IXC in the Filter field. Additional filtering capabilities are available through the Filter actions icon ( ), which is located in the far right corner of the table header.
- To reorganize the default presentation of the Message table for the current Interval page:
  - Click the Details column icon ( ) in the Messages table toolbar. The display changes to show more diagnostic columns on the right side of the Messages table.
  - Click **Actions** > **View Message Summary** or **Actions** > **View Message Full Text** to display either the Message Example or Message Summarization column.

    By default, this column displays the full message text with the column heading **Example**. To change the display to the summary view, click **Actions** > **View Message Summary**. To reset the display to the default example view, click **Actions** > **View Message Full Text**.
  - Click a column heading to resort the message entries according to the values in that specific column. If you click more than one column heading in sequence, you have to option of specifying whether this column sort takes priority over the earlier column sort. All columns in the Messages table are sortable except for the Example or Summary column.

  Any changes that you make to the Interval page apply only for the current display. If you return to the Analysis page and then select an interval to open another Interval page, IBM zAware uses the default presentation for the Messages table.
- The table footer provides a total count of table entries, and the number of selected table entries. Depending on the total count, you might need to use the vertical scroll bar to view all of the table entries.

Linking to IBM Operations Analytics for z Systems for message analysis

Through the **Actions** menu in the **Interval** view, you can link to IBM Operations Analytics for z Systems for further message analysis.

Before you begin

To link to IBM Operations Analytics for z Systems from the Messages table, your **Search Options** must be properly configured. For more information, see Chapter 18, "Configuring the Search Options," on page 177.

About this task

The Messages table provides diagnostic details about the messages that IBM z Advanced Workload Analysis Reporter (IBM zAware) collected during an analysis interval, which IBM Operations Analytics for z Systems can analyze further. If you are unfamiliar with how IBM zAware collects message data from monitored clients, review "Using the Interval page to pinpoint the causes of system anomalies" on page 161.

Procedure

1. Go to the Messages table.

   Fast path: Click **Analysis** > **Analysis Graph** > **Interval view**, select the messages, and then go to step 3.

   a. Click the **Analysis** pane, and then click the **System Group** that you want to analyze. If you need to select a different group or system, click **Change Source**.
   b. Select the anomaly score for the system that you want to analyze, and then select the interval in the **Details for System** *system_name* page. You are in the **Interval View for System** *system_name* page.
2. Select the messages that you want IBM Operations Analytics for z Systems to analyze.
   - To link all messages, select the check box in the Messages table header.
   - To link specific messages, click the check boxes next to one or more messages.
3. Click **Actions** > **Search Logs:** ***Name defined in Search Options*** to link the messages to IBM Operations Analytics for z Systems for analysis. If **Search Logs:** ***Name defined in Search Options*** is not defined, see Chapter 18, "Configuring the Search Options," on page 177.

What to do next

See the topic about "Preparing to analyze z/OS log data" with **IBM Operations Analytics for z Systems** at www.ibm.com/support/knowledgecenter/SS55JD.

Time Line Summary window

The **Time Line Summary** window provides a text-only format that indicates when a selected message ID was issued during a selected analysis interval. To display the text-only format, position your cursor over the graphic display in the Time Line column on the Interval page, and click to open the **Time Line Summary** window.

Fields in the Time Line Summary window

Table 28 describes the fields that are displayed in the Time Line Summary window.

Table 28. Fields displayed in the Time Line Summary window

Fields	Description
Message ID	Identifies the message ID selected from the Messages table in the Interval view.
System	Identifies the name of the monitored system on which the selected message ID was issued.
Time interval (UTC)	Specifies the start and end times for the selected analysis interval, in Coordinated Universal Time, using the 24-hour clock.
Date (UTC)	Specifies the date on which the selected analysis interval occurred.
System time interval (UTC offset)	Specifies the start and end times for the selected analysis interval in the local time for the central processor complex (CPC) on which the monitored system runs, using the 24-hour clock. The label for this field contains the UTC offset, if any.
System date (UTC offset)	Specifies the date on which the selected analysis interval occurred, using the local time for the CPC on which the monitored system runs. The label for this field contains the UTC offset, if any.
Appearance count (in interval)	Specifies the total number of times that the selected message ID was issued during the selected analysis interval.
Analysis interval (minutes)	Specifies the length, in minutes, of the selected analysis interval. The length of an analysis interval varies, depending on the type of monitored system.
Time Line of Appearances table	Specifies whether or not the selected message ID was issued during each minute of the selected analysis interval. Each row in the table represents one minute in the analysis interval.
Close	Closes the Time Line Summary window and returns to the Interval page.

Ignore Message Status window

You can use the Ignore Message Status window to display the current status for a specific message ID. To display the Ignore Message Status window, click the icon ( ) in the Rules Status column on the Interval page. Only those users with an ID mapped to the Administrator role can view the icon and click it to open the Ignore Message Status window.

Fields in the Ignore Message Status window

Table 29 describes the fields that are displayed in the Ignore Message Status window.

Table 29. Fields displayed in the Ignore Message Status window

Field	Description
Selected message ID	Lists the message ID selected from the Messages table on the Interval page.
Current system	Lists the name of the system on which the selected message ID was issued.

Table 29. Fields displayed in the Ignore Message Status window (continued)

Field	Description
Current ignore status	Lists the current ignore status. Possible values are: **Not Ignored** The message ID does not have any ignore status value currently applied to it. **Until next training** The message is to be ignored during analysis until the next time IBM zAware successfully builds a model of behavior for the monitored system, as the result of either a manually requested or automatically scheduled training operation. The date for the next scheduled model rebuild is listed under "Next scheduled training date" after you select **Next Training Period Model Dates** on the Manage Model Dates page. **Until manually restored** The message is to be ignored during analysis until an administrator manually changes the ignore status of the message, or removes it from the list of messages in the Ignored Messages table.
Current ignore status applied (UTC)	Indicates the date and time (in UTC, by using the 12-hour clock) at which an administrator most recently updated the ignore status for the message.
Ignore message option for future intervals	Lists the options for setting ignore status for this message ID. By default, the current status is selected. You can select one of the following options. • Do not ignore message. • Ignore message until next training occurs for the current system. • Ignore message until manually restored. Messages can be restored using the **Manage Ignored Messages** action on the Training Sets page.
OK	Sets the selected ignore message option and closes the Ignore Message Status window. If you want to go to the Manage Ignored Messages page instead of returning to the Interval page, click **Go to Manage Ignored Messages view on OK**. For information about the Manage Ignored Messages page, see "Managing ignored messages" on page 230.
Cancel	Closes the Ignore Message Status window and returns to the Interval page.
Go to "Manage Ignored Messages" view on OK.	Opens the Manage Ignored Messages page. For information about the Manage Ignored Messages page, see "Managing ignored messages" on page 230.

Verifying planned system changes with IBM zAware

Although the **Analysis** page display is useful for finding and diagnosing unexpected problems or reported incidents, you also can use it to verify behavior after a planned change. If you know when your installation made a change, you can navigate to the time period immediately following that change and view the analytical data for the affected monitored client (system). For example, you can check the intervals roughly one hour after a system IPL to verify that the system is operating normally.

A higher volume of messages might indicate a potential problem after the following types of changes:
- Software updates for the operating system, middleware, or applications

- Updated system settings
- Changed system configurations, including hardware

Software updates and configuration changes might not result in higher interval anomaly scores for the affected system. When your installation introduces new workloads or applications to a system, however, IBM zAware scores the intervals that follow these changes as highly anomalous because it detects messages that are not reflected in the system model. IBM zAware is not able to detect whether these new messages indicate problems or routine, expected behavior until enough data is available to update the system model.

After using the IBM zAware analytical data to verify that the new workload or application is operating as expected, you have two options for altering IBM zAware analysis to prevent the assignment of high anomaly scores to future intervals: Marking messages for IBM zAware to ignore during analysis, or manually requesting IBM zAware to rebuild the model of system behavior.

Marking messages for IBM zAware to ignore

You can mark messages from the new workload or application for IBM zAware to ignore during analysis. You can designate messages to be ignored until the next time IBM zAware builds a model of behavior for the monitored system, or until an administrator manually changes the ignore status of the message. This capability is available only for messages that z/OS monitored systems issue.

Ignore messages until next training occurs for the current system.

Marks the selected messages to be ignored during analysis until the next time IBM zAware successfully builds a model of behavior for the monitored system, as the result of either a manually requested or automatically scheduled training operation. The date for the next scheduled model rebuild is listed under "Next scheduled training date" after you select **Next Training Period Model Dates** on the Manage Model Dates page.

Use this value for messages that you determine to be anomalous because of a workload change on the system, but you expect them to become part of the normal behavior for this system. The next training results in a model that includes these messages, which will be subject to normal analysis after the training.

Ignore messages until manually restored.

Marks the selected messages to be ignored during analysis until an administrator manually changes the ignore status of the message, or removes it from the list of messages in the Ignored Messages table.

Use this value for messages that you determine to be normal (that is, not indicative of a problem) on this system. In subsequent analysis, these messages do not contribute to the anomaly score, and thus reduce false-positive results.

This option is the most expedient method; however, if you choose to ignore messages until the next training, note that IBM zAware includes the designated messages in analysis immediately after the successful completion of either a manually requested or automatically scheduled training operation. For more information about designating messages to be ignored during analysis, see "Managing ignored messages" on page 230.

Manually requesting IBM zAware to rebuild a model

You can manually request IBM zAware to rebuild the model of system behavior.

This option requires specific timing to be effective. The sequence of events shown in Figure 50 on page 172 illustrates how you can use this option to alter IBM zAware analysis after installing a new application on a system named SYS1.

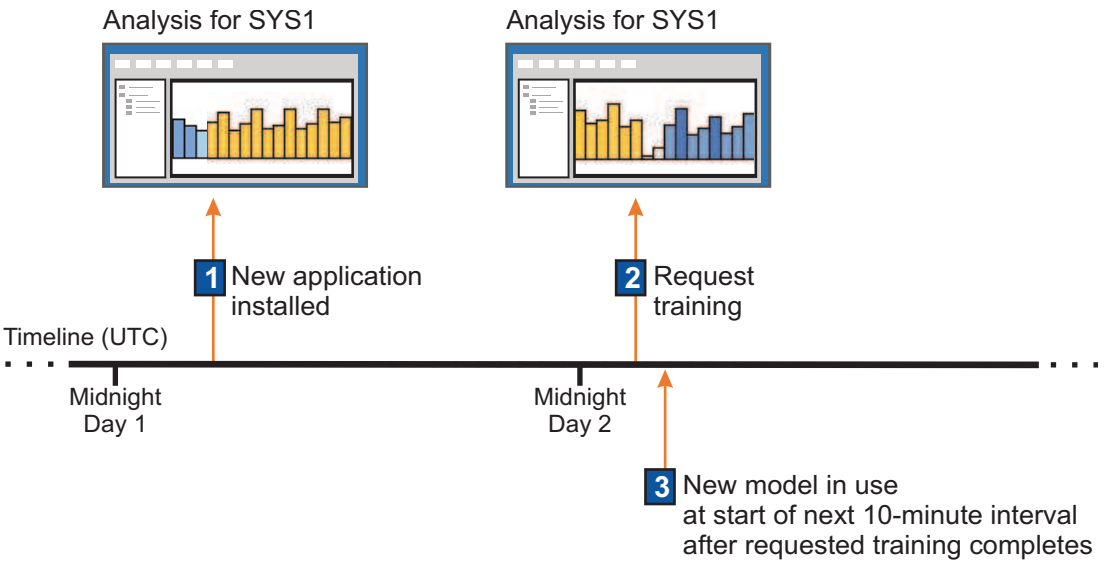

Figure 50. Sample timeline for retraining IBM zAware to analyze data from a new application

1. Several hours after midnight on Day 1, your installation installs a new application on SYS1, which is connected and sending data to the IBM zAware server.

   Before the installation occurred, the analytical data in the **Analysis** page indicates fairly normal system behavior with blue rectangles. After the installation, intervals in the **Analysis** page are dark gold because the server detects unique messages that it has not detected previously in the model for SYS1. The server assigns an interval anomaly score of 101 to each interval in which the new application issues a message.
2. Although you can request the server to rebuild a model at any time, wait until midnight of Day 2 before you request the IBM zAware server to rebuild the model for SYS1.

   This delay is necessary because the server uses only complete days of data to build models, and because it needs a significant amount of data from the new application to correctly identify and recognize message patterns. Any time after midnight on Day 2, you can request the server to rebuild the SYS1 model by using the **Request Training** action on the **Administration** > **Training Sets** page. The server uses the data it received from SYS1 during Day 1 to update the model.
3. After the requested training completes, the IBM zAware server begins to use the newly rebuilt model at the start of the next 10-minute interval.

   Using the new model, the server can more accurately detect message patterns from the new application, and intervals receive more accurate interval anomaly scores and colors. The effect of retraining depends on the stability of the system and the degree to which the new messages cluster together, so you might need to wait another day to collect additional data and rebuild the model again.

Chapter 17. Viewing the Message History page

Use the IBM z Advanced Workload Analysis Reporter (IBM zAware) **Message History** page to review the history for all occurrences of a message ID across all clients that IBM zAware monitors.

Before you begin

On the **Message History** page, you can enter a Message ID and see every occurrence of the message across all monitored clients in your environment.

About this task

The following procedure describes how to query the history of a specific message directly from the **Interval View**.

Procedure

1. Go to the specific **Interval** that you want to analyze. Click **Analysis** > **Analysis Graph** > **Interval View**.

   Fast path: Click **Analysis** > **Analysis Graph** > **Interval View**, select the specific interval for the system with the peak anomaly score that has messages that you want to analyze, and then go to Step 3.
   You are in the **Interval View** for **System *sysname*.**
2. Click the specific interval for the system with the peak anomaly score that has one or more messages to analyze. Figure 51 shows the **Interval View** table with z/OS message identifiers that are listed in the **ID** column.

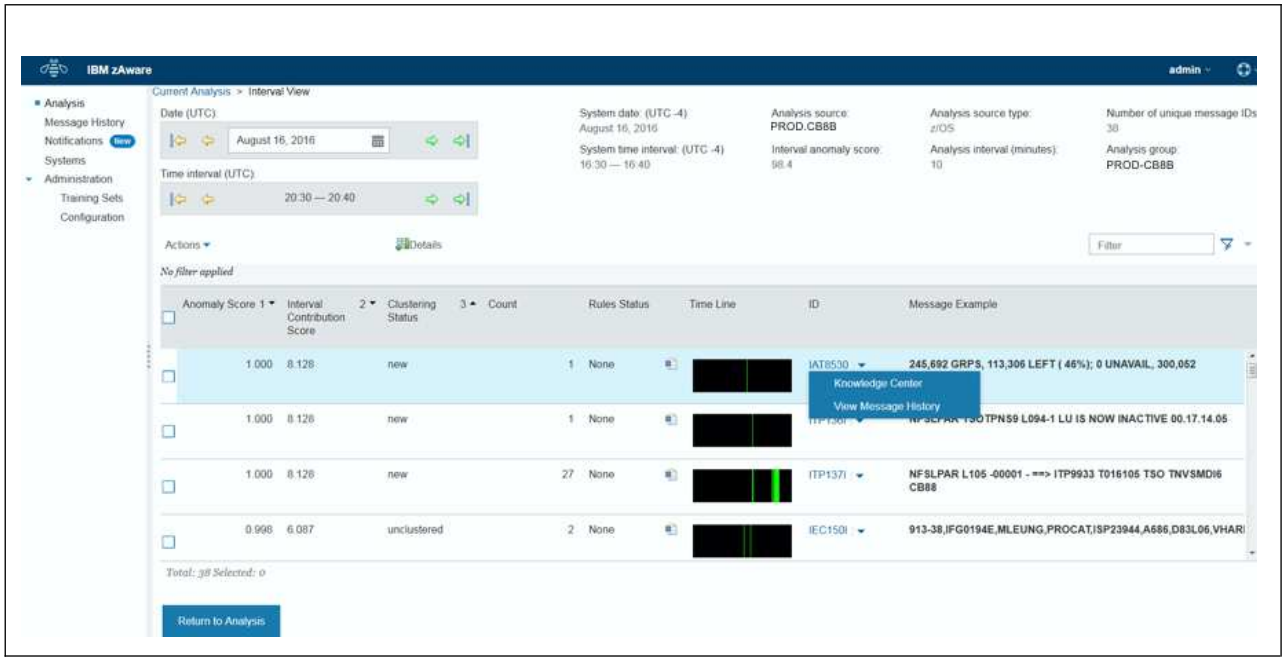

Figure 51. Interval View for z/OS with links to IBM Knowledge Center or View Message History

3. Click the message ID to link to either the **Message History** page or **Knowledge Center**. The action depends on which operating system you are analyzing.

- For Linux, click the message number to open a new browser tab with the **Message History** from IBM zAware. The list contains all occurrences of the message ID across all monitored clients.
- For z/OS, click the message number to open a new browser tab with a list of the message search results from IBM Knowledge Center. You can also use the down arrow next to the message ID to access the **Message History** page to view all occurrences of the message ID across all monitored clients.

Results

The history of the message ID displays as shown in Figure 52 on page 175. The **Message History** table contains the following fields:

System
Specifies the name of the system on which the message appeared.

Interval (Coordinated Universal Time)
Specifies the start and end times for the interval when the message appeared. You can click the **Interval** to return to the **Interval View for System** ***system name*****(z/OS) | *host name*****(Linux)** page.

Count Specifies the number of times the message appeared during the interval.

Message Text Summary
Displays a summary of the common message text that was issued for each occurrence of the same message.

Interval Anomaly Score
Indicates the anomaly score during the 10-minute interval when the message appeared.

Cluster Status
Indicates whether the message is part of a pattern that is associated with common system events. For example, starting a subsystem or a workload is a common event. The patterns or groups are called "clusters" and specific message identifiers constitute a specific cluster. When IBM zAware analyzes data from a monitored client, the server determines whether a message is expected to appear within a specific cluster. A message that is issued out of context, without the other messages in the same cluster, might indicate a problem.

New IBM zAware did not previously detect this message in the model or detected one or more messages for the first time.

Unclustered
This message is not part of a defined cluster.

In context
IBM zAware expects this message to be issued within a specific cluster, and the message was issued as expected in the analysis interval.

Out of context
IBM zAware expects this message to be issued within a specific cluster, but the message was issued in a different context during the analysis interval.

Periodicity
Indicates whether the message tends to repeat at specific times, and if it repeated as expected within the analysis interval.

The following values are for Periodicity Status:

NEW IBM zAware did not previously detect the message or it is the first time that the message is issued.

IN_SYNC
IBM zAware expects the message to be issued in a periodic pattern, and the pattern was as expected during the analysis interval.

NOT_IN_SYNC
IBM zAware expects the message to be issued in a periodic pattern, but the message was not issued as expected during the analysis interval.

NOT_PERIODIC
IBM zAware does not expect the message to be issued in a periodic pattern.

Example

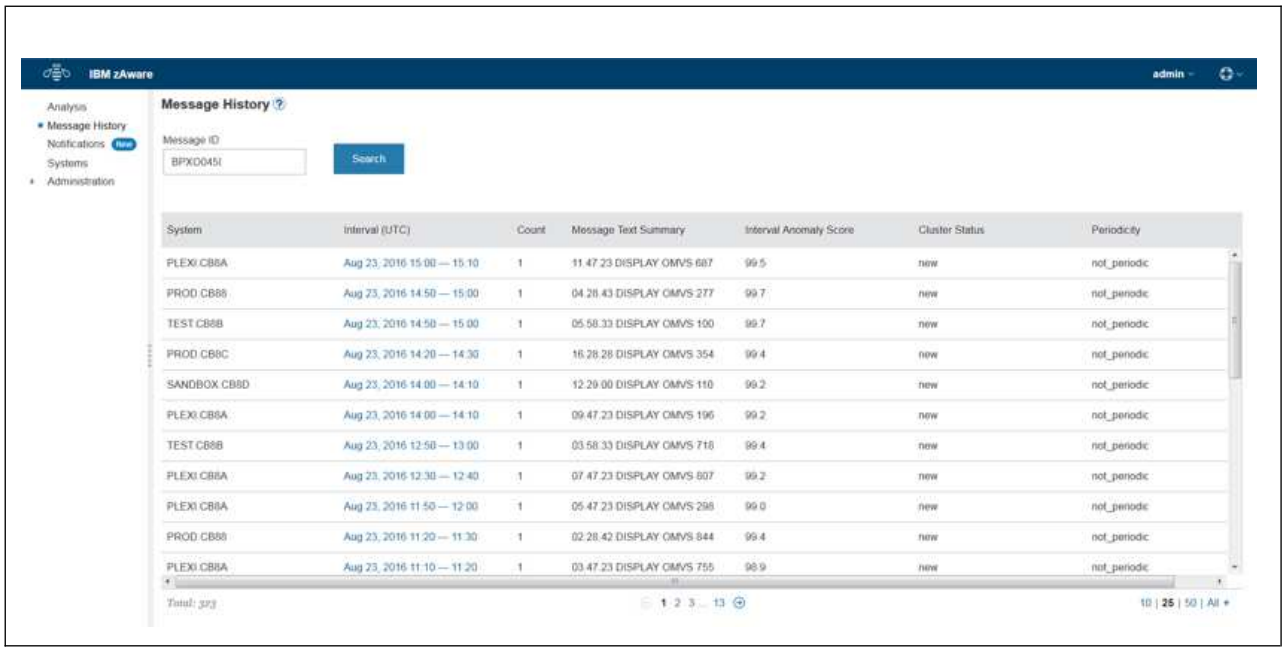

Figure 52. Viewing a z/OS message in the Message History

Related information:
"Tips for using the Interval page" on page 167

Chapter 18. Configuring the Search Options

Use this procedure to set up the **Configuration** > **Search Options** page to enable IBM z Advanced Workload Analysis Reporter (IBM zAware) to link to IBM Operations Analytics for z Systems.

Before you begin

If you are unfamiliar with how IBM zAware collects message data from z/OS clients, review "Using the Interval page to pinpoint the causes of system anomalies" on page 161.

About this task

IBM zAware collects message data from z/OS monitored clients for analysis, but does not keep the entire message log that is transmitted. IBM Operations Analytics for z Systems can help quickly identify and isolate potential problems. Use the **Configuration** > **Search Options** tab to set up the server options to enable IBM zAware to link to IBM Operations Analytics for z Systems.

Procedure

1. Click **Configuration** > **Search Options** > **Add New** to set up the server for IBM Operations Analytics for z Systems.
2. Provide a name for the IBM Operations Analytics for z Systems configuration in the **Name** field. For example, say that you select "Analytics" as the name. On the Intervals page, when you are ready to transmit messages, "Analytics" appears as **Actions** > **Search logs: Analytics**.
3. Provide the complete URL, host name, and port number, in the **Server** field. For example, `https://analytics.ibm.com:port_number` or `http://valid_server_name`.
4. Click **Save** to save the server configuration. The other options include the following choices:
   - Click **Edit** to change the **Name**, **Server**, or the information in both fields.
   - Click **Delete** to remove a server configuration.

Results

You can search the IBM Operations Analytics for z Systems knowledge base for the selected messages and system for the time frame that IBM zAware discovered the anomaly.

What to do next

For more information, see "Linking to IBM Operations Analytics for z Systems for message analysis" on page 168.

Chapter 19. Specifying security settings for the IBM zAware GUI

To access the IBM zAware graphical user interface (GUI), a user requires a valid user ID and password to authenticate with the IBM zAware GUI and users must be assigned to an IBM zAware role, which permits access to IBM zAware functions.

IBM zAware supports two methods of user authentication. You can authenticate users through the use of a Lightweight Directory Access Protocol (LDAP) repository or through the use of a local file-based repository. For simplicity, using only an LDAP repository is the preferred option. However, you might want to define one or two user IDs in a local repository so you can access the IBM zAware GUI when the LDAP server is unavailable. If you configure an LDAP repository and also define users or groups in a local repository, both sets of users or groups are available through the IBM zAware GUI. Do not define the same user ID in more than one repository; results are not predictable.

With either type of repository, the authentication process is the same:

1. A user enters a user ID and password to log in to the IBM zAware GUI.
2. The IBM zAware server verifies that the login credentials match a user ID and password stored in the LDAP repository or the local repository.
3. If there is a match, the user is considered to be authenticated. However, IBM zAware does not grant the user access to the GUI until it verifies that a user is also assigned to an IBM zAware role – *User* or *Administrator*.

When your installation configures and activates the IBM zAware partition, your installation defines a master user ID and password that you can use to initially log in to the IBM zAware GUI. By default, the master user ID is assigned to the administrator role, which has authority to perform any task that is available through the IBM zAware GUI.

To configure your security settings for the first time, log in to the GUI using the master user ID and password. IBM zAware provides the following security mechanisms that your installation can configure to limit access to the IBM zAware GUI:

Server SSL certificate
: You can, optionally, replace the automatically generated Secure Sockets Layer (SSL) certificate that is configured in the IBM zAware server with a certificate that is signed by a certificate authority. Doing so prevents the browser warning message that is displayed when you initially log in to the IBM zAware GUI. For instructions, see “Replacing the default SSL certificate” on page 180.

LDAP authentication for the IBM zAware GUI
: You can, optionally, configure the IBM zAware server to authenticate users through the use of an existing LDAP repository. For instructions, see “Enabling LDAP authentication for IBM zAware users” on page 183.

  For instructions for using the local file-based repository for authentication, see “Setting up a local repository to secure access to the IBM zAware GUI” on page 108.

Role-based access to IBM zAware GUI functions
: You can control access to IBM zAware GUI functions by assigning users to specific IBM zAware roles. For instructions, see “Assigning users or groups to a role” on page 187.

Browser session timeout setting
: You can change the browser session time out from the default setting to a value that is more appropriate for your installation. By default, browser sessions time out after 12 hours (720 minutes). For instructions, see “Specifying the duration of a browser session” on page 190.

These security mechanisms are described in more detail in the sections that follow.

Replacing the default SSL certificate

A Secure Sockets Layer (SSL) certificate is automatically generated for IBM zAware when your installation initially activates the IBM zAware partition. The certificate is not signed by a certificate authority (CA). Therefore, the first time you log in to the IBM zAware graphical user interface (GUI), the browser displays a warning message because it does not recognize the default SSL certificate. You can resolve this problem by replacing the default SSL certificate with a certificate that is signed by a certificate authority of your choice. Doing so provides secure communication between the IBM zAware server and the browsers of all authorized users.

Before you begin

If you decide to replace the automatically generated certificate, which is the recommended practice for improved security, you can use any third-party certificate authority of your choice, or your installation can provide an internal certificate authority for certificate signing tasks. IBM zAware does not renew these replacement certificates; in this case, managing replacement certificates becomes the responsibility of the security administrator.

You might need to process the CA reply before you can paste it into the appropriate field in the IBM zAware GUI.

- The required format for replacement certificates is Base64 encoded X509 certificate blocks.
- When you receive a reply from a certificate authority, the reply might contain a chain of certificates, starting with the signed server certificate. Then, it is possibly followed by certificates from one or more intermediate CAs and finally, the self-signed certificate of the CA. When you supply this information in the GUI, provide the entire certificate chain, starting with the signed server certificate through the self-signed certificate of the CA.
- In some cases, the CA reply that you receive is delivered in a public-key cryptography standards (PKCS) #7 file. You must extract the certificates from the file before pasting them into the GUI. One method of extracting certificates from a PKCS #7 file is to use the **openssl pkcs7** command; for more information, see the OpenSSL Project website at the following URL.

  www.openssl.org/

About this task

This security task is optional.

The following procedure might take several days to complete, depending on the time that the certificate authority requires to receive and process your request, and to send a reply. You can complete other tasks in the IBM zAware GUI while you wait for a reply from the certificate authority.

Procedure

1. Expand the Administration category in the navigation pane and select **Configuration**. The Configuration page is displayed.
2. Click **Security** > **SSL Settings** to display the SSL Settings tab.
3. To create a certificate signing request (CSR), in the Certificate Actions section, click **Generate Certificate Signing Request**. The Create Certificate Signing Request page is displayed.
4. Enter the appropriate information for the following fields:
   a. In the **Common name** field, verify the host name or IP address of the IBM zAware partition. IBM zAware loads this field with a value that matches the host name or IP address that is specified in the image profile of the IBM zAware partition. The common name is required.

b. In the **Organization** field, enter the name of your company. The value that you supply for this field must be a string of length 1-64. The organization name is optional.

c. In the **Organizational unit** field, enter the company organization or department name. The value that you supply for this field must be a string of length 1-64. The organizational unit is optional.

d. In the **Locality** field, enter the city in which your company is located. The value that you supply for this field must be a string of length 1-128. The city is optional.

e. In the **State or province** field, enter the state or province in which your company is located. The value that you supply for this field must be a string of length 1-128. The state or province is optional.

f. In the **Postal code** field, enter the postal code for your company address. The value that you supply for this field must be a string of length 1-16. The postal code is optional.

g. In the **Country code** field, enter the two-character abbreviation for the country in which your company is located. The value that you supply for this field must be a string of length 1-2. The country code is optional.

5. Click **Generate** to generate the certificate request. IBM zAware displays the generated request.
6. Expand the Generated Request Input section to display and verify the generated request input.
7. Copy the generated certificate signing request, which is displayed in the **Generated CSR** field.
8. Submit the request to a certificate authority using the procedures required by that entity.
9. Click **Close** to return to the main SSL Settings tab.
10. When you receive a reply from the certificate authority, do the following:

    a. Extract the certificates, if necessary.

       When you receive a reply from a certificate authority, the reply might contain a chain of certificates, starting with the signed server certificate. Then, it is possibly followed by certificates from one or more intermediate CAs and finally, the self-signed certificate of the CA. When you supply this information in the GUI, provide the entire certificate chain, starting with the signed server certificate through the self-signed certificate of the CA.

       Make sure that you do not insert any lines or spaces between the end of one certificate and the beginning of the next certificate. When you paste certificate replies in the GUI, make sure that you include all of the content, including the header `-----BEGIN CERTIFICATE-----` through and including `-----END CERTIFICATE-----`

       If you need to view a sample CA reply that contains a certificate chain, see Appendix B, "Sample certificate authority (CA) reply," on page 289.

    b. Return to the main SSL Settings tab.

    c. Click **Receive Certificate Request Reply**. The Receive Certificate Authority Reply page is displayed.

    d. Copy the reply to your clipboard and paste it into the **Certificate Authority reply** field.

    e. Click **Receive** to import the CA reply into the IBM zAware server.

Results

The main SSL Settings tab now displays information from the received CA reply.

SSL Settings tab

Use the actions provided on the SSL Settings tab to replace the default IBM zAware SSL certificate with a certificate that is signed by a certificate authority.

For instructions, see "Replacing the default SSL certificate" on page 180.

For a description of the content and controls included on the SSL Settings tab, see Table 30 on page 182.

Table 30. Items displayed on the SSL Settings tab

Item	Description
Current zAware Server Certificate	Displays information about the SSL certificate that is configured in the IBM zAware server.
Certificate Actions	Provides the following buttons: **Generate Certificate Signing Request** Allows you to specify the information required to generate a new certificate signing request (CSR) and to generate the request. **View Last Generated Request** Allows you to view the last certificate signing request that IBM zAware generated. This button is enabled after you generate the CSR. **Receive Certificate Request Reply** Allows you to input the reply you received from the certificate authority and to receive that certificate into IBM zAware. This button is enabled after you generate the CSR.

Create Certificate Signing Request page

Use the Create Certificate Signing Request (CSR) page to specify the information to include in the certificate signing request.

To display the Create Certificate Signing Request page, click **Generate Certificate Signing Request** on the main SSL Settings tab. The fields initially included on the page are described in Table 31.

After you provide the appropriate information, click **Generate** to generate the request. The fields that are included on the page after the CSR is generated are described in Table 32 on page 183.

Table 31. Fields displayed on the page before the CSR is generated

Field	Description
Common name	Verify the host name or IP address of the IBM zAware partition. IBM zAware loads this field with a value that matches the host name or IP address that is specified in the image profile of the IBM zAware partition. The common name is required.
Organization	Enter the name of your company. The value that you supply for this field must be a string of length 1-64. The organization name is optional.
Organizational unit	Enter the company organization or department name. The value that you supply for this field must be a string of length 1-64. The organizational unit is optional.
Locality	Enter the city in which your company is located. The value that you supply for this field must be a string of length 1-128. The city is optional.
State or province	Enter the state or province in which your company is located. The value that you supply for this field must be a string of length 1-128. The state or province is optional.
Postal code	Enter the postal code for your company address. The value that you supply for this field must be a string of length 1-16. The postal code is optional.
Country code	Enter the two-character abbreviation for the country in which your company is located. The value that you supply for this field must be a string of length 1-2. The country code is optional.

Table 32. Fields displayed on the page after the CSR is generated

Field	Description
Generated Request Input	Provides the input that you supplied for the request. Verify the request input. If any information is incorrect, you must create a new request. To do so, click **Close** to return to the main SSL Settings tab. Then, click **Generate Certificate Signing Request**.
Generated CSR	Provides the certificate signing request that IBM zAware generated by using your input. Copy the CSR and submit it to a certificate authority.

View Last Generated Request page

Use the View Last Generated Request page to view the last certificate signing request that IBM zAware generated.

To display the View Last Generated Request page, click **View Last Generated Request** on the main SSL Settings tab. The fields included on the page are described in Table 33.

Table 33. Fields displayed on the View Last Generated Request page

Field	Description
Last Generated Request Input	Provides the input that was supplied for the request.
Last generated CSR	Provides the certificate signing request that IBM zAware generated using the input. If necessary, copy the CSR and submit it to a certificate authority.

Receive Certificate Authority Reply page

Use the Receive Certificate Authority Reply page to input the reply you received from the certificate authority and to receive that certificate into IBM zAware.

To display the Receive Certificate Authority Reply page, click **Receive Certificate Request Reply** on the main SSL Settings tab.

In the **Certificate Authority reply** field, which is displayed on the page, enter the reply that you received from the certificate authority. Verify that:

- The format of the reply is Base64 encoded X509 certificate blocks.
- The reply contains the entire certificate chain. When you receive a reply from a certificate authority, the reply might contain a chain of certificates, starting with the signed server certificate. Then, it is possibly followed by certificates from one or more intermediate CAs and finally, the self-signed certificate of the CA. When you supply this information in the GUI, provide the entire certificate chain, starting with the signed server certificate through the self-signed certificate of the CA.

In some cases, the CA reply that you receive is delivered in a public-key cryptography standards (PKCS) #7 file. You must extract the certificates from the file before pasting them into the GUI. One method of extracting certificates from a PKCS #7 file is to use the **openssl pkcs7** command; for more information, see the OpenSSL Project website at the following URL.

Click **Receive** to import the certificate into the IBM zAware server.

Enabling LDAP authentication for IBM zAware users

Lightweight Directory Access Protocol (LDAP) is an information directory where users and groups can be defined only once and shared across multiple machines and multiple applications. IBM zAware can authenticate user login requests against an existing LDAP server. To enable LDAP authentication, specify the settings for an existing LDAP server in your installation.

Before you begin

- Obtain the master user ID and password that was provided in the image profile for the IBM zAware partition.
- Check with your LDAP administrator or network administrator to ensure that the IBM zAware server can access the LDAP server. IBM zAware cannot save LDAP settings unless it can communicate with the LDAP server when you apply the new or changed settings.
- Check with your LDAP administrator to determine the settings needed to configure your LDAP server. For a description of each setting, see "LDAP Settings tab" on page 185.
- If you are using this procedure to modify an existing LDAP configuration, make sure you verify that all mapped users and groups are still valid and appropriate in the new LDAP configuration. Before making any changes to the LDAP configuration, go to the **Role Mapping** tab and review the lists of currently mapped users and groups for both the **Administrator** and **User** roles.

  If any users or groups are not appropriate for the new LDAP configuration, select them and click **Remove**. If you need more information about using the controls on the **Role Mapping** tab, see "Assigning users or groups to a role" on page 187.

Procedure

1. Log in to the IBM zAware GUI using the master user ID and password that was provided in the image profile for the IBM zAware partition.
2. Expand the Administration category and select **Configuration**. The Configuration page is displayed.
3. Click **Security** > **LDAP Settings** to display the **LDAP Settings** tab.
4. Specify the settings required for the IBM zAware server to communicate with the LDAP server and authenticate users. For a description of the LDAP settings provided, see "LDAP Settings tab" on page 185.
5. Click **Apply** to save the LDAP settings you specified. When you click **Apply**, the GUI displays a confirmation message to indicate that the web server must be restarted for your changes to be applied. This process might take a considerable amount of time to complete. While the server is restarting, other GUI users receive *page not available* errors until the restart process is complete.
6. Click **OK** to confirm that you want to apply your changes, or click **Cancel**.

What to do next

- To verify that IBM zAware is configured to authenticate against the LDAP server, do the following:
  1. Click **Role Mapping** on the **Security** tab.
  2. Assign your user ID to the *User* or *Administrator* role:
     a. In the **Role** field, select either **Administrator** or **User** as the role to which you want to assign your user ID.
     b. Specify a search filter to use to find your user ID in the LDAP directory. You can specify an asterisk (*) as a wildcard value at any position in the filter value.
     c. Select your user ID in the **Available users** list and click **Add** to move your user ID to the **Current mapped users** list.
     d. Click **Apply** to save your changes to the role mapping.
  3. Try logging in to the IBM zAware GUI with the user ID and password that is specified for you in the LDAP directory.

  If IBM zAware logs you in, LDAP user authentication is working correctly.
- If you cannot log in with LDAP user authentication, verify that:
  – Your user ID and password are correct.
  – The IBM zAware server can access the LDAP server.
  – The LDAP settings you specified are correct. To do so, log in to the IBM zAware GUI using the master user ID and password provided in the image profile for the IBM zAware partition.

- To allow additional users or groups to access IBM zAware, map users or groups in the new LDAP repository to specific IBM zAware roles. For instructions, see "Assigning users or groups to a role" on page 187.

LDAP Settings tab

You can use the **Security** > **LDAP Settings** tab on the Configuration page to authorize users to access the IBM zAware GUI through the use of a Lightweight Directory Access Protocol (LDAP) repository. To do so, specify the appropriate settings on the **LDAP Settings** tab.

Tip: Your installation can opt to use a local file-based repository for user authentication. For instructions, see "Setting up a local repository to secure access to the IBM zAware GUI" on page 108.

The LDAP settings are described in the following sections:

- "General LDAP settings"
- "Group LDAP settings" on page 186

For instructions for enabling LDAP authentication, see "Enabling LDAP authentication for IBM zAware users" on page 183.

To save the LDAP settings you specified, click **Apply**. If necessary, click **Restore** to restore the LDAP configuration values that were in effect before you clicked **Apply**.

Note that to apply the settings, the server is automatically restarted. This process might take a considerable amount of time to complete. While the server is restarting, other GUI users receive *page not available* errors until the restart process is complete.

General LDAP settings

Table 34. General LDAP settings

Setting	Description
LDAP server hostname	Enter the resolvable host name or IP address of the LDAP server to which you want to connect. The host name is required.
LDAP server port	Enter the port on which the LDAP server listens for TCP/IP connections. The value can range from 0 - 65535. The port is required.
Follow referrals	A referral is an entity that is used to redirect a client request to another LDAP server. A referral contains the names and locations of other objects. It is sent by the server to indicate that the information the client requested can be found at another location, possibly at another server or several servers. Select one of the following options: **Follow** Indicates that referrals to other LDAP servers will be followed. **Ignore** Indicates that referrals to other LDAP servers will be ignored. This option is selected by default. A selection is required.
Bind distinguished name	Enter the distinguished name used to bind to the LDAP repository. The name must be a string of length 0-512. If no name is specified, the server binds anonymously. The name is optional.
Bind password	Enter the password used to bind to the LDAP directory. The password must be a string of length 0-512. The password is optional.
Base distinguished name	Enter the distinguished name of a base entry in the repository. The name must be a string of length 1-512. The name is required.

Table 34. General LDAP settings (continued)

Setting	Description
Login attribute	Enter the LDAP attribute of a user entity used to login to the IBM zAware GUI. The attribute must uniquely identify a user in the directory. Typical login attributes include *uid*, *mail*, *primaryuserid*, and so on. The value must be a string of length 1-512. The default value is *uid*. The login attribute is required.
User object classes	Enter the object class or classes that are associated with user entities in the LDAP repository. Delimit multiple object classes with semicolons (;). Typical object classes include *Person*, *ePerson*, *inetOrgPerson*, and so on. The value must be a string of length 1-512. At least one user object class is required.
User search bases	Specify the base object of the directory (or level of the directory) from which to start a search for user entities in the LDAP repository. The search bases must be subtrees of the base distinguished name. Delimit multiple search bases with semicolons (;). The value must be a string of length 1-512. The user search base is optional. If unspecified, the base distinguished name is used.
SSL enabled	Select this option to enable secure socket communication to the LDAP server. If selected, you must supply the SSL certificate in the **LDAP server certificate** field. By default, SSL is disabled.
LDAP server certificate	Enter the Base64 encoded certificate that is required to validate the certificate of the LDAP server. This certificate should be the signer of the server certificate for the LDAP repository. A certificate is required only when SSL is enabled.

Group LDAP settings

Table 35. Group LDAP settings

Setting	Description
User group membership attribute	Enter the LDAP attribute of a user entity that indicates the groups to which an entry belongs. If your LDAP server does not support the group membership attribute, do not specify this attribute. The value must be a string of length 1-512. The user group membership attribute is optional.
User group membership scope	Select the scope of the user group membership attribute. You can select one of the following options: **Direct** Indicates that the attribute contains only immediate members of the group without members of subgroups. This option is selected by default. **Nested** Indicates that the attribute contains direct members and members nested within subgroups of this group. **All** Indicates that the attribute contains all direct, nested, and dynamic members. A selection is required if a value is specified in the **User group membership attribute** field.
Group object classes	Enter the object class or classes that are associated with group entities in the LDAP repository. Delimit multiple object classes with semicolons (;). Typical object classes include *groupOfNames*, *groupOfUniqueNames*, and so on. The value must be a string of length 1-512. At least one group object class is required.
Group search bases	Specify the base object of the directory (or level of the directory) from which to start a search for group entities in the LDAP repository. The search bases must be subtrees of the base distinguished name. Delimit multiple search bases with semicolons (;). The value must be a string of length 1-512. The group search base is optional. If unspecified, the base distinguished name is used.

Table 35. Group LDAP settings (continued)

Setting	Description
Group member attributes	For each group object class specified in the **Group object classes** field, indicate the LDAP attribute of a group entity in the object class that contains the members of the group. Delimit multiple group member attributes with semicolons (;). A value is required, and must be a string of length 1-512. For example, if the group object classes specification is *groupOfNames;groupOfUniqueNames*, the group member attributes specification might be *member;uniqueMember*.
Group member object classes	For each attribute specified in the **Group member attributes** field, specify the object classes of the group that uses the member attribute. The value must be a string of length 1-512. The group member object classes are optional. If unspecified, the member attributes apply to all group object classes.
Group member scope	Select the scope of the group member attribute. You can select one of the following options: **Direct** Indicates that the member attribute contains only direct members. This option is selected by default. **Nested** Indicates that the member attribute contains both direct and nested members. **All** Indicates that the member attribute contains direct, nested, and dynamic members. A selection is required.

Assigning users or groups to a role

In IBM zAware, a role represents the ability to perform one or more tasks in the IBM zAware graphical user interface (GUI). To assign users or groups to an IBM zAware role, use the **Security** > **Role Mapping** tab on the Configuration page.

Before you begin

Configure user authentication for the IBM zAware GUI. You can configure IBM zAware to authenticate users against a Lightweight Directory Access Protocol (LDAP) repository or a local file-based repository.

For more details, see one of the following topics:

- "Enabling LDAP authentication for IBM zAware users" on page 183
- "Setting up a local repository to secure access to the IBM zAware GUI" on page 108.

About this task

You can map authorized users and groups to specific roles: either Administrator or User. Users or groups with Administrator authority can use any task in the GUI, while those with User authority can view only the following pages and use only the actions as noted:

- On all views of the **Analysis** page, all controls and actions are permitted.
- On the **Interval** page, all controls and actions are permitted except for modifying the non-IBM rules status for a specific message ID. Only administrators can view and change a rules status value. IBM rules cannot be changed.
- On the **Notifications** page, all actions are disabled.
- On the **Systems** > **System Status** tab, all actions are disabled.
- On the **Systems** > **Model Groups** tab, all actions except for **Search Systems** are disabled.

This procedure describes how to add a user or group to a specific role, and how to remove a user or group from a role.

- Only a person with a user ID mapped to the Administrator role can add or remove users or groups from a role.
- An administrator cannot remove his or her user ID from the Administrator role. If an administrator attempts to remove the user or group through which his or her user ID is mapped to the Administrator role, IBM zAware rejects the removal request unless a duplicate Administrator role mapping exists for this administrator's user ID, either through an additional group or an individual user ID mapping.

Procedure

1. Log in to the IBM zAware GUI using the user ID and password that was provided in the hardware definition for the IBM zAware partition.
2. Expand the Administration category in the navigation pane and select **Configuration**. The Configuration page is displayed.
3. Click **Security** > **Role Mapping** to display the **Role Mapping** tab.
4. In the **Role** field, select either **Administrator** or **User** as the role to which you want to map particular users or groups.

   The IBM zAware server populates the **Current mapped users** and **Current mapped groups** lists with all users or groups that are currently mapped to the selected role. Initially, only the default master user ID appears in the **Current mapped users** list for both the Administrator role and the User role.
5. Specify the search filter to use when selecting user and group entries from the LDAP directory or local repository. You can specify an asterisk (*) as a wildcard value at any position in the filter value. An asterisk (*) is the default filter value. A filter value is required.

   The IBM zAware server populates the **Available users** and **Available groups** lists with user and group entries in the LDAP repository or local repository that match the search filter.
6. Specify the maximum number of matching entries the IBM zAware server can display in the **Available users** and **Available groups** lists. A search limit is required and must be a whole number in the range of 1-200. The default value is 20.

   The search limit applies to both users and groups; therefore, a search limit of 20 might return 40 entries: 20 users and 20 groups.
7. To add a user or group to the selected role, complete one or more of the following actions:
   - Select one or more users from the **Available users** list and click **Add** to move the users to the **Current mapped users** list.
   - Click **Add All** to move all the users to the **Current mapped users** list.
   - Select one or more groups from the **Available groups** list and click **Add** to move the groups to the **Current mapped groups** list.
   - Click **Add All** to move all the groups to the **Current mapped groups** list.
8. To remove a user or group from the selected role, complete one or more of the following actions:
   - Select one or more users from the **Current mapped users** list and click **Remove** to move the users to the **Available users** list.
   - Click **Remove All** to move all the users to the **Available users** list.
   - Select one or more groups from the **Current mapped groups** list and click **Remove** to move the groups to the **Available groups** list.
   - Click **Remove All** to move all the groups to the **Available groups** list.
9. Click **Apply** to save your changes to the role mapping. The Apply Role Mappings window is displayed.
10. Review your role assignments. If the changes are correct, click **Apply**.

Results

The IBM zAware GUI displays a confirmation message to indicate that the server must be restarted for the new role mappings to be applied. This process might take a considerable amount of time to complete because the server is automatically restarted with the new role mapping values. While the server is restarting, other GUI users receive *page not available* errors until the restart process is complete.

Role Mapping tab

You can use the **Security** > **Role Mapping** tab on the Configuration page to assign users or groups to an IBM zAware role, which permits them access to function provided in the IBM zAware graphical user interface (GUI).

Important: Before you can map users or groups to roles, you must configure IBM zAware to authenticate users against a Lightweight Directory Access Protocol (LDAP) repository or a local file-based repository. Otherwise, the **Role Mapping** tab displays only the master user ID that was provided in the partition hardware definition.

For more details, see one of the following topics:

- "Enabling LDAP authentication for IBM zAware users" on page 183
- "Setting up a local repository to secure access to the IBM zAware GUI" on page 108.

For instructions for assigning users or groups to a role and for a description of what tasks a role authorizes users or groups to perform, see "Assigning users or groups to a role" on page 187.

For a description of the items included on the **Role Mapping** tab, see Table 36.

Table 36. Items displayed in the Role Mapping tab

Item	Description
Role	Indicates the role to which users or groups are or will be assigned. Select **User** or **Administrator**. A selection is required. The IBM zAware server populates the **Current mapped users** and **Current mapped groups** lists with all users or groups that are currently mapped to the selected role. Initially, only the default master user ID appears in the **Current mapped users** list for both the Administrator role and the User role.
Filter	Specify the search filter to use when selecting user and group entries from the LDAP directory or local repository. You can specify an asterisk (*) as a wildcard value at any position in the filter value. An asterisk (*) is the default filter value. A filter value is required. The IBM zAware server populates the **Available users** and **Available groups** lists with user and group entries in the LDAP repository or local repository that match the search filter.
Search limit	Specify the maximum number of matching entries the IBM zAware server can display in the **Available users** and **Available groups** lists. A search limit is required and must be a whole number in the range of 1-200. The default value is 20. The search limit applies to both users and groups; therefore, a search limit of 20 might return 40 entries: 20 users and 20 groups.
Available Users	Lists the users that match the filter value, up to the search limit value, who are not assigned to an IBM zAware role. You can sort the list by clicking the list header.
Current Mapped Users	Lists the users who are assigned to the selected IBM zAware role. You can sort the list by clicking the list header.
Available Groups	Lists the groups that match the filter value, up to the search limit value, that are not assigned to an IBM zAware role. You can sort the list by clicking the list header.

Table 36. Items displayed in the Role Mapping tab (continued)

Item	Description
Current Mapped Groups	Lists the groups that are assigned to the selected IBM zAware role. You can sort the list by clicking the list header.
Buttons	Provides the following buttons: **Add** Moves the selected users or groups from the **Available users or groups** list to the **Current mapped users or groups** list. **Add All** Moves all the users or groups from the **Available users or groups** list to the **Current mapped users or groups** list. **Apply** Saves the changes you made to the role mappings. **Remove** Moves the selected users or groups from the **Current mapped users or groups** list to the **Available users or groups** list. **Remove All** Moves all the users or groups from the **Current mapped users or groups** list to the **Available users or groups** list. **Reset** Resets role mappings to the values that were in effect before you made any changes, only if you have not already clicked **Apply** to save your changes. **Search** Searches for user and group entries that match the specified search filter values, and displays entries that match the search filter value, up to the search limit value.

Apply Role Mappings window

Use the Apply Role Mappings window to review the changes you made to the role mappings and to confirm that you want to apply those changes.

For a description of the fields included in the Apply Role Mappings window, see Table 37.

To store your changes, click **Apply**. The IBM zAware GUI displays a confirmation message to indicate that the server must be restarted for the new role mappings to be applied. This process might take a considerable amount of time to complete because the server is automatically restarted with the new role mapping values. While the server is restarting, other GUI users receive *page not available* errors until the restart process is complete.

Table 37. Fields displayed in the Apply Role Mappings window

Field	Description
Roles with changes	Lists the roles for which there are changes to be reviewed.
Mapping changes for role	Select the role for which you want to view changes.
Users to remove	Lists the users that will be removed from the selected role.
Users to add	Lists the users that will be added to the selected role.
Groups to remove	Lists the groups that will be removed from the selected role.
Groups to add	Lists the groups that will be added to the selected role.

Specifying the duration of a browser session

The IBM zAware graphical user interface (GUI) allows you to configure the LTPA timeout value, which determines the duration of your browser session. By default, browser sessions time out after 12 hours (720 minutes). To modify this setting, use the **Security** > **LTPA Settings** tab.

Procedure

1. Expand the Administration category in the navigation pane and select **Configuration**. The Configuration page is displayed.
2. Click **Security** > **LTPA Settings** to display the **LTPA Settings** tab.
3. In the **LTPA timeout** field, specify the duration to use for a browser session in minutes. The allowable range of values is 10 - 525600 minutes (365 days).
4. Click **Apply** to save your changes. When you click **Apply**, the GUI displays a confirmation message to indicate that the web server must be restarted for your changes to be applied. This process might take a considerable amount of time to complete. Click **OK** to confirm that you want to apply your changes, or click **Cancel**.

Chapter 20. Managing IBM zAware operation and resources

You can accomplish most operations tasks through the IBM zAware graphical user interface (GUI). The following topics provide references when other interfaces or tools are required.

Accessing your notifications

A *notification* is a message notifying you of some occurrence in the IBM zAware environment that requires your awareness or response. That is, a notification can be informational in nature, or it can be an error that requires a response from you. To view and manage your notifications for IBM zAware, use the Notifications page.

When you have unread notification messages, the New label (New) is displayed in the navigation pane, to the right of the **Notifications** link. To display the Notifications page, click **Notifications** in the navigation pane. If no unread notification messages await your attention, the New label is not displayed.

On the Notifications page, each notification is displayed as a row in the Notification Messages table. The messages that are listed can be related to an action you performed, to an action that another user performed, or to independent server processing (such as automatically scheduled retraining). The list is shared across users, and is intended to inform you of activity in the IBM zAware environment.

The following information is displayed for each message:

Message ID
Provides the identifier of the message.

Message Text
Provides the text of the message.

Message Date/Time
Provides the date and the approximate time when the message was issued.

By default, IBM zAware appends the newest messages to the end of the list and preserves all notification messages until the IBM zAware partition is deactivated. If your user ID is assigned to the Administrator role, you can select entries in the table and select **Remove** in the **Actions** list to delete messages that you no longer need. You cannot undo this action.

To change the default sort, click the column headers.

To refresh the messages list, click **Refresh**.

Assigning storage devices to IBM zAware

To provide analytical data for monitored clients, IBM zAware requires continuous access to a set of Extended Count Key Data (ECKD) direct-access storage devices (DASD).

Attention: The recommended practice is to configure the IBM zAware partition such that it has access to only those channel path identifiers (IDs), control units, and I/O devices that are required for network connectivity and storage. If this practice is not followed, an IBM zAware administrator might inadvertently assign storage devices that are in use by other partitions. To avoid the potential loss of critical system and application data on storage devices that are connected to the IBM zAware host system, make sure that you use the GUI to assign only those storage devices that are intended for IBM zAware use. Check with your storage administrator to make sure that you select the appropriate storage devices to assign to the IBM zAware server.

To assign storage devices for IBM zAware to use for storing analytical data, use the **Data Storage** tab on the **Configuration** page. Storage-related tasks and the content and controls displayed on the **Data Storage** tab are described in the following sections:

- "Estimating external storage device requirements"
- "Fields on the Data Storage tab"
- "Data Storage Devices"

To refresh the information that is displayed on the **Data Storage** tab, click **Refresh**.

Note that no storage devices are listed until you connect storage devices to the IBM zAware partition. For instructions, see Chapter 12, "Configuring network connections and storage for the IBM zAware partition," on page 95.

Estimating external storage device requirements

IBM zAware stores the following analytical data on DASD:

- Current data from each z/OS or Linux monitored client, as well as priming data, if any.
- IBM zAware models for each z/OS monitored client and for each model group, which is an administrator-defined group of Linux clients.
- Analysis results for each z/OS or Linux monitored client.

IBM zAware sets default retention times for each of these types of analytical data and, through an automated process, removes the data when the retention time has elapsed. Your installation can change these defaults through the **Administration** > **Configuration** > **Analytics** tab in the IBM zAware GUI.

Storage requirements vary depending on the retention times for each type of analytical data and on the number of monitored systems that you plan to connect to IBM zAware. Start with 500 GB of storage for IBM zAware to use, plus 4 - 5 GB of storage for each monitored system.

If you increase the number of monitored clients, you need to configure an extra 4 - 5 GB of storage for each monitored system. If you increase the retention times of instrumentation data, training models, or analysis results, you also might need to increase the amount of persistent storage that IBM zAware can use. To determine whether you need to add storage devices, periodically use the **Administration** > **Configuration** > **Data Storage** tab to monitor the list of assigned storage devices, their status, and capacity.

Fields on the Data Storage tab

Table 38 provides a description of the fields that are displayed on the **Data Storage** tab.

Table 38. Fields on the Data Storage tab

Field	Description
Total capacity (GB)	Specifies the total capacity, in gigabytes (GB), of all the storage devices that are assigned to IBM zAware.
Total storage used (GB)	Specifies the total amount of space, in gigabytes, that IBM zAware has allocated on the storage devices.
Total storage used (%)	Specifies the percentage of the total capacity that is currently allocated (in use).

Data Storage Devices

The Data Storage Devices table lists the storage devices that are available and connected to the IBM zAware partition. Table 39 on page 195 provides a description of the columns in the Data Storage Devices table.

To sort the data in the table, click the column header for the appropriate column.

For more details about the **Add and Remove Devices** and the **Apply Pending Removals** actions, see "Adding and removing storage devices."

Table 39. Columns in the Data Storage Devices table

Column	Description
Device	Provides, in hexadecimal, the device number or unit address assigned to the device on which the volume is mounted.
Status	Provides the status of the device. Possible values are: **Available** Indicates that the device is not assigned to the IBM zAware server. Although the device is identified as available, it might be in use by another operating system or another IBM zAware server. Before assigning devices to the IBM zAware server, check with your storage administrator to make sure you select the appropriate storage devices. If the device remains available after IBM zAware processes an administrator request to add it, the device status is displayed as Available (Add Failed). **Being Added** Indicates that the IBM zAware server is formatting the device and preparing it for use. **In Use** Indicates that the device is assigned to the IBM zAware server. If the device remains in use after IBM zAware processes an administrator request to remove it, the device status is displayed as In Use (Remove Failed). **Being Removed** Indicates that the IBM zAware server is in the process of removing the device. **Pending Removal** Indicates that you selected to remove the device but IBM zAware was unable to immediately remove it. To complete the removal process, click **Apply Pending Removals** on the **Data Storage** tab. When the removal operation is complete, IBM zAware changes the status to *Available*. If the device remains in Pending Removal state after IBM zAware processes an administrator request to apply pending removals, the device status is displayed as Pending Removal (Remove Failed).
Device Type	Provides the type of device on which the volume is mounted.
Capacity (GB)	Provides the size of the volume in gigabytes. If the value in the Status column is *Available*, a dash (-) is displayed in the Capacity column.

Adding and removing storage devices

To assign storage devices for the IBM zAware server to use for storing analysis results, system behavior models, and data from monitored systems, go to the **Administration** > **Configuration** > **Data Storage** tab, and select **Add and Remove Devices** from the **Actions** list in the Data Storage Devices table. You can also use **Add and Remove Devices** to unassign previously assigned storage devices.

Before you begin

Complete the following actions:

- Connect storage devices to the IBM zAware partition. For instructions, see Chapter 12, "Configuring network connections and storage for the IBM zAware partition," on page 95.

- Check with your storage administrator to make sure that you know which specific storage devices you can assign to the IBM zAware server, and the intended use (normal operations or backup) for these devices.
  - The recommended practice is to configure the IBM zAware partition such that it has access to only those channel path identifiers (IDs), control units, and I/O devices that are required for network connectivity and storage. If this practice is not followed, an IBM zAware administrator might inadvertently assign storage devices that are in use by other partitions. To avoid the potential loss of critical system and application data on storage devices that are connected to the IBM zAware host system, make sure that you use the GUI to assign only those storage devices that are intended for IBM zAware use.
  - If any devices are to be used for storing backup copies of IBM zAware data, your installation must define two physically separate but equivalent sets of storage devices:
    - One set for IBM zAware to use for normal operations.
    - Another set for storing backup copies of data.

    The number of storage devices in each set must match, and each backup device must be equivalent in size to the device from which the data is copied. These number and size requirements also apply for configurations that contain primary and alternate IBM zAware partitions.
- Check with your storage administrator to confirm that sufficient storage will remain if you remove a device. Otherwise, IBM zAware might not have sufficient capacity to store data for monitored systems.
- Click **Refresh** to ensure that you are viewing the most recent information.

Procedure

1. To display the Configuration page, expand the Administration category in the navigation pane and select **Configuration**.
2. Click the **Data Storage** tab.
3. In the Data Storage Devices table, select **Add and Remove Devices** from the **Actions** list. The Add and Remove Devices window opens.
4. To add one or more storage devices for the IBM zAware server to use, complete the following steps.
   a. If the **Preserve data** option is displayed on the Add and Remove Devices window, select it only if you are adding storage devices that already contain backup copies of IBM zAware data. This option prevents IBM zAware from overwriting data on the devices to be added.
   b. Select one or more devices in the Devices Available list.
   c. Click **Add** to move the devices to the Devices In Use list.

   As an alternative, you can use **Add All** to move all the devices to the Devices In Use list.

Attention: Do not use **Add All** if any of the available storage devices are shared. If a device is shared and in use by another application, data will be lost or overwritten if the IBM zAware server formats the device.

5. To remove a storage device that the IBM zAware server is currently using, complete one of the following actions:
   a. Select one or more devices in the Devices In Use list.
   b. Click **Remove** to move the devices to the Devices Available list.

   As an alternative, you can use **Remove All** to move all the devices to the Devices Available list.

Attention: IBM zAware rejects any attempt to reduce storage below the amount that is currently in use by the server, unless you selected all in-use devices for removal. Removing all in-use devices is a destructive operation because all stored data is deleted and IBM zAware is no longer able to provide analytical data for any monitored clients.

6. Click **OK**.

Results

The IBM zAware server completes the following actions:

- Unless you selected the **Preserve data** option, the IBM zAware server formats the devices that were moved to the Devices In Use list. While the formatting is in progress, the device status is *Being Added*.

  Depending on the number of devices that you assign, this formatting process might take some time to complete. Periodically click **Refresh** to update the information in the Data Storage Devices table. When the formatting process is complete, the device status is *In Use*. As part of the formatting process, the volume serial (VOLSER) for the device is renamed.
- For devices that were moved to the Devices Available list, the IBM zAware server changes the status to *Being Removed* or *Pending Removal*, if the device cannot be removed immediately. To complete the removal process, click **Apply Pending Removals** on the **Data Storage** tab. IBM zAware unassigns the device, moves the data currently stored on the device to a device that is in use, and changes the device status to *Available*.

  Important: Removing a device requires IBM zAware to recycle the analytics engine. While the removal is in progress, only the Systems page is available. When the removal process completes, you must reconnect your monitored systems to the IBM zAware server. For more details, see "Starting and stopping data collection for your monitored systems" on page 203.

Add and Remove Devices window

You can use the Add and Remove Devices window to assign storage devices to the IBM zAware server for use and to remove previously assigned devices.

To display the Add and Remove Devices window, go to the **Administration** > **Configuration** > **Data Storage** tab, and select **Add and Remove Devices** from the **Actions** list in the Data Storage Devices table.

For a description of the items that are displayed in the Add and Remove Devices window, see Table 40.

For more details about adding and removing devices, see "Adding and removing storage devices" on page 195.

Table 40. Items displayed in the Add and Remove Devices window

Item	Description
Preserve data option	Prevents IBM zAware from overwriting data on the device to be added. This option is available only when no storage devices in the IBM zAware storage configuration are designated as "In use".
Devices Available	Lists the storage devices that are not assigned to the IBM zAware server.
Devices In Use	Lists the storage devices that are currently assigned to the IBM zAware server.

Table 40. Items displayed in the Add and Remove Devices window (continued)

Item	Description
Buttons	**Add** Use to move one or more individual devices from the Devices Available list to the Devices In Use list. **Add All** Use with caution. This button moves all devices in the Devices Available list to the Devices In Use list. After you click **OK**, IBM zAware formats and initializes the devices to be added. Unless these storage devices have been configured in the IODF for the exclusive use of IBM zAware, initializing these devices might result in the loss of data that other applications use. This button is not recommended for adding storage devices that contain backup copies of IBM zAware data. **Remove** Use to move one or more individual devices from the Devices In Use list to the Devices Available list. After you click **OK**, IBM zAware moves data that is stored on these devices to any devices that are currently in use. **Remove All** Use with caution. This button moves all devices in the Devices In Use list to the Devices Available list. After you click **OK**, IBM zAware removes all data from these devices and is no longer able to provide analytical data for any monitored clients. **OK** Initiates the add or remove request and closes the window. **Reset** Discards your selections and keeps the window open. **Cancel** Discards your selections, cancels the action, and closes the window.

Replacing storage devices

If needed, you can replace a storage device with a faster one. The procedure to replace the storage device varies depending on whether the device is used for data storage or product image.

Replacing a storage device for data storage

To replace a storage device for data storage, replicate the data on the original device to a new device and attach the new device to the LPAR configuration after you detach the original device.

About this task

Procedure

1. Deactivate IBM zAware. It is the best that you deactivate any product that is sending data to IBM zAware. For more information, see “Deactivating the IBM zAware partition” on page 216.
2. Replicate data on the original storage device to a new storage device by using the method of your choice. The size of the new device must match the size of the original device. For more backup and replication related information, see Table 12 on page 72.
3. In HMC of the host system, activate IBM zAware with the original storage device.
4. In the IBM zAware user interface, go to **Configuration** > **Data Storage** and click **Remove All** to remove the device. After applying Pending Removals on the **Data Storage** tab, the removal operation is completed. IBM zAware changes the disk status to **Available**.
5. In HMC of the host system, deactivate the IBM zAware partition.
6. Update the I/O definition file (IODF) of the host system, attach the new storage device, and detach the original storage device.

7. In HMC of the host system, activate the IBM zAware partition. You see the new storage device in the **Data Storage** tab.
8. Use the **Add All** button with **Preserve data** option to add the new device.

Replacing a storage device for image

You can replace the image storage device with a faster one.

Procedure

1. Find the ID of the storage device on which the existing software appliance image is installed. This device must be attached to the server that hosts the new Secure Service Container partition.
2. Configure and start a Secure Service Container partition with the boot option Secure Service Container Installer selected Connect to the Secure Service Container installer through the browser of your choice.
3. On the main page, click the plus (+) icon to install image files from local media. The page display changes to the Install Software Appliance page.
4. On the Install Software Appliance page, select Attach existing disk.
5. From the disk list, select the disk where the software appliance is on, and click Apply.
6. On the confirmation dialog, click Reboot to have the installer automatically reactivate the partition.
7. Click Yes to continue with the installation.

Results

The storage device is replaced. For more information, see the *IBM z Systems Secure Service Container User's Guide*, SC28-6978-02, in the IBM Support Portal at https://ibm.biz/Bd2Kqe Chapter 15. Moving an existing software appliance into a different Secure Service Container partition on the same system.

Specifying settings for the analytics engine

IBM zAware administrators use the **Administration** > **Configuration** > **Analytics** tab to view and optionally modify the configuration values that control the analytics engine. Tabs on the left of the **Analytics** tab display provide access to settings for z/OS or Linux monitored clients.

The controls and content displayed on the **Analytics** tab are described in the following sections:

- "Specifying settings for z/OS monitored clients"
- "Specifying settings for Linux monitored clients" on page 201

Specifying settings for z/OS monitored clients

Use the **Analytics** > **z/OS** tab to view and optionally modify the configuration values that the analytics engine uses for all z/OS monitored clients.

Table 41 on page 200 describes the fields that are displayed on the **Analytics** > **z/OS** tab. Each field contains default values that represent reasonable estimates for IBM zAware analytics for all z/OS monitored clients. These estimates might not be appropriate for monitored systems at your installation, so you might need to change the default values according to your knowledge of system workloads.

If you increase the retention times of instrumentation data, training models, or analysis results, you might need to increase the amount of persistent storage that IBM zAware can use. To determine whether you need to add storage devices, periodically use the **Data Storage** tab to monitor the list of assigned storage devices, their current status, and capacity.

The values on the **Analytics** > **z/OS** tab are global settings that apply for all z/OS monitored systems. You cannot specify different date ranges for individual monitored systems but you can manage the

training dates used for each monitored system through the Training Sets page. For more details, see Chapter 21, “Managing the training for monitored clients,” on page 221.

If you modify the configuration values, click **Apply** to store them. For new training period and training interval values to take effect for currently connected clients, you need to stop and reconnect those clients. For instructions, see “Starting and stopping data collection for your monitored systems” on page 203.

To undo any changes that you have not applied, click **Restore**.

Table 41. Fields on the **Analytics** > **z/OS** *tab*

Field	Description
Instrumentation data retention time	Specifies the number of consecutive calendar days for which the IBM zAware server keeps the data that is received from monitored clients. This data provides the source for training models so the retention time must match or exceed the duration that is specified for **Training period**. For example, if you specify 90 days: • The server stores client data for 90 consecutive days, whether instrumentation data is available on each of the days between that date range. • Periodically, the server deletes data that is older than 90 days. IBM zAware uses this value when it initiates an automated database pruning process to identify and remove obsolete information from its database. IBM zAware runs this pruning process only when the IBM zAware partition is reactivated. **Default value:** 365 days **Valid range:** 1 through 730
Training models retention time	Specifies the number of consecutive calendar days for which the IBM zAware server keeps training models for all monitored clients. If you enter 0 as the retention time, the server keeps only the current training model for each monitored client. IBM zAware uses this value when it initiates an automated database pruning process to identify and remove obsolete information from its database. IBM zAware runs this pruning process only when the IBM zAware partition is reactivated. **Default value:** 365 days **Valid range:** 0 through 730
Analysis results retention time	Specifies the number of consecutive calendar days for which the IBM zAware server keeps analysis results for all monitored clients. IBM zAware uses this value when it initiates an automated daily pruning process to identify and remove obsolete analysis results. **Default value:** 365 days **Valid range:** 30 through 3650

Table 41. Fields on the **Analytics** > **z/OS** *tab (continued)*

Field	Description
Training period	Specifies the number of consecutive calendar days that the IBM zAware server uses to identify the instrumentation data to include in training models. The instrumentation data that is received on days during this time period serves as input for creating the model of normal system behavior for each monitored client. If the monitored clients at your installation process workloads in a six-month cycle, for example, you can change the value to 180 days. The server builds a training model from the instrumentation data that is received during this time period, whether or not data is available for each day. **Default value:** 90 days **Valid range:** 1 through 365 For a new training value to take effect for currently connected clients, you need to stop and reconnect those clients. For instructions, see "Starting and stopping data collection for your monitored systems" on page 203. For more information about training periods, see "Understanding training periods and intervals" on page 221.
Training interval	Specifies the number of consecutive calendar days between automatic builds of system behavior models. IBM zAware uses this value to schedule automatic builds only after the initial model is built successfully. For an automatic build to be scheduled, the client must be connected to the IBM zAware server. **Default value:** 30 days **Valid range:** 7 through 365 For a new training value to take effect for currently connected clients, you need to stop and reconnect those clients. For instructions, see "Starting and stopping data collection for your monitored systems" on page 203. For more information about training intervals, see "Understanding training periods and intervals" on page 221.

Specifying settings for Linux monitored clients

Use the **Analytics** > **Linux** tab to view and optionally modify the configuration values that the analytics engine uses for all Linux monitored clients.

Table 42 on page 202 describes the fields that are displayed on the **Analytics** > **Linux** tab. Each field contains default values that represent reasonable estimates for IBM zAware analytics for all Linux monitored clients. These estimates might not be appropriate for monitored systems at your installation, so you might need to change the default values according to your knowledge of system workloads.

If you increase the retention times of instrumentation data, training models, or analysis results, you might need to increase the amount of persistent storage that IBM zAware can use. To determine whether you need to add storage devices, periodically use the **Data Storage** tab to monitor the list of assigned storage devices, their current status, and capacity.

The values on the **Analytics** > **Linux** tab are global settings that apply for all Linux monitored systems. You cannot specify different date ranges for individual monitored systems but you can manage the training dates used for each monitored system through the Training Sets page. For more details, see Chapter 21, "Managing the training for monitored clients," on page 221.

If you modify the configuration values, click **Apply** to store them. For new training period and training interval values to take effect for currently connected clients, you need to stop and reconnect those clients. For instructions, see "Starting and stopping data collection for your monitored systems" on page 203.

To undo any changes that you have not applied, click **Restore**.

Table 42. Fields on the **Analytics** > **Linux** *tab*

Field	Description
Instrumentation data retention time	Specifies the number of consecutive calendar days for which the IBM zAware server keeps the data that is received from monitored clients. This data provides the source for training models so the retention time must match or exceed the duration that is specified for **Training period**. IBM zAware uses this value when it initiates an automated database pruning process to identify and remove obsolete information from its database. IBM zAware runs this pruning process only when the IBM zAware partition is reactivated. **Default value:** 365 days **Valid range:** The minimum setting is the value in effect for the training period; the maximum value is 365
Training models retention time	Specifies the number of consecutive calendar days for which the IBM zAware server keeps training models for all Linux model groups. If you enter 0 as the retention time, the server keeps only the current training model for each model group. IBM zAware uses this value when it initiates an automated database pruning process to identify and remove obsolete information from its database. IBM zAware runs this pruning process only when the IBM zAware partition is reactivated. **Default value:** 365 days **Valid range:** 0 through 365
Analysis results retention time	Specifies the number of consecutive calendar days for which the IBM zAware server keeps analysis results for all monitored clients. IBM zAware uses this value when it initiates an automated daily pruning process to identify and remove obsolete analysis results. **Default value:** 365 days **Valid range:** 30 through 3650
Training period	Specifies the number of consecutive calendar days that the IBM zAware server uses to identify the instrumentation data to include in training models. The instrumentation data received on days during this time period serves as input for creating the model of normal system behavior for each Linux model group. The server builds a training model from the instrumentation data that is received during this time period, whether or not data is available for each day. **Default value:** 120 days **Valid range:** 1 through 365 For a new training value to take effect for currently connected clients, you need to stop and reconnect those clients. For instructions, see "Starting and stopping data collection for your monitored systems" on page 203. For more information about training periods, see "Understanding training periods and intervals" on page 221.

Table 42. Fields on the **Analytics** > **Linux** *tab (continued)*

Field	Description
Training interval	Specifies the number of consecutive calendar days between automatic builds of system behavior models. IBM zAware uses this value to schedule automatic builds only after the initial training period has elapsed. For the initial training period: • IBM zAware automatically schedules early training every seven days, starting from the first day for which IBM zAware has data available. The seven-day early training schedule continues until at least one of the group members is connected to IBM zAware for the configured training period; at that point, IBM zAware uses the configured training interval to schedule automatic training. • If an automatic training attempt fails and a model is not available, IBM zAware automatically retries the training attempt the next day and, if necessary, every following day until a model is successfully built. **Default value:** 30 days **Valid range:** 7 through 365 For a new training value to take effect for currently connected clients, you need to stop and reconnect those clients. For instructions, see "Starting and stopping data collection for your monitored systems." For more information about training intervals, see "Understanding training periods and intervals" on page 221.

Managing system connections and model groups

Through tabs on the Systems page, IBM zAware users and administrators can view information about monitored system connections and model groups. Administrators also can stop and restart the IBM zAware analytics engine, and create or modify model group definitions.

In the navigation pane, click **Systems** to display the Systems page and its tabs:

System Status
Displays information about all monitored systems that are or were connected to the IBM zAware server, and also displays the status of the IBM zAware analytics engine.

Model Groups
Provides a list of the administrator-defined collections of Linux monitored clients. Users can view defined model groups, and search them to find the group to which a particular monitored client belongs. Administrators also can create and modify model group definitions.

The content and controls displayed on the Systems page tabs are described in the following sections:

- "Starting and stopping data collection for your monitored systems"
- "Viewing the status of monitored clients" on page 205
- "System Status tab" on page 206
- "Managing groups of Linux monitored clients" on page 210
- "Model Groups tab" on page 212
- "Search Systems window" on page 215

Starting and stopping data collection for your monitored systems

Administrators can use controls on the **Systems** > **System Status** tab to instruct the analytics engine to start or stop collecting data from monitored systems. The process of stopping and restarting the analytics engine is called *recycling*.

Before you begin

To perform the actions described in this topic, your user ID must be assigned to the Administrator role.

About this task

You must explicitly recycle the analytics engine only after you modify the analytics configuration values, such as the training period or training interval. If you or another administrator performs an action that requires IBM zAware to stop the analytics engine, such as assigning priming data or modifying the topology, IBM zAware automatically restarts the engine.

When you click **Stop** ( ), the following occurs:

- IBM zAware stops the analytics engine for all monitored systems. No new data is collected or analyzed for the monitored systems. Previously collected data and analysis results are preserved.
- IBM zAware disconnects all monitored systems from the server by closing the connection that exists between the systems and the server.
- IBM zAware does not accept any new connections from monitored systems.

When you click **Start** ( ), IBM zAware starts the analytics engine for all monitored systems and accepts new connections from monitored systems. The systems might not be automatically reconnected.

- When z/OS monitored clients are disconnected from the server, they automatically attempt to reconnect and continue to buffer data for approximately 10 minutes. If the analytics engine is restarted within that time, the z/OS system reconnects and sends the buffered data to the IBM zAware server. If the z/OS system times out before the analytics engine is restarted, any buffered data is lost and you must manually reconnect the system. To reconnect a z/OS system, issue the SETLOGR command.

  ```
  SETLOGR FORCE,ZAICONNECT,LSN=SYSPLEX.OPERLOG
  ```

- When Linux monitored clients are disconnected from the server, they normally attempt to reconnect to the server; if they do not reconnect, you must manually reconnect them. To reconnect a Linux system, restart the syslog daemon, using the appropriate command for the type of syslog daemon and for the type of initialization (init) process that is in use on the Linux system.

When the system is reconnected, the analytics engine resumes collecting and analyzing the data that is received from the monitored systems.

If the system was disconnected for an extended period of time and buffered data was lost, IBM zAware might not have enough available data to create a model, or the model might not be representative of normal system behavior. For z/OS monitored systems only, you can use the z/OS bulk load client for IBM zAware to provide the missing data to the IBM zAware server. If you do so, keep in mind that the server does not analyze priming data, so analytical data is not available for the time period during which the monitored client was disconnected. Also, the process of assigning priming data results in automatic recycling of the analytics engine and the disconnection of all monitored clients, so you need to determine whether the missing message data is worth this disruption to your IBM zAware environment.

Procedure

1. To display the **System Status** tab, select **Systems** in the navigation pane.
2. To ensure that you are viewing the most recent status information, click **Refresh**.
3. Check the status value shown in the “Analytics engine status” field, which indicates whether or not you can start or stop the engine. Table 43 on page 205 lists the possible status values and administrator actions. If the suggested action is to wait for an in-progress operation to complete, click **Refresh** at least once to determine whether the current operation has completed. If an in-progress operation completes successfully, IBM zAware automatically restarts the analytics engine, and its status changes to Running.

Table 43. Analytics engine status values and administrator actions

Analytics engine status value	Administrator actions
Quiesced because the engine is being recycled	Wait for the current operation to complete.
Running	Click **Stop** to stop the analytics engine.
Status cannot be determined	Click **Refresh** at least once to obtain the latest status information. If this status value persists, check the Notifications page for error messages that might indicate potential problems, and take the appropriate corrective action.
Stopped because migration is in progress	Wait for the current operation to complete.
Stopped because migration is required	Go to the **Configuration** > **Migration** tab to start the automated migration.
Stopped because of a storage error	Go to the Notifications page to look for error messages that indicate problems with data storage devices. When the problem is corrected, IBM zAware automatically restarts the analytics engine.
Stopped because the storage configuration is incomplete	Go to the **Configuration** > **Data Storage** tab to assign storage devices. After devices are successfully added to the storage configuration, IBM zAware automatically restarts the analytics engine.
Stopped by an administrator	Click **Start** to restart the analytics engine.
Unavailable because database configuration is in progress	Wait for the current operation to complete.
Unavailable because storage configuration is in progress	Wait for the current operation to complete.

Viewing the status of monitored clients

A *monitored client* is a z/OS or Linux system that is configured to send data to the IBM zAware server for analysis. To view the status of the systems that IBM zAware is monitoring, use the **System Status** tab on the Systems page.

Before you begin

To ensure that you are viewing the most recent status information, click **Refresh**.

Procedure

1. To display the System Status tab, select **Systems** in the navigation pane.
2. In the "Analytics engine status" field, verify that the analytics engine is running. If the status is not Running, IBM zAware is not collecting data from or analyzing data for the monitored systems.
3. In the IBM zAware Monitored System Data Suppliers table, do the following:
   - View the list of systems that are or were connected to the IBM zAware server.
   - Use the Type column to determine whether the monitored client is a z/OS or Linux system.
   - Use the Status column and the Connect Start Time column to determine which systems are connected to (Active) or disconnected from (Inactive) the IBM zAware server, and to determine when the current or last connection started.
   - Use the Instrumentation Data Type column to identify the type of data that the monitored system is supplying to the IBM zAware server.
     – z/OS systems can send two types of data:
       - Operations log (OPERLOG) data, when the system is sending current message traffic through the z/OS system logger. OPERLOG data contains all messages and commands from all z/OS systems in a sysplex.

- System log (SYSLOG) data, when the system sending priming data by running the z/OS bulk load client for IBM zAware. SYSLOG data contains:
  - All messages issued through WTL macros
  - All messages entered by LOG operator commands
  - Usually, the hardcopy log
  - Any messages routed to the SYSLOG from any system component or program

– Linux systems send system log (syslog) data, which is well-known, standardized, UNIX syslog data; for example, the contents of the `/var/log/messages` directory.

System Status tab

You can use the **System Status** tab to view information about the monitored systems that are or were previously connected to the IBM zAware server, to verify that the analytics engine is running, and, if your user ID is assigned to the Administrator role, to stop or start the analytics engine.

To display the **System Status** tab, click **Systems** in the navigation pane. The controls and content displayed on the **System Status** tab are described in the following sections:

- "About the analytics engine"
- "About monitored clients" on page 208

To display the most recent status for the analytics engine and the monitored systems, click **Refresh**.

About the analytics engine

The *analytics engine* is the component of IBM zAware that manages the data the server receives from each monitored system. Management actions include reading, storing, processing, and analyzing the data, as well as determining when to build new models for each monitored system or model group.

The **System Status** tab provides controls that IBM zAware administrators can use to start or stop the analytics engine. **Start** is enabled when the engine is stopped. **Stop** is enabled when the engine is running. Both buttons are disabled when the engine is quiesced or unavailable.

IBM zAware recycles the engine when you perform actions that require the engine to update the database where it stores the data. Such actions include assigning priming data, modifying the sysplex topology, or removing storage devices. You must explicitly recycle the analytics engine only after you modify the analytics configuration values, such as the training period or training interval. As an alternative to explicitly recycling the analytics engine, you can reconnect all monitored clients for your configuration changes to take effect.

When the analytics engine is stopped or recycled, IBM zAware disconnects all monitored systems from the server.

- When z/OS monitored clients are disconnected from the server, they automatically attempt to reconnect and continue to buffer data for approximately 10 minutes. If the analytics engine is restarted within that time, the z/OS system reconnects and sends the buffered data to the IBM zAware server. If the z/OS system times out before the analytics engine is restarted, any buffered data is lost and you must manually reconnect the system. To reconnect a z/OS system, issue the SETLOGR command.

  ```
  SETLOGR FORCE,ZAICONNECT,LSN=SYSPLEX.OPERLOG
  ```

- When Linux monitored clients are disconnected from the server, they normally attempt to reconnect to the server; if they do not reconnect, you must manually reconnect them. To reconnect a Linux system, restart the syslog daemon, using the appropriate command for the type of syslog daemon and for the type of initialization (init) process that is in use on the Linux system.

For more details about the consequences of starting or stopping the analytics engine, see "Starting and stopping data collection for your monitored systems" on page 203.

The following list describes the possible states displayed for the analytics engine.

Quiesced because the engine is being recycled
Indicates that the analytics engine is temporarily not receiving data from monitored systems, and that all connections between the server and its monitored systems have been disconnected, because a request that requires the engine to be recycled is being processed.

If the analytics engine was running before you issued the request, IBM zAware restarts the engine when your request completes. After the restart, you might need to reconnect the monitored systems to the IBM zAware server.

Running
Indicates that the analytics engine is managing the data that the monitored systems are transferring to the IBM zAware server.

Status cannot be determined
Indicates that the IBM zAware server is unable to determine the current state of the analytics engine.

Stopped because migration is in progress
Indicates that the analytics engine is not receiving data from monitored systems, and that all connections between the server and its monitored systems have been disconnected, because an administrator started the automated migration process through the **Administration** > **ConfigurationMigration** tab.

Stopped because migration is required
Indicates that the analytics engine is not receiving data from monitored systems, and that all connections between the server and its monitored systems have been disconnected, because IBM zAware detected that migration is required for this IBM zAware environment.

Stopped because of a storage error
Indicates that the analytics engine is not receiving data from monitored systems, and that all connections between the server and its monitored systems have been disconnected, because IBM zAware detected an error in the storage configuration.

Stopped because the storage configuration is incomplete
Indicates that the analytics engine is not receiving data from monitored systems, and that all connections between the server and its monitored systems have been disconnected, because the storage configuration is not complete.

Stopped by an administrator
Indicates that the analytics engine is not receiving data from monitored systems, and that all connections between the server and its monitored systems have been disconnected, because an administrator explicitly stopped the engine.

Unavailable because database configuration is in progress
Indicates that the analytics engine is temporarily not managing the data transferred by monitored systems, and that all connections between the server and its monitored systems have been disconnected, because a database configuration operation is in progress.

If the analytics engine was running before the database configuration operation began, IBM zAware restarts the engine when the request completes. After the restart, you might need to reconnect the monitored systems to the IBM zAware server.

Unavailable because storage configuration is in progress
Indicates that the analytics engine is temporarily not managing the data transferred by monitored systems, and that all connections between the server and its monitored systems have been disconnected, because a storage configuration operation is in progress.

If the analytics engine was running before the storage configuration operation began, IBM zAware restarts the engine when the request completes. After the restart, you might need to reconnect the monitored systems to the IBM zAware server.

About monitored clients

A *monitored client* is a z/OS or Linux system that is configured to send data to the IBM zAware server for analysis. To detect problems, IBM zAware compares the system and application messages in these log files to a model of normal behavior, and highlights anomalous results through the IBM zAware graphical user interface (GUI).

The IBM zAware Monitored System Data Suppliers table lists the z/OS and Linux systems that IBM zAware is monitoring, and provides the connection status for each system. For a description of the columns in the table, see Table 44 on page 209.

You can change the sort order of table entries by clicking any of the column headings. The table footer provides a total count of table entries, and the number of selected table entries. Depending on the total count, you might need to use the vertical scroll bar to view all of the table entries, or use the Filter field in the table header to limit the entries. To make sure the display contains the latest information, click **Refresh**.

To quickly determine the total number of active systems, you can use Filter functions:

1. Click the down arrow to the right of the Filter actions icon ( ), which is located in the far right corner of the table header.
2. Select **Build Filter** from the Filter actions list to open the Build Filter window.
3. On the Build Filter window, create a rule by selecting the following options: Status, equal, and Active. Then click **Filter** to create and apply this filter rule.

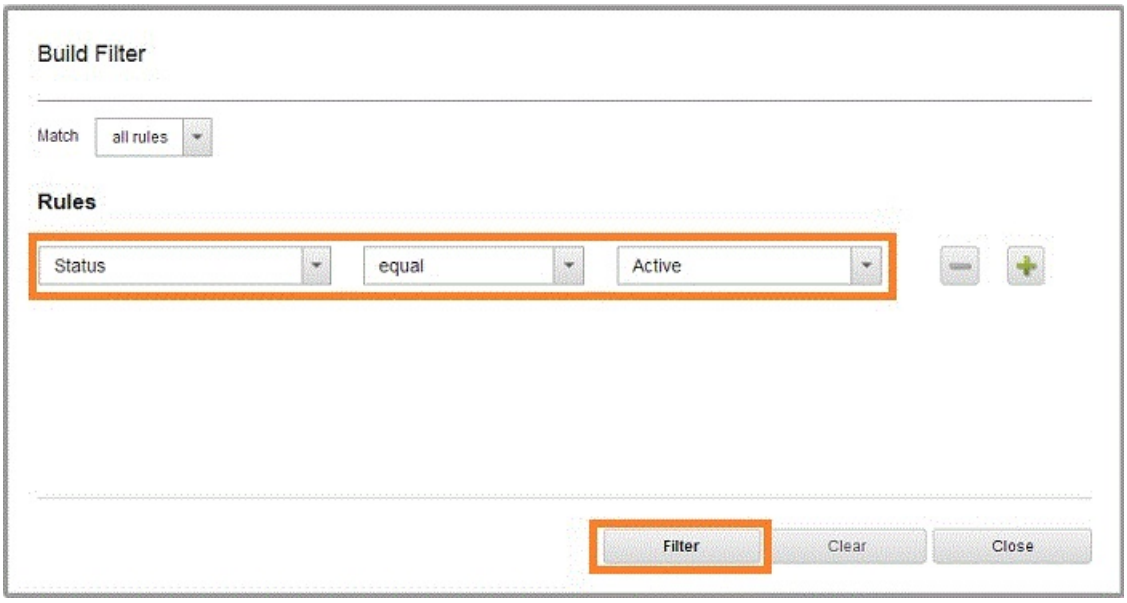

Figure 53. Build Filter window

4. In the row above the column headings in the IBM zAware Monitored System Data Suppliers table, the display indicates the filter result: only 43 active systems in the total list of 81 systems.

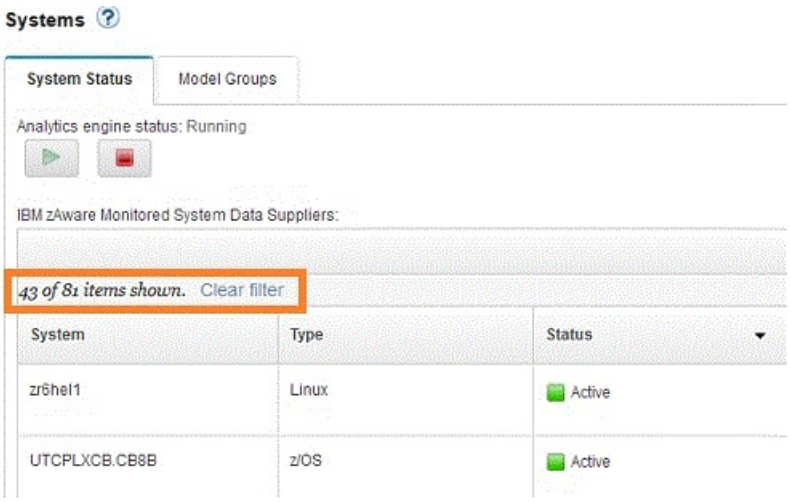

Figure 54. Example filter results

Note: This table provides you with a history of all the systems that have been connected to the server unless an administrator explicitly removes a system through the **Remove Selected Systems** action on the **Administration** > **Configuration** > **Topology** tab. If you remove a monitored system from your installation or if you move a system to another sysplex, the systems remain in this table unless an administrator explicitly removes them.

Table 44. Columns in the IBM zAware Monitored System Data Suppliers table

Column	Description
System	Provides the name of the monitored system. • For a z/OS monitored client, the name has the format *sysplex-name.system-name*, where *sysplex-name* is the name of the sysplex to which the system belongs and *system-name* is the name of the system • For a Linux monitored client, the name is a fully qualified domain name, a host name, or an IP address.
Type	Indicates whether the monitored client is a z/OS or Linux monitored client.
Status	Indicates whether the system is connected to the IBM zAware server. The system can have one of the following status values. **Active** Indicates that the system is connected to the IBM zAware server. The system might or might not be transmitting data to the server. **Inactive** Indicates that the system was previously connected to the IBM zAware server but is disconnected.
Instrumentation Data Type	Identifies the type of data that the monitored system is supplying to the IBM zAware server. • z/OS systems can send two types of data: – Operations log (OPERLOG) data, when the system is sending current message traffic through the z/OS system logger. OPERLOG data contains all messages and commands from all z/OS systems in a sysplex. – System log (SYSLOG) data, when the system sending priming data by running the z/OS bulk load client for IBM zAware. SYSLOG data contains: - All messages issued through WTL macros - All messages entered by LOG operator commands - Usually, the hardcopy log - Any messages routed to the SYSLOG from any system component or program • Linux systems send system log (syslog) data, which is well-known, standardized, UNIX syslog data; for example, the contents of the `/var/log/messages` directory.

Table 44. Columns in the IBM zAware Monitored System Data Suppliers table (continued)

Column	Description
Connect Start Time	Provides the date and time when the current or last connection started.

Managing groups of Linux monitored clients

Linux systems are often activated and deactivated frequently to meet various operational needs, such as additional resources, system availability, and system maintenance. For IBM zAware to evaluate message traffic for such systems, they must belong to a model group. Before IBM zAware can provide analysis results for a Linux system, that system must belong to a model group other than the UNASSIGNED model group, and IBM zAware must have successfully built a model of system behavior for that model group. Use this procedure to create, modify, and delete model group definitions, and to determine the model group to which a specific Linux system belongs.

Before you begin

To perform the steps in this procedure, you need an ID that is mapped to the IBM zAware Administrator role.

About this task

A model group definition consists of a name for the group, an optional description, a membership rule that is based on Linux system naming conventions, and a membership evaluation order. If no administrator-defined model groups exist, a Linux monitored system automatically becomes part of the UNASSIGNED model group. Otherwise, to assign a Linux system to an administrator-defined model group, IBM zAware compares the Linux system name to the membership rule for each model group, according to the specified membership evaluation order. When IBM zAware finds the first matching membership rule, it assigns the Linux system to the model group associated with that rule. IBM zAware performs this comparison when:

- A Linux system connects to the IBM zAware server and starts sending data.
- When an administrator modifies the membership rule or evaluation order for a model group, and saves the changes.
- When an administrator deletes a model group definition.

To find the model group to which a specific Linux system belongs, complete the following steps.

1. Navigate to the **Systems** > **Model Groups** tab.
2. In the header of the Model Groups table, click **Actions** to open the Actions list.
3. Select **Search Systems** to open the Search Systems window.
4. Enter the name of the monitored system in the Name field and click **OK**.

   IBM zAware returns you to the **Model Groups** tab and, in the Model Groups table, identifies the group by displaying a checkmark next to the model group name.
5. To verify that the system belongs to the indicated model group, click the link for that model group in the Name column of the Model Groups table. In the Model Group Details pane, check the "Known matching member systems" table for the system name.

Because IBM zAware builds one model for a group of Linux systems with similar workloads, and uses that model to compare to current syslog data from each system in the group, modifying or deleting a model group has a direct effect on the analysis results for Linux monitored systems. For best results, administrators need to follow a system naming convention that reflects the function of or workload supported by different Linux systems.

For information about building models to accurately reflect normal system behavior, see Chapter 11, "Planning to create IBM zAware models," on page 87.

Procedure

1. To work with Linux model groups, navigate to the **Systems** > **Model Groups** tab. The Model Groups table lists the UNASSIGNED model group and all administrator-defined model groups, if any. The table footer provides a total count of table entries, and the number of selected table entries. Depending on the total count, you might need to use the vertical scroll bar to view all of the table entries, , or use the Filter field in the table header to limit the entries.
2. To create a new model group, you have two choices: using the **New Group** action, or creating a group by selecting one or more systems from the UNASSIGNED group. To use the **New Group** action, complete the following steps; otherwise, skip to step 3.
   a. In the header of the Model Groups table, click **Actions** to open the Actions list.
   b. Select **New Group** to open the Model Group Details pane. If that pane was already open, the field values are cleared and NEWGROUP is displayed in the Name field. For field descriptions, see "Fields in the Model Group Details pane" on page 214.
   c. Type a new value in the Name field. The name can contain alphanumeric characters (A through Z, a through z, and 0 through 9), underscores (_), and blanks.
   d. Optional: Provide a description in the Description field.
   e. Required: In the "Membership rule" field, specify the text string for IBM zAware to use when assigning Linux systems to this model group. The text string is a full or partial Linux system name, which can be a fully qualified domain name, a hostname, or an IP address. The text string can contain alphanumeric characters (A through Z, a through z, and 0 through 9), periods (.), colons (:), dashes (–), and forward slashes (/). To specify a partial name or IP address, use an asterisk (*) or question mark (?) as a wildcard for any one character in the text string (for example, `LNXVM5*`).
   f. Required: Specify a value for the "Membership evaluation order" field. When you specify an evaluation order, make sure that more specific membership rules are evaluated before more generic rules; otherwise, a Linux system might be assigned to the wrong group. For example, suppose that you have several systems with names that range from LNXVM50 to LNXVM59. If you define a group for them with a rule of LNXVM5*, that rule must be moved higher in the evaluation order than a more general rule, such as LNXVM*.
   g. Click **Evaluate Membership** and check the results presented in the "Known matching member systems" table. If the membership rule for a different model group definition contains an error, IBM zAware prompts you to correct the error before you can successfully evaluate the membership for the model group that you are creating or modifying.
   h. When you are satisfied with the membership results, click **Save** to save your changes.
3. To create a model group by selecting one or more systems from the UNASSIGNED group, complete the following steps.
   a. Click the link for the UNASSIGNED group in the Name column in the Model Groups table. The Model Group Details pane opens.
   b. Select one or more systems listed in the "Known matching member systems" table, and click **Create Model Groups**. IBM zAware creates one model group definition for each selected system, and the Model Group Details pane content changes to display the details for the first selected system. IBM zAware fills in the Name and "Membership rule" fields with the system name. If any characters in the system name are not allowed in the Name field, IBM zAware strips out those characters.
   c. Edit fields in the Model Group Details pane as necessary, and click **Save** to save the new model group definition.
   d. If you selected more than one system, find the name of the new model group for that system in the Model Groups table, and click the link in the Name column to the open Model Group Details pane. Edit the appropriate fields and click **Save**. Repeat as necessary for any other systems that you selected in the "Known matching member systems" table.
4. To modify an existing model group, complete the following steps.

a. Select the model group by clicking the link in the Name column in the Model Groups table. The Model Group Details pane opens.

   Note that you cannot explicitly modify the UNASSIGNED model group, but the list of systems that belong to it can change after you create, modify, or remove any administrator-defined model groups.

b. To change the membership evaluation order, use the **Move Up** or **Move Down** actions in the Actions list, or change the value of the "Membership evaluation order" field in the Model Group Details pane.

c. To change any other elements of the model group definition, use the fields in the Model Group Details pane. For field descriptions, see "Fields in the Model Group Details pane" on page 214.

d. If you change the membership rule or evaluation order, click **Evaluate Membership** and check the results presented in the "Known matching member systems" table. If the membership rule for a different model group definition contains an error, IBM zAware prompts you to correct the error before you can successfully evaluate the membership for the model group that you are creating or modifying.

e. When you are satisfied with the membership results, click **Save** to save your changes.
   - After you successfully save your changes, you can re-edit the model group definition, but you cannot use **Reset** to restore the model group definition to its previously saved values.
   - If another administrator has removed or modified the same model group after you began editing it, IBM zAware cannot save your changes.
     – If the model group that you edited has been removed, you are returned to a refreshed view of the Model Groups tab.
     – Otherwise, you are returned to a refreshed view of the Details pane, where you can make changes using the latest model group information.

What to do next

- For a new or modified model group, navigate to the **Administration** > **Training Sets** > **Model Groups** tab to request training. Note that removing a model group results in changes to the system membership of one or more administrator-defined model groups; in this case, you might also want to request new models to be built. For more information, see "Training sets for Linux model groups" on page 238.
- If you need to remove an existing model group, complete the following steps.
  1. Select the model group by clicking the link in the Name column in the Model Groups table.
  2. In the header of the Model Groups table, click **Actions** to open the Actions list.
  3. Select **Remove Group**.

  This action does not remove any associated analysis results but does cause current member systems to be assigned to a different model group.

Model Groups tab

You can use the **Model Groups** tab to view information about administrator-defined collections of Linux monitored clients. If your user ID is assigned to the Administrator role, you can create new or modify existing model group definitions.

To display the **Model Groups** tab, click **Systems** in the navigation pane, then click the **Model Groups** tab. To view more details about a model group, click the link in the Name column of the Model Groups table to expand the Model Group Details pane. Only administrators can enter or edit information in the Model Group Details pane.

The controls and content displayed on the **Model Groups** tab are described in the following sections:

- "About Linux model groups" on page 213
- "Fields in the Model Groups table" on page 213
- "Fields in the Model Group Details pane" on page 214

About Linux model groups

A model group is a collection of one or more systems that handle the same type of workload, and thus can be expected to exhibit similar behavior. When a Linux system is first connected to the IBM zAware server and starts sending data, IBM zAware assigns the new monitored system to the UNASSIGNED model group or, if any exist, to an administrator-defined model group. Before IBM zAware can provide analysis results for a Linux system, that system must belong to a model group other than the UNASSIGNED model group, and IBM zAware must have successfully built a model of system behavior for that model group.

A model group definition consists of a name for the group, an optional description, a membership rule that is based on Linux system naming conventions, and a membership evaluation order. To assign a newly connected Linux system to an administrator-defined model group, IBM zAware compares the Linux system name to the membership rule for each model group, according to the specified membership evaluation order.

- IBM zAware first compares the system name to the membership rule of the model group with the membership evaluation order of 1; if the name does not match the rule, IBM zAware, then compares the name to the rule of the group model with the evaluation order of 2, and so on, until it finds a match.
- When IBM zAware finds the first matching membership rule, it assigns the Linux system to the model group associated with that rule.
- If the name does not match any membership rule, the new monitored system becomes part of the UNASSIGNED model group. The UNASSIGNED model group contains any Linux systems that do not belong to a defined model group.

"Managing groups of Linux monitored clients" on page 210 contains instructions for administrators to use when creating or modifying model groups.

Fields in the Model Groups table

Table 45 on page 214 lists the fields and controls in the Model Groups table. The table footer provides a total count of table entries, and the number of selected table entries. Depending on the total count, you might need to use the vertical scroll bar to view all of the table entries, , or use the Filter field in the table header to limit the entries.

Table 45. Fields in the Model Groups table

Field	Description
Actions	Contains the list of management tasks for model groups. **Move Up** Changes the membership evaluation order of the selected model group by subtracting one. For example, if an administrator selects a model group with the membership evaluation order of 3 and selects **Move Up**, the membership evaluation order changes from 3 to 2. Only administrators can select this action. **Move Down** Changes the membership evaluation order of the selected model group by adding one. For example, if an administrator selects a model group with the membership evaluation order of 3 and selects **Move Down**, the membership evaluation order changes from 3 to 4. Only administrators can select this action. **Remove Group** Deletes the selected model group. This action does not remove any associated analysis results but does cause current member systems to be assigned to a different model group. Only administrators can select this action. **New Group** Opens the Model Group Details pane, or clears the fields in that pane, so an administrator can define a new model group. Only administrators can select this action. **Search Systems** Opens the Search Systems window, which is described in "Search Systems window" on page 215.
Name	Indicates the name of an administrator-defined collection of Linux monitored systems.
Membership Evaluation Order	Indicates the position of this model group in the search order that IBM zAware uses when it compares a Linux system name to a membership rule. The model group with the membership evaluation order of 1 is searched first, the group with the evaluation order of 2 is searched next, and so on.
Membership Rule	Specifies the text string that IBM zAware uses to assign a Linux system to a model group. The text string is a full or partial Linux system name, which can be a fully qualified domain name, a hostname, or an IP address. The text string can contain alphanumeric characters (A through Z, a through z, and 0 through 9), periods (.), colons (:), dashes (–), and forward slashes (/).
Description	Provides an optional description of the selected model group.

Fields in the Model Group Details pane

Table 46 on page 215 lists the fields and controls in the Model Group Details pane, which opens in view-only mode for users, and in edit mode for administrators. To display the Model Group Details pane, click the link in the Name column of the Model Groups table; administrators can display this pane also by selecting the **New Group** action from the **Actions** list.

Table 46. Fields in the Model Group Details pane

Field	Description
Name	Indicates the name of an administrator-defined collection of Linux monitored systems. This field contains the value NEWGROUP when an administrator selects **New Group** from the **Actions** list. An administrator can replace this name by typing a new value in the Name field. The Name field is a required field for a model group definition. The name can contain alphanumeric characters (A through Z, a through z, and 0 through 9), underscores (_), and blanks.
Description	Provides an optional description of a selected or a new model group.
Membership Rule	Specifies the text string that IBM zAware uses to assign a Linux system to a model group. The text string is a full or partial Linux system name, which can be a fully qualified domain name, a hostname, or an IP address. The text string can contain alphanumeric characters (A through Z, a through z, and 0 through 9), periods (.), colons (:), dashes (–), and forward slashes (/). The Membership Rule field is a required field for a model group definition. To specify a partial name or IP address, use an asterisk (*) or question mark (?) as a wildcard for any one character in the text string (for example, `LNXVM5*`).
Membership Evaluation Order	Indicates the position of this model group in the search order that IBM zAware uses when it compares a Linux system name to a membership rule. The model group with the membership evaluation order of 1 is searched first, the group with the evaluation order of 2 is searched next, and so on. The Membership Evaluation Order field is a required field for a model group definition. When you specify an evaluation order, make sure that more specific membership rules are evaluated before more generic rules; otherwise, a Linux system might be assigned to the wrong group. For example, suppose that you have several systems with names that range from LNXVM50 to LNXVM59. If you define a group for them with a rule of LNXVM5*, that rule must be moved higher in the evaluation order than a more general rule, such as LNXVM*.
Known matching member systems table	Lists the names of Linux systems with names that match the membership rule. The table footer provides a total count of table entries, and the number of selected table entries. Depending on the total count, you might need to use the vertical scroll bar to view all of the table entries, , or use the Filter field in the table header to limit the entries.
Evaluate Membership	Provides a method of populating or refreshing the contents of the “Known matching member systems” table.
Create Model Groups	Creates one model group definition for each selected system in the “Known matching member systems” table.
Save	Saves the information for a new model group, or saves the changes to an existing model group definition.
Reset	Restores an existing model group definition to its previously saved values, by undoing any unsaved edits.

Search Systems window:

You can use the Search Systems window to find the model group to which a particular system belongs. To display the Search Systems window, open the Actions list in the Model Groups table on the **Systems** > **Model Groups** tab, and select **Search Systems**.

Fields in the Search Systems window

Table 47 describes the fields that are displayed in the Search Systems window.

Table 47. Fields displayed in the Search Systems window

Field	Description
Name	Specifies the search string for IBM zAware to use when searching the member systems in model group definitions.
OK	Starts the search for the model group that contains the system identified in the Name field. If IBM zAware finds a system name that matches the search string, it returns to the **Model Groups** tab and, in the Model Groups table, identifies the group by displaying a checkmark next to the model group name.
Cancel	Closes the Search Systems window and returns to the **Model Groups** tab, without searching for a system.

Monitoring processor, memory, and storage resources

The IBM zAware server monitors its own memory and storage use, and issues notification messages if usage exceeds specific thresholds. IBM zAware administrators can periodically check the **Notifications** page for these messages. In addition, administrators can view information about the IBM zAware partition and its use of processor and memory resources through several tools.

- The System Activity display for the IBM zAware partition provides information about processor resources. To access the System Activity display, use the **Monitors Dashboard** task in the HMC for the IBM zAware host system.
- The Partition Data Report section of the CPU Activity report, which is available through z/OS Resource Management Facility (RMF™), provides information about processor resources. When shared processors are used, this report might provide data on processor usage.

  For more information about the Partition Data Report, see *z/OS Resource Measurement Facility™ Report Analysis*, SC33-7991.

For information about changing the processor and memory resources that are defined for a logical partition, see *z Systems PR/SM Planning Guide*, SB10-7162.

For information about HMC and SE tasks, see IBM Knowledge Center at http://www.ibm.com/support/knowledgecenter/.

Deactivating the IBM zAware partition

The following steps describe the formal procedure for deactivating the IBM zAware partition.

1. Stop the IBM zAware analytics engine and the data transmission from monitored clients.
   a. Go to the **Systems** > **System Status** tab, and click **Stop** ( ) to stop the analytics engine. This action prevents the IBM zAware server from accepting any data transmission from clients.
   b. To prevent monitored systems from experiencing communication errors, stop them from transmitting data.
      - For monitored z/OS clients, use the **SETLOGR** command on each z/OS system to prevent the systems from attempting to reestablish the TCP/IP connection to the IBM zAware server.

        ```
        SETLOGR FORCE,ZAIQUIESCE,ALL
        ```

      - For monitored Linux systems, stop the syslog daemon, by using the appropriate command for the type of syslog daemon that is in use on the Linux system.

2. Deactivate the IBM zAware partition. Use the **Deactivate** task in the Hardware Management Console (HMC). For authorization requirements and other information about the **Deactivate** task, see HMC/SE topics in IBM Knowledge Center, at http://www.ibm.com/support/knowledgecenter/

Part 6. Advanced topics for managing IBM zAware

Topics in this part describe specialized management tasks for IBM zAware.

Topics covered in this part are:

Chapter 21. Managing the training for monitored clients

Training is the process of using data from monitored clients to build a model of normal system behavior for analysis. Training, combined with the use of pattern recognition techniques, is how IBM zAware learns about the typical behavior of monitored systems and their workloads. To detect differences that might indicate a problem, IBM zAware compares current data from monitored systems to the model that was created during training.

Administrators can manage the training schedule for monitored systems. Management actions include:

- Viewing the training schedule. For example, you can view the next scheduled training date and the dates to be included in the next period.
- Selecting the dates to exclude from future training because of very abnormal behavior, or to re-include in a future training.
- Determining for which dates in the training period IBM zAware has or does not have data.
- Requesting training.
- Viewing the status of a training request.
- Canceling a training request.
- Identifying specific messages that are to be ignored during analysis for the selected monitored system. This capability is available only for messages that z/OS monitored systems issue.

Understanding training periods and intervals

IBM zAware uses two concepts, training periods and training intervals, to manage the training schedule for a single monitored system or for a group of monitored systems, which is called a model group. The *training period* is the number of consecutive calendar days that the IBM zAware server uses to identify the monitored system data to include in training models. The *training interval* is the number of consecutive calendar days between automatic builds of system behavior models.

That is, the training period is how many days of data IBM zAware needs to include in a model and the training interval is how often IBM zAware automatically recreates the model. To modify the default training period and training interval for a specific type of monitored client, go to the **Administration** > **Configuration** > **Analytics** tab.

The following examples depict the relationship between the training period and the training interval, and describe the training schedule for different types of monitored clients.

- "z/OS example 1: Allowing IBM zAware to collect the data for the initial model"
- "z/OS example 2: Using priming data to build the initial model" on page 222
- "Linux example: Allowing IBM zAware to collect the data for the initial group model" on page 223

Note that the low training values used in these examples are for illustration only, and might not yield enough data to build a representative model for the monitored systems in your installation. The default training values vary depending on the type of monitored system, are based on IBM experience with building models for different system types, and are more likely to yield high quality models during the initial training. For system-specific guidance about building models, see the appropriate planning topic for the type of monitored client in Chapter 11, "Planning to create IBM zAware models," on page 87.

z/OS example 1: Allowing IBM zAware to collect the data for the initial model

Figure 55 on page 222 depicts the calendar days for a sample training schedule in which IBM zAware collects the data for the initial z/OS model and schedules automatic training requests, based on the

values set for the training period and the training interval. For illustration purposes only, this example uses values other than the default values for the training period and the training interval for z/OS monitored clients:

- The training period is set to 10 days.
- The training interval is set to seven days.

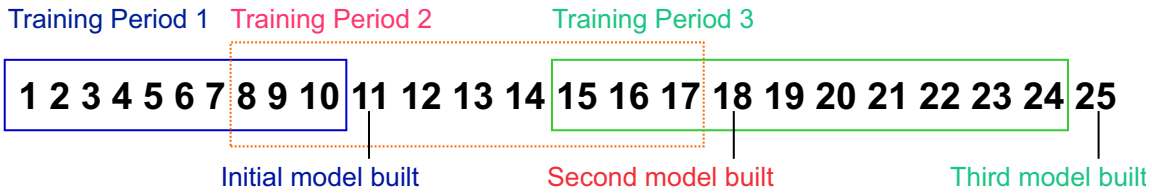

Figure 55. Training schedule for z/OS example 1

In this example:

- On day 1, the z/OS monitored system is initially connected to the IBM zAware server, and the training period begins.

  For the initial model, IBM zAware does not consider the training interval setting. Instead, it uses the training period value to determine when it has collected enough data to build a model. In this example, IBM zAware schedules the first model build for day 11, the day after the first training period ends.

  – For IBM zAware scheduling purposes, a day begins at UTC midnight and ends at 23:59:59 UTC.
  – IBM zAware server does not include data from the current day in a model, for either automatically scheduled training or an administrator request for training.
- On day 11, IBM zAware builds the initial model for the z/OS monitored system from data collected from day 1 through and including day 10.
- After the model is successfully built, IBM zAware begins analyzing current data from the z/OS monitored system, and calculates the date for the next automatic build, using the setting for the training interval. IBM zAware schedules the second training (model build) for day 18, which is seven days (the training interval) after the initial training.
- On day 18, IBM zAware builds a new model that includes the data that was collected during the past 10 days, from day 8 through and including day 17, which represent the second training period. Also on day 18, IBM zAware schedules the third training for day 25.
- On day 25, IBM zAware builds a new model that includes the data that was collected during the third training period, from day 15 through and including day 24. Assuming 31 days in the current month, IBM zAware uses the configured training interval value to schedule the next training for day 1 of the next month.
- For future training periods, IBM zAware repeats this pattern of calculation for scheduling automatic builds, using both the training period and training interval settings that are in effect for z/OS monitored systems.

  If an administrator manually requests training and a model is successfully built, IBM zAware recalculates the scheduled date for the next automatic build, using the date on which the model was created and the configured training interval value. For example, if an administrator manually requests training on day 26 and a new model is created, IBM zAware changes the next scheduled date for an automatic build from day 1 of the next month to day 2.

z/OS example 2: Using priming data to build the initial model

Figure 56 on page 223 depicts the calendar days for a sample training schedule in which IBM zAware uses priming data to build the initial z/OS model. Priming data is data from the hardcopy log or system log for the z/OS monitored client. For illustration purposes only, this example uses values other than the default values for the training period and the training interval for z/OS monitored clients:

- The training period is set to 10 days.
- The training interval is set to seven days.

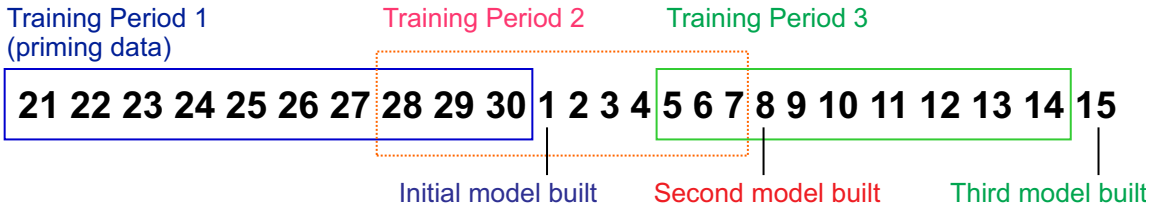

Figure 56. Training schedule for z/OS example 2

In this example:

- On day 1, the z/OS monitored system is initially connected to the IBM zAware server. Also on that day:
  1. An IBM zAware administrator sends priming data to IBM zAware.
  2. The administrator assigns priming data to the sysplex to which the z/OS monitored system belongs.
  3. The administrator submits a training request for IBM zAware to build a model for the z/OS monitored system for which priming data was sent. In this example, the priming data contains data for day 21 through day 30 of the previous month; these days constitute the first training period. Note that an administrator can provide more than 10 days of priming data for the z/OS monitored system.

     Because the administrator requests training on day 1 of the following month, IBM zAware includes all of the priming data in the initial model. If the administrator had submitted the training request on day 30, the priming data for day 30 would not be included in the model. The IBM zAware server does not include data from the current day in a model.
- After the model is successfully built, IBM zAware begins analyzing current data from the z/OS monitored system, and calculates the date for the first automatic build, using the setting for the training interval. IBM zAware schedules the second training for day 8, which is seven days (the training interval) after the initial training (model build).
- On day 8, IBM zAware builds a new model that includes data from day 28 through day 30 from the priming data, as well as data that IBM zAware collected from the z/OS monitored system, from day 1 through and including day 7. These dates represent the second training period. Also on day 8, IBM zAware schedules the third training for day 15.
- On day 15, IBM zAware builds a new model that includes the data that was collected during the third training period, from day 5 through and including day 14. IBM zAware uses the configured training interval value to schedule the third training for day 22.
- For future training periods, IBM zAware repeats this pattern of calculation for scheduling automatic builds, using both the training period and training interval settings that are in effect.

  If an administrator manually requests training and a model is successfully built, IBM zAware recalculates the scheduled date for the next automatic build, using the date on which the model was created and the configured training interval value. For example, if an administrator manually requests training on day 16 and a new model is created, IBM zAware changes the next scheduled date for an automatic build from day 22 to day 23.

Linux example: Allowing IBM zAware to collect the data for the initial group model

For Linux model groups, IBM zAware automatically schedules early training every seven days, starting from the first day for which IBM zAware has data available. The seven-day early training schedule continues until at least one of the group members is connected to IBM zAware for the configured training period; at that point, IBM zAware uses the configured training interval to schedule automatic training.

The following example contains two diagrams:
- Figure 57 depicts the calendar days in the early training schedule.
- Figure 58 on page 225 depicts the early training schedule along with the first training period and training interval that use the configured training values.

For illustration purposes only, this example uses values other than the default values for the training period and the training interval for Linux monitored clients:
- The training period is set to 20 days.
- The training interval is set to 10 days.

In this example:
- On day 1, an IBM zAware administrator defines a model group for a collection of Linux systems with similar names and workloads. The administrator also connects at least one of the member systems to the IBM zAware server.

  When the first system is connected and starts sending data, IBM zAware assigns the system to the model group, and detects how many days of data are available for building a model for the group. For IBM zAware scheduling purposes, a day begins at UTC midnight and ends at 23:59:59 UTC.

  IBM zAware schedules the first automatic training for day 8, as shown in Figure 57.

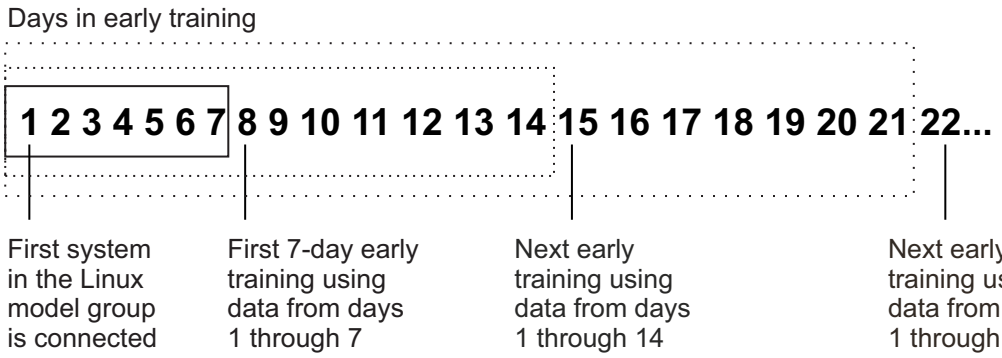

Figure 57. Early training schedule for Linux

- On day 8, IBM zAware attempts to build a model for the group of systems, using all data that it collected from member systems on day 1 through and including day 7. The IBM zAware server does not include data from the current day in a model.
  - If a model is successfully built, IBM zAware begins analyzing current data from the monitored systems in the group, and calculates the date for the next automatic training.
  - If an automatic training attempt fails and a model is not available, IBM zAware automatically retries the training attempt the next day and, if necessary, every following day until a model is successfully built. Analysis cannot begin until a model is successfully built.

  For this example, assume that IBM zAware successfully builds a model on day 8. Because member systems in the group have been connected for only eight days, which is less than the configured training period of 20 days, IBM zAware continues to follow the early training schedule, setting the next early training for day 15.
- On day 15, IBM zAware successfully completes an early training attempt, using the data that it collected from member systems on day 1 through and including day 14. Because member systems in the group have been connected for less than the configured training period of 20 days, IBM zAware continues to follow the early training schedule, setting the next training for day 22.
- On day 22, IBM zAware successfully completes an early training attempt, using the data that it collected from member systems on day 1 through and including day 21.

  Because member systems have been connected for more days than the configured training period of 20 days, IBM zAware uses the configured training interval value (10) to schedule the next training for day

1 of the next month, as shown in Figure 58.

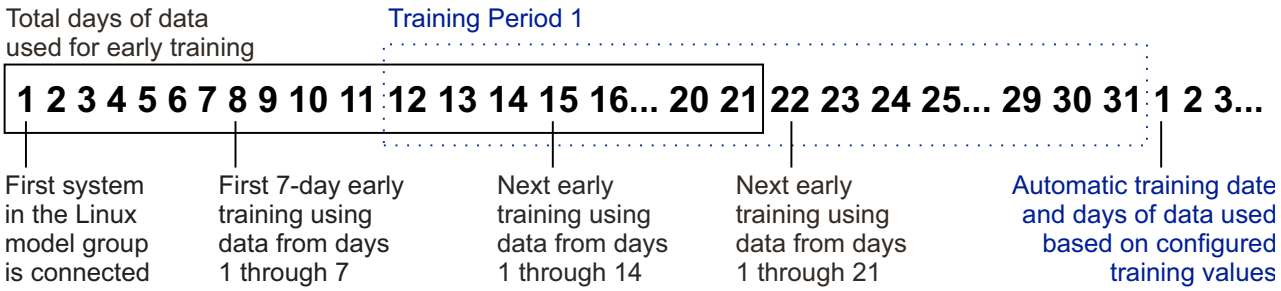

Figure 58. Early and configured training schedule for Linux example

- On day 1 of the next month, IBM zAware builds a new model, using data that it collected from the monitored systems in the group, from day 12 through and including day 31. These days represent the first full training period according to the configured training period value (20), as shown in Figure 58. IBM zAware uses the configured training interval value (10) to schedule the next training for day 11 of the next month.
- For future training periods, IBM zAware repeats this pattern of calculation for scheduling automatic builds, using both the training period and training interval settings that are in effect.

  If an administrator manually requests training and a model is successfully built, IBM zAware recalculates the scheduled date for the next automatic build, using the date on which the model was created and either the early seven-day schedule or the configured training interval, whichever is in effect. For example, if an administrator manually requests training on day 29 and a new model is created, IBM zAware uses the configured interval value to calculate and change the next scheduled date for an automatic build from day 1 to day 8 of the next month.

Note that an administrator can provide priming data, which is prior data from system logs, for one or more of the Linux systems in a model group. Priming data allows for earlier training and better quality models. For example, suppose an administrator provided priming data for day 25 through and including day 31 of the prior month. With this seven days of data, IBM zAware can attempt the first early training on day 1 of this month, after one or more of the systems in the model group are first connected.

Viewing model dates

To view model dates for a z/OS or Linux monitored client, use the **Manage Model Dates** action on the Training Sets page.

About this task

Through the Manage Model Dates page, you can view the information using the Summary View or the Calendar View. Through either view, you can:
- Determine the training period begin and end dates.
- View the days in the training period for which data is available, unavailable, or excluded.
- Display the date when the current model was built or when the next model is scheduled to be built.

This topic describes how to view model dates; see “Excluding dates from a model” on page 226 to learn how to exclude specific dates.

Procedure

1. To display the Training Sets page, click **Administration** in the navigation pane and select **Training Sets**.
2. On the Training Sets page, click either the **z/OS Systems** tab or the **Model Groups** tab.

3. Depending on the tab you clicked, select either a z/OS system or Linux model group.
   - For a z/OS monitored client, select the system for which you want to view the model dates from the Monitored z/OS Systems table.

     You can select only one system and the value in the Last Training Result column cannot be Never Connected. The system does not have to be currently connected to the IBM zAware server, but it must have been connected to the server at least once.
   - For a Linux monitored client, select the model group to which the Linux system belongs from the Monitored Model Groups table. You can select only one model group.
4. From the **Actions** list, select **Manage Model Dates**. This action is disabled if the value in the Training Progress column is In Progress or In Queue, or if the value in the Last Training Result column is Empty.
5. For a Linux system only, use the "Training system" field to type the name of the Linux system or to select it from the list of member systems in the model group.
6. Review the information that is provided on the Manage Model Dates page. You can view the information using the Summary View or the Calendar View. To switch between the views, click **Switch to Calendar View** or **Switch to Summary View**.

Excluding dates from a model

Ideally, a model represents a predictable, stable workload that generates the same artifacts when the monitored system, its subsystems, hardware, and applications are working as your installation expects them to function. If unexpected errors or activity occurred during the training period for a z/OS or Linux monitored client, you can specify for that data to be excluded from the next model. To do so, use the **Manage Model Dates** action on the Training Sets page.

About this task

You can exclude data for dates that fall in the range of yesterday minus the training period, including the first day and yesterday. For example, if the training period is 90 days and the current date is August 8th, you can exclude dates that fall in the following range: May 9th - August 7th.

You cannot modify the current model dates. You can exclude dates only when working with the next training period model dates. After you select the dates to exclude, you can either manually request training or wait for the next scheduled training. In either case, IBM zAware builds and uses a model that excludes the selected dates.

Procedure

1. To display the Training Sets page, click **Administration** in the navigation pane and select **Training Sets**.
2. On the Training Sets page, click either the **z/OS Systems** tab or the **Model Groups** tab.
3. Depending on the tab you clicked, select either a z/OS system or Linux model group.
   - For a z/OS monitored client, select the system for which you want to manage model dates from the Monitored z/OS Systems table.

     You can select only one system and the value in the Last Training Result column cannot be Never Connected. The system does not have to be currently connected to the IBM zAware server, but it must have been connected to the server at least once.
   - For a Linux monitored client, select the model group to which the Linux system belongs from the Monitored Model Groups table. You can select only one model group.
4. From the **Actions** list, select **Manage Model Dates**. This action is disabled if the value in the Training Progress column is In Progress or In Queue, or if the value in the Last Training Result column is Empty.

5. For a Linux system only, use the “Training system” field to type the name of the Linux system or to select it from the list of member systems in the model group.
6. Complete one of the following steps:
   - To use the Manage Model Dates: Summary View to exclude dates, do the following:
     – In the “Model dates” field, select **Next Training Period Model Dates**. The content of the Manage Model Dates page changes from current model dates to model dates that apply for the next training period. You cannot exclude dates from the current model.
     – Using the “Excluded dates” field, specify the date to be excluded. You can type the value; select it using the left and right arrow icons, which allow you to scroll through the dates one day or one month at a time; or select it from the calendar, which is displayed when you click the calendar icon.
     – Click **Add** to add the date to the “Excluded dates” list.
     – Repeat this process until all the dates you want to exclude are listed in the “Excluded dates” list. If you previously excluded a date that you want to re-add to the model, in the “Excluded dates” list, select the date and click **Remove**.
   - To use the Manage Model Dates: Calendar View to exclude dates, do the following:
     – Click **Switch to Calendar View**.
     – In the “Model dates” field, select **Next Training Period Model Dates**. You cannot exclude dates from the current model.
     – In the “Training calendar” field, click the dates to be excluded. You can click any date that is not marked as unavailable. If the date will be excluded, it is surrounded by a burgundy square. If you previously excluded a date that you want to re-add to the model, click it again. The burgundy square will be removed.

       Tip: To scroll backwards and forwards through the calendar by month, click the left and right arrow icons, which are displayed in the month header. To scroll backwards and forwards by year, click the first and last year link.

Results

IBM zAware excludes the selected dates when it builds the next model.

Requesting training automatically or manually

For the initial model for a z/OS monitored system, the IBM zAware server requests an automatic build when the training period elapses. Meanwhile, you can manually raise a request to build a model of normal system behavior for a z/OS monitored system or for a group of Linux monitored systems by using the **Request Training** action.

Before you begin

Verify that the IBM zAware server has received data for the number of days in the training period, which is specified on the **Administration** > **Configuration** > **Analytics** tab.

About this task

Any data that monitored system is currently sending to the IBM zAware server does not become part of the model until you request training or the IBM zAware server automatically builds the model.

- For the initial model for a z/OS monitored system, the IBM zAware server requests an automatic build when the training period elapses. For subsequent builds, the server requests an automatic build when the training interval elapses.

  By default, the training period is 90 days and the training interval is 30 days for z/OS OPERLOG. If you choose Option 1: Waiting for the server to build a z/OS model, you don't need to manually

request training. IBM zAware will automatically start after 90 days. For more information about Option 1 and Option 2, see “Planning to create IBM zAware models for z/OS monitored clients” on page 87.

This behavior does not apply to the training that is manually requested. The result of the manually requested training depends on whether there are required message IDs in the log. The message traffic for the system must contain a minimum of 250 message IDs. Each message ID must be issued at least once in three different 10-minute intervals during the training period.

If a training fails, IBM zAware automatically retries the training the next day. The daily retry continues until a model is set up successfully. After the training model is set up, it will start to analyze the log.

- For the initial model for a Linux model group, IBM zAware automatically schedules early training every seven days, starting from the first day for which IBM zAware has data available. The seven-day early training schedule continues until at least one of the group members is connected to IBM zAware for the configured training period; at that point, IBM zAware uses the configured training interval to schedule automatic training. For subsequent builds, the server requests an automatic build when the training interval elapses.

  To build a robust model of Linux system behavior, IBM zAware generally needs a minimum of 120 days of message data. Although you can request training for a Linux model group at any time before the default 120-day training period, your request (and the seven-day early training attempts) might either fail or produce a limited model. For more information about limited models, see “Understanding how IBM zAware calculates and displays anomaly scores ” on page 142.
- Consider requesting a build if you modify the software, hardware, or network configuration for a monitored system or if the workload increases on the system. When these changes occur, anomalies or new message patterns might be introduced for the system. Any differences between the model and the current data are identified on the Analysis page. After you verify that these differences are normal for the system, you might want to build a new model that includes the new message patterns or exclude this data from the model so that other differences that might indicate a potential problem are easier to identify.
- For z/OS monitored systems only, if you used the z/OS bulk load client to collect the data to include in the model, you can wait for the next scheduled training where the IBM zAware server will use the priming data to build the model or you can request a build. The latter option is recommended because analysis can start shortly after the model is built.
- If you are using JES3 on a z/OS monitored system, requesting a build is also recommended after moving the JES3 global function from one system to another, through either an IPL or the dynamic system interchange (DSI) facility. After the JES3 global is moved and restarted, the message traffic on the new system contains messages that are not reflected in the existing model but are not necessarily indicative of problems. After moving the JES3 global and viewing the Analysis page results for the new system, rebuild the system model to include the JES3 messages.
- For Linux model groups, consider requesting a build if you change the membership of one or more Linux model groups. Note that IBM zAware does not provide analysis results for Linux monitored systems that belong to the UNASSIGNED model group, nor does it build a model for that group. To assign or move Linux systems to a model group, administrators create or edit a model group through the **Model Groups** tab on the Systems page; for more information, see “Managing groups of Linux monitored clients” on page 210.

If a training request completes without errors, the following occurs:

- IBM zAware creates a new model that contains all the data that was collected from the beginning of the training period to yesterday at 23:59:59 UTC. Data for the current day is not included in the model.
- At the start of the next analysis interval, IBM zAware begins using the new model to identify changes in message patterns for the monitored client or clients that are associated with that model.
- IBM zAware calculates the new date for the next automatic build. For example, if the training interval is 30 days, the next automatic build is scheduled for 30 days from the current date.

  Shortly after it begins to use the new model, IBM zAware updates the next training period model dates on the Manage Model Dates page.

Note: IBM zAware does not create a model when a training request fails. If a model existed before the failure, IBM zAware continues to use that model for analysis and attempts to retry the failed training request on the next day and, if necessary, every following day until a model is successfully built.

Procedure

1. To display the Training Sets page, expand the Administration category in the navigation pane and select **Training Sets**.
2. Select either the **z/OS Systems** tab or the **Model Groups** tab.
   - On the **z/OS Systems** tab, the Monitored z/OS Systems table lists all of the z/OS monitored systems that are or were connected to the IBM zAware server.
   - On the **Model Groups** tab, the Monitored Model Groups table lists all of the groups of Linux monitored systems that administrators defined through the **Model Groups** tab on the Systems page.
3. In the table, select the system or model group for which you want to build a new model.
   - You can select only one system or model group.
   - If you select a z/OS system, the value in the Last Training Result column cannot be Never Connected. The system does not have to be currently connected to the IBM zAware server, but it must have been connected to the server at least once.
4. From the **Actions** list, select **Request Training**. This action is disabled if the value in the Training Progress column is In Progress or In Queue, or if the value in the Last Training Result column is Empty.
5. Click **OK** to confirm that you want to rebuild the model. IBM zAware displays a message indicating whether the request was successfully submitted. If the request was successfully submitted, the status in the Training Progress column is changed to either *In Progress* or *In Queue*. To determine the result of the training request, periodically click **Refresh** to update the Training Sets page display.

Results

When the request completes, IBM zAware updates the Last Training Result column and the Last Training Result Time column. If the model was built, the value displayed in the Last Training Result column is Complete; in this case, IBM zAware also updates the Current Model Built column.

If the training request fails, the value displayed in the Last Training Result column is Failed. Typically, training requests fail for one of the following reasons:

- An I/O error occurred. For example, IBM zAware might be unable to read or write to a file or file system or a file system might be unavailable.
- IBM zAware does not have enough log data to build a model for a system or group of systems.

In the case of a failure, check the **Notifications** page for additional messages that provide more information about the reason for the failure.

- To resolve an I/O error, verify that the storage devices that are managed by IBM zAware are online and try adding storage devices for IBM zAware to use. To view and manage the storage devices, use the **Administration** > **Configuration** > **Data Storage** tab.
- To resolve a failure caused by insufficient data, you might only need to load more data or wait until the system has sent more data, and then retry the training request. Other possible corrective actions include the following:
  - If an administrator is priming IBM zAware with prior data and a model does not exist yet, consider sending additional days of log data for this system or systems in the group, and retrying the training request for the system or group. For information about sending more log data, see the appropriate configuration topic for the type of monitored client in Part 4, “Configuring IBM zAware and its monitored clients,” on page 93.

– If an administrator modified the default training set such that days that represent normal system activity are excluded from the training period, use **Manage Model Dates** to restore excluded dates to the training set, and retry the training request. If you need more information, see "Excluding dates from a model" on page 226.
– If an administrator modified the default training period but more days are required to build the model, go to the appropriate **Configuration** > **Analytics** tab, increase the "Training period" value, and apply the change. After reconnecting monitored clients as necessary, retry the training request. If you need more information about reconnecting monitored clients, see "Starting and stopping data collection for your monitored systems" on page 203.

Canceling training

To cancel an automatic training request or a request that you submitted for a monitored system or model group, use the **Cancel Training** action.

Procedure

1. To display the Training Sets page, expand the Administration category in the navigation pane and select **Training Sets**.
2. Select either the **z/OS Systems** tab or the **Model Groups** tab.
   - On the **z/OS Systems** tab, the Monitored z/OS Systems table lists all of the z/OS monitored systems that are or were connected to the IBM zAware server.
   - On the **Model Groups** tab, the Monitored Model Groups table lists all of the groups of Linux monitored systems that administrators defined through the **Model Groups** tab on the Systems page.
3. If necessary, click the Training Progress column heading in the table to sort the entries by the column value, and review the entries with a value of In Queue.
4. Select the system or model group for which you want to cancel training. You can select only one system or model group.
5. From the **Actions** list, select **Cancel Training**. This action is enabled only if the value in the Training Progress column is In Queue. In other words, you cannot cancel training that is already in progress.

Results

If the training request was canceled, the value in the Last Training Result column is Cancelled and the Last Training Result Time column is updated.

IBM zAware does not create a model when a training request is canceled. If a model existed prior to the cancellation, IBM zAware continues to use that model.

Managing ignored messages

To designate specific message IDs for IBM zAware to ignore during analysis of message data from a specific monitored system, use the **Manage Ignored Messages** action provided in the Monitored Systems table on the Training Sets page. This capability is available only for messages that z/OS monitored systems issue.

About this task

You can designate specific message IDs for IBM zAware to ignore during analysis. This capability is especially useful when you recently changed a monitored system, such as adding a new workload. When your installation intentionally makes a significant change to a monitored system, IBM zAware might detect the resulting change in message traffic as anomalous behavior, and assign high anomaly scores to

the intervals during and after the change. You can use this analysis to confirm that the anomalies are a result of expected message traffic rather than problems, and then mark the new or unusual messages as messages for IBM zAware to ignore.

You can designate messages to be ignored until the next time IBM zAware builds a model of behavior for the monitored system, or until an administrator manually changes the ignore status of the message. The ignore status for any message applies only on a per-system basis. If you want IBM zAware to ignore the same message on several monitored systems, you must repeat this procedure for setting the ignore status on each system.

You cannot identify a message to be ignored when an IBM rule already applies to that message. An IBM rule takes precedence, even when an administrator already successfully marked the message to be ignored. This situation can occur when an administrator successfully designated a message to be ignored, but IBM zAware later assigns an IBM rule to that message, as a result of the analysis of training data. In this case, the following conditions apply:

- For this message, the IBM rule is displayed in the Rules Status column on the Interval page.
- This message is displayed in the Ignored Messages table on the Manage Ignored Messages window, even though the IBM rule takes precedence. An administrator can remove this message from the table, but cannot successfully reapply an ignore status value to this message.

Procedure

1. To display the Training Sets page, expand the Administration category in the navigation pane and select **Training Sets**.
2. In the Monitored Systems table, select the system for which you want to manage messages. You can select only one z/OS system and the value in the Last Training Result column cannot be *Never Connected*. The system does not have to be currently connected to the IBM zAware server, but it must have been connected to the server at least once.
3. From the **Actions** list, select **Manage Ignored Messages**. The Manage Ignored Messages page opens, displaying the Ignored Messages table for the system that you selected.

   The Ignored Messages table lists any messages that an administrator has requested IBM zAware to ignore while analyzing current data from a specific monitored client. The table is empty if an administrator did not previously designate any messages for IBM zAware to ignore during analysis.
4. To add one or more messages for IBM zAware to ignore for the selected system, complete the following steps.
   a. Click **Add Messages** from the **Actions** list in the Ignored Messages table. The Add Ignored Messages window opens.
   b. In the "Message IDs to ignore" field, enter one or more message IDs, separating each message ID with a comma. The message identifier (ID) must be a well-formed ID. IBM zAware recognizes messages IDs that conform to the z/OS standard, which consists of a component identifier, a message number, and an action code, in that order. IBM zAware also is capable of recognizing message IDs that do not completely conform to this z/OS standard.
   c. Select one of the ignore status values to apply to all of the messages that you entered.

      Ignore messages until next training occurs for the current system.
      Marks the selected messages to be ignored during analysis until the next time IBM zAware successfully builds a model of behavior for the monitored system, as the result of either a manually requested or automatically scheduled training operation. The date for the next scheduled model rebuild is listed under "Next scheduled training date" after you select **Next Training Period Model Dates** on the Manage Model Dates page.

      Use this value for messages that you determine to be anomalous because of a workload change on the system, but you expect them to become part of the normal behavior for this system. The next training results in a model that includes these messages, which will be subject to normal analysis after the training.

Ignore messages until manually restored.
Marks the selected messages to be ignored during analysis until an administrator manually changes the ignore status of the message, or removes it from the list of messages in the Ignored Messages table.

Use this value for messages that you determine to be normal (that is, not indicative of a problem) on this system. In subsequent analysis, these messages do not contribute to the anomaly score, and thus reduce false-positive results.

d. Click **OK** to apply your changes, or click **Cancel** to return to the Manage Ignored Messages page. When you click **OK**, IBM zAware indicates whether it successfully processed the messages you entered. You can verify the results by clicking **Refresh** on the Manage Ignored Messages page, and viewing the updated list in the Ignored Messages table.

5. To change the ignore status of one or more messages in the Ignored Messages table, complete the following steps.
   a. Select one or more messages listed in the Ignored Messages table.
   b. Click one of the following actions from the **Actions** list in the Ignored Messages table.

      These actions are equivalent to the status values described in step 4 on page 231.
      - **Ignore Until Next Training**
      - **Ignore Until Manually Restored**

   IBM zAware indicates whether it successfully processed the action for the messages that you selected. You can verify the results by clicking **Refresh** on the Manage Ignored Messages page, and viewing the updated list in the Ignored Messages table.
6. To remove one or more messages from the Ignored Messages table, complete the following steps.
   a. Select one or more messages listed in the Ignored Messages table.
   b. Click **Remove** from the **Actions** list in the Ignored Messages table. The Remove Ignored Messages window opens. This window displays a list of the messages that you selected for removal for this system.
   c. Click **OK** to apply your changes, or click **Cancel** to return to the Manage Ignored Messages page.

   When you click **OK**, IBM zAware indicates whether it successfully removed the messages that you selected. You can verify the results by clicking **Refresh** on the Manage Ignored Messages page, and viewing the updated list in the Ignored Messages table.

Training sets for z/OS systems

You can use the **Training Sets** > **z/OS Systems** tab to request training for a z/OS monitored system, to display the current training status for each system, and to view the current and future training dates for each system.

To display the **z/OS Systems** tab, expand the Administration category in the navigation pane and select **Training Sets**. The controls and content displayed on the **z/OS Systems** tab are described in the following sections:

- "Monitored z/OS Systems table"
- "Actions list in the Monitored z/OS Systems table" on page 235
- "Current Training Status Details section" on page 235

Monitored z/OS Systems table

The Monitored z/OS Systems table provides the current training status for all of the z/OS monitored systems that are or were connected to the IBM zAware server. For a description of the columns in the Monitored z/OS Systems table, see Table 48 on page 233. To display the most recent training status, click **Refresh**.

Table 48. Columns in the Monitored z/OS Systems table

Column	Description
System	Provides the name of the monitored system.
Sysplex	Provides the name of the sysplex to which the monitored system belongs.
Training Progress	Indicates the current training activity for the monitored system. One of the following values is displayed: **In Progress** Indicates that the training request is being processed. **In Queue** *(n)* Indicates that the training request is in the training queue and is waiting to be processed, where *(n)* represents the position of the request in the queue. The queue position is incremented for each subsequent request that is added to the queue. Because training is resource intensive, IBM zAware processes only one training request at a time. All subsequent training requests are added to the training queue and are processed in the order in which they were placed in the queue. **"-" (dash)** Indicates that there are no active or queued training requests for the monitored system. With the exception of the dash (-), the value in the Training Progress column is a hyperlink. Click it to expand the Current Training Status Details section, which provides additional information about the training request that was submitted for the selected monitored system.

Table 48. Columns in the Monitored z/OS Systems table (continued)

Column	Description
Last Training Result	Provides the outcome of the last training activity. One of the following values is displayed: **Cancelled** Indicates that the training request was canceled while it was in the training queue. IBM zAware does not create a model when a training request is canceled. If a model existed prior to the cancellation, IBM zAware continues to use that model. **Complete** Indicates that the last training request completed and a new model was built. **Failed** Indicates that the last training request failed. Typically, training requests fail for one of the following reasons: • An I/O error occurred. For example, IBM zAware might be unable to read or write to a file or file system or a file system might be unavailable. • IBM zAware does not have enough log data to build a model for a system or group of systems. To resolve an I/O error, verify that the storage devices that are managed by IBM zAware are online and try adding storage devices for IBM zAware to use. To view and manage the storage devices, use the **Administration** > **Configuration** > **Data Storage** tab. To resolve a log data problem, go to the **Notifications** tab and look for messages that are related to the training attempt, and follow any corrective actions that are suggested in the message descriptions. **Note:** IBM zAware does not create a model when a training request fails. If a model existed before the failure, IBM zAware continues to use that model for analysis and attempts to retry the failed training request on the next day and, if necessary, every following day until a model is successfully built. **Never Connected** Indicates that the monitored system has not been connected to the IBM zAware server. This result occurs under the following circumstances: • When you use the **Administration** > **Configuration** > **Priming Data** tab to assign priming data for multiple systems to a sysplex, but only one of the sysplex members has been connected to the server. • When you use the **Administration** > **Configuration** > **Topology** tab to move a system to a different sysplex. In both cases, the IBM zAware server does not recognize the system because it identifies systems by the system and sysplex name combination. To make the IBM zAware server aware of this new system and sysplex name combination, connect the system to the IBM zAware server. If you used the **Topology** tab to move a system, and a model already exists for the old system and sysplex name combination, that model is preserved and is associated with the new system and sysplex name combination. **Not Trained** Indicates that no training has been requested for the monitored system. **"-" (dash)** Indicates that the result of the last training request is not available because the request is either being processed or is in the training queue. With the exception of the dash (-), the value in the Last Training Result column is a hyperlink. Click it to expand the Current Training Status Details section, which provides additional information about the training request that was submitted for the selected monitored system.
Last Training Result Time	Provides the date and time when the last training result was obtained. That is, the date and time the last training completed, failed, or was canceled. A dash (-) is displayed if the value in the Last Training Result column is *Never Connected* or *Not Trained* or if the value in the Training Progress column is *In Progress* or *In Queue*.

Table 48. Columns in the Monitored z/OS Systems table (continued)

Column	Description
Current Model Built	Provides the date and time when the current model was built. A dash (-) is displayed if no model is available for the system.

Actions list in the Monitored z/OS Systems table

The **Actions** list provided in the Monitored z/OS Systems table lists the actions that you can take against a monitored system. Table 49 provides a description of each action and provides links to additional information that explains how to perform the action.

Table 49. Actions for z/OS monitored systems

Action	Description	Additional Information
Manage Model Dates	Specify the dates to exclude from the training for the selected monitored system. View the training dates that were used to build the current model and that will be used to build the next model.	"Excluding dates from a model" on page 226 "Viewing model dates" on page 225
Request Training	Build a new model of normal system behavior for the selected monitored system.	"Requesting training automatically or manually" on page 227
Cancel Training	Cancel the training request for the selected monitored system.	"Canceling training" on page 230
Manage Ignored Messages	Identify specific messages that are to be ignored during analysis for the selected monitored system.	"Managing ignored messages" on page 230

Current Training Status Details section

The Current Training Status Details section provides additional information about a training request. To display this information, click the link that is provided in either the Last Training Result column or the Training Progress column in the Monitored z/OS Systems table. Doing so expands the Current Training Status Details section and populates the fields with information about the training request that was submitted for the selected monitored system.

You can also expand and collapse the section by clicking the column header. If you expand the section before selecting a training request, the value for each field will be blank.

See Table 50 for a description of each field that is displayed in the Current Training Status Details section.

Table 50. Fields in the Current Training Status Details section

Field	Description
System name	Provides the name of the monitored system.

Table 50. Fields in the Current Training Status Details section (continued)

Field	Description
Training progress	Indicates the current training activity for the monitored system. One of the following values is displayed: **In Progress** Indicates that the training request is being processed. **In Queue** *(n)* Indicates that the training request is in the training queue and is waiting to be processed, where *(n)* represents the position of the request in the queue. The queue position is incremented for each subsequent request that is added to the queue. Because training is resource intensive, IBM zAware processes only one training request at a time. All subsequent training requests are added to the training queue and are processed in the order in which they were placed in the queue. **"-" (dash)** Indicates that there are no active or queued training requests for the monitored system.

Table 50. Fields in the Current Training Status Details section (continued)

Field	Description
Last training result	Provides the outcome of the last training activity. One of the following values is displayed: **Cancelled** Indicates that the training request was canceled while it was in the training queue. IBM zAware does not create a model when a training request is canceled. If a model existed prior to the cancellation, IBM zAware continues to use that model. **Complete** Indicates that the last training request completed and a new model was built. **Failed** Indicates that the last training request failed. Typically, training requests fail for one of the following reasons: • An I/O error occurred. For example, IBM zAware might be unable to read or write to a file or file system or a file system might be unavailable. • IBM zAware does not have enough log data to build a model for a system or group of systems. To resolve an I/O error, verify that the storage devices that are managed by IBM zAware are online and try adding storage devices for IBM zAware to use. To view and manage the storage devices, use the **Administration** > **Configuration** > **Data Storage** tab. To resolve a log data problem, go to the **Notifications** tab and look for messages that are related to the training attempt, and follow any corrective actions that are suggested in the message descriptions. **Note:** IBM zAware does not create a model when a training request fails. If a model existed before the failure, IBM zAware continues to use that model for analysis and attempts to retry the failed training request on the next day and, if necessary, every following day until a model is successfully built. **Never Connected** Indicates that the monitored system has not been connected to the IBM zAware server. This result occurs under the following circumstances: • When you use the **Administration** > **Configuration** > **Priming Data** tab to assign priming data for multiple systems to a sysplex, but only one of the sysplex members has been connected to the server. • When you use the **Administration** > **Configuration** > **Topology** tab to move a system to a different sysplex. In both cases, the IBM zAware server does not recognize the system because it identifies systems by the system and sysplex name combination. To make the IBM zAware server aware of this new system and sysplex name combination, connect the system to the IBM zAware server. If you used the **Topology** tab to move a system, and a model already exists for the old system and sysplex name combination, that model is preserved and is associated with the new system and sysplex name combination. **Not Trained** Indicates that no training has been requested for the monitored system. **"-" (dash)** Indicates that the result of the last training request is not available because the request is either being processed or is in the training queue.
Training start time	Specifies the date and time the training started. The start date and time is provided only when the value in the Training Progress column is *In Progress*. Otherwise, a dash (-) is displayed.
Time in training (h:m:s)	Specifies the total number of hours, minutes, and seconds that IBM zAware has been processing the training request. The training length is provided only when the value in the Training Progress column is *In Progress*. Otherwise, a dash (-) is displayed.

Table 50. Fields in the Current Training Status Details section (continued)

Field	Description
Last training result time	Provides the date and time when the last training result was obtained. That is, the date and time the last training completed, failed, or was canceled. A dash (-) is displayed if the value in the Last Training Result column is *Never Connected* or *Not Trained* or if the value in the Training Progress column is *In Progress* or *In Queue*.
Entered queue time	Specifies the date and time the training request was added to the training queue. The queue time is provided only when the value in the Training Progress column is *In Progress* or *In Queue*. Otherwise, a dash (-) is displayed.
Time in queue (h:m:s)	Specifies the total number of hours, minutes, and seconds that the training request was or has been in the queue. The total time in the queue is provided only when the value in the Training Progress column is *In Progress* or *In Queue*. Otherwise, a dash (-) is displayed.

Training sets for Linux model groups

You can use the **Training Sets** > **Model Groups** tab to request training for an administrator-defined group of Linux monitored systems, to display the current training status for each model group, and to view the current and future training dates for each model group.

To display the **Model Groups** tab, expand the Administration category in the navigation pane and select **Training Sets**, then click the **Model Groups** tab. The controls and content displayed on the **Model Groups** tab are described in the following sections:

- "Monitored Model Groups table"
- "Actions list in the Monitored Model Groups table" on page 240
- "Current Training Status Details section" on page 241

Monitored Model Groups table

The Monitored Model Groups table provides the current training status for all of the groups of Linux monitored systems that administrators defined through the **Model Groups** tab on the Systems page. To determine which Linux monitored systems belong to a particular model group:

1. Go to the **Systems** > **Model Groups** tab.
2. Select one model group in the Model Groups table.
3. In the Model Group Details page, review the Linux systems in the "Known matching member systems" list.

For a description of the columns in the Monitored Model Groups table, see Table 51. To display the most recent training status, click **Refresh**.

Table 51. Columns in the Monitored Model Groups table

Column	Description
Model Group	Indicates the name of an administrator-defined collection of Linux monitored systems. The UNASSIGNED model group is not displayed in this table because IBM zAware does not provide analysis results for Linux monitored systems that belong to the UNASSIGNED model group, nor does it build a model for that group.

Table 51. Columns in the Monitored Model Groups table (continued)

Column	Description
Training Progress	Indicates the current training activity for the model group. One of the following values is displayed: **In Progress** Indicates that the training request is being processed. **In Queue** *(n)* Indicates that the training request is in the training queue and is waiting to be processed, where *(n)* represents the position of the request in the queue. The queue position is incremented for each subsequent request that is added to the queue. Because training is resource intensive, IBM zAware processes only one training request at a time. All subsequent training requests are added to the training queue and are processed in the order in which they were placed in the queue. **"-" (dash)** Indicates that there are no active or queued training requests for the model group. With the exception of the dash (-), the value in the Training Progress column is a hyperlink. Click it to expand the Current Training Status Details section, which provides additional information about the training request that was submitted for the selected model group.

Table 51. Columns in the Monitored Model Groups table (continued)

Column	Description
Last Training Result	Provides the outcome of the last training activity. One of the following values is displayed: **Cancelled** Indicates that the training request was canceled while it was in the training queue. IBM zAware does not create a model when a training request is canceled. If a model existed prior to the cancellation, IBM zAware continues to use that model. **Complete** Indicates that the last training request completed and a new model was built. **Empty** Indicates that the last training request failed because the model group does not contain any member systems. In this case, IBM zAware has no data to use for training. **Failed** Indicates that the last training request failed. Typically, training requests fail for one of the following reasons: • There is not enough data or message traffic to build a model. In this case, you might need to wait until the member systems in the group have sent more data. • An I/O error occurred. For example, IBM zAware might be unable to read or write to a file or filesystem or a filesystem might be unavailable. To resolve this issue, verify that the storage devices that are managed by IBM zAware are online and try adding storage devices for IBM zAware to use. To view and manage the storage devices, use the **Administration** > **Configuration** > **Data Storage** tab. **Note:** IBM zAware does not create a model when a training request fails. If a model existed before the failure, IBM zAware continues to use that model for analysis and attempts to retry the failed training request on the next day and, if necessary, every following day until a model is successfully built. **Not Trained** Indicates that no training has been requested for the model group. **"-" (dash)** Indicates that the result of the last training request is not available because the request is either being processed or is in the training queue. With the exception of the dash (-), the value in the Last Training Result column is a hyperlink. Click it to expand the Current Training Status Details section, which provides additional information about the training request that was submitted for the selected model group.
Last Training Result Time	Provides the date and time when the last training result was obtained. That is, the date and time the last training completed, failed, or was canceled. A dash (-) is displayed if the value in the Last Training Result column is *Empty* or *Not Trained* or if the value in the Training Progress column is *In Progress* or *In Queue*.
Current Model Built	Provides the date and time when the current model was built. A dash (-) is displayed if no model is available for the model group.

Actions list in the Monitored Model Groups table

The **Actions** list provided in the Monitored Model Groups table lists the actions that you can take against a model group. Table 52 on page 241 provides a description of each action and provides links to additional information that explains how to perform the action. All actions are disabled for a model group with a value of Empty in the Last Training Result column.

Table 52. Actions for model groups

Action	Description	Additional Information
Manage Model Dates	Specify the dates to exclude from the training for the selected model group. View the training dates that were used to build the current model and that will be used to build the next model.	"Excluding dates from a model" on page 226 "Viewing model dates" on page 225
Request Training	Build a new model of normal system behavior for the selected model group.	"Requesting training automatically or manually" on page 227
Cancel Training	Cancel the training request for the selected model group.	"Canceling training" on page 230

Current Training Status Details section

The Current Training Status Details section provides additional information about a training request. To display this information, click the link that is provided in either the Last Training Result column or the Training Progress column in the Monitored Model Groups table. Doing so expands the Current Training Status Details section and populates the fields with information about the training request that was submitted for the selected model group.

You can also expand and collapse the section by clicking the column header. If you expand the section before selecting a training request, the value for each field will be blank.

See Table 53 for a description of each field that is displayed in the Current Training Status Details section.

Table 53. Fields in the Current Training Status Details section

Field	Description
Model group name	Indicates the name of an administrator-defined collection of Linux monitored systems.
Training progress	Indicates the current training activity for the model group. One of the following values is displayed: **In Progress** Indicates that the training request is being processed. **In Queue** *(n)* Indicates that the training request is in the training queue and is waiting to be processed, where *(n)* represents the position of the request in the queue. The queue position is incremented for each subsequent request that is added to the queue. Because training is resource intensive, IBM zAware processes only one training request at a time. All subsequent training requests are added to the training queue and are processed in the order in which they were placed in the queue. **"-" (dash)** Indicates that there are no active or queued training requests for the model group.

Table 53. Fields in the Current Training Status Details section (continued)

Field	Description
Last training result	Provides the outcome of the last training activity. One of the following values is displayed: **Cancelled** Indicates that the training request was canceled while it was in the training queue. IBM zAware does not create a model when a training request is canceled. If a model existed prior to the cancellation, IBM zAware continues to use that model. **Complete** Indicates that the last training request completed and a new model was built. **Empty** Indicates that the last training request failed because the model group does not contain any member systems. In this case, IBM zAware has no data to use for training. **Failed** Indicates that the last training request failed. Typically, training requests fail for one of the following reasons: • There is not enough data or message traffic to build a model. In this case, you might need to wait until the member systems in the group have sent more data. • An I/O error occurred. For example, IBM zAware might be unable to read or write to a file or filesystem or a filesystem might be unavailable. To resolve this issue, verify that the storage devices that are managed by IBM zAware are online and try adding storage devices for IBM zAware to use. To view and manage the storage devices, use the **Administration** > **Configuration** > **Data Storage** tab. **Note:** IBM zAware does not create a model when a training request fails. If a model existed before the failure, IBM zAware continues to use that model for analysis and attempts to retry the failed training request on the next day and, if necessary, every following day until a model is successfully built. **Not Trained** Indicates that no training has been requested for the model group. **"-" (dash)** Indicates that the result of the last training request is not available because the request is either being processed or is in the training queue.
Training start time	Specifies the date and time the training started. The start date and time is provided only when the value in the Training Progress column is *In Progress*. Otherwise, a dash (-) is displayed.
Time in training (h:m:s)	Specifies the total number of hours, minutes, and seconds that IBM zAware has been processing the training request. The training length is provided only when the value in the Training Progress column is *In Progress*. Otherwise, a dash (-) is displayed.
Last training result time	Provides the date and time when the last training result was obtained. That is, the date and time the last training completed, failed, or was canceled. A dash (-) is displayed if the value in the Last Training Result column is *Empty* or *Not Trained* or if the value in the Training Progress column is *In Progress* or *In Queue*.
Entered queue time	Specifies the date and time the training request was added to the training queue. The queue time is provided only when the value in the Training Progress column is *In Progress* or *In Queue*. Otherwise, a dash (-) is displayed.
Time in queue (h:m:s)	Specifies the total number of hours, minutes, and seconds that the training request was or has been in the queue. The total time in the queue is provided only when the value in the Training Progress column is *In Progress* or *In Queue*. Otherwise, a dash (-) is displayed.

Manage Model Dates page

You can use the Manage Model Dates page provided in IBM zAware to view the dates that are associated with the current model and the next training period, and you can specify which dates to exclude from the model and which dates to re-include.

To display the Manage Model Dates page, expand the Administration category in the navigation pane and select **Training Sets**. Then, invoke the **Manage Model Dates** action for a monitored system.

The Manage Model Dates page contains a Summary view and a Calendar view. The Summary view is the default view. For more details about the controls and content that are displayed in each view, see the sections that follow.

Summary view

The Summary view provides a text-based version of the dates that are associated with a model of system behavior. The Summary view is the default view when you go to the **Training Sets** > **Manage Model Dates** page.

The controls and content displayed in the Summary view depend on whether you select *Next Training Period Model Dates* or *Current Model Dates* in the **Model dates** field. For a description of the items that are displayed for each option, see the following sections:

- "Fields displayed for Next Training Period Model Dates"
- "Fields displayed for Current Model Dates" on page 245

To switch to a pictorial view of the model dates, click **Switch to Calendar View**. To display the Training Sets page, click **Return to Training Sets** or click the **Training Sets** breadcrumb.

Fields displayed for Next Training Period Model Dates

When you select *Next Training Period Model Dates* in the **Model dates** field, IBM zAware displays information about the next model it will automatically build for the monitored system. For a description of the fields displayed for the next model, see Table 54.

Table 54. Fields displayed in the Summary view for the next model

Field	Description
Training system	Identifies the z/OS or Linux system for which data is displayed. • For a z/OS monitored client, the name displayed is the name of the system selected on the **Training Sets** > **z/OS Systems** tab. To select a different z/OS system, click **Return to Training Sets** to select another system from the Monitored z/OS Systems table. • For a Linux monitored client, use the "Training system" field to type the name of the Linux system or select it from the list of member systems in the model group that you selected on the **Training Sets** > **Model Groups** tab. A Linux system name is a fully qualified domain name, a host name, or an IP address.
Model dates	Allows you to select the dates you want to view. You can select the Next Training Period Model Dates or the Current Model Dates. Depending on the value that you select for Model dates, the content of the Manage Model Dates page changes. When you display the Manage Model Dates page, Next Training Period Model Dates is selected by default.
Today's date (UTC)	Provides the current date in Coordinated Universal Time (UTC).
Manual training period begin date (UTC)	Provides the earliest date in UTC for which IBM zAware will include data when building a model in response to a training request you submit today. That is, the model will include data for dates that occurred on or between UTC midnight on the begin date through 23:59:59 UTC yesterday.

Table 54. Fields displayed in the Summary view for the next model (continued)

Field	Description
Next training period begin date (UTC)	Provides the earliest date in UTC for which IBM zAware will include data when automatically building the next model. That is, the model will include data for dates that occurred on or between UTC midnight on the begin date through 23:59:59 UTC on the day before the next scheduled training date. If a dash (–) is displayed in this field, the corresponding monitored system is most likely not connected to the IBM zAware server.
Next scheduled training date (UTC)	Provides the date in UTC when IBM zAware is scheduled to submit a request to automatically build a model. For an automatic build to be scheduled, the system must be connected to the IBM zAware server. • If a model has not been built for this system, the next training date is based on the first date for which data is available and the training period value. • If a model has been built for this system, the next training date is based on the training interval value. IBM zAware uses the training interval value to determine the schedule for automatic builds only after the initial client model is built successfully. If a dash (–) is displayed in this field, the corresponding monitored system is most likely not connected to the IBM zAware server.
Excluded dates (UTC)	Allows you to select dates to exclude from the model. You can type the date or select it from the calendar widget. Use the features depicted in Figure 59 on page 245 to select the available dates to exclude from the model. After you specify the date to exclude, click **Add** to add it to the **Excluded Days** list. For more information about excluding dates, see "Excluding dates from a model" on page 226.
Unavailable dates (UTC)	Lists the dates in UTC for which IBM zAware will not include the data in the next model. To remove a date from the list, select it and click **Remove**.

The callout labels in Figure 59 on page 245 correspond to the following controls:
1. Scroll by day
2. Display the month list
3. Display the calendar
4. Scroll by month
5. Dates for which data is not available
6. Dates for which data is available
7. Scroll by year

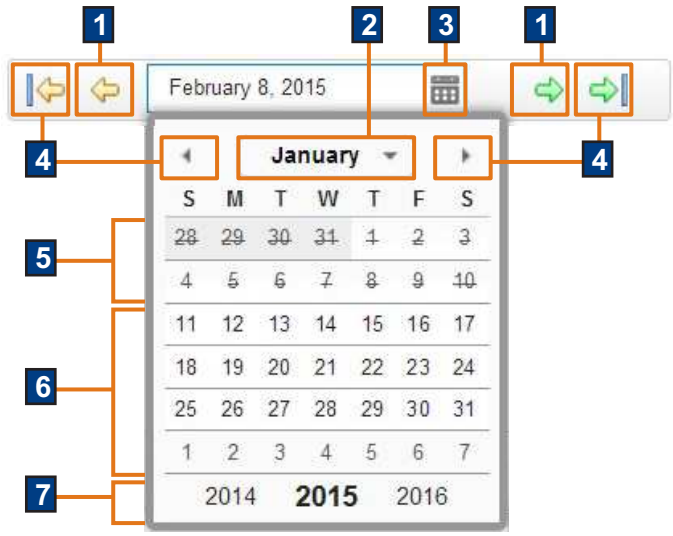

Figure 59. Features in the calendar widget

Fields displayed for Current Model Dates

When you select *Current Model Dates* in the **Model dates** field, IBM zAware displays information about the model it is currently using for the monitored system. For a description of the fields displayed for the current model, see Table 55.

Table 55. Fields displayed in the Summary view for the current model

Field	Description
Training system	Identifies the z/OS or Linux system for which data is displayed.
Model dates	Allows you to select the dates you want to view. You can select the Next Training Period Model Dates or the Current Model Dates. Depending on the value that you select for Model dates, the content of the Manage Model Dates page changes. When you display the Manage Model Dates page, Next Training Period Model Dates is selected by default.
Today's date (UTC)	Provides the current date in Coordinated Universal Time (UTC).
Current model trained date (UTC)	Provides the date in UTC when IBM zAware built the current model.
Current model begin date (UTC)	Provides the training period begin date in UTC that IBM zAware used when building the current model. That is, the current model includes data for dates that occurred on or between UTC midnight on the begin date through 23:59:59 UTC on the day before the current model trained date.
Excluded Days (UTC)	Lists the dates in UTC for which IBM zAware did not include the data in the current model. For more information about excluding dates, see "Excluding dates from a model" on page 226.

Calendar view

The Calendar view provides a pictorial view of the dates that are associated with the model. To display this view, click **Switch to Calendar View** on the Summary view of the Manage Model Dates page.

The controls and content displayed in the Calendar view are described in the following sections:

- "Fields displayed in the Calendar view" on page 246
- "Understanding the training calendar" on page 246

To switch to a text-based view of the model dates, click **Switch to Summary View**. To display the Training Sets page, click **Return to Training Sets** or click the **Training Sets** breadcrumb.

Fields displayed in the Calendar view

Table 56 describes the fields that are displayed in the Calendar view.

Table 56. Fields displayed in the Calendar view

Field	Description
Training system	Identifies the z/OS or Linux system for which data is displayed. • For a z/OS monitored client, the name displayed is the name of the system selected on the **Training Sets** > **z/OS Systems** tab. To select a different z/OS system, click **Return to Training Sets** to select another system from the Monitored z/OS Systems table. • For a Linux monitored client, use the "Training system" field to type the name of the Linux system or select it from the list of member systems in the model group that you selected on the **Training Sets** > **Model Groups** tab. A Linux system name is a fully qualified domain name, a host name, or an IP address.
Model dates	Allows you to select the dates you want to view. You can select the Next Training Period Model Dates or the Current Model Dates. Depending on the value that you select for Model dates, the content of the Manage Model Dates page changes. When you display the Manage Model Dates page, Next Training Period Model Dates is selected by default.
Training calendar	Uses a calendar to display the dates that are associated with the model. For more details about the training calendar, see "Understanding the training calendar."

Understanding the training calendar

The training calendar displays the important dates for a model using squares of different colors. For example, a green square represents the date when the next training is scheduled. Dates that are shown in a white square (without a colored outline) represent days for which IBM zAware has system data available for training purposes. For more details about the calendar, see Figure 60 on page 247 and Table 57 on page 247.

If Next Training Period Model Dates is selected in the **Model dates** field, you can use the training calendar to select the dates to exclude from or re-include in the next model. To exclude a date, select an available date. A burgundy box will surround the date. To re-include a date, select it again. The box will be removed.

You cannot modify the excluded dates for the current model, and you cannot use the Calendar view to modify any of the other dates for either model.

For more information about excluding days, see "Excluding dates from a model" on page 226.

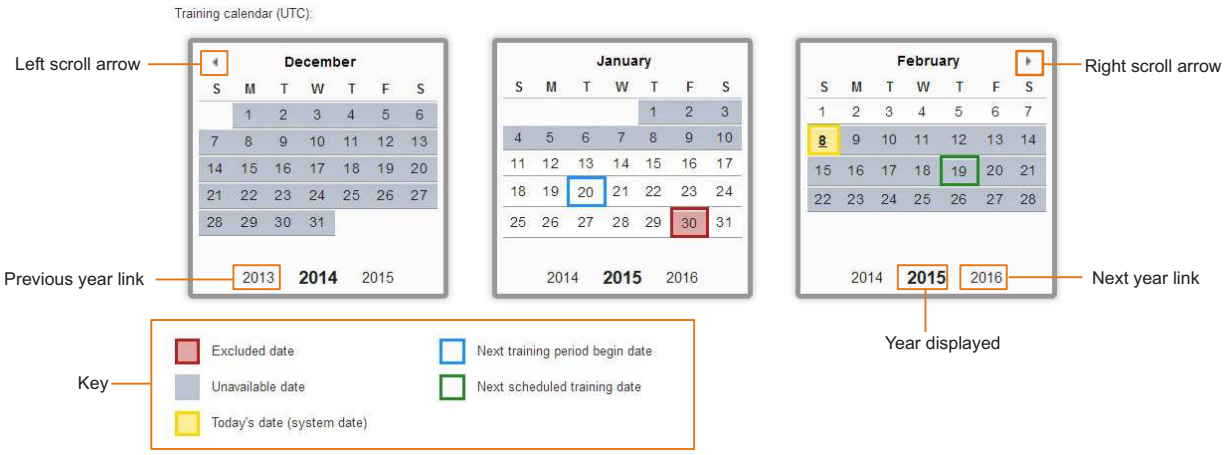

Figure 60. Training Calendar

Table 57. Items displayed in the training calendar

Field	Description
Left scroll arrow	Scrolls to the previous month.
Right scroll arrow	Scrolls to the next month.
Previous year link	Scrolls to the previous year.
Next year link	Scrolls to the next year.
Year displayed	Provides the year for which data is displayed.
Key	Explains what the different color boxes in the calendar represent. Dates that are shown in a white square (without a colored outline) represent days for which IBM zAware has system data available for training purposes. The colors can represent one of the following dates: **Current model begin date (purple square)** The training period begin date that IBM zAware used when building the current model. This date is displayed only when Current Model Dates is selected in the **Model dates** field. **Current model trained date (orange square)** The date when IBM zAware built the current model. This date is displayed only when Current Model Dates is selected in the **Model dates** field. **Excluded days (burgundy square)** The dates that IBM zAware excluded from the current model or that IBM zAware will exclude from the next model it builds. **Next scheduled training date (green square)** The date when IBM zAware is scheduled to submit a request to automatically build the next model. This date is displayed only when Next Training Period Model Dates is selected in the **Model dates** field. **Next training period begin date (blue square)** The earliest date for which IBM zAware will include data when automatically building the next model. This date is displayed only when Next Training Period Model Dates is selected in the **Model dates** field. **Today's date (yellow square)** The current date on the system. **Unavailable days (gray square)** The dates for which IBM zAware did not or has not received data for the monitored system.

Manage Ignored Messages page

The Manage Ignored Messages page displays the Ignored Messages table for a particular system. The Ignored Messages table lists any messages that an administrator has requested IBM zAware to ignore while analyzing current data from a specific monitored client. The table is empty if an administrator did not previously designate any messages for IBM zAware to ignore during analysis. An administrator can identify only z/OS messages to be ignored during analysis.

Figure 61 shows a sample view of the Manage Ignored Messages page.

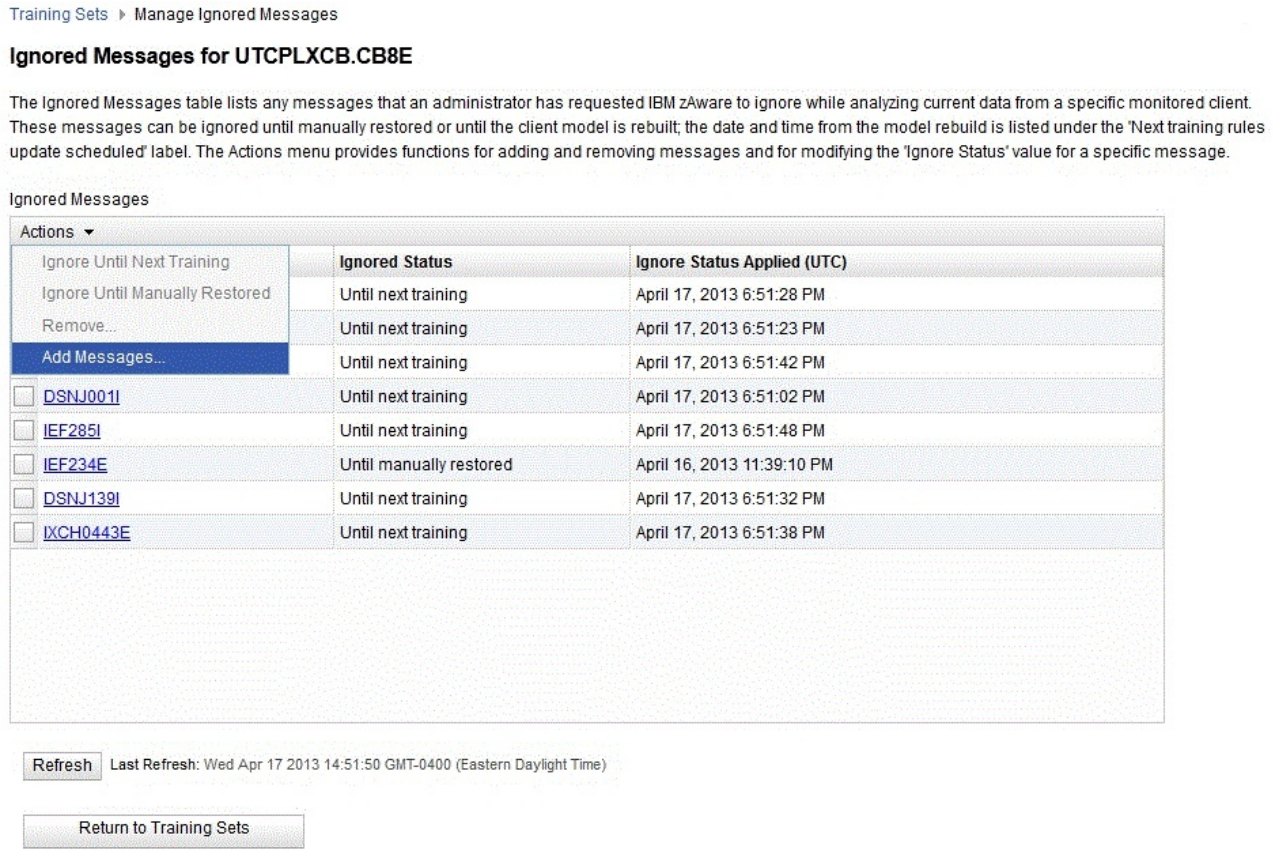

Figure 61. The Manage Ignored Messages page

The Actions menu provides functions for adding and removing messages, and for modifying the ignore status value for a specific message. The Manage Ignored Messages page also includes the following controls:

- Click **Refresh** to update the display contents in the Ignored Messages table.
- Click **Return to Training Sets** to close the Manage Ignored Messages page and return to the Training Sets page.

Fields displayed in the Ignored Messages table

Table 58 describes the fields that are displayed in the Ignored Messages table.

Table 58. Fields displayed in the Ignored Messages table

Field	Description
Actions	The **Actions** list provides functions for adding and removing messages, and for modifying the ignore status value for a specific message. All actions except for **Add Messages** require that you first select one or more messages in the Ignored Messages table. **Ignore Until Next Training** Marks the selected messages to be ignored during analysis until the next time IBM zAware successfully builds a model of behavior for the monitored system, as the result of either a manually requested or automatically scheduled training operation. The date for the next scheduled model rebuild is listed under "Next scheduled training date" after you select **Next Training Period Model Dates** on the Manage Model Dates page. Use this value for messages that you determine to be anomalous because of a workload change on the system, but you expect them to become part of the normal behavior for this system. The next training results in a model that includes these messages, which will be subject to normal analysis after the training. **Ignore Until Manually Restored** Marks the selected messages to be ignored during analysis until an administrator manually changes the ignore status of the message, or removes it from the list of messages in the Ignored Messages table. Use this value for messages that you determine to be normal (that is, not indicative of a problem) on this system. In subsequent analysis, these messages do not contribute to the anomaly score, and thus reduce false-positive results. **Remove** Opens the Remove Ignored Messages window, which displays a list of the messages that you selected for removal for this system. Click **OK** to apply your changes, or click **Cancel** to return to the Manage Ignored Messages page. **Add Messages** Opens the Add Ignored Messages window, through which you can enter message IDs of messages to ignore for the selected system. "Add Ignored Messages window" on page 250 describes the items in this window.
Message ID	Provides the message identifier.
Ignore Status	Lists the ignore status for the message, which can be one of the following values: **Until next training** The message is ignored during analysis until the next time IBM zAware successfully builds a model of behavior for the monitored system, as the result of either a manually requested or automatically scheduled training operation. The date for the next scheduled model rebuild is listed under "Next scheduled training date" after you select **Next Training Period Model Dates** on the Manage Model Dates page. **Until manually restored** The message is ignored during analysis until an administrator manually changes the ignore status of the message, or removes it from the list of messages in the Ignored Messages table.
Ignore Status Applied (UTC)	Indicates the date and time (in UTC, by using the 12-hour clock) at which an administrator most recently updated the ignore status for the message.

Add Ignored Messages window

Use the Add Ignored Messages window to enter the message IDs of messages to ignore for the selected system. This capability is available only for messages that z/OS monitored systems issue.

Table 59. Fields displayed in the Add Ignored Messages window

Field	Description
Current system	Identifies the system for which IBM zAware is to ignore specific messages during its analysis of current data.
Message IDs to ignore	The text entry field into which you can enter one or more message IDs for IBM zAware to ignore. If you enter more than one message ID, separate each ID with a comma.The message identifier (ID) must be a well-formed ID. IBM zAware recognizes messages IDs that conform to the z/OS standard, which consists of a component identifier, a message number, and an action code, in that order. IBM zAware also is capable of recognizing message IDs that do not completely conform to this z/OS standard.
Ignore messages until next training occurs for the current system	Marks the selected messages to be ignored during analysis until the next time IBM zAware successfully builds a model of behavior for the monitored system, as the result of either a manually requested or automatically scheduled training operation. The date for the next scheduled model rebuild is listed under "Next scheduled training date" after you select **Next Training Period Model Dates** on the Manage Model Dates page. Use this value for messages that you determine to be anomalous because of a workload change on the system, but you expect them to become part of the normal behavior for this system. The next training results in a model that includes these messages, which will be subject to normal analysis after the training.
Ignore messages until manually restored	Marks the selected messages to be ignored during analysis until an administrator manually changes the ignore status of the message, or removes it from the list of messages in the Ignored Messages table. Use this value for messages that you determine to be normal (that is, not indicative of a problem) on this system. In subsequent analysis, these messages do not contribute to the anomaly score, and thus reduce false-positive results.
OK	Click **OK** to apply your changes.
Cancel	Click **Cancel** to return to the Manage Ignored Messages page.

Chapter 22. Viewing and modifying the topology of IBM zAware monitored systems

The IBM zAware server dynamically discovers and provides topology information for all of the systems that it monitors. Administrators can use the **Topology** tab on the **Administration** > **Configuration** page to view and modify topology information.

The **Topology** tab displays a list of system groups and their member systems, along with systems that have not been assigned to a particular group. Through this display, administrators can select and move all z/OS members or individual z/OS systems from one sysplex to another, and remove any type of monitored system from the IBM zAware topology.

For additional information about using the **Topology** tab, see the following topics:
- "The Topology tab"
- "Modifying the z/OS sysplex topology" on page 252
- "Removing systems from the IBM zAware topology" on page 254

To assign or move Linux systems to a model group, administrators create or edit a model group through the **Model Groups** tab on the Systems page; for more information, see "Managing groups of Linux monitored clients" on page 210.

The Topology tab

The **Topology** tab displays a list of system groups and their member systems, along with systems that have not been assigned to a particular group. Through this display, administrators can select and move all z/OS members or individual z/OS systems from one sysplex to another, and remove any type of monitored system from the IBM zAware topology.

- To move a z/OS system from one sysplex node in the topology to another sysplex node, use the instructions in "Modifying the z/OS sysplex topology" on page 252.
- To assign or move Linux systems to a model group, administrators create or edit a model group through the **Model Groups** tab on the Systems page; for more information, see "Managing groups of Linux monitored clients" on page 210.
- To remove one or more systems from the topology, use the instructions in "Removing systems from the IBM zAware topology" on page 254.

Table 60 on page 252 describes the fields that are displayed in the Systems Topology table.
- To make sure the display contains the latest information, click **Refresh**.
- To filter the display by group or system name, type a value in the Filter field.
- To change the sort order, click any of the column headings.
- To view individual systems in a particular group, click the plus sign next to the group name to expand the list of member systems.

Table 60. Items in the Systems Topology table

Column	Description
Actions	An administrator must select one or more systems before choosing an action. **Move Selected Systems** Opens the Move Selected Systems window, through which an administrator can move selected z/OS systems into an available sysplex. If you have inadvertently selected a model group or a Linux system, the **Move Selected Systems** action is not enabled. **Remove Selected Systems** Opens the Remove Selected Systems window, through which an administrator can confirm the selected z/OS or Linux systems to be removed from the IBM zAware topology.
System Group / System	Identifies the name of a system group or monitored system.
Type	Identifies the type of system group or the type of monitored system. • Values for type of system group: Sysplex or Model Group • Values for type of monitored system: z/OS or Linux
Status	Indicates whether the system is connected to the IBM zAware server. The system can have one of the following status values. **Active** Indicates that the system is connected to the IBM zAware server. The system might or might not be transmitting data to the server. **Inactive** Indicates that the system was previously connected to the IBM zAware server but is disconnected.

Modifying the z/OS sysplex topology

The IBM zAware server dynamically discovers and provides sysplex topology information for all the z/OS systems it monitors. If you change your sysplex topology, such as moving a system to a different sysplex, you need to modify the IBM zAware sysplex topology accordingly. To modify the sysplex topology for your monitored systems, use the **Topology** tab on the Configuration page.

Before you begin

Ensure that at least one z/OS system is or was previously connected to the IBM zAware server. Otherwise, the topology display will not contain any sysplexes.

About this task

The IBM zAware server receives OPERLOG or SYSLOG data from the z/OS system logger running on a z/OS monitored system. The server processes the log data and extracts the name of the z/OS system and the name of the sysplex to which the monitored system belongs. The server uses the system and sysplex name to:

- Uniquely identify a z/OS monitored system.
- Associate log data, models, and analytic data with the correct z/OS system.
- Build the sysplex topology.

The z/OS system name and sysplex name must uniquely identify the system to be monitored. IBM zAware identifies each monitored client by sysplex and system name, in the format

sysplex_name.*system_name*; for example: SYSPLEX1.SYSA. IBM zAware cannot monitor more than one system with the same sysplex and system name combination.

Note that the priming data from the z/OS bulk load client does not include the sysplex name; therefore, you must use the **Priming Data** tab on the Configuration page to assign the priming data, if any, to the correct sysplex. After you assign the priming data, the IBM zAware server updates the sysplex topology.

If your installation moves a z/OS system to another sysplex, you must do the following:

- Update the sysplex topology through the IBM zAware GUI. For example, if your installation moved SYS1 from SYSPLEXB to SYSPLEXC, you also need to move SYS1 from SYSPLEXB to SYSPLEXC in the IBM zAware sysplex topology. Doing so instructs IBM zAware to associate the model and training sets information that it collected for SYSPLEXB.SYS1 with SYSPLEXC.SYS1. However, the analysis results for prior dates are not associated with SYSPLEXC.SYS1.
- Connect the z/OS system to the IBM zAware server.

To move a z/OS system from one sysplex node in the topology to another sysplex node, complete the steps that follow.

Procedure

1. Expand the Administration category in the navigation pane and select **Configuration**. The Configuration page is displayed.
2. Click **Topology** tab to display the Systems Topology table.
3. In the Systems Topology table, select one or more z/OS systems to be moved.

   If necessary, click the Type column heading to sort the system groups such that sysplexes are the first entries in the table. Select only the z/OS systems to be moved into the same target sysplex.

   - If you are moving all of the systems in one sysplex to another sysplex, you can select the containing sysplex to automatically select all of its system members, rather than selecting each system individually.
   - To select individual systems that are currently in different sysplexes, expand the sysplex nodes, as necessary, and select only those systems to be moved into the same target sysplex.
4. From the Actions list, click **Move Selected Systems**. The Move Selected Systems window opens.

   If you have inadvertently selected a model group or a Linux system, the **Move Selected Systems** action is not enabled. To assign or move Linux systems to a model group, administrators create or edit a model group through the **Model Groups** tab on the Systems page; for more information, see "Managing groups of Linux monitored clients" on page 210.
5. In the "Available sysplexes" field, select the target sysplex to which you want to move the systems identified in the "Selected systems" list. You can select only one sysplex.
6. Click **OK** to update the topology.

Results

IBM zAware recycles the analytics engine so that your changes take effect. When the analytics engine is stopped or recycled, IBM zAware disconnects all monitored systems from the server.

- When z/OS monitored clients are disconnected from the server, they automatically attempt to reconnect and continue to buffer data for approximately 10 minutes. If the analytics engine is restarted within that time, the z/OS system reconnects and sends the buffered data to the IBM zAware server. If the z/OS system times out before the analytics engine is restarted, any buffered data is lost and you must manually reconnect the system.

  To determine whether any z/OS systems require reconnection, click **Refresh** to display updated status information for systems in the topology. If Inactive continues to be displayed in the Status column after a reasonable amount of time, you need to manually reconnect the system.

  To reconnect a z/OS system, issue the SETLOGR command.

```
SETLOGR FORCE,ZAICONNECT,LSN=SYSPLEX.OPERLOG
```

- When Linux monitored clients are disconnected from the server, they normally attempt to reconnect to the server; if they do not reconnect, you must manually reconnect them. To reconnect a Linux system, restart the syslog daemon, using the appropriate command for the type of syslog daemon and for the type of initialization (init) process that is in use on the Linux system.

Move Selected Systems window

Use the Move Selected Systems window to select the sysplex to which you want to move the selected z/OS systems.

In the Move Selected Systems window, the following information is displayed:

Selected systems
: Lists the systems that you selected on the **Topology** tab. These systems will be moved to the sysplex you select in the "Available sysplexes" field.

Available sysplexes
: Lists the sysplexes to which one or more monitored systems belong. Select the sysplex to associate with the systems specified in the "Selected systems" list.

Click **OK** to update the sysplex topology.

The IBM zAware server will recycle the analytics engine so that your changes can take affect. When the server restarts the engine, you might need to reconnect monitored systems to the server.

Removing systems from the IBM zAware topology

The IBM zAware server dynamically discovers and provides topology information for all of the systems it monitors. If you want to remove one or more systems from the IBM zAware topology, use the **Topology** tab on the Configuration page.

About this task

When an administrator selects one or more systems to remove from the topology, IBM zAware also removes the data that is associated with the selected system, including:

- Current instrumentation data and priming data, if any.
- The system model, or the group model only when no other systems remain in the group.
- Analysis results for the selected system.

A sysplex or model group remains in the topology, even after all of the systems in the sysplex or model group are removed.

Procedure

1. To prevent the possible rediscovery of a system, disconnect each monitored system that you want to remove.
   - For monitored z/OS clients, use the **SETLOGR** command on each z/OS system to prevent the systems from attempting to reestablish the TCP/IP connection to the IBM zAware server.

     ```
     SETLOGR FORCE,ZAIQUIESCE,ALL
     ```

   - For monitored Linux systems, stop the syslog daemon, by using the appropriate command for the type of syslog daemon that is in use on the Linux system.
2. Expand the Administration category in the navigation pane and select **Configuration**. The Configuration page is displayed.
3. Click **Topology** tab to display the Systems Topology table.
4. In the Systems Topology table, select one or more monitored systems to be removed. From the Actions list, click **Remove Selected Systems**. The Remove Selected Systems window opens.

5. Verify that the listed systems are the monitored clients that you want to remove. If the Status value indicates that a system is still active, use the appropriate command to manually disconnect the system, as described in step 1 on page 254.
6. Click **OK** to remove the selected systems from the topology.

Results

IBM zAware disconnects any currently connected systems, and recycles the analytics engine so that your changes take effect. When the remove request for each system completes, IBM zAware issues message AIFB0018I, which identifies the removed system by name.

What to do next

After IBM zAware restarts the analytics engine, you might need to reconnect some monitored systems to the server. For more details, see "Starting and stopping data collection for your monitored systems" on page 203.

Remove Selected Systems window

Use the Remove Selected Systems window to verify the list of systems for IBM zAware to remove from its topology.

In the Remove Selected Systems window, the following information is displayed:

System Name
Specifies the name of a system that is to be removed from the topology.
- For a z/OS monitored client, the name has the format *sysplex-name.system-name*, where *sysplex-name* is the name of the sysplex to which the system belongs and *system-name* is the name of the system
- For a Linux monitored client, the name is a fully qualified domain name, a host name, or an IP address.

Type Indicates whether the monitored client is a z/OS or Linux monitored client.

Status Indicates whether the system is connected to the IBM zAware server. The system can have one of the following status values.

Active Indicates that the system is connected to the IBM zAware server. The system might or might not be transmitting data to the server.

Inactive
Indicates that the system was previously connected to the IBM zAware server but is disconnected.

Click **OK** to remove the listed systems and their associated data, or click **Cancel** to return to the Topology tab.

Chapter 23. Collecting priming data for z/OS system models

To provide analytical data for a monitored client, IBM zAware requires a model of normal system behavior to use for comparison. For z/OS monitored clients, IBM zAware builds a model for each z/OS system, using message data from that system. You have two options for building a model: waiting for the server to build a model from data collected over a specific time period, or priming the server with prior data. This priming option is recommended because analysis can start shortly after the model is built.

The following topics discuss the two options for building a model for a z/OS monitored client:

- "Waiting for the server to build a z/OS model"
- "Transferring priming data to build a z/OS model"

Waiting for the server to build a z/OS model

After your installation configures z/OS monitored systems to send data to the IBM zAware server, no additional configuration is required because the server automatically starts receiving current data from the z/OS system logger running on the z/OS monitored systems. When the training period elapses, the IBM zAware server automatically builds the IBM zAware model.

The estimated amount of data for building the most accurate z/OS models is 90 days of data for each system. Therefore, with this option, you have to wait for the IBM zAware server to collect data for 90 days before the IBM zAware server can build a model and use that model to start detecting system problems. Your installation can modify the number of days required for this training period, based on your knowledge of the workloads running on z/OS monitored systems. For more details, see "Specifying settings for the analytics engine" on page 199.

Transferring priming data to build a z/OS model

Instead of waiting for the IBM zAware server to collect data over the course of the training period, you can prime the server by transferring prior data from the hardcopy or system logs of z/OS monitored systems, and request the server to build a model for each z/OS system from the transferred data. To do so, configure and run the z/OS bulk load client for IBM zAware on the z/OS priming system. For instructions, see "Creating an IBM zAware model for new z/OS monitored clients" on page 118.

In contrast to data that the IBM zAware server receives from the z/OS system logger running on a z/OS monitored system, the priming data from the z/OS bulk load client does not include the name of the sysplex to which the z/OS monitored system belongs. Without the sysplex name, the IBM zAware server cannot associate the priming data with the appropriate sysplex and cannot include the data in a model.

The sections that follow explain how to associate the z/OS priming data with a sysplex.

Assigning z/OS priming data to a sysplex

If you used the z/OS bulk load client to transfer priming data to the IBM zAware server, use the **Priming Data** tab on the Configuration page to assign the received priming data to the appropriate sysplex.

Before you begin

Ensure that your installation has completed the following actions:

1. Configured storage, security, and analytics for the IBM zAware server.

2. Configured z/OS monitored systems to send data to the IBM zAware server. Verify that at least one z/OS monitored system is connected to the IBM zAware server for each sysplex to which you want to assign data.
3. Configured and run the z/OS bulk load client on the z/OS priming system.

About this task

The following steps explain how to assign priming data from z/OS monitored systems by moving those systems from the "Priming message data by systems" list to the Sysplex Topology list on the **Priming Data** tab. You do not have to assign all systems in the list until you are ready to do so. Unassigned systems remain in the "Priming message data by systems" list until you add them to the sysplex topology, or use **Delete** to remove them.

Procedure

1. Expand the Administration category in the navigation pane, and select **Configuration**. The Configuration page is displayed.
2. Click **Priming Data** to display the **Priming Data** tab.
3. In the Sysplex Topology list, select the sysplex to which you want to assign systems. You can select only one sysplex. If the sysplex to which you want to assign the system is not listed, to add it to the topology, you must configure a z/OS monitored system in that sysplex and connect it to the IBM zAware server.
4. Complete one of the following actions:
   a. Select one or more z/OS systems in the "Priming message data by systems" list and click **Add** to move those systems to the selected sysplex node in the Sysplex Topology list.
   b. Click **Add All** to move all the z/OS systems listed in the "Priming message data by systems" list to the selected sysplex node in the Sysplex Topology list.

   The selected systems are displayed under the selected sysplex, with the parenthetical phrase *to be assigned* displayed after the system name. If necessary, expand the sysplex node to see the list of systems for the selected sysplex.

   If the selected sysplex already contains a system with the same name as the selected system, you can still add the system to the sysplex topology. In this case, during the assign process, the IBM zAware server merges the priming data with the data that already exists for the system.

   If the selected sysplex does not contain the selected system, during the assign process, the IBM zAware server moves the data to that sysplex.
5. Repeat step 4, as needed, to move each z/OS system in the "Priming message data by systems" list to the appropriate sysplex in the Sysplex Topology list.

   If you add one or more systems to the incorrect sysplex, you can move them back to the "Priming message data by systems" list by selecting the sysplex and clicking **Remove**. To move all *to be assigned* systems from the Sysplex Topology list, click **Remove All**.
6. When you have finished moving z/OS systems from the "Priming message data by systems" list to the appropriate sysplex, click **Assign** to apply your changes. The Assign Priming Data window is displayed.
7. Review and confirm your changes by clicking **OK**.

Results

IBM zAware recycles the analytics engine so that your changes take effect. When the analytics engine is stopped or recycled, IBM zAware disconnects all monitored systems from the server.

- When z/OS monitored clients are disconnected from the server, they automatically attempt to reconnect and continue to buffer data for approximately 10 minutes. If the analytics engine is restarted within that time, the z/OS system reconnects and sends the buffered data to the IBM zAware server. If

the z/OS system times out before the analytics engine is restarted, any buffered data is lost and you must manually reconnect the system. To reconnect a z/OS system, issue the SETLOGR command.

```
SETLOGR FORCE,ZAICONNECT,LSN=SYSPLEX.OPERLOG
```

- When Linux monitored clients are disconnected from the server, they normally attempt to reconnect to the server; if they do not reconnect, you must manually reconnect them. To reconnect a Linux system, restart the syslog daemon, using the appropriate command for the type of syslog daemon and for the type of initialization (init) process that is in use on the Linux system.

If the IBM zAware server is processing a training request when the analytics engine must be restarted, the training request is canceled and replaced in the queue so it is the first request to be processed when the analytics engine is available again.

What to do next

To verify that the transferred data is available for the IBM zAware server to use, complete the following steps:

1. Expand the Administration category in the navigation pane and select **Training Sets** to display the Training Sets page.
2. Select the z/OS monitored system for which you transferred priming data.
3. From the **Actions** list, select **Manage Model Dates**. The Manage Model Dates page is displayed.
4. In the Model dates field, select **Current Model Dates**.
5. Click **Switch to Calendar View** to use the calendar to determine days for which transferred data is available. Calendar days that are not marked as *Excluded* or *Unavailable* identify the dates for which the IBM zAware server has data to use.
6. Click **Return to Training Sets** to return to the Training Sets page.
7. Repeat these steps, as needed, for each system for which data was transferred.

After you verify that the data is available for the server to use, request training for the monitored system. For instructions, see "Requesting training automatically or manually" on page 227.

If you want to delete any unassigned systems, select one or more systems in the "Priming message data by system list" and click **Delete** to remove them. The Delete Priming Data window is displayed so you can review the systems to be deleted, and confirm or cancel this request. If you click **OK** to confirm the delete request, IBM zAware recycles the analytics engine so that your changes take effect.

Priming Data tab

When the IBM zAware server receives priming data from the z/OS bulk load client, the data does not include the name of the sysplex to which the z/OS monitored system belongs. Without the sysplex name, the IBM zAware server cannot associate the priming data with the appropriate sysplex and cannot include the data in a model. You can use the **Priming Data** tab on the Configuration page to assign the z/OS priming data to the appropriate sysplex.

For a description of the items that are displayed on the **Priming Data** tab, see Table 61.

Table 61. Items displayed on the Priming Data tab

Item	Description
Priming message data by system	Lists the z/OS systems for which priming data is available that is not assigned to a sysplex. Because the z/OS bulk load client can send data for more than one monitored system at a time, several systems might be listed.
Sysplex Topology	Lists the z/OS monitored systems that are or were previously connected to the IBM zAware server and organizes that list by sysplex. If you want to make changes to the sysplex topology, complete the steps provided in "Modifying the z/OS sysplex topology" on page 252.

Table 61. Items displayed on the Priming Data tab (continued)

Item	Description
Buttons	**Add** Moves the selected z/OS systems in the "Priming message data by systems" list to the selected sysplex in the Sysplex Topology list. **Add All** Moves all the z/OS systems in the "Priming message data by systems" list to the selected sysplex in the Sysplex Topology list. **Assign** Allows you to confirm and proceed with the assignments you specified. **Delete** Deletes one or more unassigned systems from the "Priming message data by systems" list. **Remove** Moves all the z/OS systems in the selected sysplex that are marked as *to be assigned* to the "Priming message data by systems" list. **Remove All** Moves all the z/OS systems in the Sysplex Topology list that are marked as *to be assigned* to the "Priming message data by systems" list.

Assign Priming Data window

Use the Assign Priming Data window to verify that the priming data for each z/OS monitored system is assigned to the correct sysplex and to start the priming data assignment process.

The Assign Priming Data window provides the following information:

Priming Data by System
Name of the z/OS system associated with the priming data.

Sysplex to Assign
Name of the sysplex to which the priming data will be assigned.

Review and confirm your assignments by clicking **OK**. IBM zAware restarts the analytics engine so that your changes take affect.

Chapter 24. Setting up a local repository to secure access to the IBM zAware GUI

During installation, you can provide user authentication to the IBM zAware graphical user interface (GUI) through either an existing Lightweight Directory Access Protocol (LDAP) repository or a local file-based repository. For simplicity, using only an LDAP repository is the preferred option.

However, you might want to define one or more user IDs in a local repository, so you can access the IBM zAware GUI when the LDAP server is not available. If you configure an LDAP repository and also define users or groups in a local repository, both sets of users or groups are available through the IBM zAware GUI. Use the following procedures to add users or groups to a local file-based repository.

To configure an existing LDAP repository for user authentication, see Chapter 13, "Configuring storage, security, and analytics for the IBM zAware server," on page 99 for instructions.

Before you begin

- Use the IBM zAware GUI to define users or groups to the local repository. You can also delete users and groups, delete group members, and change a user password using the GUI.
- Log in to the GUI with a user ID that has the appropriate authority to add or define users or groups. This user ID can be the default master user ID that is defined in the image profile for the IBM zAware partition on the host system, or another user ID that is assigned to the IBM zAware Administrator role.
- Make sure that you have reviewed the planning considerations in Chapter 9, "Planning for security," on page 75.
- Prepare a list of user IDs or groups to define in the local repository. Do not define the same user ID in more than one repository; results are not predictable.

Roadmap of this chapter

To define users or groups to a local repository, see the following steps:

- "Defining new local users"
- "Defining new local groups" on page 262
- "Adding one or more local users or members to a local group" on page 262

To delete a user or group from the local repository, remove a local member from a local group, or change a user password, see the following steps:

- "Deleting local users or groups" on page 262
- "Deleting one or more local users or members from a local group" on page 263
- "Changing a user password" on page 263

To assign each user ID or group in the local repository to a specific IBM zAware role, see "Assigning users or groups to a role" on page 187.

Defining new local users

Procedure

1. Log in to the IBM zAware GUI and provide a user ID and password for a user with the IBM zAware Administrator role.
2. On the navigation panel, select **Configuration**>**Security**>**Users and Groups**.

3. Define one new user at a time to the local repository. Under the **New User** section of the page, follow these steps:
   a. Enter the **New User Name** in the provided space.

      Note: zAware does not allow duplicate user names. Each name must be unique.
   b. Enter the **New User Password** in the provided space.
   c. Reenter the **New User Password** in the provided space.
   d. Click **Create New User**.
4. To create more users, repeat these steps.

Results

The user IDs and passwords that you add are defined to the local repository.

Defining new local groups

Procedure

1. Log in to the IBM zAware GUI and provide a user ID and password for a user with the IBM zAware Administrator role.
2. On the navigation panel, select **Configuration**>**Security**>**Users and Groups**.
3. Define one new group at a time to the local repository. Under the **New Group** section of the page, follow these steps:
   a. Enter the **New Group Name** in the provided space.

      Note: zAware does not allow duplicate group names. Each name must be unique.
   b. Click **Create New Group**.
4. To create more groups, repeat these steps.

Adding one or more local users or members to a local group

Procedure

1. Log in to the IBM zAware GUI and provide a user ID and password for a user with the IBM zAware Administrator role.
2. On the navigation panel, select **Configuration**>**Security**>**Users and Groups**.
3. Add one or more local users to a local group. Under the **Manage Group Membership** section of the page, follow these steps:
   a. From the drop-down menu under **Select a group to manager members**, select a group that you want to add one or more members to. This can produce a list of available users, the actions users can take, and a list of currently mapped users.
   b. From the **Available Users** list, select the user or users that you want to add to the group. Then, select **add** action to add one user or **add all** action to add multiple users.
   c. Click **Apply**. The users that you add to the group now appear under **Current Mapped Users**.

Deleting local users or groups

Procedure

1. Log in to the IBM zAware GUI and provide a user ID and password for a user with the IBM zAware Administrator role.
2. On the navigation panel, select **Configuration**>**Security**>**Users and Groups**.

3. Delete one local user or group at a time from the local repository. Under the **Delete Members** section of the page, follow these steps:
   a. From the drop-down menu, select a member to delete (either a user or a group).
   b. Click **Delete Member**.
4. To delete more users or groups, repeat these steps.

Deleting one or more local users or members from a local group

Procedure

1. Log in to the IBM zAware GUI and provide a user ID and password for a user with the IBM zAware Administrator role.
2. On the navigation panel, select **Configuration**>**Security**>**Users and Groups**.
3. Delete one or more local users from a local group. Under the **Manage Group Membership** section of the page, follow these steps:
   a. From the drop-down menu under **Select a group to manager members**, select a group that you want to remove one or more members from. This can produce a list of available users, the actions that users can take, and a list of currently mapped users.
   b. From the **Current Mapped Users** list, select the user or users that you want to remove from the group. Then, select the **remove** action to remove one user or the **remove all** action to remove multiple users.
   c. Click **Apply**. The users that you remove from the group now appear under **Available Users**.

Changing a user password

In the IBM zAware GUI, the administrator cannot change user passwords. Users must change their own passwords.

Procedure

1. Log in to the IBM zAware GUI with the user ID whose password you want to change.
2. In the GUI where the user name appears, open the drop-down menu and select **User Profile**.
3. On the user profile page, under the **Password Management** section, enter your current password, your new password, and confirm the new password.
4. Click **Update Password**. If the password is updated, the following message is displayed: `AIFF0035I Your password has been changed successfully`.

Chapter 25. Restoring IBM zAware configuration data

You can restore the configuration data from IBM zAware on the storage devices that you are currently using for the IBM zAware partition on a host system. The restore utility is available through the **Configuration** > **Utilities** tab.

About this task

The following IBM zAware configuration data can be restored:

- Current analytics settings for each type of monitored system listed under the **Configuration** > **Analytics** tab.
- Current Secure Sockets Layer (SSL) certificate information listed under the **Configuration** > **Security** > **SSL Settings** tab.
- Depending on the security mechanisms in use for this IBM zAware server, current user authentication details from one or both of the following:
  - An existing Lightweight Directory Access Protocol (LDAP) repository, the values for which are specified in the **Configuration** > **Security** > **LDAP Settings** tab
  - A local file-based repository that is defined through the zAware GUI.
- Current users and groups that are defined to the IBM zAware Administrator or User role, as shown under the **Configuration** > **Security** > **Role Mapping** tab.
- Current values listed under the **Configuration** > **Security** > **LTPA Settings** tab.

Procedure

1. Log in to the IBM zAware GUI.

   If you plan to use the restore utility to populate a new IBM zAware server with saved configuration data, you need to log in with the master user ID and password that was provided in the image profile for the new IBM zAware partition. Otherwise, you can log in with user ID that is mapped to an Administrator role.
2. Expand the Administration category in the navigation pane and select **Configuration**. The Configuration page is displayed.
3. Click **Utilities** to display the restore utility control.
   - To restore a previously saved copy of configuration data, click **Restore**. This utility can take a considerable amount of time to complete, partly because the IBM zAware server is automatically restarted.

     The restore utility is not available under the following conditions:
     - When a saved copy of configuration data is not available for use. Make sure that the appropriate MCLs have been applied to the host system.
     - When IBM zAware is processing another operation, and cannot run the restore utility. In this case, IBM zAware displays text that indicates an operation is in progress.

   IBM zAware presents a confirmation message before proceeding to run the utility. Click **OK** to submit the restore utility request, or **Cancel** to return to the **Utilities** tab.
   - When you submitted a request for the restore utility but another operation has already started, IBM zAware issues a message to indicate that you can retry the request at a later time. Otherwise, IBM zAware overwrites any configuration changes that were made to IBM zAware after the last saved copy was created, and restarts the server.
     - If the saved configuration data does not contain the user ID that you used to log in to the GUI for this session, you will be logged out, and will have to use a different administrator ID to log in again.

– If the saved configuration data contains an SSL certificate that has not been accepted into your browser, the restart operation might time out. In this situation, IBM zAware presents an error message indicating that the server has not restarted in the expected time. To resolve the problem:
  a. Click **OK** to dismiss the error message.
  b. Using your browser controls, refresh the page to force the browser to prompt you to accept the certificate.
  c. Accept the certificate.

Results

IBM zAware issues a message to indicate successful completion: AIFF0030I for the restore utility.

What to do next

- If you used the restore utility and want to verify that configuration data has been restored, go to the **Analytics** or **Security** tabs and check the displayed content.
- Because the IBM zAware server is automatically restarted as part of the restore utility processing, any connected monitored clients are disconnected. After the utility completes, go to the **Systems** > **System Status** tab to determine whether any systems need to be reconnected. For more information about reconnecting monitored clients, see "Starting and stopping data collection for your monitored systems" on page 203.

Chapter 26. Setting up multiple IBM zAware partitions for switchover situations

Your installation can configure more than one IBM zAware partition, with one for normal operations and another reserved for switchover situations. This type of configuration enables you to quickly restore IBM zAware operations after a failure. The primary and alternate partitions can reside on the same IBM zAware host system or on different host systems. Use the instructions in this topic to configure an IBM zAware environment that contains primary and alternate IBM zAware partitions.

Before you begin

- Read the planning information in "Planning persistent storage configuration and capacity" on page 59 before you configure a primary or alternate IBM zAware partition. The storage administrator at your installation needs to plan for and provide a list of storage devices that are reserved for use by the primary for normal operations, and a list of equivalent storage devices for the alternate to use, when necessary. The alternate set is a backup set that contains replicated data from the primary set of storage devices.
- Note that both host systems must be in the same IBM product family and have the same machine type. For example, switching from a z13 to an IBM zEnterprise host system (zEC12) is not supported.

About this task

- To correctly configure the partition in which the alternate server runs, use the same IP address as you defined for the primary partition. Doing so guarantees that you cannot have multiple IBM zAware servers running simultaneously, and also eliminates the need to reconfigure the TCP/IP settings of monitored clients if you have to switch from using the primary server to the alternate server.
- To correctly configure persistent storage for primary and alternate IBM zAware partitions, your installation must define physically separate but equivalent sets of storage devices for each partition, and also set up replication to copy the content of the primary storage devices to the alternate storage devices. For data replication to be successful, the number of storage devices in the primary set must match the number of devices in the alternate set. Additionally, each alternate device must be equivalent in size to the primary device.

The primary and alternate partitions can reside on the same IBM zAware host system or on different host systems. The alternate host system might have the IBM zAware disaster recovery (DR) feature installed, but this feature is not required.

Procedure

1. Use the instructions in Chapter 12, "Configuring network connections and storage for the IBM zAware partition," on page 95 to set up the network connections and physical storage devices for the IBM zAware environment.

   Use the Hardware Configuration Definition (HCD) to define network connections and storage devices for the primary and alternate IBM zAware partitions in the input/output configuration data set (IOCDS) for the appropriate host system.

   - For network connections, make sure that both the primary and alternate IBM zAware partitions have access to the same networks.
   - For physical storage devices, the storage administrator can use image access and candidate lists for channel path definitions to allow only the IBM zAware partition to access specific devices. Using the explicit device candidate list is an alternative method of restricting access to specific devices.
     a. Use the explicit candidate list to allow the primary IBM zAware partition to access only the set of storage devices that are intended for normal operations.

b. Use the explicit candidate list to allow the alternate IBM zAware partition to access only the set of storage devices that are to contain backup copies of IBM zAware data.
c. For the purposes of replication only, allow only one z/OS partition to access both sets of storage devices.

2. Use the instructions in "Configuring the IBM z Systems Secure Service Container for IBM zAware" on page 27 to create an activation profile for the primary IBM zAware partition. Through the Hardware Management Console (HMC), make sure you select the appropriate host system for the primary IBM zAware partition.
3. Use the instructions in Chapter 13, "Configuring storage, security, and analytics for the IBM zAware server," on page 99 to configure the primary IBM zAware partition.

   When you use the IBM zAware graphical user interface (GUI) to assign storage devices for the primary server, these devices become the in-use set.
4. To complete the configuration of the primary IBM zAware environment, use the instructions in "Configuring z/OS monitored clients to send data to the IBM zAware server" on page 111 and "Creating an IBM zAware model for new z/OS monitored clients" on page 118.

   When you complete this step, the primary IBM zAware server is analyzing data for its connected monitored clients, and is storing information related to its operation in the in-use set of physical storage devices.
5. Set up replication to copy the contents of the in-use set of storage devices to the backup set.

   For replication, your installation can consider using IBM FlashCopy or one of several Data Facility Storage Management Subsystem (DFSMS) copy services, including Extended Remote Copy (XRC) and Peer-to-Peer Remote Copy (PPRC) solutions. Another possible alternative is using DFSMShsm to copy data. In contrast to real-time replication solutions, DFSMShsm requires deactivating the IBM zAware partition before copying data, then reactivating the partition after the copy operation completes. Non-IBM replication products also are available for use.

   If you change the set of in-use devices by adding or removing devices through the GUI, make sure that you adjust replication accordingly. To successfully replace an in-use device with its equivalent backup device, the set of in-use devices must match the set of backup devices in number of devices, size of devices, and content.
6. After the contents of the in-use set of storage devices have been replicated to the backup set at least once, configure the alternate IBM zAware partition.
   a. Disconnect the monitored clients that are sending data to the primary IBM zAware server.
   b. Deactivate the primary IBM zAware partition.
   c. Using the image profile for the primary IBM zAware partition as a model, create an activation profile for the alternate IBM zAware partition.
      - Select the appropriate host system for the alternate IBM zAware partition.
      - Make sure that you use the same IP address as the one you defined for the primary partition.
   d. Activate the alternate IBM zAware partition.
7. Through the IBM zAware GUI, configure storage, security, and analytics for the alternate IBM zAware server.

   The instructions in Chapter 13, "Configuring storage, security, and analytics for the IBM zAware server," on page 99 are essentially the same for both the primary and alternate servers, with the exception of assigning storage devices. When assigning storage devices for the alternate server:
   a. Select only those devices in the backup set that are equivalent to the devices that are currently in use by the primary server.
   b. When adding the selected devices, use the **Preserve data** option to ensure that IBM zAware does not format or initialize these devices, so the replicated data that they contain is preserved and usable.

   After the selected devices are added, they constitute the in-use set for the alternate IBM zAware server.

8. Deactivate the alternate IBM zAware partition.
9. Reactivate the primary IBM zAware partition and reconnect its monitored clients.

Results

Your installation has two partitions of IBM zAware configured for use: the primary for normal operations, and the alternate for switchover situations, should any occur.

What to do next

- When a failure occurs and the primary partition is no longer available, activate the alternate IBM zAware partition. This switchover operation is successful only if the number of devices in the primary in-use set and in the alternate backup set match.
  - When z/OS monitored clients are disconnected from the server, they automatically attempt to reconnect and continue to buffer data for approximately 10 minutes. If the analytics engine is restarted within that time, the z/OS system reconnects and sends the buffered data to the IBM zAware server. If the z/OS system times out before the analytics engine is restarted, any buffered data is lost and you must manually reconnect the system. To reconnect a z/OS system, issue the SETLOGR command.

    ```
    SETLOGR FORCE,ZAICONNECT,LSN=SYSPLEX.OPERLOG
    ```
  - When Linux monitored clients are disconnected from the server, they normally attempt to reconnect to the server; if they do not reconnect, you must manually reconnect them. To reconnect a Linux system, restart the syslog daemon, using the appropriate command for the type of syslog daemon and for the type of initialization (init) process that is in use on the Linux system.

  Depending on the timing of the CPC failure and the replication schedule for backing up IBM zAware data, the data on the backup set might be back-level. In this case, the alternate IBM zAware cannot provide analytical data for the dates between the last day of replication and the date and time when the administrator activated the alternate IBM zAware.
- During normal operations, an administrator might need to change the set of in-use devices for the primary server by adding or removing devices through the GUI. Because the in-use set and backup set of devices must be equivalent for a switchover to be successful, the administrator also must adjust replication and the set of storage devices for the alternate server so both of the primary and alternate sets match in number of devices, size of devices, and content. If the number of devices in the in-use set and in the backup set do not match, you cannot successfully switch over to using the alternate partition.
- If you need to add, replace, or remove storage devices from the host system after initially configuring the primary and alternate partitions of IBM zAware, use the following procedures.
  - **To add a storage device**
    1. Use HCD to add the storage devices for the primary and alternate IBM zAware partitions in the input/output configuration data set (IOCDS) for the appropriate host system.
    2. Through the IBM zAware GUI, assign the new storage device to the in-use set of storage for the primary server.
    3. Update the replication method in use at your installation to copy data from the newly added storage device to its equivalent backup device.
    4. After replication occurs at least once during normal operations, update the configuration for the alternate partition of IBM zAware:
       a. Disconnect the monitored clients that are sending data to the primary IBM zAware server.
       b. Deactivate the primary IBM zAware partition.
       c. Activate the alternate IBM zAware partition.
       d. Through the IBM zAware GUI, assign the new backup device to the in-use set of storage for the alternate server.
       e. Deactivate the alternate IBM zAware partition.
       f. Activate the primary IBM zAware partition.
       g. Reconnect the monitored clients to the primary IBM zAware partition.

– **To remove a storage device**

1. Update the replication method in use at your installation to remove the storage device and its equivalent backup device.
2. Through the IBM zAware GUI, remove the storage device from the in-use set of storage for the primary server.
3. After replication occurs at least once during normal operations, update the configuration for the alternate partition of IBM zAware:
   a. Disconnect the monitored clients that are sending data to the primary IBM zAware server.
   b. Deactivate the primary IBM zAware partition.
   c. Activate the alternate IBM zAware partition.
   d. Through the IBM zAware GUI, remove the backup device from the in-use set of storage for the alternate server.
   e. Deactivate the alternate IBM zAware partition.
   f. Activate the primary IBM zAware partition.
   g. Reconnect the monitored clients to the primary IBM zAware partition.
4. Use HCD to remove the storage devices from the IOCDS for the appropriate host system.

Chapter 27. Enabling system management products to use IBM zAware data

IBM zAware provides an application programming interface (API) through which existing alerting products can be enhanced by including IBM zAware data into their alerting framework. For example, if IBM Tivoli OMEGAMON XE for z/OS detects a service level agreement (SLA) violation, it can use IBM zAware anomaly information to confirm that the SLA violation needs immediate attention. Through the IBM zAware API, system management products can request and receive IBM zAware analytical data in XML format. This data is equivalent to the information that is available through the Analysis views and Interval page in the IBM zAware GUI.

For additional information about using IBM zAware with various system management products, see the following topics:

- Starting with Tivoli OMEGAMON on z/OS V5.1.1, IBM zAware data is consolidated with performance and other information to support diagnoses of problems and to include in OMEGAMON XE on z/OS situations. OMEGAMON XE on z/OS provides a workspace through which users can display, manage, and customize IBM zAware data.

  Use the following IBM Knowledge Center link to find instructions for configuring OMEGAMON XE for z/OS to connect to an IBM zAware server, as well as additional topics related to OMEGAMON XE and IBM zAware. If you are using a different version of OMEGAMON XE for z/OS, use IBM Knowledge Center search and product filters to find the appropriate set of instructions for the product version. https://www.ibm.com/support/knowledgecenter/en/SS2JNN_5.5.0/com.ibm.omegamon_xezos.doc_5.5.0/configuration/complete_zaware.htm
- To use IBM zAware data with other system management products, use the instructions in "Integrating IBM zAware data into monitoring and alerting products."

  For an example of a system management program that uses the IBM zAware API, see the IBM Tivoli NetView for z/OS topic in the IBM Redbooks publication *Extending z/OS System Management Functions with IBM zAware*, SG24-8070. This Redbooks publication is available at the following URL:

  `http://www.redbooks.ibm.com/`
- To use IBM zAware through the z/OS Management Facility (z/OSMF), see "Viewing the IBM zAware GUI through z/OS Management Facility" on page 273.

Integrating IBM zAware data into monitoring and alerting products

Your installation can modify system management products to request and receive IBM zAware analytical data in XML format by using the IBM zAware application programming interface (API). Use this procedure as an overview for programming a system management product on z/OS to issue the API and process the returned XML data.

Before you begin

- Before your program can establish a connection with the IBM zAware server, you must enable the z/OS system on which the program runs for Application Transparent Transport Layer Security (AT-TLS). For more information, see the AT-TLS topic in IBM Knowledge Center: https://www.ibm.com/support/knowledgecenter/SSLTBW_2.1.0/com.ibm.zos.v2r1.halz002/attls.htm?lang=en
- To establish a connection, you need to know the IP address of the IBM zAware server. You can code your program to use the IXGQUERY service to retrieve the IP address. For more information, see the IXGQUERY topic in IBM Knowledge Center: https://www.ibm.com/support/knowledgecenter/SSLTBW_2.1.0/com.ibm.zos.v2r1.ieaa600/que.htm?lang=en

  You also can code your program to use the ENFREQ service to listen for ENF event code 48, which is issued for z/OS system logger configuration changes, including changes to the IBM zAware server IP

address. For more information, see the ENFREQ topic in IBM Knowledge Center: https://www.ibm.com/support/knowledgecenter/SSLTBW_2.1.0/com.ibm.zos.v2r1.ieaa200/enf.htm?lang=en

- The z/OS user ID under which the program runs must be added to the IBM zAware user authentication mechanism: an existing Lightweight Directory Access Protocol (LDAP) directory or a local file-based repository. This user ID must be assigned to either the User or the Administrator role through the **Administration** > **Configuration** > **Role Mapping** page in the IBM zAware GUI.

  Because authentication relies on the use of cookies, your program must be configured to accept, save, and send cookies.

Procedure

1. Connect and authenticate to the IBM zAware server. Issue a POST request to the following URL. In the code example, the variable *server_IP_address* is the IP address of the IBM zAware server.

   ```
   http://server_IP_address/zAware/j_security_check
   ```

   With the POST request, supply the user ID and password for the z/OS user ID under which the program runs.

   ```
   j_username=username
   j_password=password
   ```

2. Issue an HTTP GET request to retrieve analytical data from IBM zAware for a specific monitored client. Version 1 GET requests are supported for downward compatibility. Whenever possible, use Version 2 syntax and use the **version** parameter to specify which version of XML that you want the IBM zAware server to return. The syntax for a Version 2 request is:

   ```
   GET https://server_ip_address/zAware/authuser/Analysis?reqtype=request_type&time=time
   &sysname=system_name&clienttype=client_type&version=version
   ```

   The request type indicates the type of analytical data that you are requesting from the IBM zAware server. Valid types are:

   analysis (or LPAR)
   Requests analytical data for one day for one monitored client. (LPAR is the Version 1 parameter.)

   INTERVAL
   Requests analytical data for a specific analysis interval for one monitored client.

   For the GET request parameter descriptions and examples, see "Syntax and description of a GET request for IBM zAware data" on page 294.

3. Receive and process the XML data that the IBM zAware server returns in response to the GET request. On successful completion, the response is an XML document that matches the request type: analysis or INTERVAL.

   For a Version 1 request

   - "XML for a Version 1 LPAR or ANALYSIS request" on page 298 describes the returned XML for an `LPAR` or `ANALYSIS` request.
   - "XML for a Version 1 INTERVAL request" on page 302 describes the returned XML for an `INTERVAL` request.

   For a Version 2 request

   - "XML for a Version 2 ANALYSIS request" on page 307 describes the returned XML for an **ANALYSIS** request.
   - "XML for a Version 2 INTERVAL request" on page 313 describes the returned XML for an `INTERVAL` request.

   For a Version 1 request, IBM zAware returns Version 1 XML for a successfully processed GET request. For a successfully processed Version 2 request, however, the returned XML might be Version 1 format if Version 2 data is not available. To determine the version of the returned XML, see "Content-Type header values for HTTP responses" on page 297.

4. Periodically repeat steps 2 on page 272 and 3 on page 272 to provide the system management functions that you consider necessary for each monitored system.

   Consider providing notifications for the following analytical data:
   - An interval anomaly score of 101.
   - Multiple intervals with an interval anomaly score of 101.
   - A significant change in interval anomaly score from one interval to the next.
   - Intervals that contain no analytical data.
   - Intervals that contain new messages.

Viewing the IBM zAware GUI through z/OS Management Facility

Use this procedure as an overview for configuring z/OS Management Facility (z/OSMF) so that users can access the IBM zAware graphical user interface (GUI) through the z/OSMF navigation area. The procedure for configuring z/OSMF to link to the IBM zAware GUI is the same procedure as for any other external link that you define in the z/OSMF navigation area. For details about defining external links, use the z/OSMF online help for the **Links** page.

Before you begin

- You need to know the URL for the IBM zAware GUI.

  The URL includes the IP address or host name that is assigned to the IBM zAware partition:

  `https://`*`ip_address`*`/zAware/` **or** `https://`*`host_name`*`/zAware/`

  The "zAware" portion of the URL is case-sensitive.
- To define links for z/OSMF, you must be authorized to do so. By default, only the z/OSMF Administrator can define links.

  How users are authorized to links, and whether the authorization is performed in your security management product or through the Links task, depends on the authorization mode that is in effect for your installation. For more information about authorization modes, see the topic about setting up security in *IBM z/OS Management Facility Configuration Guide*, SA38-0652.

  To access this security topic in the product documentation for z/OSMF, go to IBM Knowledge Center at the following URL, and select the set of topics for the version of z/OS and the version of z/OSMF that you are using.

  http://www.ibm.com/support/knowledgecenter/.

Procedure

To display the **Links** page, expand the z/OSMF **Administration** category in the navigation area and select **Links** to begin a sequence of steps for defining links for z/OSMF.
The following list provides a summary of the steps that are described in detail in the z/OSMF online help for the **Links** page.

- Specify the link name and its location (a URL).
- Provide a system authorization facility (SAF) resource name to be used for managing user authorizations to the link.
- Select a z/OSMF category for the link. Suggested categories are "Links" or "Problem Determination".
- Specify how the link opens in the user's browser session (the launch behavior). The recommended behavior is to open the link in a new browser window.
- Manage access to the link for z/OSMF users.

Results

Users can launch the IBM zAware GUI through the z/OSMF **Links** page.

Chapter 28. Troubleshooting problems in the IBM zAware environment

The topics in this information describes potential problems and provides suggested corrective actions.

If you experience any problems that are related to the IBM zAware environment, make sure that you check the following sources of diagnostic information:

- Through the **Notifications** page in the GUI, you can view messages that IBM zAware issues to notify you of some activity or condition that requires your awareness or response. When you have unread notification messages, the New label (New) is displayed in the navigation pane, to the right of the **Notifications** link.

  On the Notifications page, each notification is displayed as a row in the Notification Messages table. The messages that are listed can be related to an action you performed, to an action that another user performed, or to independent server processing (such as automatically scheduled retraining). The list is shared across users, and is intended to inform you of activity in the IBM zAware environment.
- For problems with the IBM zAware partition, you can check the Service Container UI. Click the **Help icon** > **Container Settings**, and then look at the Service Container UI for log messages.

If you are unable to correct the problem, use the instructions in Chapter 29, "Reporting IBM z Advanced Workload Analysis Reporter (IBM zAware) problems to IBM," on page 281 to request support from IBM.

Troubleshooting problems with the IBM zAware partition

Table 62 lists problems and fixes that you might encounter with IBM zAware partition.

Table 62. Troubleshooting tips for the IBM zAware partition

Problem	Explanation or fix
From a correctly configured z/OS system, I attempted to start sending OPERLOG data to the IBM zAware server, but the communication failed.	The communication between a monitored client and IBM zAware might fail for one of the following reasons: • The IBM zAware partition might have been deactivated or is in the process of being activated. • The IBM zAware partition has been activated but the IBM zAware software is still initializing. • The IBM zAware partition has been activated, IBM zAware is fully initialized and accessible through its GUI, but storage devices either have not yet been assigned or are still being formatted for IBM zAware use. To check the status of the partition, use the HMC or SE for the IBM zAware host system. To check the status of storage devices assigned to IBM zAware, go to the **Data Storage** tab in the GUI.

Troubleshooting problems with the IBM zAware server and GUI

Table 63 on page 276 lists some problems and fixes that you might encounter with the IBM zAware server or the GUI.

Table 63. Troubleshooting tips for browser or GUI page displays

Problem	Explanation or fix
I cannot successfully log in with the master user ID or password for the GUI.	The answer depends on the server you are using. • On IBM z13 and above, the master user ID and password can be changed or reset through container partition that is set up in the IBM zAware LPAR profile. • On IBM zEnterprise Systems, use the bootstrap configuration utility to update the settings. For more information, see "Modifying the Bootstrap Configuration for IBM zAware" on page 33.
I have been using the IBM zAware GUI for a while but now some pages are not displaying properly.	To ensure that the GUI displays the current page and content: • Make sure that the browser you are using meets requirements by using the environment checker: click the down arrow next to the Help icon () on the IBM zAware header, and select **Environment Checker**. The Environment Checker window opens in a new browser tab, and presents the evaluation results. If the current setting for a particular option does not meet requirements, the display includes a warning icon for that setting. • Clear your browser cache after you apply service for IBM zAware, and periodically during normal use.
When I click on a link on a GUI page, such as the icon for help, nothing seems to happen.	Make sure that the browser you are using is correctly configured to enable pop-up windows to open.

I connected a new monitored client and the IXG messages that are issued on the z/OS system indicate that the connection was successful. However, the system does not appear on the **Systems** > **System Status** tab, or the value that is shown in the Instrumentation Data Type column is not correct.	1. On the newly connected system, enter the **SETLOGR** command to stop data transmission:
	`SETLOGR FORCE,ZAIQUIESCE,ALL`
	2. Wait a few minutes, then enter the **SETLOGR** command to reconnect the client:
	`SETLOGR FORCE,ZAICONNECT,LSNAME=SYSPLEX.OPERLOG`

I expected to see analysis data for a particular monitored client on a specific date, but the **Analysis** page display does not contain any data for that client and date.	Analytical data might not be available for all systems for the date and time that you select for the Analysis page display. Data is not available under the following circumstances: • The monitored system was added to the topology after the date you select for the Analysis page display. • The monitored system is not connected to the IBM zAware server. • The monitored system and the applications that run on it did not issue any messages. The Analysis display indicates the "not connected" or "no data" conditions.
I have been using the IBM zAware GUI but it now appears to be hanging.	IBM zAware might have lost access to one or more of its in-use storage devices. When an in-use storage device becomes unavailable, IBM zAware operations stop. To a user of the IBM zAware GUI, the GUI appears to hang. On the SE for the IBM zAware host system, hardware messages indicate input/output (I/O) problems that are related to the loss of access to physical storage devices. See the response to message "AIFP0013E" on page 337 for instructions to diagnose and correct this condition.

Troubleshooting problems with the z/OS bulk load client for IBM zAware

Table 64 on page 277 lists some problems and fixes that you might encounter with running the z/OS bulk load client to transfer priming data to the IBM zAware server.

Table 64. Troubleshooting tips for running the z/OS bulk load client for IBM zAware

Problem	Explanation or fix
The REXX call to run the z/OS bulk load client is failing with an invalid character on line number 1.	Check the TSO profile to make sure that the PACK option is set to OFF and resubmit the z/OS bulk load client job.

Troubleshooting problems with z/OS monitored clients

Table 65 lists some problems and fixes that you might encounter with configuring and managing z/OS monitored clients. For a more comprehensive list of possible errors and fixes, see the topic on resolving z/OS IBM zAware log stream client errors in *z/OS MVS Diagnosis: Reference*, GA22-7588.

Table 65. Troubleshooting tips for z/OS monitored clients

Problem	Explanation or fix
When configuring a z/OS system as an IBM zAware monitored client, I attempted to define or update an existing OPERLOG log stream with the ZAI and ZAIDATA parameters but the request failed with reason code "839"x.	The active primary LOGR couple data set is not formatted at the level that is required to process the request. You can specify the ZAI and ZAIDATA keywords for a log stream only when the LOGR CDS format level is at least HBB7705. To determine what format level is in use for a sysplex, enter the following command and check the resulting message display. `D XCF,COUPLE,TYPE=LOGR` If the LOGR CDS format level is not HBB7705, your installation needs to run the format CDS utility IXCL1DSU with the DATA TYPE(LOGR) and ITEM NAME(SMDUPLEX) NUMBER(1) options. For more information, see the topic about LOGR parameters for the format utility in *z/OS MVS Setting Up a Sysplex*.
From a correctly configured z/OS system, I attempted to start sending OPERLOG data to the IBM zAware server, but the communication failed.	The communication between a monitored client and IBM zAware might fail for one of the following reasons: • The IBM zAware partition might have been deactivated or is in the process of being activated. • The IBM zAware partition has been activated but the IBM zAware software is still initializing. • The IBM zAware partition has been activated, IBM zAware is fully initialized and accessible through its GUI, but storage devices either have not yet been assigned or are still being formatted for IBM zAware use. To check the status of the partition, use the HMC or SE for the IBM zAware host system. To check the status of storage devices assigned to IBM zAware, go to the **Data Storage** tab in the GUI.

Table 65. Troubleshooting tips for z/OS monitored clients (continued)

Problem	Explanation or fix
On a z/OS system that is established as an IBM zAware monitored client, I see repeated socket error messages. What happened to the connection between the z/OS system and the IBM zAware server?	The IBM zAware analytics engine might be stopped or recycling. When the analytics engine is stopped or recycled, IBM zAware disconnects all monitored systems from the server. When z/OS monitored clients are disconnected from the server, they automatically attempt to reconnect and continue to buffer data for approximately 10 minutes. During this time, the socket error messages are issued until the z/OS system successfully reconnects or stops retrying. If the analytics engine is restarted within that time, the z/OS system reconnects and sends the buffered data to the IBM zAware server. If the z/OS system times out before the analytics engine is restarted, any buffered data is lost and you must manually reconnect the system. To reconnect a z/OS system, issue the SETLOGR command. `SETLOGR FORCE,ZAICONNECT,LSN=SYSPLEX.OPERLOG` To check the status of the IBM zAware analytics engine, go to the **System Status** page in the GUI.

Troubleshooting problems with Linux on z Systems monitored clients

Table 66 lists some problems and fixes that you might encounter with configuring and managing Linux on z Systems monitored clients.

Table 66. Troubleshooting tips for Linux monitored clients

Problem	Explanation or fix

My installation uses Security-Enhanced Linux (SELinux), and another process is already using the port that is required for IBM zAware.	You might need to install the **semanage** command before you can complete this procedure to correct the problem.
	1. To determine which object type is assigned to port 2003, enter the following **semanage** command:
	`semanage port -l | grep 2003`
	
	The system response indicates the object type that is assigned to port 2003; for example:
	`lmtp_port_t tcp 24, 2003`
	2. To assign port 2003 to the syslog daemon, enter the following **semanage** command:
	`semanage port -m -t syslogd_port_t -p tcp 2003`
	3. To verify the change, enter the command `semanage port -l | grep 2003` again. The system response indicates which object types are assigned to port 2003; for example:
	`lmtp_port_t tcp 24, 2003`
	`syslogd_port_t tcp 2003,` `601`

Table 66. Troubleshooting tips for Linux monitored clients (continued)

Problem	Explanation or fix

My installation uses SELinux, and I cannot successfully configure the syslog daemon to connect to the IBM zAware.	The SELinux policy might be stopping a normally started syslog daemon from using TCP. Running the daemon in the foreground and in debug mode might bypass the restriction.
	1. Check the audit log to determine whether the syslog daemon is being denied. For example, if you are using rsyslog, search the audit log `/var/log/audit/audit.log` for the word "rsyslog"; if rsyslog is being denied, the audit log has an entry similar to the following:
	`type=SYSCALL msg=audit(1434389426.850:821):`
	`arch=80000016 syscall=102 per=400000`
	`success=no exit=-13`
	`a0=3 a1=3fffbba52c8 a2=0 a3=3fffbba6a18 items=0`
	`ppid=1 pid=14708 auid=0 uid=0 gid=0 euid=0`
	`suid=0 fsuid=0 egid=0 sgid=0 fsgid=0`
	`tty=(none) ses=90`
	`comm=72733A6D61696E20513A526567`
	`exe="/sbin/rsyslogd"`
	`subj=unconfined_u:system_r:syslogd_t:s0 key=(null)`
	2. If the syslog daemon is being denied, add an exception to the SELinux policy.
	`semanage port -a -t syslogd_port_t -p tcp 2003`
	3. Start the syslog daemon by using the appropriate command, For example, `service rsyslog start`

Chapter 29. Reporting IBM z Advanced Workload Analysis Reporter (IBM zAware) problems to IBM

Use the following topic to determine whether you need to learn how to report a problem and prepare for your call with IBM Support.

Before you begin

If you encounter a problem that appears to be directly related to a particular IBM z Advanced Workload Analysis Reporter (IBM zAware) component, you can open a Software problem management record (PMR). If you are not sure which component is the cause of the problem, open the software PMR. After the problematic component is identified, IBM Support transfers the problem record to the appropriate support center.

About this task

If an IBM Support representative requests dump data, use this procedure to generate the dump and send it to IBM.

Procedure

1. Log in to the IBM zAware, and then click **Help** > **Service Container**. You are redirected to the **Service Container** in IBM zAware. To access the **Help**, click the down arrow next to the Help icon ().
2. Select **Dumps** from the left navigation pane, and then click the type of dump you were asked to send. Depending on the type of problem, IBM Support might ask you to send one of the following dump types.
   - Concurrent
   - Disruptive
3. Enter any comments about the dump. If you opened a PMR, include the PMR number. Next, click **Create Dump**.

Results

After the dump data is generated, it is encrypted, and ready to send to IBM Support.

What to do next

Prepare for your call with IBM Support. If you have not sent a dump for the IBM zAware partition, IBM Support might request you to do so. Depending on the type of problem, IBM Support also might request screen captures of your IBM zAware configuration settings.

IBM Support also might ask you to take further action or answer one or more of the following questions:

- View the message logs that are associated with the IBM zAware.
- For the IBM zAware partition:
  - Are you experiencing a problem with a new IBM zAware partition that is not active?
  - What is the name of the IBM zAware partition?
  - Provide details about the network interface card (NIC) and virtual local area network (VLAN).
  - Are all IBM zAware stream microcode load levels (MCLs) installed and activated?
- For an IBM zAware monitored client:

  - Is the z/OS system that is running in a partition on the IBM zAware host system on the same CPC as the IBM zAware partition?
  - What release of the z/OS operating system is installed?
  - What brand, version, and release of Linux is installed?
- For the problem itself, you might also need the following information:
  - What are the problem symptoms? For example, LPAR activation or training requests fail.
  - Can you easily reproduce the problem or is it intermittent?
  - What is the exact date and time of the last failure?

Part 7. Appendixes

Appendix A. Summary of IBM zAware tasks and required IT skills, tools and information

- The IBM zEnterprise System (zEnterprise) product library, which includes z Systems hardware books, is available through **Resource Link** at http://www.ibm.com/servers/resourcelink.
- The z/OS product library is available in IBM Knowledge Center at http://www.ibm.com/support/knowledgecenter/. From the IBM Knowledge Center welcome page:
  1. Use the Table of Contents to go to **IBM Operating Systems**, expand the appropriate platform.
  2. Click **z/OS** and select the appropriate z/OS version and release that you are using.
  3. Then navigate to one of the following book collections for each z/OS element:
     - z/OS Communications Server
     - z/OS MVS
     - z/OS Security Server RACF
     - z/OS UNIX System Services

Table 67. Summary of IBM zAware tasks and required IT skills, tools and information

Task	IT role / skill	Tools / interfaces	Information resources
Plan to use IBM zAware	System planners and installation managers	—	• Chapter 3, "Project plan for configuring and using IBM zAware," on page 17 • Part 3, "Planning to configure IBM zAware," on page 37 • *IBM z14 Technical Guide*, SG24-8451 • *IBM z13 Technical Guide*, SG24-8251 • *IBM z13s Technical Guide*, SG24-8294
Configure IBM hardware, networking and storage devices that support the IBM zAware partition	• System planners and installation managers • Network administrators • Storage administrators	Hardware Management Console (HMC) user interface (UI)	• Part 3, "Planning to configure IBM zAware," on page 37 and Part 4, "Configuring IBM zAware and its monitored clients," on page 93 • *z Systems PR/SM Planning Guide*, SB10-7162
Configure the IBM zAware partition	• Systems programmers • Network administrators • Storage administrators	HMC UI	• Part 4, "Configuring IBM zAware and its monitored clients," on page 93 • HMC information can be found on the console help system, or on IBM Knowledge Center at http://www.ibm.com/support/knowledgecenter/

Table 67. Summary of IBM zAware tasks and required IT skills, tools and information (continued)

Task	IT role / skill	Tools / interfaces	Information resources
Prepare the IBM zAware server for operation	• Systems programmers • Storage administrators • Security administrators	IBM zAware graphical user interface (GUI) or operating system interfaces	• Part 4, "Configuring IBM zAware and its monitored clients," on page 93 • Books listed in Table 68
Configure operating systems to send data to the IBM zAware server	• Systems programmers • Network administrators	IBM zAware GUI or operating system interfaces	• Part 4, "Configuring IBM zAware and its monitored clients," on page 93 • Books listed in Table 68
Manage the use and operation of the IBM zAware server	• Systems programmers • Network administrators • Storage administrators • Security administrators	IBM zAware GUI or operating system interfaces	Part 5, "Managing and using the IBM zAware server," on page 137 and Part 6, "Advanced topics for managing IBM zAware," on page 219
View and interpret analytical data and resolve potential system problems	• Systems programmers • Experienced application programmers	IBM zAware GUI	• Part 5, "Managing and using the IBM zAware server," on page 137 and Part 6, "Advanced topics for managing IBM zAware," on page 219 • *z/OS MVS Diagnosis Reference*
Connect the IBM zAware GUI to other system management or monitoring products	• Systems programmers • Experienced application programmers	Various, depending on system management or monitoring product	Part 3, "Planning to configure IBM zAware," on page 37 and Part 5, "Managing and using the IBM zAware server," on page 137
Update vendor programs that duplicate IBM zAware functions	Experienced application programmers		Appendix C, "Application Programming Interface (API) for monitoring products," on page 293

Table 68. IBM zAware information in the z/OS product library

z/OS title and order number	IBM zAware-related content
z/OS MVS System Commands	• Table 7. MVS Commands, RACF Access Authorities, and Resource Names • DISPLAY LOGGER command • DISPLAY MSGFLD command • SET command (IXGCNF parameter) • SETLOGR command
z/OS Planning for Installation	The APAR requirement for system logger is listed in the entry for the BCP component in *Table 39. Hardware requirements for z/OS V1R13 elements and features.*

Table 68. IBM zAware information in the z/OS product library (continued)

z/OS title and order number	IBM zAware-related content
z/OS Setting up a Sysplex	• Planning the IXGCNF system parameter • Planning for system logger applications – Define authorization for the system logger address space – Updating a log stream's attributes – Preparing for z/OS IBM zAware log stream client usage • LOGR parameters for format utility • LOGR keywords and parameters for the administrative data utility
z/OS MVS Planning: Operations	Exploiting the IBM z Advanced Workload Analysis Reporter (IBM zAware) for OPERLOG
z/OS MVS Programming: Assembler Services Reference, Volume 2	• IXGINVNT — Managing the LOGR inventory couple data set • IXGQUERY — Query a log stream or system logger information
z/OS MVS Programming: Assembler Services Guide	Using system logger services: • IXGINVNT: Managing the LOGR policy • IXGQUERY: Get information about a log stream or system logger
z/OS MVS Programming: Authorized Assembler Services Guide	• Using system logger services • Setting up the system logger configuration • Writing an ENF event 48 listen exit
z/OS Diagnosis: Reference	System logger: • Resolving system logger z/OS IBM zAware log stream client errors
z/OS MVS Diagnosis: Tools and Service Aids	Updated topic: SYSLOGR component trace
z/OS System Messages, Vol 1	New AIZ message descriptions for messages that the z/OS bulk load client for IBM zAware issues
z/OS System Messages, Vol 10	New and updated IXG message descriptions for messages that the z/OS system logger issues
z/OS MVS System Management Facility	Updated topic: Record Type 88 (58) — System Logger Data
z/OS Initialization and Tuning Reference	• IEASYSxx (system parameter list): IXGCNF parameter • IXGCNFxx (system logger initialization parameters)

Appendix B. Sample certificate authority (CA) reply

Your installation has the option of replacing the IBM zAware default SSL certificate with a certificate signed by a certificate authority of your choice. This topic provides sample certificate blocks to illustrate the content that you might receive from a certificate authority.

When you receive a reply from a certificate authority, the reply might contain a chain of certificates, starting with the signed server certificate. Then, it is possibly followed by certificates from one or more intermediate CAs and finally, the self-signed certificate of the CA. Figure 62 on page 290 provides sample certificate blocks to illustrate the content that you might receive from a certificate authority. In this sample, the reply chain consists of the following:

1. The first block of certificate is the IBM zAware server certificate, as returned by the certificate authority. This certificate is signed by the next signer in the chain.

   ```
   subject=/C=US/ST=NY/L=Poughkeepsie/O=ibm.com/OU=SysTest/CN=198.xx.xx.xx/UID=xxx97/mail=emplye@us.ibm.com
   issuer=/C=US/O=International Business Machines Corporation/CN=IBM INTERNAL INTERMEDIATE CA
   ```

2. The next block of certificate text is the signer of the IBM zAware server certificate; in this sample, the signer is an intermediate certificate authority. This certificate is signed by the next signer in the chain. Note that the reply can contain one or more blocks for intermediate signers.

   ```
   subject=/C=US/O=International Business Machines Corporation/CN=IBM INTERNAL INTERMEDIATE CA
   issuer=/C=US/O=International Business Machines Corporation/CN=IBM Internal Root CA
   ```

3. The final block of certificate text is the self-signed certificate for the certificate authority itself. It is self signed. This certificate must be added to the browser's trust store to authenticate the IBM zAware server.

   ```
   subject=/C=US/O=International Business Machines Corporation/CN=IBM Internal Root CA
   issuer=/C=US/O=International Business Machines Corporation/CN=IBM Internal Root CA
   ```

```
subject=/C=US/ST=NY/L=Poughkeepsie/O=ibm.com/OU=SysTest/CN=198.xx.xx.xx/UID=xxx97/mail=emplye@us.ibm.com
issuer=/C=US/O=International Business Machines Corporation/CN=IBM INTERNAL INTERMEDIATE CA
-----BEGIN CERTIFICATE-----
MIIF0DCCBLigAwIBAgICD6kwDQYJKoZIhvcNAQEFBQAwajELMAkGA1UEBhMCVVMx
NDAyBgNVBAoTK0ludGVybmF0aW9uYWwgQnVzaW5lc3MgTWFjaGluZXMgQ29ycG9y
⋮
5i5BozwFbvxCDmC2INWzEaejdmejdCSDkgDGqgVtXXZnZeCtREOGME99nm3fHW7h
QkyXkg==
-----END CERTIFICATE-----

subject=/C=US/O=International Business Machines Corporation/CN=IBM INTERNAL INTERMEDIATE CA
issuer=/C=US/O=International Business Machines Corporation/CN=IBM Internal Root CA
-----BEGIN CERTIFICATE-----
MIID7TCCAtWgAwIBAgIBAjANBgkqhkiG9w0BAQUFADBiMQswCQYDVQQGEwJVUzE0
MDIGA1UEChMrSW50ZXJuYXRpb25hbCBCdXNpbmVzcyBNYWNoaW5lcyBDb3Jwb3Jh
⋮
irUCKeSX1o3HGZFhMYwlKsYwog470qbYqDIqP+JM2N161GaNHilDcW49qKvQTkV5
fg==
-----END CERTIFICATE-----

subject=/C=US/O=International Business Machines Corporation/CN=IBM Internal Root CA
issuer=/C=US/O=International Business Machines Corporation/CN=IBM Internal Root CA
-----BEGIN CERTIFICATE-----
MIIDxDCCAqygAwIBAgIBADANBgkqhkiG9w0BAQUFADBiMQswCQYDVQQGEwJVUzE0
MDIGA1UEChMrSW50ZXJuYXRpb25hbCBCdXNpbmVzcyBNYWNoaW5lcyBDb3Jwb3Jh
⋮
bwnogYppATaH1z2PpMC3nqhyMv6B+NfAen1iMVbAFERrDRUuPD+RtO9s8ayEwVqp
3+HY0FBqhlI=
-----END CERTIFICATE-----
```

Figure 62. Sample reply from a third-party certificate authority

When you supply this information in the GUI, provide the entire certificate chain, starting with the signed server certificate through the self-signed certificate of the CA. Make sure that you do not insert any lines or spaces between the end of one certificate and the beginning of the next certificate. Figure 63 on page 291 illustrates the correct format for pasting certificate content. In the figure, the ellipses represent certificate content that has been removed only for publication in this book. When you paste certificate replies in the GUI, make sure that you include all of the content, including the header `-----BEGIN CERTIFICATE-----` through and including `-----END CERTIFICATE-----`

```
-----BEGIN CERTIFICATE-----
MIIF0DCCBLigAwIBAgICD6kwDQYJKoZIhvcNAQEFBQAwajELMAkGA1UEBhMCVVMx
NDAyBgNVBAoTK0ludGVybmF0aW9uYWwgQnVzaW5lc3MgTWFjaGluZXMgQ29ycG9y
⋮
5i5BozwFbvxCDmC2INWzEaejdmejdCSDkgDGqgVtXXZnZeCtREOGME99nm3fHW7h
QkyXkg==
-----END CERTIFICATE-----
-----BEGIN CERTIFICATE-----
MIID7TCCAtWgAwIBAgIBAjANBgkqhkiG9w0BAQUFADBiMQswCQYDVQQGEwJVUzE0
MDIGA1UEChMrSW50ZXJuYXRpb25hbCBCdXNpbmVzcyBNYWNoaW5lcyBDb3Jwb3Jh
⋮
irUCKeSX1o3HGZFhMYwlKsYwog470qbYqDIqP+JM2N161GaNHilDcW49qKvQTkV5
fg==
-----END CERTIFICATE-----
-----BEGIN CERTIFICATE-----
MIIDxDCCAqygAwIBAgIBADANBgkqhkiG9w0BAQUFADBiMQswCQYDVQQGEwJVUzE0
MDIGA1UEChMrSW50ZXJuYXRpb25hbCBCdXNpbmVzcyBNYWNoaW5lcyBDb3Jwb3Jh
⋮
bwnogYppATaH1z2PpMC3nqhyMv6B+NfAen1iMVbAFERrDRUuPD+RtO9s8ayEwVqp
3+HY0FBqhlI=
-----END CERTIFICATE-----
```

Figure 63. Illustration of required format for pasting into the GUI

Appendix C. Application Programming Interface (API) for monitoring products

IBM zAware provides an application programming interface (API) that system management products can use to request analytical data to display through their own graphic user interfaces. Through this API, system management products, such as IBM Tivoli OMEGAMON, can request and receive IBM zAware analytical data in XML format.

This data is equivalent to the information that is available through the **Analysis** page and **Interval view** in the IBM zAware GUI.

API versioning

Because the IBM zAware API might be modified to match changes to the functions provided through the IBM zAware GUI, each functional level of the API is identified by a version number. Each version number corresponds to a specific engineering change (EC) or microcode control level (MCL) for IBM zAware V1 and V2 or a software version number for IBM zAware V3.1 and above.

To inter-operate with multiple versions of the IBM zAware API, system management products that use the API must be designed to ignore, without error, the following possible modifications in XML responses:

- Any field that is not recognized by the application.
- Any header or body field that is not recognized by the application.

Table 69 lists each API version number, its corresponding MCL number, and a summary of modifications to the IBM zAware API for each version.

Table 69. Summary of API version updates for SE-ZAWARE MCLs

API version	Machine type	SE-ZAWARE MCL	Description
—	2827 or 2828	SE-ZAWARE MCL H09126.006	Initial version of the IBM zAware API, delivered with the IBM zEnterprise EC12 (zEC12) or IBM zEnterprise BC12 (zBC12).
1	2827 or 2828	SE-ZAWARE MCL H09126.021	Updated version of the IBM zAware API, available with the zEC12 or zBC12. • New XML fields returned for an `LPAR` request type. – **version** – **gmt_offset** • New XML field returned for an `INTERVAL` request type: **version** • Removed the **intervalid** parameter from the GET syntax for an INTERVAL request. This parameter is obsolete.
2	2964	SE-ZAWARE MCL N98812.012	New version of the IBM zAware API, available with the IBM z13 (z13).
2	2964	SE-ZAWARE MCL N98812.022	Updated version that contains a new XML field returned for an `ANALYSIS` request type: **limited_model**

Table 69. Summary of API version updates for SE-ZAWARE MCLs (continued)

API version	Machine type	SE-ZAWARE MCL	Description
2	2964 or 2965	SE-ZAWAREMCL P08444.002 or later	Version that is available with the following host systems: • z13 • z13s
2	—	—	IBM zAware V3.1 software appliance

Syntax and description of a GET request for IBM zAware data

Use an HTTP GET request to retrieve analytical data from IBM zAware for a specific monitored client. Depending on the request type, you can request the interval anomaly scores for one day or details for a specific 10-minute interval.

HTTP method and URI

In the following request syntax examples:

- The URI variable *server_ip_address* is the IP address of the IBM zAware server.
- The remaining variables are parameters that are described in "Request contents."

For a Version 1 GET request

GET https://*server_ip_address***/zAware/authuser/Analysis?reqtype=***request_type***&time=***time*
&plexname=*sysplex_name***&lparname=***system_name*

For a Version 2 GET request

GET https://*server_ip_address***/zAware/authuser/Analysis?reqtype=***request_type***&time=***time*
&sysname=*system_name***&clienttype=***client_type***&version=***version*

Examples

For a Version 1 GET request

The following example shows the syntax of a GET request for the analytical data available on 14 August 2016 for the z/OS monitored client named "z4", which is a member of the sysplex named ZPLEX2.

```
GET https://198.51.100.00/zAware/authuser/Analysis?reqtype=LPAR
&time=20160814&lparname=z4&plexname=ZPLEX2
```

For a Version 2 GET request

The following example shows the syntax of a GET request for the analytical data available on 6 February 2016 for the z/OS monitored client named SY04, which is a member of the sysplex named PLEX2.

```
GET https://198.51.100.24/zAware/authuser/Analysis?reqtype=ANALYSIS
&time=20160206&sysname=PLEX2.SY04&clienttype=zOS&version=V2
```

Request contents

Version 1 GET requests are supported for downward compatibility. Whenever possible, use Version 2 syntax and use the **version** parameter to specify which version of XML that you want the IBM zAware server to return.

For a Version 1 GET request

All parameters are required, and can be specified in any order.

reqtype

Indicates what type of analytic data that you are requesting from the IBM zAware server.

LPAR
Requests analytical data for one day for one monitored client. The returned analytical data is equivalent to the information in the **Analysis** page display in the IBM zAware graphical user interface (GUI). The returned data provides interval anomaly scores for each interval since UTC midnight on the date indicated by the *time* variable. If *time* specifies the current date, the IBM zAware returns interval anomaly scores for every interval that has occurred. If *time* specifies a prior date, the IBM zAware returns interval anomaly scores for 144 intervals.

INTERVAL
Requests analytical data for a specific 10-minute time interval for one monitored client in a specific sysplex. The returned analytical data is equivalent to the information in the **Interval view** display in the IBM zAware GUI. The returned data provides details about each unique message ID that was issued during the 10-minute interval indicated by the *time* variable.

time
Indicates the date and time period for which analytical data is requested, specified in the following format:

*YYYYMMDDhhmm*00

The *hhmm*00 portion is required for an INTERVAL request, and optional for an LPAR request. If you specify the time period, use 24-hour clock time.

plexname
Specifies the name of the sysplex to which the monitored client belongs.

lparname
Specifies the name of the monitored client for which analytical data is requested.

For a Version 2 GET request
Parameters can be specified in any order.

reqtype
Indicates what type of analytic data that you are requesting from the IBM zAware server. The **reqtype** parameter values can be specified in either upper or lower case.

ANALYSIS
Requests analytical data for one day for one monitored client. The returned analytical data is equivalent to the information in the **Analysis** page display in the IBM zAware graphical user interface (GUI). The returned data provides interval anomaly scores for each analysis snapshot since UTC midnight on the date indicated by the *time* variable. If *time* specifies the current date, the IBM zAware returns interval anomaly scores for every analysis snapshot that has been taken. If *time* specifies a prior date, the IBM zAware returns interval anomaly scores for 144 intervals.

INTERVAL
Requests analytical data for a specific analysis interval for one monitored client. The returned analytical data is equivalent to the information in the **Interval** page display in the IBM zAware GUI. The returned data provides details about each unique message ID that was issued during the analysis interval indicated by the *time* variable.

time
Indicates the UTC date and time for which analytical data is requested, specified in the following format:

*YYYYMMDDhhmm*00

The *hhmm*00 portion is required for an INTERVAL request, and optional for an ANALYSIS request. If you specify the time period, use 24-hour clock time.

sysname
Indicates is the name of the monitored client for which analytical data is requested. For z/OS systems it is in the format *sysplex.system_name*. For Linux systems, the name is the system hostname.

clienttype
Indicates the type of system for which data is being requested. Acceptable values are `zOS` or `linux`.

version
Indicates the highest version of XML for IBM zAware to return for a successfully processed GET request. Acceptable values are `v1` or `v2`. IBM zAware returns Version 1 XML under the following conditions:
- The value coded for the **version** parameter is `v1`.
- The **version** parameter is omitted from the GET request.
- Version 2 XML is not available.

For information about determining the version of returned XML, see "Content-Type header values for HTTP responses" on page 297.

Response contents

On successful completion, the response is an XML document that matches the request type. To determine the version of the returned XML, see "Content-Type header values for HTTP responses" on page 297.

For a Version 1 request
- "XML for a Version 1 LPAR or ANALYSIS request" on page 298 describes the returned XML for an `LPAR` or `ANALYSIS` request.
- "XML for a Version 1 INTERVAL request" on page 302 describes the returned XML for an `INTERVAL` request.

For a Version 2 request
- "XML for a Version 2 ANALYSIS request" on page 307 describes the returned XML for an **ANALYSIS** request.
- "XML for a Version 2 INTERVAL request" on page 313 describes the returned XML for an `INTERVAL` request.

Authorization requirements

Before you can send a GET request to retrieve analysis results, you must authenticate to the IBM zAware server with a user ID that is defined in the repository that your installation configured for user authentication to the IBM zAware server. The user ID must be assigned to either the User or the Administrator role through the **Administration** > **Configuration** > **Role Mapping** page in the IBM zAware GUI.

To authenticate to the IBM zAware server, issue a POST request to the server IP address:

POST https://*server_ip_address***/zAware/j_security_check**

Through the following parameters, supply your user ID and password with your POST request:

j_username=*username*
j_password=*password*

Because authentication relies on the use of cookies, your user agent must be configured to accept, save, and send cookies.

HTTP status and reason codes

On successful completion, the HTTP status code 200 (OK) is returned and the response body is provided as described in "Response contents" on page 296. The following HTTP status codes can be returned for the indicated errors; the response body is a standard error response body providing an associated error message.

HTTP error status code	Description
200 (OK)	The IBM zAware server returned XML data for the request. For information about determining the version of returned XML, see "Content-Type header values for HTTP responses."
204 (No content)	The IBM zAware server did not return any XML data for the request. Check the parameter values that you supplied in the request. • For an ANALYSIS request, this error code is returned if interval data is not available for the date specified on the **time** parameter. • For an INTERVAL request, this error code is returned if interval data is not available for the date and time specified on the **time** parameter.
400® (Bad Request)	The IBM zAware server did not return any XML data for the request. Check the parameter types that you supplied in the request.
403 (Forbidden)	The IBM zAware server did not return any XML data for the request. Make sure that your program meets the authorization requirements in "Authorization requirements" on page 296.
404 (Not Found)	No XML data was returned because the request specified a *server_ip_address* that is not a valid address for an IBM zAware server.
503 (Server Environmental Error)	Storage devices are not configured for the IBM zAware server so no analytical data is available.

Content-Type header values for HTTP responses

For a Version 1 request, IBM zAware returns Version 1 XML for a successfully processed GET request. For a successfully processed Version 2 request, however, the returned XML might be Version 1 format even if you specified the **version** parameter with a value of `v2`. IBM zAware returns Version 1 XML under the following conditions:

- The value coded for the **version** parameter is `v1`.
- The **version** parameter is omitted from the GET request.
- Version 2 XML is not available.

Version 2 XML is not available for analysis results produced prior to IBM zAware Version 2.0. For example, if the request is for IBM zAware data that has been produced by an IBM zAware server on a zEnterprise host system, and that data has been migrated to a server on a z13, z13s, or z14 host system, IBM zAware can return data only in the Version 1 format.

To determine the version of the XML that IBM zAware returned for a GET request, you can query the HTTP Content-Type header. Table 70 lists the possible Content-Type header values and XML versions associated with those values.

Table 70. Possible Content-Type header values and XML versions for GET request types

Request type	Version returned	Content-Type header value
LPAR or ANALYSIS	Version 1 XML	`application/vnd.ibm.zaware-results-V1+xml`
	Version 2 XML	`application/vnd.ibm.zaware-results-V2+xml`
INTERVAL	Version 1 XML	`application/vnd.ibm.zaware-interval-V1+xml`
	Version 2 XML	`application/vnd.ibm.zaware-interval-V2+xml`

Requesting an XML response document through a supported browser

Use this procedure to request IBM zAware data in an XML response document that you can view or save. Saving the XML document as a file is useful if you need to send XML to IBM for diagnostic purposes.

Before you begin

- To log in to the IBM zAware GUI, you need to know the URL.

  The URL includes the IP address or host name that is assigned to the IBM zAware partition:

  `https://`*`ip_address`*`/zAware/` **or** `https://`*`host_name`*`/zAware/`

  The "zAware" portion of the URL is case-sensitive.
- To submit the request, you need to know the API syntax for a GET request. Use the same syntax as noted in "HTTP method and URI" on page 294, but remove the word GET from the syntax and replace variables as appropriate.

Procedure

1. Log in to the GUI with a user ID that is mapped to either an IBM zAware Administrator or User role. If you successfully log on, the GUI displays the default Analysis view page, or the last page you viewed during your previous session under this user ID.
2. In the browser address bar, enter a modified GET request, remembering to remove the GET parameter and to provide the appropriate values for the request parameters. The syntax for a Version 2 request is:

   https://*server_ip_address***/zAware/authuser/Analysis?reqtype=***request_type***&time=***time***&sysname=***system_name***&clienttype=***client_type***&version=***version*

   The following example shows the syntax of a GET request for the analytical data available on 1 June 2015 for the z/OS monitored client named D0, which is a member of the sysplex named SVPLEX3.

   ```
   https://192.12.18.65/zAware/authuser/Analysis?reqtype=analysis&time=20150601022000&sysname=SVPLEX3.D0
   &clienttype=zOS&version=V2
   ```

   The GUI page display changes to a blank screen, and a separate browser window opens to display your options for the returned XML file:
   - Open the file with a text editor, or
   - Download and save the file to the location of your choice.
3. If you want to request additional intervals, edit the time parameter as necessary, and view or save each file separately. Otherwise, to return to the original GUI page display, use the browser back button to reload the previous page.

Version 1 API

Through Version 1 of the IBM zAware API, system management products can request and receive IBM zAware analytical data that is equivalent to the information available through the Analysis page and Interval page in the IBM zAware GUI.

You can use the Version 1 API to request and receive data from an IBM zAware server that is running on any of the supported host systems, which are listed in Chapter 2, "Prerequisites for configuring and using IBM zAware," on page 13. If you use Version 1 on any supported host system other than an IBM zEnterprise EC12 (zEC12) or IBM zEnterprise BC12 (zBC12), however, the XML does not contain all of the analytical data that is available with the Version 2 API.

XML for a Version 1 LPAR or ANALYSIS request

This topic provides the XML structure, XML element descriptions, and a sample XML response that the IBM zAware server returns in response to one of the following HTTP GET method request types: an **LPAR** request type, or an **ANALYSIS** request type for which Version 1 XML is explicitly or implicitly

specified. This XML response contains information that is equivalent to the interval anomaly scores that the server displays through the **Analysis** page in the IBM zAware graphical user interface (GUI).

The following code illustrates the XML structure of the response to an HTTP GET method with an **LPAR** or an **ANALYSIS** request type. The major element is the **systems** element, which identifies the specific date and monitored client (system) for which analytical data was requested. The **systems** element also identifies the number and size of intervals returned in the XML document. The XML also contains one **interval** element for each analysis snapshot since UTC midnight on the requested date. The **interval** element provides the interval anomaly score and number of unique message IDs that were issued during the specific analysis snapshot.

"XML element descriptions for a Version 1 LPAR or ANALYSIS request" provides additional information about each element in the XML response.

```
<xs:schema xmlns:xs="http://www.w3.org/2001/XMLSchema"
 targetNamespace="http://www.example.org/MelodyCorePlex" xmlns="http://www.example.org/MelodyCorePlex"
 elementFormDefault="qualified">

 <xs:element name="systems" >
 <xs:complexType>
  <xs:sequence>
   <xs:element name="version" type="xs:int" />
   <xs:element name="start_time" type="xs:dateTime" />
   <xs:element name="end_time" type="xs:dateTime" />
   <xs:element name="gmt_offset" type="xs:string" />
   <xs:element name="number_intervals" type="xs:int" />
   <xs:element name="interval_size" type="xs:int" />
   <xs:element name="system" type="systems_system_type"
    maxOccurs="unbounded" />
  </xs:sequence>
 </xs:complexType>
 </xs:element>

 <xs:complexType name="systems_system_type">
  <xs:sequence>
   <xs:element name="interval" type="systems_interval_type"
    minOccurs="0" maxOccurs="unbounded" />
  </xs:sequence>
  <xs:attribute name="sys_id" type="xs:string" use="required" />
 </xs:complexType>

 <xs:complexType name="systems_interval_type">
  <xs:sequence>
   <xs:element name="num_unique_msg_ids" type="xs:int" />
   <xs:element name="anomaly_score" type="xs:double" />
  </xs:sequence>
 </xs:complexType>

</xs:schema>
```

XML element descriptions for a Version 1 LPAR or ANALYSIS request

The following list describes the major elements in the **systems** element.

version
An integer that identifies the version of the IBM zAware application programming interface (API). For information about specific API versions, see "API versioning" on page 293.

start_time
Indicates the beginning of the first interval for which data is available for the specified system on the date in the LPAR request. The start time is indicated in the XML `dateTime` data type format in Coordinated Universal Time (UTC).

*YYYY-MM-DD***T***hh:mm:ss.ttt***Z**

end_time
Indicates the beginning of the first interval *after* the date specified in the LPAR request. The end time is indicated in the XML `dateTime` data type format in Coordinated Universal Time (UTC).

*YYYY-MM-DD***T***hh:mm:ss.ttt***Z**

gmt_offset
An integer that indicates the difference in hours and minutes from Coordinated Universal Time (UTC) for the requested start time.

number_intervals
An integer that indicates the number of intervals for which analytical data is available for the system and the date specified on the LPAR request. Analytical data might not be available for this system for all intervals during the date specified on the request. Data is not available under the following circumstances:

- The monitored client was not connected and sending current data to the IBM zAware server on the specified date.
- The monitored client was added to the IBM zAware topology after the specified date.

Note that analytical data is not available for the dates for which you supplied priming data, unless the z/OS monitored client was connected and sending data to the IBM zAware server on those dates. The server uses priming data only for creating the model of system behavior.

interval_size
An integer that indicates the number of seconds in an interval.

system
An element that provides additional details about intervals for the system specified on the LPAR request. The `sys_id` attribute for this element provides the name of the system that was specified on the LPAR request, and the name of the system group to which the system belongs.

interval
An element that provides additional details about a specific interval. For the system and the date specified on the LPAR request, the XML response contains one interval element for each element for which analytical data is available.

num_unique_msg_ids
An integer that provides the number of unique message IDs that were issued during this analysis interval. If the same message ID was issued more than once during the interval, the message ID is counted only once.

anomaly_score
A double value that provides the anomaly score for this interval. The interval anomaly score is the percentile of the sum of each anomaly score for individual message IDs within an interval. When the IBM zAware server uses priming data and current data to create a model of system behavior, a process that is called "training", the server captures the distribution of interval anomaly scores for all intervals that are represented in the training data. The server uses the distribution results and uses them to establish the range of values for each percentile.

The possible interval anomaly scores are:

0 through 99.4
The analysis interval contains messages and message clusters that match or exhibit relatively insignificant differences in expected behavior, as defined in the IBM zAware model. A score of 0 is possible because the server eliminates all expected, in-context messages from its scoring calculation. A score of 0 indicates intervals that exhibit no difference in behavior when compared to the system or group model. The analysis snapshots for these analysis intervals are colored with the lightest blue shade.

Analysis intervals with scores that are greater than 0 but less than 99.5 contain some messages that are unexpected or issued out of context. Scores in this range indicate

intervals that do not vary significantly from the system model. The analysis snapshots for these analysis intervals are colored with varying shades of blue.

99.5 Analysis intervals with this score contain rarely seen, unexpected, or out-of-context messages. Generally speaking, this score indicates analysis intervals with some differences from the system or group model but do not contain messages of much diagnostic value. The analysis snapshots for these analysis intervals are colored with the darkest blue shade.

99.6 - 100
Analysis intervals with this score contain rarely seen messages (these messages appear in the model only once or twice), or many messages that are unexpected or issued out of context. This score indicates analysis intervals with more differences from the system or group model; these intervals can contain messages that might help you diagnose anomalous system behavior. The analysis snapshots for these analysis intervals are the color gold.

101 Analysis intervals with this score exhibit the most significant differences from the system or group model; these intervals contain messages that merit investigation. The analysis snapshots for these analysis intervals are the color orange. IBM zAware assigns this score to analysis intervals that contain:
- Unusual or unexpected messages.
- Messages that IBM rules define as critical.
- A much higher volume of messages than expected.

Sample XML response for a Version 1 LPAR or ANALYSIS request

This sample provides the returned XML in response to a GET request for the anomaly scores for a z/OS system named "C05" in sysplex "ABCPLEX" on 12 September 2012. The sample response shown is formatted for publication and contains information for only a few of the 144 returned intervals.

GET method:

```
GET https://xx.xx.xx.xxx/zAware/authuser/Analysis?reqtype=analysis
&time=20120912120000&sysname=ABCPLEX.C05&clienttype=zOS
```

XML response:

```
<?xml version="1.0" encoding="UTF-8"?>
<?xml-stylesheet href='./xslt/MelodyCorePlex.xsl' type='text/xsl'?>
<systems xmlns="http://www.example.org/MelodyCorePlex"
xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
xsi:noNamespaceSchemaLocation="xslt/MelodyCorePlex.xsd">
  <version>1</version>
  <start_time>2012-09-12T00:00:00.000Z</start_time>
  <end_time>2012-09-13T00:00:00.000Z</end_time>
  <gmt_offset>GMT-05:00</gmt_offset>
  <number_intervals>144</number_intervals>
  <interval_size>600</interval_size>
  <system sys_id="ABCPLEX-C05">
   <interval>
      <num_unique_msg_ids>71</num_unique_msg_ids>
     <anomaly_score>82.0</anomaly_score>
   </interval><
   <interval>
     <num_unique_msg_ids>51</num_unique_msg_ids>
     <anomaly_score>52.0</anomaly_score>
   </interval>

⋮

   <interval>
     <num_unique_msg_ids>99</num_unique_msg_ids>
     <anomaly_score>98.7</anomaly_score>
```

```
  </interval>
  <interval>
    <num_unique_msg_ids>0</num_unique_msg_ids>
    <anomaly_score>0.0</anomaly_score>
  </interval>
 </system>
 </systems>
```

XML for a Version 1 INTERVAL request

This topic provides the XML structure, XML element descriptions, and a sample XML response that the IBM zAware server returns in response to an HTTP GET method with an `INTERVAL` request type. This XML response contains information that is equivalent to the interval and message details that the server displays through the **Interval view** in the IBM zAware graphical user interface (GUI).

The following code illustrates the XML structure of the response to an HTTP GET method with an `INTERVAL` request type. The major element is the **interval** element, which contains information about a specific analysis interval for a specific system that is established as an IBM zAware monitored client. The **interval** element also contains one **interval_message** element for each unique message issued during the interval. If the same message ID was issued more than once during the selected interval, the XML contains only one **interval_message** element for that unique message ID.

"XML element descriptions for a Version 1 INTERVAL request" on page 303 provides additional information about each element in the XML response.

```
<xs:schema xmlns:xs="http://www.w3.org/2001/XMLSchema"
  targetNamespace="http://www.example.org/MelodyCoreInterval"
 xmlns="http://www.example.org/MelodyCoreInterval"
  elementFormDefault="qualified">

  <xs:element name="interval">
     <xs:complexType>
        <xs:sequence>
      <xs:element name="version" type="xs:int" />
        <xs:element name="sys_id" type="xs:string"/>
        <xs:element name="start_time" type="xs:dateTime" />
        <xs:element name="end_time" type="xs:dateTime" />
        <xs:element name="anomaly_score" type="xs:double"/>
        <xs:element name="model_internal_id" type="xs:int"/>
        <xs:element name="melody_version" type="xs:int"/>
        <xs:element name="interval_message" type="interval_message_type"
          maxOccurs="unbounded" minOccurs="0"/>
      </xs:sequence>
    </xs:complexType>
 </xs:element>

     <xs:complexType name="interval_message_type">
      <xs:sequence>
        <xs:element name="num_instances" type="xs:int"/>
         <xs:element name="bernoulli" type="xs:double"/>
          <xs:element name="cluster_id" type="xs:int"/>
        <xs:element name="poisson" type="xs:double"/>
        <xs:element name="intCont" type="xs:double"/>
        <xs:element name="normIntCont" type="xs:double"/>
        <xs:element name="anomaly" type="xs:double"/>
        <xs:element name="cluster_status" type="xs:string"/>
        <xs:element name="critical_words" type="xs:double"/>
        <xs:element name="text_sum" type="xs:string"/>
        <xs:element name="text_smp" type="xs:string"/>
        <xs:element name="time_vec" type="interval_time_vector_type"/>
        <xs:element name="active_rules" type="active_rules_type" maxOccurs="1" minOccurs="0"/>
      </xs:sequence>
      <xs:attribute name="msg_id" type="xs:string" use="required" />
    </xs:complexType>
```

```
    <xs:complexType name="interval_time_vector_type">
       <xs:sequence>
         <xs:element name="occ" maxOccurs="unbounded" minOccurs="0" type="xs:int"/>
       </xs:sequence>
    </xs:complexType>

    <xs:complexType name="active_rules_type">
      <xs:sequence>
         <xs:element name="rule" type="rule_type" maxOccurs="unbounded" minOccurs="0"/>
       </xs:sequence>
   </xs:complexType>

   <xs:complexType name="rule_type">
      <xs:sequence>
           <xs:element name="name" type="xs:string" />
           <xs:element name="action" type="xs:string" />
      </xs:sequence>
   </xs:complexType>
</xs:schema>
```

XML element descriptions for a Version 1 INTERVAL request

The following list describes the major elements in the **interval** element.

version
An integer that identifies the version of the IBM zAware application programming interface (API). For information about specific API versions, see "API versioning" on page 293.

sys_id
A string that provides the name of the system that was specified on the INTERVAL request, and the name of the sysplex to which the system belongs.

start_time
Indicates the beginning of the first interval for which data is available for the specified system on the date in the LPAR request. The start time is indicated in the XML `dateTime` data type format in Coordinated Universal Time (UTC).

*YYYY-MM-DD***T***hh:mm:ss.ttt***Z**

end_time
Indicates the beginning of the first interval *after* the date specified in the LPAR request. The end time is indicated in the XML `dateTime` data type format in Coordinated Universal Time (UTC).

*YYYY-MM-DD***T***hh:mm:ss.ttt***Z**

anomaly_score
A double value that provides the anomaly score for this interval. The interval anomaly score is the percentile of the sum of each anomaly score for individual message IDs within an interval. When the IBM zAware server uses priming data and current data to create a model of system behavior, a process that is called "training", the server captures the distribution of interval anomaly scores for all intervals that are represented in the training data. The server uses the distribution results and uses them to establish the range of values for each percentile.

The possible interval anomaly scores are:

0 through 99.4
The analysis interval contains messages and message clusters that match or exhibit relatively insignificant differences in expected behavior, as defined in the IBM zAware model. A score of 0 is possible because the server eliminates all expected, in-context messages from its scoring calculation. A score of 0 indicates intervals that exhibit no difference in behavior when compared to the system or group model. The analysis snapshots for these analysis intervals are colored with the lightest blue shade.

Analysis intervals with scores that are greater than 0 but less than 99.5 contain some messages that are unexpected or issued out of context. Scores in this range indicate intervals

that do not vary significantly from the system model. The analysis snapshots for these analysis intervals are colored with varying shades of blue.

99.5 Analysis intervals with this score contain rarely seen, unexpected, or out-of-context messages. Generally speaking, this score indicates analysis intervals with some differences from the system or group model but do not contain messages of much diagnostic value. The analysis snapshots for these analysis intervals are colored with the darkest blue shade.

99.6 - 100
Analysis intervals with this score contain rarely seen messages (these messages appear in the model only once or twice), or many messages that are unexpected or issued out of context. This score indicates analysis intervals with more differences from the system or group model; these intervals can contain messages that might help you diagnose anomalous system behavior. The analysis snapshots for these analysis intervals are the color gold.

101 Analysis intervals with this score exhibit the most significant differences from the system or group model; these intervals contain messages that merit investigation. The analysis snapshots for these analysis intervals are the color orange. IBM zAware assigns this score to analysis intervals that contain:
- Unusual or unexpected messages.
- Messages that IBM rules define as critical.
- A much higher volume of messages than expected.

model_internal_id
An integer that the IBM zAware server uses to identify this system model.

melody_version
An integer that represents the version of the analytics engine that the IBM zAware server is using.

interval_message
The XML response contains one **interval_message** element for each unique message ID that was issued within the interval specified on the LPAR request. The attribute `msg_id` on each **interval_message** element contains a string that identifies the unique message ID.

Each **interval_message** contains the following attributes for the message.

num_instances
An integer that specifies the number of times that this message was issued within this 10-minute interval.

bernoulli
A double value that indicates how frequently the message ID is issued within a sampled set of 10-minute intervals in the system model. Values range from 1 to 101:
- A value of 1 indicates that the message is issued in almost all analysis intervals in the model.
- A value of 100 indicates that the message is issued in almost none of the analysis intervals in the model.
- A value of 101 indicates that this message ID was not issued in any analysis interval in the model.

cluster_id
An integer that represents the identifier of the cluster to which this message belongs. When the message is not part of a recognized cluster, the cluster ID is `-1`.

poisson
A double value that indicates how closely the message ID distribution in current data matches the Poisson distribution of that message ID in data during the training period for the system model. This value is provided only for message IDs that are not part of a cluster. The higher the **poisson** value, the greater the difference from expected behavior.

intCont
A double value that indicates the relative contribution of this message to the interval anomaly

score for the 10-minute interval. This interval score is a function of the message anomaly score, the number of times that the message appears within this interval, and whether the message appeared in context.

anomaly
A double value that indicates the rarity of this specific message ID within the selected interval. The anomaly score is a combination of the interval contribution score for this message and the rule, if any, that is in effect for this message. Higher scores indicate greater anomaly so messages with high anomaly scores are more likely to indicate a problem.

cluster_status
A string that indicates whether or not this message is part of an expected pattern of messages associated with a routine system event (for example, starting a subsystem or workload). IBM zAware identifies and recognizes these patterns or groups, which are called "clusters", and the specific message IDs that constitute a specific cluster. When analyzing data from a monitored client, the server determines whether a specific message is expected to be issued within a specific cluster. A message that is issued out of context (without the other messages in the same cluster) might indicate a problem.

Values for **cluster_status** are:

New IBM zAware did not previously detect this message in the model or detected one or more messages for the first time.

Unclustered
This message is not part of a defined cluster.

In context
IBM zAware expects this message to be issued within a specific cluster, and the message was issued as expected in the analysis interval.

Out of context
IBM zAware expects this message to be issued within a specific cluster, but the message was issued in a different context during the analysis interval.

critical_words
A double value that indicates whether the message contains specific words that indicate potential problems. Critical words include "abend", "failure", and "warning".

text_sum
A string that contains a summary of the common message text that was issued for each occurrence of the same message.

text_smp
A string that contains the full message text for the first occurrence of this message within the interval.

time_vec
The XML response contains one **time_vec** element for each unique message ID that was issued within the interval specified on the LPAR request.

occ
The XML response contains one **occ** element for each time that this message ID was issued within the interval specified on the LPAR request.

active_rules
The XML response contains one **active_rules** element for each unique message ID that was issued within the interval specified on the LPAR request.

rule
The XML response contains one **rule** element for each rule that is in effect for this message ID.

name
A string that contains the name of the rule that was applied for this message. The rule can be one of the following types:
- Predefined by IBM.
- Assigned by IBM zAware as a result of the analysis of training data.
- Assigned by IBM zAware when an administrator has identified the message as one that IBM zAware is to ignore during analysis, either until the next model is built or until an administrator manually resets the ignore status.

action
A string that contains the status value associated with the applied rule. Possible values are:

CRITICAL
An IBM rule identifies this message as critical for diagnosing a potential system problem. For example, message IXC101I, which indicates that a system is being removed from a sysplex, is classified as critical.

IMPORTANT
An IBM rule identifies this message as likely to indicate a problem. For example, message IEA911E, which indicates that an SVC dump was taken, is classified as important.

INTERESTING
An IBM rule identifies this message as indicative of a diagnostically useful event, such as a health check exception.

NONE
No rule is applied for this message.

NON-INTERESTING
One of the following conditions is true for this message:
- A predefined IBM rule or an IBM zAware-assigned rule identifies this message as one with little or no diagnostic value.
- An administrator identified the message as one that IBM zAware is to ignore during analysis, either until the next model is built or until an administrator manually resets the ignore status.

Sample XML response for a Version 1 INTERVAL request

This sample provides the returned XML in response to a GET request for the details for a system named "CB8E" in sysplex "ABCPLEX" during the 10-minute interval that started on 12 September 2012 at 14:20:00 UTC. The sample response shown is formatted for publication and contains information for only a few of the returned messages issued within this interval.

GET method:

```
GET https://xx.xx.xx.xx/zAware/authuser/Analysis?
reqtype=INTERVAL&amp;amp;time=20120912142000&amp;amp;sysname=ABCPLEX.CB8E&amp;amp;clienttype=zOS&amp;amp;version=v1
```

XML response:

```
<?xml version='1.0' encoding='UTF-8' ?>
<?xml-stylesheet href='./xslt/MelodyCoreInterval.xsl' type='text/xsl' ?>
<interval xsi:noNamespaceSchemaLocation="xslt/MelodyCoreInterval.xsd"
  xmlns="http://www.example.org/MelodyCoreInterval"
    xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance">
<version>1</version>
<sys_id>ABCPLEX-CB8E</sys_id>
<start_time>2016-09-12T14:20:00.000Z</start_time>
<end_time>2016-09-12T14:30:00.000Z</end_time>
<anomaly_score>95.3</anomaly_score>
```

```
<model_internal_id>2</model_internal_id>
<melody_version>170</melody_version>
<interval_message msg_id="BLWH0001E">
     <num_instances>1</num_instances>
     <bernoulli>58.0</bernoulli>
     <cluster_id>120</cluster_id>
     <poisson>0.0</poisson>
     <intCont>0.0</intCont>
     <normIntCont>0.0</normIntCont>
     <anomaly>0.0</anomaly>
     <cluster_status>IN_CONTEXT</cluster_status>
     <critical_words>0.0</critical_words>
     <text_sum>AutoIPL policy is not active.</text_sum>
     <text_smp>AutoIPL policy is not active.</text_smp>
     <time_vec>
        <occ>0</occ>
     </time_vec>
</interval_message>
<interval_message msg_id="HZS0002E">
     <num_instances>12</num_instances>
     <bernoulli>12.0</bernoulli>
     <cluster_id>-1</cluster_id>
     <poisson>3.245</poisson>
     <intCont>3.245</intCont>
     <normIntCont>8.947</normIntCont>
     <anomaly>0.96</anomaly>
     <cluster_status>UNCLUSTERED</cluster_status>
     <critical_words>0.0</critical_words>
     <text_sum>CHECK(*,*):</text_sum>
     <text_smp>CHECK(IBMSVA,SVA_AUTOIPL_DEFINED):</text_smp>
     <time_vec>
        <occ>0</occ>
        <occ>1</occ>
        <occ>4</occ>
        <occ>16</occ>
        <occ>24</occ>
        <occ>25</occ>
      </time_vec>
</interval_message>

⋮
</interval>
```

Version 2 API

Through Version 2 of the IBM zAware API, system management products can request and receive IBM zAware analytical data that is equivalent to the information available through the Analysis Graph view and Interval page in the IBM zAware GUI.

You can use the Version 2 API to request and receive data from an IBM zAware server that is running on any of the supported host systems, which are listed in Chapter 2, "Prerequisites for configuring and using IBM zAware," on page 13, but the returned results might be in Version 1 format if the Version 2 format is not available. For example, if the request is for IBM zAware data that has been produced by an IBM zAware server on a zEnterprise host system, and that data has been migrated to a server on a z13, z13s, or z14 host system, IBM zAware can return data only in the Version 1 format. For information about determining the version of returned data, see "Content-Type header values for HTTP responses" on page 297.

XML for a Version 2 ANALYSIS request

The information here provides the XML element descriptions and sample XML responses that the IBM zAware server returns in response to an HTTP GET method with an **ANALYSIS** request type. This XML

response contains information that is equivalent to the interval anomaly scores that the server displays through the **Analysis** page in the IBM zAware graphical user interface (GUI).

The following code illustrates the XML structure of the response to an HTTP GET method with an **ANALYSIS** request type. The major element is the **systems** element, which identifies the specific date and monitored client (system) for which analytical data were requested. The **systems** element also identifies the number and size of intervals that are returned in the XML document. The XML also contains one **interval** element for each analysis snapshot since Coordinated Universal Time (UTC) midnight on the requested date. The **interval** element provides the interval anomaly score and number of unique message IDs that were issued during the specific analysis snapshot.

See "XML element descriptions for a Version 2 ANALYSIS request" on page 309 for additional information about each element in the XML response.

```
<xs:schema xmlns:xs="http://www.w3.org/2001/XMLSchema"
    targetNamespace="http://www.ibm.com/zAware/MelodyCorePlexV2"
    xmlns="http://www.ibm.com/zAware/MelodyCorePlexV2"
    elementFormDefault="qualified">

    <xs:element name="systems" >
    <xs:complexType>
        <xs:sequence>
            <xs:element name="version" type="xs:int"/>
            <xs:element name="start_time" type="xs:dateTime" />
            <xs:element name="end_time" type="xs:dateTime" />
            <xs:element name="gmt_offset" type="xs:string" />
            <xs:element name="number_intervals">
                <xs:complexType>
                    <xs:simpleContent>
                        <xs:extension base="xs:int">
                            <xs:attribute name="analysis_snapshot_size" type="xs:int" use="required"/>
                        </xs:extension>
                    </xs:simpleContent>
                </xs:complexType>
            </xs:element>
            <xs:element name="interval_size" type="xs:int" />
            <xs:element name="model_info">
                <xs:complexType>
                    <xs:attribute name="model_creation_date" type="xs:dateTime" use="required"/>
                    <xs:attribute name="training_period" type="xs:int" use="required"/>
                    <xs:attribute name="analysis_group" type="xs:string" use="required"/>
                </xs:complexType>
            </xs:element>
            <xs:element name="system" type="systems_system_type"
                minOccurs="1" maxOccurs="1" />
        </xs:sequence>
    </xs:complexType>
    </xs:element>

    <xs:complexType name="systems_system_type">
        <xs:sequence>
            <xs:element name="interval" type="systems_interval_type"
                minOccurs="0" maxOccurs="unbounded" />
        </xs:sequence>
        <xs:attribute name="sys_id" type="xs:string" use="required" />
        <xs:attribute name="log_type" type="xs:string" use="required"/>
    </xs:complexType>

    <xs:complexType name="systems_interval_type">
        <xs:sequence>
            <xs:element name="num_unique_msg_ids" type="xs:int" />
            <xs:element name="anomaly_score" type="xs:double" />
```

```
        </xs:sequence>
     <xs:attribute name="num_never_seen_before_messages" type="xs:int" use="required" />
        <xs:attribute name="num_new_messages" type="xs:int" use="required" />
        <xs:attribute name="num_new_messages_first_reported" type="xs:int" use="required" />
        <xs:attribute name="index" type="xs:int" use="required"/>
        <xs:attribute name="missing" type="xs:boolean" use="required"/>
        <xs:attribute name="missing_reason" type="xs:string" use="optional"/>
        <xs:attribute name="limited_model" use="optional">
            <xs:simpleType>
                <xs:restriction base="xs:string">
                    <xs:enumeration value="Yes" />
                    <xs:enumeration value="No" />
                    <xs:enumeration value="Unknown" />
                </xs:restriction>
            </xs:simpleType>
        </xs:attribute>
    </xs:complexType>

</xs:schema>
```

XML element descriptions for a Version 2 ANALYSIS request

The following list describes the major elements in the **systems** element.

version
An integer that identifies the version of the IBM zAware application programming interface (API). For information about specific API versions, see "API versioning" on page 293.

start_time
Indicates the beginning of the first interval for which data is available for the specified system on the date in the **ANALYSIS** request. The start time is indicated in the XML `dateTime` data type format in Coordinated Universal Time (UTC).

*YYYY-MM-DD***T***hh:mm:ss.ttt***Z**

end_time
Indicates the beginning of the first interval *after* the date specified in the **ANALYSIS** request. The end time is indicated in the XML `dateTime` data type format in Coordinated Universal Time (UTC).

*YYYY-MM-DD***T***hh:mm:ss.ttt***Z**

gmt_offset
An integer that indicates the difference in hours and minutes from Coordinated Universal Time (UTC) for the requested start time.

number_intervals
An integer that indicates the number of intervals for which analytical data is available for the system and the date that is specified on the **ANALYSIS** request. The attribute **analysis_snapshot_size** provides the size of an analysis snapshot in seconds.

analysis_snapshot_size
An integer that indicates the amount of time in an analysis snapshot.

interval_size
An integer that indicates the number of seconds in an interval.

model_info
Provides information about the model that is associated with the specified system.

model_creation_date
An element that provides the date and time when IBM zAware successfully built the most recent model of system behavior.

training_period
An integer that indicates the number of consecutive calendar days that the IBM zAware server uses to identify the instrumentation data to include in training models.

analysis_group
An element that provides the name of a z/OS sysplex or Linux model group in the IBM zAware topology.

system
An element that provides more details about intervals for the system that is specified on the **ANALYSIS** request.

interval
An element that provides more details about a specific interval. For the system and the date that is specified on the **ANALYSIS** request, the XML response contains one interval element for each element for which analytical data is available.

num_unique_msg_ids
An integer that provides the number of unique message IDs that were issued during this analysis interval. If the same message ID was issued more than once during the interval, the message ID is counted only once.

anomaly_score
A double value that provides the anomaly score for this interval. The interval anomaly score is the percentile of the sum of each anomaly score for individual message IDs within an interval. When the IBM zAware server uses priming data and current data to create a model of system behavior, a process that is called "training", the server captures the distribution of interval anomaly scores for all intervals that are represented in the training data. The server uses the distribution results and uses them to establish the range of values for each percentile.

The possible interval anomaly scores are:

0 through 99.4
The analysis interval contains messages and message clusters that match or exhibit relatively insignificant differences in expected behavior, as defined in the IBM zAware model. A score of 0 is possible because the server eliminates all expected, in-context messages from its scoring calculation. A score of 0 indicates intervals that exhibit no difference in behavior when compared to the system or group model. The analysis snapshots for these analysis intervals are colored with the lightest blue shade.

Analysis intervals with scores that are greater than 0 but less than 99.5 contain some messages that are unexpected or issued out of context. Scores in this range indicate intervals that do not vary significantly from the system model. The analysis snapshots for these analysis intervals are colored with varying shades of blue.

99.5 Analysis intervals with this score contain rarely seen, unexpected, or out-of-context messages. Generally speaking, this score indicates analysis intervals with some differences from the system or group model but do not contain messages of much diagnostic value. The analysis snapshots for these analysis intervals are colored with the darkest blue shade.

99.6 - 100
Analysis intervals with this score contain rarely seen messages (these messages appear in the model only once or twice), or many messages that are unexpected or issued out of context. This score indicates analysis intervals with more differences from the system or group model; these intervals can contain messages that might help you diagnose anomalous system behavior. The analysis snapshots for these analysis intervals are the color gold.

101 Analysis intervals with this score exhibit the most significant differences from the

system or group model; these intervals contain messages that merit investigation. The analysis snapshots for these analysis intervals are the color orange. IBM zAware assigns this score to analysis intervals that contain:
- Unusual or unexpected messages.
- Messages that IBM rules define as critical.
- A much higher volume of messages than expected.

num_never_seen_before_messages
An integer that indicates the number of message IDs that were issued during the new analysis interval for "the first time." These messages were never seen in any of the previous analysis intervals or in the current model.

num_new_messages
An integer that indicates the number of message IDs that are not in the model and also meets one of the following conditions:

1. The message was issued during the current analysis interval "for the first time".
2. The message was issued before the current model was started.
3. The message was issued before the current model was created.

For more information, see Table 71 on page 312.

num_new_messages_first_reported
An integer that indicates the number of message IDs that are not in the model and also meets one of the following conditions:

1. The message was issued during the current analysis interval "for the first time".
2. The message was issued before the current model was started.
3. The message was issued before the current model was created.

The difference between **num_new_messages** and **num_new_messages_first_reported** can be seen only on a Linux system with its overlapping 60-minute analysis intervals (as compared to z/OS with its 10-minute analysis intervals that do not overlap). For an example, see Table 71 on page 312.

index
An integer that indicates the sequence number of this interval within the date that is specified on the LPAR request.

missing
A Boolean value that identifies whether analytical data is available for this interval.

missing_reason
An element that indicates why analytical data is not available; this element has a value only when the value returned for `missing` is `true`.

limited_model
Indicates whether IBM zAware used a limited model to calculate the anomaly score for the interval. Valid values are Yes, No, or Unknown, which indicates temporary conditions under which IBM zAware cannot determine whether the model is limited.

sys_id
Provides the name of the system that was specified on the **ANALYSIS** request.

log_type
An element that identifies the type of data that the monitored system is supplying to the IBM zAware server.

Sample XML response for a Version 2 ANALYSIS request

A sample ANALYSIS request and response for a z/OS system
The request:

```
https://198.12.18.06/zAware/authuser/Analysis?reqtype=
analysis&time=20150206142000&sysname=PLEX2.SY06&clienttype=zOS&version=v2
```

The response:

```
<?xml version='1.0' encoding='UTF-8' ?>
<?xml-stylesheet href='./xslt/MelodyCorePlexV2.xsl' type='text/xsl' ?>
<systems xsi:noNamespaceSchemaLocation="/xml/MelodyCorePlexV2.xsd"
xmlns="http://www.ibm.com/zAware/MelodyCorePlexV2"
xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance">
<version>2</version>
<start_time>2016-02-06T00:00:00.000Z</start_time>
<end_time>2016-02-07T00:00:00.000Z</end_time>
<gmt_offset>GMT-05:00</gmt_offset>
<number_intervals analysis_snapshot_size="600">144</number_intervals>
<interval_size>600</interval_size>
<model_info model_creation_date="2016-02-05T19:26:10.815Z"
training_period="120" analysis_group="PLEX2-SY06"/>
<system sys_id="PLEX2-SY06" log_type="z/OS operlog">
<interval
   num_never_seen_before_messages="10"
   num_new_messages="10"
   num_new_messages_first_reported="10"
   index="0" missing="false" limited_model="No">
   <num_unique_msg_ids>48</num_unique_msg_ids>
   <anomaly_score>99.9</anomaly_score>
</interval>
<interval
   num_never_seen_before_messages="0"
   num_new_messages="0"
   num_new_messages_first_reported="0"
   index="1" missing="false" limited_model="No">
   <num_unique_msg_ids>32</num_unique_msg_ids>
   <anomaly_score>101.0</anomaly_score>
</interval>
.
.
.
</system>
</systems>
```

Table 71. Difference between new messages and first reported new messages for a Linux system

Interval	num_new_messages	num_new_messages_first_reported
11:00-12:00	0	0
11:10-12:10	0	0
11:20-12:20	**1 (msg3000)** msg3000 first appears at 12:14, which means that the number of new messages is 1 for msg3000 and for the next several intervals that include 12:14.	**1 (msg3000)** msg3000 is only counted as first reported during the first instance that an interval includes 12:14.
11:30-12:30	**1 (msg3000)** **num_new_messages** interval includes 12:14.	0
11:40-12:40	**1 (msg3000)** **num_new_messages** interval includes 12:14.	0
11:50-12:50	**1 (msg3000)** **num_new_messages** interval includes 12:14.	0

Table 71. Difference between new messages and first reported new messages for a Linux system (continued)

Interval	num_new_messages	num_new_messages_first_reported
12:00-13:00	**1 (msg3000)** **num_new_messages** interval includes 12:14.	0
12:10 - 13:10	**1 (msg3000)** **num_new_messages** interval includes 12:14.	0
12:20-13:20	**0** msg3000 appears at 13:11, but not considered a new message because it is not the first interval that the message was seen during the analysis period.	**0** msg3000 appears at 13:11, but not considered a first reported message because it is not the first time the message appeared during the analysis period.
13:20-13:40	0	0

XML for a Version 2 INTERVAL request

This topic provides the XML element descriptions and sample XML that the IBM zAware server returns in response to an HTTP GET method with an `INTERVAL` request type. This XML response contains information that is equivalent to the interval and message details that the server displays through the Interval page in the IBM zAware graphical user interface (GUI).

The following code illustrates the XML structure of the response to an HTTP GET method with an `INTERVAL` request type. The major element is the **interval** element, which contains information about a specific analysis interval for a specific system that is established as an IBM zAware monitored client. The **interval** element also contains one **interval_message** element for each unique message issued during the interval. If the same message ID was issued more than once during the selected interval, the XML contains only one **interval_message** element for that unique message ID.

"XML element descriptions for a Version 2 INTERVAL request" on page 315 provides additional information about each element in the XML response.

```
<xs:schema xmlns:xs="http://www.w3.org/2001/XMLSchema"
    targetNamespace="http://www.ibm.com/zAware/MelodyCoreIntervalV2"
    xmlns="http://www.ibm.com/zAware/MelodyCoreIntervalV2" elementFormDefault="qualified">
    <xs:element name="interval">
        <xs:complexType>
            <xs:sequence>
                <xs:element name="version" type="xs:int" />
                <xs:element name="sys_id" type="xs:string" />
                <xs:element name="start_time" type="xs:dateTime" />
                <xs:element name="end_time" type="xs:dateTime" />
                <xs:element name="anomaly_score" type="xs:double" />
                <xs:element name="model_internal_id" type="xs:int" />
                <xs:element name="melody_version" type="xs:int" />
                <xs:element name="gmt_offset" type="xs:string" />
                <xs:element name="model_info">
                    <xs:complexType>
                        <xs:attribute name="model_creation_date" type="xs:dateTime" use="required" />
                        <xs:attribute name="training_period" type="xs:int" use="required" />
                        <xs:attribute name="interval_size_in_sec" type="xs:long" use="required" />
                        <xs:attribute name="analysis_group" type="xs:string" use="required" />
                    </xs:complexType>
                </xs:element>
                <xs:element name="msg_summary">
                    <xs:complexType>
                        <xs:attribute name="num_new_msg" type="xs:int" use="required" />
                    </xs:complexType>
```

```
            </xs:element>
            <xs:element name="interval_message" type="interval_message_type" maxOccurs="unbounded"
                minOccurs="0" />
        </xs:sequence>
    </xs:complexType>
</xs:element>

<xs:complexType name="interval_message_type">
    <xs:sequence>
        <xs:element name="num_instances" type="xs:int" />
        <xs:element name="bernoulli">
            <xs:complexType>
                <xs:simpleContent>
                    <xs:extension base="xs:double">
                        <xs:attribute name="frequency" type="xs:double" use="required" />
                    </xs:extension>
                </xs:simpleContent>
            </xs:complexType>
        </xs:element>
        <xs:element name="cluster_id" type="xs:int" />
        <xs:element name="periodicity">
            <xs:complexType>
                <xs:attribute name="status" use="required">
                    <xs:simpleType>
                        <xs:restriction base="xs:string">
                            <xs:enumeration value="IN_SYNC" />
                            <xs:enumeration value="NOT_IN_SYNC" />
                            <xs:enumeration value="NOT_PERIODIC" />
                            <xs:enumeration value="NEW" />
                        </xs:restriction>
                    </xs:simpleType>
                </xs:attribute>
                <xs:attribute name="last_issued" type="xs:dateTime" use="optional" />
                <xs:attribute name="score" type="xs:double" use="optional" />
            </xs:complexType>
        </xs:element>
        <xs:element name="poisson">
            <xs:complexType>
                <xs:simpleContent>
                    <xs:extension base="xs:double">
                        <xs:attribute name="mean" type="xs:double" use="required" />
                    </xs:extension>
                </xs:simpleContent>
            </xs:complexType>
        </xs:element>
        <xs:element name="intCont" type="xs:double" />
        <xs:element name="normIntCont" type="xs:double" />
        <xs:element name="anomaly" type="xs:double" />
        <xs:element name="cluster_status" type="xs:string" />
        <xs:element name="critical_words" type="xs:double" />
        <xs:element name="text_sum" type="xs:string" />
        <xs:element name="text_smp" type="xs:string" />
        <xs:element name="time_vec" type="interval_time_vector_type" />
        <xs:element name="active_rules" type="active_rules_type" maxOccurs="1" minOccurs="0" />
    </xs:sequence>
    <xs:attribute name="msg_id" type="xs:string" use="required" />
</xs:complexType>

<xs:complexType name="interval_time_vector_type">
    <xs:sequence>
        <xs:element name="occ" maxOccurs="unbounded" minOccurs="0" type="xs:int" />
    </xs:sequence>
</xs:complexType>



<xs:complexType name="active_rules_type">
```

```
        <xs:sequence>
            <xs:element name="rule" type="rule_type" maxOccurs="unbounded" minOccurs="0" />
        </xs:sequence>
    </xs:complexType>

    <xs:complexType name="rule_type">
        <xs:sequence>
            <xs:element name="name">
                <xs:complexType>
                    <xs:simpleContent>
                        <xs:extension base="xs:string">
                            <xs:attribute name="affected_score" type="xs:boolean" />
                        </xs:extension>
                    </xs:simpleContent>
                </xs:complexType>
            </xs:element>
            <xs:element name="action" type="xs:string" />
        </xs:sequence>
    </xs:complexType>
</xs:schema>
```

XML element descriptions for a Version 2 INTERVAL request

The following list describes the major elements in the **interval** element.

version
An integer that identifies the version of the IBM zAware application programming interface (API). For information about specific API versions, see "API versioning" on page 293.

sys_id
A string that provides the name of the system that was specified on the INTERVAL request, and the name of the system group to which the system belongs.

start_time
Indicates the beginning of the first interval for which data is available for the specified system on the date in the INTERVAL request. The start time is indicated in the XML `dateTime` data type format in Coordinated Universal Time (UTC).

*YYYY-MM-DD***T***hh:mm:ss.ttt***Z**

end_time
Indicates the beginning of the first interval *after* the date specified in the INTERVAL request. The end time is indicated in the XML `dateTime` data type format in Coordinated Universal Time (UTC).

*YYYY-MM-DD***T***hh:mm:ss.ttt***Z**

anomaly_score
A double value that provides the anomaly score for this interval. The interval anomaly score is the percentile of the sum of each anomaly score for individual message IDs within an interval. When the IBM zAware server uses priming data and current data to create a model of system behavior, a process that is called "training", the server captures the distribution of interval anomaly scores for all intervals that are represented in the training data. The server uses the distribution results and uses them to establish the range of values for each percentile.

The possible interval anomaly scores are:

0 through 99.4
The analysis interval contains messages and message clusters that match or exhibit relatively insignificant differences in expected behavior, as defined in the IBM zAware model. A score of 0 is possible because the server eliminates all expected, in-context messages from its scoring calculation. A score of 0 indicates intervals that exhibit no difference in behavior when compared to the system or group model. The analysis snapshots for these analysis intervals are colored with the lightest blue shade.

Analysis intervals with scores that are greater than 0 but less than 99.5 contain some messages that are unexpected or issued out of context. Scores in this range indicate intervals that do not vary significantly from the system model. The analysis snapshots for these analysis intervals are colored with varying shades of blue.

99.5
Analysis intervals with this score contain rarely seen, unexpected, or out-of-context messages. Generally speaking, this score indicates analysis intervals with some differences from the system or group model but do not contain messages of much diagnostic value. The analysis snapshots for these analysis intervals are colored with the darkest blue shade.

99.6 - 100
Analysis intervals with this score contain rarely seen messages (these messages appear in the model only once or twice), or many messages that are unexpected or issued out of context. This score indicates analysis intervals with more differences from the system or group model; these intervals can contain messages that might help you diagnose anomalous system behavior. The analysis snapshots for these analysis intervals are the color gold.

101
Analysis intervals with this score exhibit the most significant differences from the system or group model; these intervals contain messages that merit investigation. The analysis snapshots for these analysis intervals are the color orange. IBM zAware assigns this score to analysis intervals that contain:
- Unusual or unexpected messages.
- Messages that IBM rules define as critical.
- A much higher volume of messages than expected.

model_internal_id
An integer that the IBM zAware server uses to identify this system model.

melody_version
An integer that represents the version of the analytics engine that the IBM zAware server is using.

gmt_offset
An integer that indicates the difference in hours and minutes from Coordinated Universal Time (UTC) for the requested start time.

model_info
Provides information about the model associated with the specified system.

model_creation_date
An element that provides the date and time when IBM zAware successfully built the most recent model of system behavior.

training_period
An integer that indicates the number of consecutive calendar days that the IBM zAware server uses to identify the instrumentation data to include in training models.

interval_size_in_sec
An integer that indicates the number of seconds in an interval.

analysis_group
An element that provides the name of a z/OS sysplex or Linux model group in the IBM zAware topology.

msg_summary
An element that contains summary information about messages in the interval.

num_new_msg
An integer that provides the total number of new messages issued by the monitored system during this interval.

interval_message
The XML response contains one **interval_message** element for each unique message ID that was issued within the interval specified on the INTERVAL request. Each **interval_message** contains the following attributes for the message.

num_instances
An integer that specifies the number of times that this message was issued within this 10-minute interval.

bernoulli
A double value that indicates how frequently the message ID is issued within a sampled set of 10-minute analysis snapshots in the system model. Values range from 1 to 101:
- A value of 1 indicates that the message is issued in almost all analysis intervals in the model.
- A value of 100 indicates that the message is issued in almost none of the analysis intervals in the model.
- A value of 101 indicates that this message ID was not issued in any analysis interval in the model.

frequency
An integer that indicates the average number of analysis intervals in which the message is issued each day, according to analysis of the message data that IBM zAware uses for training.

cluster_id
An integer that represents the identifier of the cluster to which this message belongs. When the message is not part of a recognized cluster, the cluster ID is -1.

periodicty_status
An element that indicates whether or not this message has a tendency to recur at specific times, and whether the message recurred as expected within the analysis interval. Valid values are:
NEW IBM zAware did not previously detect this message.
IN_SYNC
IBM zAware expects this message to be issued in a periodic pattern, and the message was issued as expected during the analysis interval.
NOT_IN_SYNC
IBM zAware expects this message to be issued in a periodic pattern, but the message was not issued as expected during the analysis interval.
NOT_PERIODIC
IBM zAware does not expect this message to be issued in a periodic pattern.

last_issued
An element that provides the UTC date and time when this message was last issued on the monitored system.

score
An integer that indicates how the periodicity status of this message contributed to the message anomaly score for the analysis interval. Higher scores generally indicate greater contribution to the message anomaly score

poisson
A double value that indicates how closely the message ID distribution in current data matches the Poisson distribution of that message ID in data during the training period for the system model. This value is provided only for message IDs that are not part of a cluster. The higher the **poisson** value, the greater the difference from expected behavior.

intCont
A double value that indicates the relative contribution of this message to the interval anomaly score for the analysis interval. This interval score is a function of the message anomaly score, the number of times that the message appears within this interval, and whether the message appeared in context.

normIntCont
A double value that indicates the normalized contribution of this message to the interval anomaly score for the analysis interval.

anomaly
A double value that indicates the rarity of this specific message ID within the selected interval. The anomaly score is a combination of the interval contribution score for this message and the rule, if any, that is in effect for this message. Higher scores indicate greater anomaly so messages with high anomaly scores are more likely to indicate a problem.

cluster_status
A string that indicates whether or not this message is part of an expected pattern of messages associated with a routine system event (for example, starting a subsystem or workload). IBM zAware identifies and recognizes these patterns or groups, which are called "clusters", and the specific message IDs that constitute a specific cluster. When analyzing data from a monitored client, the server determines whether a specific message is expected to be issued within a specific cluster. A message that is issued out of context (without the other messages in the same cluster) might indicate a problem.

Values for **cluster_status** are:

New IBM zAware did not previously detect this message in the model or detected one or more messages for the first time.

Unclustered
This message is not part of a defined cluster.

In context
IBM zAware expects this message to be issued within a specific cluster, and the message was issued as expected in the analysis interval.

Out of context
IBM zAware expects this message to be issued within a specific cluster, but the message was issued in a different context during the analysis interval.

critical_words
A double value that indicates whether the message contains specific words that indicate potential problems. Critical words include "abend", "failure", and "warning".

text_sum
A string that contains a summary of the common message text that was issued for each occurrence of the same message.

text_smp
A string that contains the full message text for the first occurrence of this message within the interval.

time_vec
The XML response contains one **time_vec** element for each unique message ID that was issued within the interval specified on the INTERVAL request.

occ
The XML response contains one **occ** element for each time that this message ID was issued within the interval specified on the INTERVAL request.

active_rules
The XML response contains one **active_rules** element for each unique message ID that was issued within the interval specified on the INTERVAL request.

rule
The XML response contains one **rule** element for each rule that is in effect for this message ID.

name
A string that contains the name of the rule that was applied for this message. The rule can be one of the following types:
- Predefined by IBM.
- Assigned by IBM zAware as a result of the analysis of training data.
- Assigned by IBM zAware when an administrator has identified the message as one that IBM zAware is to ignore during analysis, either until the next model is built or until an administrator manually resets the ignore status.

action
A string that contains the status value associated with the applied rule. Possible values are:

CRITICAL
An IBM rule identifies this message as critical for diagnosing a potential system problem. For example, message IXC101I, which indicates that a system is being removed from a sysplex, is classified as critical.

IMPORTANT
An IBM rule identifies this message as likely to indicate a problem. For example, message IEA911E, which indicates that an SVC dump was taken, is classified as important.

INTERESTING
An IBM rule identifies this message as indicative of a diagnostically useful event, such as a health check exception.

NONE
No rule is applied for this message.

NON-INTERESTING
One of the following conditions is true for this message:
- A predefined IBM rule or an IBM zAware-assigned rule identifies this message as one with little or no diagnostic value.
- An administrator identified the message as one that IBM zAware is to ignore during analysis, either until the next model is built or until an administrator manually resets the ignore status.

msg_id
A string that identifies the unique message ID.

Sample XML response for a Version 2 INTERVAL request

A sample INTERVAL request and response for a z/OS system

The request:

```
https://198.xx.xx.xx/zAware/authuser/Analysis?clienttype=zOS
&reqtype=INTERVAL&sysname=ZOSPLEX.ZOSSYS&time=20160131013000&version=V2
```

The response:

```
<?xml version='1.0' encoding='UTF-8' ?>
<?xml-stylesheet href='./xslt/MelodyCoreIntervalV2.xsl' type='text/xsl' ?>
<interval xsi:noNamespaceSchemaLocation="/xml/MelodyCoreIntervalV2.xsd"
xmlns="http://www.ibm.com/zAware/MelodyCoreIntervalV2"
xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance">
<version>2</version>
<sys_id>ZOSPLEX-ZOSSYS</sys_id>
<start_time>2016-01-31T01:30:00.000Z</start_time>
<end_time>2016-01-31T01:40:00.000Z</end_time>
<anomaly_score>0.0</anomaly_score>
<model_internal_id>1</model_internal_id>
<melody_version>321</melody_version>
```

```
<gmt_offset>GMT-05:00</gmt_offset>
<model_info model_creation_date="2015-01-31T01:00:03.382Z" training_period="90"
interval_size_in_sec="600" analysis_group="ZOSPLEX-ZOSSYS"/>
<msg_summary num_new_msg="2"/>
<interval_message msg_id="ZAI00280">
<num_instances>1</num_instances>
<bernoulli frequency="1.4498881259162102">0.9899313324589152</bernoulli>
<cluster_id>73</cluster_id>
<periodicity status="IN_SYNC" last_issued="2016-01-31T01:22:00.000Z" score="5.560681631015528"/>
<poisson mean="1.0">0.0</poisson>
<intCont>-0.0</intCont>
<normIntCont>-0.0</normIntCont>
<anomaly>0.0</anomaly>
<cluster_status>IN_CONTEXT</cluster_status>
<critical_words>0.0</critical_words>
<text_sum>This is my unique message 280</text_sum>
<text_smp>This is my unique message 280</text_smp>
<time_vec>
<occ>0</occ>
</time_vec>
<active_rules/>
</interval_message>
<interval_message msg_id="ZAI00281">
<num_instances>1</num_instances>
<bernoulli frequency="1.4498881259162102">0.9899313324589152</bernoulli>
<cluster_id>73</cluster_id>
<periodicity status="IN_SYNC" last_issued="2016-01-31T01:22:00.000Z" score="5.560681631015528"/>
<poisson mean="1.0">0.0</poisson>
<intCont>-0.0</intCont>
<normIntCont>-0.0</normIntCont>
<anomaly>0.0</anomaly>
<cluster_status>IN_CONTEXT</cluster_status>
<critical_words>0.0</critical_words>
<text_sum>This is my unique message 281</text_sum>
<text_smp>This is my unique message 281</text_smp>
<time_vec>
<occ>0</occ>
</time_vec>
<active_rules/>
</interval_message>
</interval>
```

A sample INTERVAL request and response for a Linux system

The request:

```
https://198.xx.xx.xx/zAware/authuser/Analysis?reqtype=INTERVAL
&clienttype=linux&sysname=zlinux-client-01&time=20160130215000&version=V2
```

The response:

```
<?xml version='1.0' encoding='UTF-8' ?>
<?xml-stylesheet href='./xslt/MelodyCoreIntervalV2.xsl' type='text/xsl' ?>
<interval xsi:noNamespaceSchemaLocation="/xml/MelodyCoreIntervalV2.xsd"
xmlns="http://www.ibm.com/zAware/MelodyCoreIntervalV2"
xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance">
<version>2</version>
<sys_id>zlinux-client-01</sys_id>
<start_time>2016-01-30T21:00:00.000Z</start_time>
<end_time>2016-01-30T22:00:00.000Z</end_time>
<anomaly_score>95.7</anomaly_score>
<model_internal_id>1</model_internal_id>
<melody_version>321</melody_version>
<gmt_offset>GMT-05:00</gmt_offset>
<model_info model_creation_date="2016-01-30T21:29:10.133Z"
training_period="90" interval_size_in_sec="3600"
analysis_group="zLinux Group"/>
<msg_summary num_new_msg="2"/>
<interval_message msg_id="(NO_COMPONENT)_274">
```

```
<num_instances>714</num_instances>
<bernoulli frequency="3.028849794290618E-4">0.9999873797925238</bernoulli>
<cluster_id>-1</cluster_id>
<periodicity status="NEW"/>
<poisson mean="2.0">0.0</poisson>
<intCont>9.024158912343449</intCont>
<normIntCont>9.024158912343449</normIntCont>
<anomaly>0.9998795359165771</anomaly>
<cluster_status>NEW</cluster_status>
<critical_words>0.0</critical_words>
<text_sum>syslog-ng[9417]: EOF occurred while idle;</text_sum>
<text_smp>syslog-ng[9417]: EOF occurred while idle;</text_smp>
<time_vec>
<occ>73</occ>
<occ>99</occ>
<occ>100</occ>
<occ>101</occ>
<occ>110</occ>
</time_vec>
<active_rules/>
</interval_message>
<interval_message msg_id="(NO_COMPONENT)_275">
<num_instances>714</num_instances>
<bernoulli frequency="3.028849794290618E-4">0.9999873797925238</bernoulli>
<cluster_id>-1</cluster_id>
<periodicity status="NEW"/>
<poisson mean="2.0">0.0</poisson>
<intCont>9.024158912343449</intCont>
<normIntCont>9.024158912343449</normIntCont>
<anomaly>0.9998795359165771</anomaly>
<cluster_status>NEW</cluster_status>
<critical_words>0.0</critical_words>
<text_sum>syslog-ng[9417]: Connection broken; time_reopen='60'</text_sum>
<text_smp>syslog-ng[9417]: Connection broken; time_reopen='60'</text_smp>
<time_vec>
<occ>71</occ>
<occ>100</occ>
<occ>119</occ>
</time_vec>
<active_rules/>
</interval_message>
</interval>
```

Appendix D. IBM zAware operational messages

This topic provides message text and descriptions of the operational messages that the IBM zAware server issues. The message text matches the text that is used in engineering change (EC) and microcode control level (MCL) SE-ZAWARE MCL N98812.022, and later levels. Use the **Notifications** page in the IBM zAware graphical user interface (GUI) to view and manage these messages. User IDs assigned to either the Administrator or User role can view the **Notifications** page but only an administrator can remove a message from the display on the page.

AIFB0001E **IBM zAware encountered an internal error with reason code *code*. Search problem reporting databases for a fix for this error. If no fix exists, report this message ID and reason code to IBM Support.**

Explanation: IBM zAware encountered an internal error indicated by reason code *code*.

System action: For most internal errors, Licensed Internal Code automatically gathers and sends diagnostic information to IBM if the "Call Home" feature is enabled on the IBM zAware host system.

Response: On the Hardware Management Console (HMC) for the IBM zAware host system, check for outstanding hardware messages that indicate a "Call Home" call associated with the IBM zAware partition. If a service call has not been sent automatically, search problem reporting databases for a fix for this error. If no fix exists, request IBM support by generating a Type V Viewable PMH (PMV) record to report this message ID and reason code.

For additional information, see Chapter 29, "Reporting IBM z Advanced Workload Analysis Reporter (IBM zAware) problems to IBM," on page 281. If you do not have a maintenance contract with IBM, use the questions at the end of that topic to collect the information that you might need to report the problem to your hardware maintenance provider.

AIFB0003I **IBM zAware successfully assigned priming data from system *system_name* to sysplex *sysplex_name*.**

Explanation: Through the **Priming Data** tab of the IBM zAware graphical user interface (GUI), an administrator submitted a request to assign priming data for the monitored client identified by *system_name* to the sysplex *sysplex_name*. IBM zAware successfully completed the assignment request.

System action: IBM zAware recycles the analytics engine so that your changes take effect. IBM zAware is ready to use this priming data to build a model of behavior for the monitored client that is uniquely identified by its sysplex and system name.

Response: None required. To build a model for this system:

1. Navigate to the Training Sets page and select the system.
2. From the Actions list, select **Request Training** to submit a request to build a model of system behavior.

AIFB0004E **While attempting to assign priming data from system *system_name* to sysplex *sysplex_name*, IBM zAware encountered an error with reason code *code*. Search problem reporting databases for a fix for this error. If no fix exists, report this message ID and reason code to IBM Support.**

Explanation: Through the **Priming Data** tab of the IBM zAware graphical user interface (GUI), an administrator submitted a request to assign priming data for the monitored client identified by *system_name* to the sysplex *sysplex_name*. IBM zAware could not successfully complete the request because it encountered an error indicated by the reason code *code*.

System action: On the **Priming Data** tab of the IBM zAware GUI, the system remains in the **Priming data by systems** list. IBM zAware retains the priming data for this system but cannot use it to build a model.

Response: Search problem reporting databases for a fix for this error. If no fix exists, request IBM support by generating a Type V Viewable PMH (PMV) record to report this message ID and reason code.

For additional information, see Chapter 29, "Reporting IBM z Advanced Workload Analysis Reporter (IBM zAware) problems to IBM," on page 281. If you do not have a maintenance contract with IBM, use the questions at the end of that topic to collect the information that you might need to report the problem to your hardware maintenance provider.

AIFB0005I **IBM zAware successfully modified the system topology by moving system *system_name* from sysplex *source_sysplex_name* to sysplex *target_sysplex_name*.**

Explanation: Through the **Topology** tab of the IBM zAware graphical user interface (GUI), an administrator submitted a request to move the monitored client identified by *system_name* from the sysplex of which it is currently a member (*source_sysplex_name*) to a different sysplex in the topology (*target_sysplex_name*). IBM zAware successfully completed the request.

System action: IBM zAware recycles the analytics engine so that your changes take effect. When the **Topology** tab is refreshed, the system is listed under the target sysplex (*target_sysplex_name*).

Response: None required.

AIFB0006E **While attempting to move system *system_name* from sysplex *source_sysplex_name* to sysplex *target_sysplex_name*, IBM zAware encountered an error with reason code *code*. Search problem reporting databases for a fix for this error. If no fix exists, report this message ID and reason code to IBM Support.**

Explanation: Through the **Topology** tab of the IBM zAware graphical user interface (GUI), an administrator submitted a request to move the monitored client identified by *system_name* from the sysplex of which it is currently a member (*source_sysplex_name*) to a different sysplex in the topology (*target_sysplex_name*). IBM zAware could not successfully complete the request because it encountered an error indicated by the reason code *code*.

System action: In the **Topology** tab of the IBM zAware GUI, the system remains listed under the source sysplex (*source_sysplex_name*).

Response: Search problem reporting databases for a fix for this error. If no fix exists, request IBM support by generating a Type V Viewable PMH (PMV) record to report this message ID and reason code.

For additional information, see Chapter 29, "Reporting IBM z Advanced Workload Analysis Reporter (IBM zAware) problems to IBM," on page 281. If you do not have a maintenance contract with IBM, use the questions at the end of that topic to collect the information that you might need to report the problem to your hardware maintenance provider.

AIFB0007E **IBM zAware could not successfully create a backup copy of its database. Contact IBM Support.**

Explanation: IBM zAware attempted to but could not successfully create a backup copy of its database. This failure might be caused by a database integrity problem.

System action: Licensed Internal Code automatically gathers and sends diagnostic information to IBM if the "Call Home" feature is enabled on the IBM zAware host system.

Response: Contact IBM Support to diagnose the cause of this problem.

For additional information, see Chapter 29, "Reporting IBM z Advanced Workload Analysis Reporter (IBM zAware) problems to IBM," on page 281. If you do not have a maintenance contract with IBM, use the questions at the end of that topic to collect the information that you might need to report the problem to your hardware maintenance provider.

AIFB0008W **IBM zAware could not run a database backup operation because of a storage shortage. IBM zAware periodically attempts to rerun the backup operation. If the problem persists, contact IBM Support.**

Explanation: IBM zAware could not run a database backup operation because the IBM zAware partition does not have a sufficient amount of available storage. This storage shortage might be a temporary condition.

System action: IBM zAware periodically attempts to rerun the backup operation.

Response: If possible, increase the amount of storage that is available to the IBM zAware partition. If the problem persists, contact IBM Support for advice and possible corrective actions.

AIFB0009I **IBM zAware has detected integrity problems in its database, and is attempting to restore it by using a backup copy. No analytics processing can take place until this restore operation completes successfully.**

Explanation: IBM zAware detected integrity problems in its database, and started a synchronous restore operation. No analytics processing can take place until this restore operation completes successfully.

System action: The duration of this operation varies, depending on the size and complexity of the database. When IBM zAware finishes the restore operation, it issues either message AIFB0010I to report successful completion, or message AIFB0011E to report that the database could not be restored.

Response: To determine the result of the restore operation, check the **Notifications** page for either message AIFB0010I or message AIFB0011E.

AIFB0010I **IBM zAware successfully restored its database, which contains analysis data up to *timestamp*. Analytics processing can resume.**

Explanation: IBM zAware successfully resolved a database integrity problem by restoring the database

from a backup copy. The timestamp indicates, in Coordinated Universal Time (UTC), the date and time when the last restored data transaction was originally recorded in the database. Some instrumentation data and training models might have been lost during the restore operation.

System action: IBM zAware can resume its operations.

Response: None required. If the timestamp indicates that some data might have been lost during the restore operation, and you consider that missing data to be critical for analyzing monitored clients, consider resending log data from the monitored clients and rebuilding models. If you need instructions for sending log data and rebuilding models, see the appropriate configuration topic for the type of monitored client in Part 4, "Configuring IBM zAware and its monitored clients," on page 93.

AIFB0011E IBM zAware attempted but failed to restore its database from a backup copy. No analytics processing can take place until database integrity can be restored. Contact IBM Support.

Explanation: IBM zAware attempted but failed to restore its database from a backup copy. No analytics processing can take place until database integrity can be restored.

System action: Licensed Internal Code automatically gathers and sends diagnostic information to IBM if the "Call Home" feature is enabled on the IBM zAware host system.

Response: Contact IBM Support to diagnose the cause of this problem.

For additional information, see Chapter 29, "Reporting IBM z Advanced Workload Analysis Reporter (IBM zAware) problems to IBM," on page 281. If you do not have a maintenance contract with IBM, use the questions at the end of that topic to collect the information that you might need to report the problem to your hardware maintenance provider.

AIFB0012W IBM zAware could not run a restore operation after detecting integrity problems in its database. No analytics processing can take place until database integrity can be restored. Contact IBM Support.

Explanation: IBM zAware detected integrity problems in its database but could not run a restore operation, possibly because it could not find a backup copy of its database, or the backup copy is damaged. No analytics processing can take place until database integrity can be restored.

System action: Licensed Internal Code automatically gathers and sends diagnostic information to IBM if the "Call Home" feature is enabled on the IBM zAware host system.

Response: Contact IBM Support to diagnose the cause of this problem.

For additional information, see Chapter 29, "Reporting IBM z Advanced Workload Analysis Reporter (IBM zAware) problems to IBM," on page 281. If you do not have a maintenance contract with IBM, use the questions at the end of that topic to collect the information that you might need to report the problem to your hardware maintenance provider.

AIFB0013I IBM zAware database pruning is in progress.

Explanation: IBM zAware has started to remove obsolete data from its database. This data includes instrumentation data received from monitored systems, and models for monitored systems or groups.

System action: To identify and remove obsolete data for each type of monitored system, IBM zAware uses the retention times in effect for instrumentation data and training models, which IBM zAware administrators can view or modify on the **Administration** > **Configuration** > **Analytics** tab for each system type.

Response: None.

AIFB0014I IBM zAware successfully completed the removal of obsolete instrumentation and training data from its database.

Explanation: IBM zAware successfully completed a database pruning process, through which it removed obsolete instrumentation data received from monitored systems, and obsolete models for monitored systems or groups.

System action: After successful completion of this process, IBM zAware creates a backup copy of the pruned database. IBM zAware issues message AIFB0010I to report the successful completion of the database backup, or other AIFB messages if errors were encountered.

Response: None.

AIFB0015I IBM zAware successfully completed the removal of obsolete analysis results for monitored systems.

Explanation: IBM zAware successfully completed a daily pruning process, through which it removed obsolete analysis results that were generated for monitored systems or groups.

System action: To identify and remove obsolete data for each type of monitored system, IBM zAware used the retention times in effect for analysis results, which IBM zAware administrators can view or modify on the

Administration > **Configuration** > **Analytics** tab for each system type.

Response: None.

AIFB0016E IBM zAware encountered an unexpected error while attempting to complete a pruning process to remove obsolete data. Contact IBM Support.

Explanation: IBM zAware attempted but failed to complete a pruning process through which obsolete data is removed, according to the retention settings in effect. IBM zAware initiated one of the following pruning processes:

- A database pruning process, through which IBM zAware removes obsolete instrumentation data received from monitored systems, and obsolete models for monitored systems or groups. In this case, the error is most likely the result of a database integrity problem.
- A daily pruning process, through which IBM zAware removes obsolete analysis results that were generated for monitored systems or groups.

System action: Licensed Internal Code automatically gathers and sends diagnostic information to IBM if the "Call Home" feature is enabled on the IBM zAware host system. If the error occurred during the database pruning process, no analytics processing can take place until the error can be resolved.

Response: To determine which type of pruning process was attempted, go to the **Notifications** page and look for recent AIFB messages; then contact IBM Support.

- If you do not find a recent AIFB0013I message that indicates the start of a database pruning process, the attempted process is the daily pruning process.
- If you find a recent AIFB0013I message, look for additional AIFB messages that are related to database integrity problems.

For additional information, see Chapter 29, "Reporting IBM z Advanced Workload Analysis Reporter (IBM zAware) problems to IBM," on page 281. If you do not have a maintenance contract with IBM, use the questions at the end of that topic to collect the information that you might need to report the problem to your hardware maintenance provider.

AIFB0017W IBM zAware did not run a database pruning process because the preconditions for that process were not met.

Explanation: IBM zAware did not run a database pruning process because required preconditions were not met.

System action: Licensed Internal Code automatically gathers and sends diagnostic information to IBM if the "Call Home" feature is enabled on the IBM zAware host system.

Response: Contact IBM Support.

For additional information, see Chapter 29, "Reporting IBM z Advanced Workload Analysis Reporter (IBM zAware) problems to IBM," on page 281. If you do not have a maintenance contract with IBM, use the questions at the end of that topic to collect the information that you might need to report the problem to your hardware maintenance provider.

AIFB0018I IBM zAware removed system *system_name* at the request of an administrator.

Explanation: Through the **Administration** > **Configuration** > **Topology** tab, an administrator selected one or more monitored systems, and submitted a request to remove them from the IBM zAware topology. IBM zAware processes the removal request asynchronously, and issues this message for each system, identified by *system_name*, when the removal operation has completed.

System action: IBM zAware successfully removed the systems from the topology, and no longer includes them in GUI displays. IBM zAware also removes the data that is associated with each removed system, including:

- Current instrumentation data and priming data, if any.
- The system model, or the group model only when no other systems remain in the group.
- Analysis results for the selected system.

A sysplex or model group remains in the topology, even after all of the systems in the sysplex or model group are removed.

Response: None.

AIFF0001W The storage configuration is not complete. For its operation, IBM zAware requires continuous access to a set of storage devices. To assign storage devices, click Add and Remove Devices.

Explanation: An administrator has successfully logged in to the IBM zAware GUI and navigated to the **Administration** > **Configuration** > **Data Storage** tab. This warning message is displayed because no administrator has added storage devices for IBM zAware to use.

System action: IBM zAware analytics and other operations are not available until an administrator successfully adds storage devices through the **Data Storage** tab.

Response: On the **Data Storage** tab, use **Add and Remove Devices** to add storage devices. To avoid the potential loss of critical system and application data on

storage devices that are connected to the IBM zAware host system, make sure that you use the GUI to assign only those storage devices that are intended for IBM zAware use. Check with your storage administrator to make sure you select the appropriate storage devices to assign to the IBM zAware server.

AIFF0002W **When you click OK, IBM zAware formats and initializes any storage devices to be added; this process overwrites any existing data stored on those devices. To avoid the potential loss of critical system and application data on storage devices connected to the IBM zAware host system, make sure that you use the GUI to assign only those storage devices that are intended for IBM zAware use.**

Explanation: An administrator has submitted a request to add selected storage devices through the **Administration** > **Configuration** > **Data Storage** tab. This warning message is displayed as a reminder that data loss is possible if the IBM zAware server and other partitions on the host system have not been configured correctly. The recommended practice is to configure the IBM zAware partition such that it has access to only those channel path identifiers (IDs), control units, and I/O devices that are required for network connectivity and storage. If this practice is not followed, an IBM zAware administrator might inadvertently assign storage devices that are in use by other partitions.

System action: IBM zAware waits for an administrator response before processing the request.

Response: Check with your storage administrator to make sure you select the appropriate storage devices to assign to the IBM zAware server.

AIFF0003W **When you click OK, IBM zAware formats and initializes any storage devices to be added; this process overwrites any existing data stored on those devices. If you are adding a device that contains a backup copy of IBM zAware data for an alternate instance of the IBM zAware server, use the Preserve data option to ensure that IBM zAware does not overwrite the data when adding the device.**

Explanation: An administrator has submitted a request to add selected storage devices through the **Administration** > **Configuration** > **Data Storage** tab. This warning message is displayed as a reminder that, if the storage devices to be added contain backup data, the data will be lost unless the administrator selects the **Preserve data** option.

System action: IBM zAware waits for an administrator response before processing the request.

Response: Use the **Preserve data** option only when assigning a storage device that contains a backup copy of IBM zAware data.

AIFF0004W **When you are adding a device to replace a missing device, the size of the device to be added must match the size of the missing device, and also must contain a backup copy of the data that was stored on the missing device. Make sure that you select the Preserve data option when adding any replacement devices.**

Explanation: Ordinarily, IBM zAware formats and initializes devices to be added to its persistent storage, thus overwriting the data on those devices; using the **Preserve data** option ensures that IBM zAware does not overwrite the replicated data when adding the device.

System action: IBM zAware waits for an administrator response before processing the request.

Response: Use the **Preserve data** option when assigning a storage device that contains a backup copy of IBM zAware data.

AIFF0005E **A storage error has occurred. Go to the Notifications page to look for error messages that indicate problems with data storage devices. Depending on the problem, an IBM zAware administrator might need to remove the device and replace it by adding another device to the storage configuration.**

Explanation: IBM zAware detected a storage error for one or more storage devices. An IBM zAware administrator might need to resolve the error by using the **Add and Remove Devices** action on the **Configuration** > **Data Storage** tab.

System action: Depending on the type of storage error, IBM zAware analytics and other operations are not available until an administrator successfully corrects the problem.

Response: Go to the Notifications page to look for error messages that indicate problems with data storage devices. To avoid the potential loss of critical system and application data on storage devices connected to the IBM zAware host system, make sure that you use the GUI to assign only those storage devices that are intended for IBM zAware use. Check with your storage administrator to make sure you select the appropriate storage devices to assign to the IBM zAware server.

AIFF0006E **The storage configuration is not complete. For its operation, IBM zAware requires continuous access to a set of storage devices. An administrator must go to the Data Storage page to assign storage devices.**

Explanation: IBM zAware detected that the storage configuration is not complete. An IBM zAware administrator must add storage devices to correct this condition.

System action: IBM zAware analytics and other operations are not available until an administrator successfully adds storage devices.

Response: Go to the **Data Storage** tab and use **Add and Remove Devices** to add storage devices. To avoid the potential loss of critical system and application data on storage devices that are connected to the IBM zAware host system, make sure that you use the GUI to assign only those storage devices that are intended for IBM zAware use. Check with your storage administrator to make sure you select the appropriate storage devices to assign to the IBM zAware server.

AIFF0007E **Storage has not been configured. Contact an IBM zAware administrator to configure storage.**

Explanation: For its operation, IBM zAware requires continuous access to a set of storage devices. To add available devices to the IBM zAware storage configuration, you must use the master user ID or a user ID that is assigned to the IBM zAware Administrator role.

System action: IBM zAware analytics and other operations are not available until an administrator successfully adds storage devices.

Response: Contact an IBM zAware administrator to configure storage.

AIFF0008E **An error occurred while IBM zAware was processing a request to start the analytics engine. Retry the request; if the error persists, contact IBM Support.**

Explanation: An error occurred while IBM zAware was processing an administrator's request to start the analytics engine. This error might be the result of a temporary condition.

System action: None.

Response: Try submitting the start request again. If the error persists, contact IBM Support.

AIFF0009E **An error occurred when IBM zAware was processing a request to stop the analytics engine. Retry the request; if the error persists, contact IBM Support.**

Explanation: An error occurred while IBM zAware was processing an administrator's request to stop the analytics engine. This error might be the result of a temporary condition.

System action: IBM zAware analytics processing continues.

Response: Try submitting the stop request again. If the error persists, contact IBM Support.

AIFF0010E **An error occurred while IBM zAware was processing a request to remove selected messages from the Notifications page. Retry the request; if the error persists, you do not have to take any immediate action.**

Explanation: An error occurred while IBM zAware was processing an administrator's request to remove selected messages from the **Notifications** page. This error might be the result of a temporary condition.

System action: IBM zAware does not remove the selected messages.

Response: Retry the request; if the error persists, you do not have to take any immediate action. You can either contact IBM Support or deactivate the IBM zAware partition; deactivation results in the removal of all notification messages. For instructions, see "Deactivating the IBM zAware partition" on page 216.

AIFF0011W **IBM zAware has quiesced the analytics engine while processing a request that requires an update to its database. If a training request was in progress, that training request is canceled and replaced in the queue so it is the first request to be processed when the analytics engine is available again.**

Explanation: The *analytics engine* is the component of IBM zAware that manages the data the server receives from each monitored system. Management actions include reading, storing, processing, and analyzing the data, as well as determining when to build new models for each monitored system or model group. IBM zAware recycles the engine when you perform actions that require the engine to update the database where it stores the data. Such actions include assigning priming data, modifying the sysplex topology, or removing storage devices. Currently, IBM zAware has quiesced the analytics engine.

System action: Until the in-progress database update is completed, IBM zAware presents this message when a user attempts to use a GUI function that requires the use of the database.

Response: Depending on the length of time that the analytics engine is quiesced, you might need to reconnect monitored systems to the IBM zAware server after the engine is restarted. For more information about reconnecting monitored clients, see "Starting and stopping data collection for your monitored systems" on page 203.

AIFF0012E The login request is not valid. Enter a valid user ID and password and resubmit the request.

Explanation: The user ID and password combination did not match any user ID that is authorized to use IBM zAware.

System action: IBM zAware does not allow the user to access the GUI.

Response: Resubmit the request. If the problem persists, contact an IBM zAware administrator or, if you have the authority to do so, log in with the default master user ID and password.

AIFF0013E The login request failed and the LDAP server is not accessible. IBM zAware detected at least one of several login errors. Verify the credentials that you supplied and retry the login request. If the error persists, see the message description for more details.

Explanation: IBM zAware detected at least one of several possible login errors. If your installation is using an LDAP server for authenticating users, the LDAP server is not accessible. Other possible login errors include a user ID and password combination that does not match, or a user ID that is not defined as an authorized user of IBM zAware.

System action: IBM zAware does not allow the user to access the GUI.

Response: Resubmit the request with valid credentials. If the problem persists, contact an IBM zAware administrator or, if you have the authority to do so, log in with the default master user ID and password. If your installation is using an LDAP server for authenticating users, contact the LDAP administrator to resolve the problem.

AIFF0014E The LDAP server is not accessible.

Explanation: Your installation is using an LDAP server for authenticating users, and the LDAP server is not accessible. Any user IDs associated with the LDAP server cannot be authenticated until the server is accessible.

System action: IBM zAware does not allow the user to access the GUI.

Response: If you have the authority to do so, log in with the IBM zAware master user ID or with a user ID that is defined in a local file-based repository. Otherwise, contact your LDAP administrator to resolve the problem.

AIFF0017I Your changes were stored successfully. For new training period and training interval values to take effect for currently connected clients, you need to stop and reconnect those clients.

Explanation: Through one of the **Configuration** > **Analytics** tabs, an IBM zAware administrator changed one or more of the analytics configuration values, and clicked **Apply** to store the new values. Because each **Analytics** tab contains values that apply only to one particular type of monitored client, these new values apply to only monitored clients of that type.

System action: IBM zAware stores the new values and uses them for monitored clients of the type for which these new values apply.

Response: For new training period and training interval values to take effect for currently connected clients, you need to stop and reconnect any clients of the type for which these new values apply. For more information about reconnecting monitored clients, see "Starting and stopping data collection for your monitored systems" on page 203.

AIFF0018E An automatic or user-initiated training request failed for *name*.

Explanation: IBM zAware could not successfully process a training request for the system or group identified by *name*.

System action: If a model already exists, IBM zAware continues to use that model for analysis and attempts to retry the training request on the next day.

Response: If a model does not exist, retry the request. In either case, if the error persists, contact IBM Support.

AIFF0019E A user-initiated request failed to cancel training for *name*. Retry the request; if the error persists, you do not have to take any immediate action.

Explanation: An administrator attempted to cancel a training request for the system or group identified by *name*. IBM zAware could not successfully process the cancel request.

System action: The training request remains in the queue for processing. If a model for this group or system already exists, IBM zAware continues to use that model for analysis.

Response: Retry the request. If subsequent attempts to cancel training for this system also fail, contact IBM Support.

AIFF0020E **IBM zAware could not apply changes to LDAP settings because it could not connect to the LDAP server. Verify or correct the listed LDAP settings, then retry the request.** *list of one or more LDAP settings*

Explanation: Through the **Configuration** > **Security** > **LDAP Settings** tab, an IBM zAware administrator submitted a request to apply new or changed settings for the Lightweight Directory Access Protocol (LDAP) server. IBM zAware could not apply changes to LDAP settings because it could not connect to the LDAP server.

System action: The previous LDAP settings, if any, remain in effect.

Response: Verify or correct the LDAP settings listed at the end of the message text; for setting descriptions, see the online help for the **LDAP Settings** tab. Retry the request. If the error persists, contact the LDAP administrator to resolve the problem.

AIFF0021E **IBM zAware could not apply changes to LDAP settings because the SSL certificate specified for the LDAP server is not valid. Verify the value specified for the SSL certificate, then retry the request.**

Explanation: Through the **Configuration** > **Security** > **LDAP Settings** tab, an IBM zAware administrator submitted a request to apply new or changed settings for the Lightweight Directory Access Protocol (LDAP) server. IBM zAware could not apply changes to LDAP settings because the Secure Sockets Layer (SSL) certificate specified for the LDAP server is not valid.

System action: The previous LDAP settings, if any, remain in effect.

Response: Retry the request. If the error persists, contact the LDAP administrator to resolve the problem.

AIFF0022E **Persistent storage is not available because one or more devices (***total_missing***) are no longer available. IBM zAware operations cannot continue until corrective action is taken for each of the unavailable devices.**

Explanation: One or more storage devices that were previously in use by IBM zAware are no longer available. The variable in the message text, *total_missing*, indicates how many devices are no longer available.

System action: If one or more individual devices become unavailable through any method other than removal through the GUI, IBM zAware effectively loses access to all of its stored data and operations stop.

Response: See the message "AIFP0013E" on page 337 description for further details and corrective actions.

AIFF0023E **Duplicate storage devices are currently in use. IBM zAware operations cannot continue until the duplicate devices are removed.**

Explanation: An IBM zAware administrator added one or more storage devices that contain the same data as a device that is currently in use. Any duplicate devices are likely to be storage devices that are used to back up IBM zAware data.

System action: IBM zAware operations cannot continue until one of the duplicate devices is removed.

Response: Check the status values in the Data Storage Devices table to identify any duplicate storage devices, and use **Add and Remove Devices** to remove them from the IBM zAware storage configuration.

AIFF0024E **IBM zAware could not restore the excluded date (***date***) to the list of dates included in the next training period. Retry the request; if the error persists, you do not have to take any action.**

Explanation: Through the **Manage Model Dates** action on one of the **Administration** > **Training Sets** tabs, an IBM zAware administrator successfully submitted a request to exclude *date* from the next training for a system; then, at a later time, attempted to include the date by removing it from the "Excluded dates" list. IBM zAware could not restore the excluded date.

System action: IBM zAware continues to exclude data for that date from future training.

Response: Retry the request; if the error persists, you do not have to take any action. Based on your understanding of the system activities that occurred on the excluded date, and the number of days in the training period, determine whether the data from that date is necessary to accurately represent normal behavior for this system. If data from the excluded date is not required, no action is necessary. Otherwise, contact IBM Support.

AIFF0025E **IBM zAware did not successfully process the request to exclude the date (***date***) from the list of dates included in the next training period. Because the request failed, IBM zAware does not include message data for this date in the next or any subsequent training period.**

Explanation: Through the **Manage Model Dates** action on one of the **Administration** > **Training Sets** tabs, an IBM zAware administrator attempted to exclude *date* from the next training for a system by adding it to the "Excluded dates" list. IBM zAware did not successfully process the add request.

System action: Because the add request failed,IBM zAware does not include message data for this date in the next or any subsequent training period.

Response: If the add request failed because of a temporary condition, you can retry the request to exclude this date. If a subsequent add request succeeds, you have the option of restoring the date for use in future training periods. If the condition that caused the initial add request to fail is a permanent condition, IBM zAware cannot use any message data for this date in the next or any subsequent training period.

AIFF0026E **The login request failed because your user ID has not been assigned to an IBM zAware role. Ask an IBM zAware administrator to assign your user ID to the appropriate role.**

Explanation: To successfully log in to the IBM zAware GUI, your user ID must be defined in a supported repository and assigned to an IBM zAware role.

System action: IBM zAware found the user ID in the repository, but did not find it in the list of user IDs assigned to a role.

Response: Ask an IBM zAware administrator to assign your user ID to the appropriate role.

AIFF0027E **An administrator has not started the required migration process, or the migration process has not progressed to the point at which IBM zAware is fully functional. Go to the Migration tab to start or check on the progress of the automated migration.**

Explanation: For this IBM zAware server to become fully functional, an administrator must start the automated migration process, which must advance to the point at which a new database has been successfully created and activated. Until that point in the process, users who log in to the IBM zAware GUI are directed to a specific page, depending on the user's authorization. Administrators who log on to the GUI are directed to the **Migration** tab; other users are directed to the Systems page, which displays an informational message indicating that the migration is in progress. Except for the **Data Storage** tab and the **Notifications** page, other GUI pages cannot be used.

System action: Monitoring, analysis, and most administrative functions are not available until the automated migration process has advanced to the point at which a new database has been activated.

Response: Go to **Administration** > **Configuration** > **Migration** to start or check on the progress of the automated migration.

AIFF0028E **An administrator has not started the required migration process, or the migration process has not progressed to the point at which IBM zAware is fully functional.**

Explanation: For this IBM zAware server to become fully functional, an administrator must start the automated migration process, which must advance to the point at which a new database has been successfully created and activated. Until that point in the process, only the **Systems** and **Notifications** pages are available to user IDs that are mapped to the User role.

System action: Monitoring, analysis, and most administrative functions are not available until the automated migration process has advanced to the point at which a new database has been activated.

Response: None.

AIFF0029I **One or more systems or system groups no longer exist in the topology, and have been removed from the Analysis page display and the Change Source filter options. The removed systems or groups are:** *list_of_names*

Explanation: While updating an Analysis page display that a user has filtered by selecting specific systems or system groups, IBM zAware detected that one or more of the selected systems or groups no longer exist in the topology, and removed those systems or groups from the display and **Change Source** filter options. In the message text, the names of removed systems or groups are substituted for the *list_of_names* variable.

Administrators can remove systems or groups from the topology through the following actions:

- For individual z/OS systems, an administrator selects one or more systems on the **Configuration** > **Topology** tab and selects **Remove Systems** from the **Actions** list.
  - Moving a z/OS system from one sysplex to another has the same effect on the Analysis display as removing that z/OS system.
  - If an administrator removes all of the z/OS systems in a sysplex, IBM zAware also removes that sysplex from the topology.
- For z/OS sysplexes, an administrator selects one or more sysplexes on the **Configuration** > **Topology** tab and selects **Remove Systems** from the **Actions** list.
- For individual Linux systems, an administrator selects one or more systems on the **Configuration** > **Topology** tab and selects **Remove Systems** from the **Actions** list. Removing all of the Linux systems in a model group does not remove the model group itself.
- For an individual Linux model group, an administrator selects the model group on the

Systems > **Model Groups** tab and selects **Remove Group** from the **Actions** list. Renaming an existing model group has the same effect on the Analysis display as removing that model group.

System action: IBM zAware detects that a system or group no longer exists in the topology, it removes data associated with the system or group from the display on the Analysis page, and removes the name of the system or group from the **Change Source** filter options. If all of the systems or groups that the user selected to filter the display have been removed, IBM zAware refreshes the Analysis page display with the default content of all systems and groups.

Response: None.

AIFF0030I **IBM zAware successfully restored the previously saved configuration data. If you want to verify the data, go to the Analytics or Security tabs and check the displayed content.**

Explanation: Through the **Administration** > **Configuration** > **Utilities** tab, an administrator submitted a request to restore a saved copy of IBM zAware configuration data.

System action: The saved configuration data is used to repopulate the appropriate **Configuration** tabs, and the IBM zAware server is restarted.

Response: None.

AIFF0031I **IBM zAware successfully saved a copy of configuration data on in-use devices in the IBM zAware storage configuration.**

Explanation: Through the **Administration** > **Configuration** > **Utilities** tab, an administrator submitted a request to save a copy of IBM zAware configuration data.

System action: IBM zAware stored the copy on in-use devices in the IBM zAware storage configuration.

Response: None.

AIFF0032E **The system** *system_name* **does not exist. Refresh the page to display the most recent system listing.**

Explanation: On the Manage Model Dates page, an administrator selected a system that is displayed in the "Training field" list but is no longer a member of the model group. The name of the selected system is identified by the variable *system_name*.

System action: None.

Response: To display a current list of systems for the model group, refresh the Manage Model Dates page either by refreshing the browser display, or by returning to the **Administration** > **Training Sets** > **Model Groups** tab and reselecting the **Manage Model Dates** action.

AIFG0001E **IBM zAware encountered an internal error. Search problem reporting databases for a fix for this error. If no fix exists, report this message ID to IBM Support.**

Explanation: IBM zAware encountered an internal error.

System action: For most internal errors, Licensed Internal Code automatically gathers and sends diagnostic information to IBM if the "Call Home" feature is enabled on the IBM zAware host system.

Response: On the Hardware Management Console (HMC) for the IBM zAware host system, check for outstanding hardware messages that indicate a "Call Home" call associated with the IBM zAware partition. If a service call has not been sent automatically, search problem reporting databases for a fix for this error. If no fix exists, request IBM support by generating a Type V Viewable PMH (PMV) record to report this message ID.

For additional information, see Chapter 29, "Reporting IBM z Advanced Workload Analysis Reporter (IBM zAware) problems to IBM," on page 281. If you do not have a maintenance contract with IBM, use the questions at the end of that topic to collect the information that you might need to report the problem to your hardware maintenance provider.

AIFM0001W **Additional storage is required for the automated migration process to complete successfully. Go to the Data Storage tab and add more storage devices, then return to the Migration tab and click Retry.**

Explanation: IBM zAware has detected an insufficient amount of in-use storage for the automated migration process to complete successfully.

System action: The automated migration process stops until an administrator responds.

Response:

1. Go to the **Data Storage** tab to add enough storage devices. IBM testing experience indicates that you need three times the amount of in-use storage to successfully complete the migration.

   Attention: Do not use **Add All** if any of the available storage devices are shared. If a device is shared and in use by another application, data will be lost or overwritten if the IBM zAware server formats the device.

   To avoid the potential loss of critical system and application data on storage devices that are

connected to the IBM zAware host system, make sure that you use the GUI to assign only those storage devices that are intended for IBM zAware use.

2. When the "Total capacity (GB)" value meets or exceeds the required amount, return to the **Migration** tab and click **Retry**.

AIFM0002E The database consistency check failed. Contact IBM Support for advice and possible corrective actions.

Explanation: IBM zAware has detected a consistency error with the database on the in-use set of storage devices for the IBM zAware server on the zEC12 or zBC12 host system.

System action: The automated migration process stops until an administrator responds.

Response: Contact IBM Support for advice and possible corrective actions. Depending on the outcome of the discussion, take corrective action and click **Retry**, or click **Skip** to proceed with the next migration process.

AIFM0003I IBM zAware has successfully activated the new database, which is primed with migrated operations data. Now you can connect monitored clients to the new IBM zAware server.

Explanation: IBM zAware has successfully created the new database and primed it with migrated topology and training set management data. Administrators and users can use the Analysis page to view migrated analysis results for any monitored system in the topology. Current analysis results, however, are not available until that system is connected to the new IBM zAware server on the new host system, and the model is successfully rebuilt.

System action: The new IBM zAware server is ready to receive message data from monitored clients.

Response: Connect z/OS monitored clients to the new IBM zAware server on the new host system. To connect the z/OS monitored clients, issue the SETLOGR command. For example:

```
SETLOGR FORCE,ZAICONNECT,LSN=SYSPLEX.OPERLOG
```

To verify that the z/OS system has successfully connected to the new server, go to the **Status** tab on the Systems page, and look for the system name in the IBM zAware Monitored System Data Suppliers table. As you connect z/OS systems in the topology to the new server after the automated migration completes, IBM zAware detects whether a model exists for that connected system. If a model does not exist but a sufficient amount of instrumentation data is available for the connected system, IBM zAware automatically queues a training request and uses migrated data to rebuild the model.

AIFM0004I IBM zAware has successfully completed the migration cleanup. The new IBM zAware server is primed with migrated operations data, and is ready for use.

Explanation: IBM zAware has successfully cleaned up any unnecessary artifacts of the automated migration process.

System action: The new IBM zAware server is ready to receive message data from monitored clients.

Response: None required.

AIFM0005E The database pruning operation failed. Contact IBM Hardware Support for advice and possible corrective actions.

Explanation: During the automated migration process, IBM zAware could not successfully remove outdated information from the database. IBM zAware uses analytics retention times to determine which information is to be removed.

System action: The automated migration process stops.

Response: Contact IBM Support for advice and possible corrective actions.

AIFM0006I IBM zAware has found an issue with your existing database and is currently working to fix it through database repairs. These repairs might take some time to complete.

Explanation: During the automated migration process, IBM zAware detected one or more problems with the existing database, and is attempting to repair it before proceeding with the migration.

System action: On the Migration tab, IBM zAware indicates progress for this step.

Response: None required.

AIFM0099E The database migration failed with an unexpected error. Contact IBM Support.

Explanation: While performing a database migration, IBM zAware encountered an unexpected error.

System action: The automated migration process stops.

Response: Contact IBM Support.

AIFP0001E **IBM zAware encountered an internal error with reason code *code*. Search problem reporting databases for a fix for this error. If no fix exists, report this message ID and reason code to IBM Support.**

Explanation: IBM zAware encountered an internal error indicated by reason code *code*.

System action: For most internal errors, Licensed Internal Code automatically gathers and sends diagnostic information to IBM if the "Call Home" feature is enabled on the IBM zAware host system.

Response: On the Hardware Management Console (HMC) for the IBM zAware host system, check for outstanding hardware messages that indicate a "Call Home" call associated with the IBM zAware partition. If a service call has not been sent automatically, search problem reporting databases for a fix for this error. If no fix exists, request IBM support by generating a Type V Viewable PMH (PMV) record to report this message ID and reason code.

For additional information, see Chapter 29, "Reporting IBM z Advanced Workload Analysis Reporter (IBM zAware) problems to IBM," on page 281. If you do not have a maintenance contract with IBM, use the questions at the end of that topic to collect the information that you might need to report the problem to your hardware maintenance provider.

AIFP0002E **IBM zAware could not initialize the physical volume *device_id* because of a disk or I/O error. If necessary, go to the Data Storage page in the IBM zAware GUI to select another device to add.**

Explanation: Through the IBM zAware graphical user interface (GUI), an administrator used the **Add and Remove Devices** function on the Data Storage page to add one or more devices to the IBM zAware storage configuration. IBM zAware could not add the device identified by *device_id* because of a disk or input/output (I/O) error.

System action: IBM zAware does not add *device_id* to the list of in-use devices. If the add request included multiple storage devices, some of which remain to be processed, IBM zAware continues to process the request for those devices.

Response: On the Support Element (SE) for the IBM zAware host system, check for hardware messages that indicate I/O or other problems related to this storage device. If possible, correct the error and retry the add request. Otherwise, select another device to add to the IBM zAware storage configuration.

AIFP0003I **IBM zAware started processing a request to add storage device *device_id*.**

Explanation: Through the IBM zAware graphical user interface (GUI), an administrator used the **Add and Remove Devices** function on the Data Storage page to add one or more devices to the IBM zAware storage configuration. This message indicates that IBM zAware has started to process the request to add the storage device identified by *device_id*.

System action: IBM zAware continues to process the add request, and issues message AIFP0004I after successfully adding the device.

Response: None.

AIFP0004I **IBM zAware successfully added storage device *device_id*.**

Explanation: Through the IBM zAware graphical user interface (GUI), an administrator used the **Add and Remove Devices** function on the Data Storage page to add one or more devices to the IBM zAware storage configuration. This message indicates that IBM zAware successfully processed the request to add the storage device identified by *device_id*.

System action: IBM zAware uses the device for processing and storage of analysis results. If the add request included multiple storage devices, some of which remain to be processed, IBM zAware continues to process the request for those remaining devices.

Response: None.

AIFP0005I **IBM zAware rejected a request to add or remove a storage device because another storage operation is in progress. When the operation completes, retry the request.**

Explanation: Through the IBM zAware graphical user interface (GUI), an administrator used the **Add and Remove Devices** function on the Data Storage page to add or remove one or more devices. IBM zAware could not process this request because a prior request to add or remove storage devices is still in progress.

System action: IBM zAware continues to process the prior request.

Response: Resubmit the request after IBM zAware completes processing the prior request.

AIFP0006E **IBM zAware detected that at least one in-use persistent storage device, *device_id*, is missing. Try reattaching the device identified by *device_id* and reactivating the IBM zAware partition. See the AIFP0013E message description to determine whether additional corrective actions are necessary.**

Explanation: IBM zAware detected that an in-use persistent storage device, identified by *device_id*, is no longer attached to the IBM zAware partition. It might have been removed or disconnected through storage operations that are not provided through the Data Storage page of the IBM zAware GUI, such as:

- Replacing the I/O definition file (IODF) for the host system with an IODF that does not contain the in-use storage devices for IBM zAware.
- Using the Support Element (SE) to take offline one or more channel paths (CHPIDs) for storage devices or for the network through which those devices are connected to the IBM zAware partition.

Additional in-use storage devices also might be missing, as indicated by message "AIFP0013E" on page 337.

System action: When an in-use storage device becomes unavailable, IBM zAware operations stop. To a user of the IBM zAware GUI, the GUI appears to hang. On the SE for the IBM zAware host system, hardware messages indicate input/output (I/O) problems that are related to the loss of access to physical storage devices.

Response: If possible, try to reattach the device identified by *device_id* and reactivate the IBM zAware partition.

If you cannot reattach the device, see the "AIFP0013E" on page 337 message description for additional information and instructions for correcting the IBM zAware storage configuration.

AIFP0007I **During activation, IBM zAware was not able to reconnect to the previously configured LDAP server *hostname*:*port*. Only the master user ID and locally defined users are able to log in to the GUI until the LDAP server is available. An administrator must reconnect IBM zAware by reapplying the LDAP configuration through the LDAP Settings (Configuration > Security > LDAP Settings).**

Explanation: When the IBM zAware partition is reactivated at any time after its initial configuration, IBM zAware attempts to reconnect to the previously configured Lightweight Directory Access Protocol (LDAP) server. This reconnection attempt fails if the LDAP server is not available or the LDAP configuration has changed. The previously configured LDAP server is identified by *hostname* and the port number (*port*) used for communication between the LDAP and IBM zAware servers. Without access to the LDAP repository, you can log in to the IBM zAware graphical user interface (GUI) using the master user ID and password, which are specified through the activation profile for the IBM zAware partition. If you have set up a local file-based repository of user IDs, you can use one of those locally defined IDs only if you previously assigned an IBM zAware "Administrator" role for them through the Role Mapping page in the GUI.

System action: Activation of IBM zAware continues.

Response: To log in to the IBM zAware GUI after activation completes, you must use either the master user ID, or a locally defined user ID that previously was mapped to an "Administrator" role.

To enable access for users that are defined in the LDAP server:

1. Determine why the LDAP server was not accessible and correct the problem.
2. Through the IBM zAware, click the **LDAP Settings** tab on the Security page and verify the values shown for the general and group LDAP settings. Click **Apply** to reconnect IBM zAware to the LDAP server.

AIFP0008E **IBM zAware was unable to remove storage device *device_id*. Report this message ID to IBM Support.**

Explanation: Through the IBM zAware graphical user interface (GUI), an administrator used the **Add and Remove Devices** function on the Data Storage page to remove one or more devices to the IBM zAware storage configuration. IBM zAware encountered an internal error while processing the request to remove the storage device identified by the variable *device_id*.

System action: IBM zAware does not remove *device_id* from the list of in-use devices. If the remove request included multiple storage devices, some of which remain to be processed, IBM zAware continues to process the request for those devices.

Response: On the Hardware Management Console (HMC) for the IBM zAware host system, check for outstanding hardware messages that indicate a "Call Home" call associated with the IBM zAware partition. If a service call has not been sent automatically, request IBM support by generating a Type V Viewable PMH (PMV) record to report this message ID and reason code.

For additional information, see Chapter 29, "Reporting IBM z Advanced Workload Analysis Reporter (IBM zAware) problems to IBM," on page 281. If you do not have a maintenance contract with IBM, use the questions at the end of that topic to collect the information that you might need to report the problem to your hardware maintenance provider.

AIFP0009I **IBM zAware started processing a request to remove storage device *device_id* from the set of in-use devices.**

Explanation: Through the IBM zAware graphical user interface (GUI), an administrator used the **Add and**

Remove Devices function on the Data Storage page to remove one or more devices to the IBM zAware storage configuration. This message indicates that IBM zAware has started to process the request to remove the storage device identified by *device_id*.

System action: IBM zAware continues to process the remove request. IBM zAware issues message AIFP0010I when it has successfully removed the device. If it cannot remove the device at this time, IBM zAware places the device in "Pending Removal" status.

Response: Go to the **Data Storage** tab in the IBM zAware GUI to check the current status of the device. Because IBM zAware processes storage requests asynchronously, you need to click **Refresh** at least once to display the status of the device in the **Data Storage Devices** table. When IBM zAware is processing a remove request, the device status is "Being Removed"; when the process completes, the status is either "Available" or "Pending Removal". If IBM zAware placed the device in "Pending Removal" status, click **Apply Pending Removals** on the **Data Storage** tab to remove the device.

AIFP0010I **IBM zAware successfully removed storage device** *device_id* **from the set of in-use devices.**

Explanation: Through the IBM zAware graphical user interface (GUI), an administrator used the **Add and Remove Devices** function on the Data Storage page to remove one or more devices to the IBM zAware storage configuration. This message indicates that IBM zAware successfully processed the request to remove the storage device identified by *device_id*.

System action: IBM zAware moves data from *device_id* to the remaining in-use storage devices, and no longer uses *device_id* for processing and storage of analysis results. If the remove request included multiple storage devices, some of which remain to be processed, IBM zAware continues to process the request for those remaining devices.

Response: None required. At this point, a storage administrator can detach *device_id* from the IBM zAware partition without disrupting IBM zAware operations.

AIFP0011E **During activation, IBM zAware encountered an internal error while attempting to reconnect to the previously configured LDAP server** *hostname*:*port*. **Only the master user ID and locally defined users are able to log in to the GUI until the LDAP server is available. An administrator must reconnect IBM zAware by reapplying the LDAP configuration through the LDAP Settings page.**

Explanation: When the IBM zAware partition is reactivated at any time after its initial configuration, IBM zAware attempts to reconnect to the previously configured Lightweight Directory Access Protocol (LDAP) server. During a reconnection attempt, IBM zAware encountered an internal error. The previously configured LDAP server is identified by *hostname* and the port number (*port*) used for communication between the LDAP and IBM zAware servers. Without access to the LDAP repository, you can log in to the IBM zAware graphical user interface (GUI) using the master user ID and password, which are specified through the activation profile for the IBM zAware partition. If you have set up a local file-based repository of user IDs, you can use one of those locally defined IDs only if you previously assigned an IBM zAware "Administrator" role for them through the Role Mapping page in the GUI.

System action: Activation of IBM zAware continues.

Response: To log in to the IBM zAware GUI after activation completes, you must use either the master user ID, or a locally defined user ID that previously was mapped to an "Administrator" role. Through the IBM zAware, click the **LDAP Settings** tab on the Security page and verify the values shown for the general and group LDAP settings. Click **Apply** to reconnect IBM zAware to the LDAP server.

If the problem persists:

- Search the problem reporting databases for a fix for this error.
- If no fix exists, contact IBM support by generating a problem management record (PMR) to report this message ID.

For additional information, see Chapter 29, "Reporting IBM z Advanced Workload Analysis Reporter (IBM zAware) problems to IBM," on page 281. If you do not have a maintenance contract with IBM, use the questions at the end of that topic to collect the information that you might need to report the problem to your hardware maintenance provider.

AIFP0012I **Persistent storage device** *device_id* **was configured and in use by IBM zAware, but is no longer required and has been removed from the list of in-use storage devices.**

Explanation: When your installation configures a primary IBM zAware partition and an alternate IBM zAware partition for switchover situations, only one IBM zAware server can be active at a given time but both servers must have access to the same data. To correctly configure persistent storage for primary and alternate IBM zAware partitions, your installation must define physically separate but equivalent sets of storage devices for each partition, and also set up replication to copy the content of the primary storage devices to the alternate storage devices. For data replication to be successful, the number of storage devices in the

primary set must match the number of devices in the alternate set. Additionally, each alternate device must be equivalent in size to the primary device.

Message AIFP0012I is issued when the currently active IBM zAware server takes corrective action to resolve a mismatch that results from the removal of a device from either the primary or the alternate set of storage devices. Consider the following example, for which a storage administrator has defined two sets of 3390 DASD for IBM zAware:

- Devices 3001 through 3005 are intended for the exclusive use of the primary server of IBM zAware.
- Equivalent devices 9111 through 9115 are intended for the exclusive use of the alternate server.
- Through the IBM zAware graphical user interface (GUI), the administrator initially configures the primary server to use only devices 3002 and 3004.
- Through other tools or interfaces, the administrator also sets up replication such that the content of device 3002 is periodically copied to device 9112, and the content of device 3004 is copied to device 9114. For successful replication, device 9112 must be the same size as device 3002; similarly, device 9114 must be equivalent to device 3004.
- To successfully use the alternate server of IBM zAware when the primary server is no longer available, an administrator must use the GUI to initially configure the alternate server to use both devices 9112 and 9114, the devices that contain the data copied from their equivalent primary in-use devices. To configure the alternate server, the administrator:
  1. Deactivates the primary partition of IBM zAware, and activates the alternate partition.
  2. Uses the IBM zAware GUI to add devices 9112 and 9114 to the storage configuration for the alternate server.
  3. Deactivates the alternate partition of IBM zAware, and activates the primary partition.

After the initial configuration of both the primary and alternate IBM zAware servers is complete, suppose that an administrator uses the GUI to remove device 3004 from the primary set because device 3002 alone has sufficient capacity for the data that IBM zAware needs to store. This action causes a mismatch between the primary set and alternate set of storage devices that IBM zAware detects and corrects only when the alternate partition is activated:

1. When the alternate partition is activated, the alternate IBM zAware server detects that device 9114, which contains a copy of the data from removed device 3004, is still listed as an in-use device in the alternate set of storage.
2. To correct this mismatch, the alternate IBM zAware server automatically removes device 9114 from its storage configuration and issues message AIFP0012I with device ID 9114.

Although this example assumes that the alternate server detected the removal of a device from the primary set of storage devices, the reverse is also possible; when the primary partition is activated, the primary IBM zAware server can detect and correct a mismatch that results from the removal of a device from the alternate set of storage devices.

System action: The active IBM zAware server does not use the removed device identified by *device_id*.

Response: None. If you need additional information about configuring storage for primary and alternate IBM zAware partitions, see "Example: Storage configuration for multiple IBM zAware partitions" on page 66.

AIFP0013E **Persistent storage is not available because one or more devices (*total_missing*) are no longer available. IBM zAware operations cannot continue until corrective action is taken for each of the unavailable devices.**

Explanation: One or more storage devices that were previously in use by IBM zAware are no longer available. The variable in the message text, *total_missing*, indicates how many devices are no longer available. In-use storage devices can become unavailable for these reasons:

- The device is no longer attached to the IBM zAware partition. It might have been removed or disconnected through storage operations that are not provided through the Data Storage page of the IBM zAware GUI, such as:
  - Replacing the I/O definition file (IODF) for the host system with an IODF that does not contain the in-use storage devices for IBM zAware.
  - Using the Support Element (SE) to take offline one or more channel paths (CHPIDs) for storage devices or for the network through which those devices are connected to the IBM zAware partition.

  Otherwise, the storage device might be disconnected or damaged such that the host system cannot connect to it.
- The number of devices that are in use by the primary IBM zAware server no longer matches the number of devices that have been added to the alternate IBM zAware server. When your installation configures primary and alternate IBM zAware servers, the content of the primary storage devices are copied to the alternate storage devices through the replication method of your choice. For replication to be successful, the number of storage devices in the primary set must match the number of devices in the alternate set.
- An administrator is in the process of migrating an existing IBM zAware server to a new host system. In this case, this message can be ignored.

System action: When an in-use storage device becomes unavailable, IBM zAware operations stop. To a user of the IBM zAware GUI, the GUI appears to hang. On the SE for the IBM zAware host system, hardware messages indicate input/output (I/O) problems that are related to the loss of access to physical storage devices.

Response: When access to in-use storage devices is lost and the GUI appears to hang, use the System Activity display for the IBM zAware partition to check processor utilization and processor weights for the IBM zAware partition. To access the System Activity display, use the **Monitors Dashboard** task in the Hardware Management Console (HMC) for the IBM zAware host system. If the partition is not using any cycles, deactivate and reactivate the partition.

- If the partition does not successfully reactivate or still appears to be unresponsive, report the problem to IBM Support by a generating a Type V Viewable PMH (PMV) record.
- If you can successfully log in to the IBM zAware GUI as an administrator, the GUI displays the Data Storage page with message AIFP0013E. Corrective measures depend on which circumstance caused the device to become unavailable. Before you take any action, however, you need to determine which device is no longer available and, if your installation is periodically backing up IBM zAware data, which equivalent device contains the replicated data.
  - If the device is no longer attached to the partition:
    - If possible, try to reattach the device and reactivate the IBM zAware partition.
    - If the device cannot be reattached, you must replace it with an equivalent device containing a backup copy of the data that was stored on the unavailable device. Use the **Add and Remove Devices** function on the Data Storage tab to replace the unavailable devices with their equivalents.
    - If your installation does not have backup copies of IBM zAware data, you must deactivate the IBM zAware partition and reconfigure the IBM zAware environment.
  - If IBM zAware detected a mismatch between the primary and alternate sets of storage devices:
    - If the mismatch resulted from the removal of a device, you do not need to take any corrective action.
    - If the mismatch resulted from the addition of a device, you need to determine which set— the primary or alternate— has the additional device, and add its equivalent to the other set. The size of a primary device and the size of its alternate device must match.

AIFP0015I IBM zAware rejected a request to add a persistent storage device, *device_id*, to the list of in-use devices. The request is rejected because the current IBM zAware storage configuration is missing one or more storage devices, as indicated by message AIFP0013E, and the requested device does not contain a backup copy of the data that was stored on one of those missing devices.

Explanation: When IBM zAware issues message AIFP0013E to indicate that one or more in-use storage devices are missing, an administrator cannot successfully add previously unused storage devices until all of the missing devices either have been reattached to the IBM zAware partition, or have been replaced with equivalent devices containing replicated data. The request to add the persistent storage device *device_id* is rejected because that device does not contain a backup copy of the data that was stored on one of the missing devices.

System action: IBM zAware does not add *device_id* to the list of in-use devices, and reissues message AIFP0013E. IBM zAware operations cannot continue until corrective action is taken for each of the missing devices.

Response: If you were attempting to replace a missing device with an equivalent device that contains replicated data, check the device ID to verify that you are adding the correct replacement device and retry the request. Otherwise, follow the instructions for message AIFP0013E to correct the IBM zAware storage configuration.

AIFP0016E Persistent storage device *device_id* is a duplicate of storage device *device_id*. IBM zAware operations cannot continue until one of these devices is removed through the Data Storage page of the IBM zAware GUI.

Explanation: When your installation sets up replication to have a backup copy of IBM zAware data available, a storage administrator must define physically separate but equivalent sets of storage devices:

- One set for IBM zAware to use during normal operations.
- A backup set to contain replicated data.

Although both sets of devices are displayed through the IBM zAware graphical user interface (GUI) as available for use, initially an administrator can add only the set of devices that are designated for normal operations. If the backup devices are added during initial configuration of IBM zAware, replication cannot be successful because IBM zAware will use the backup devices to store its data.

After successful replication, however, an administrator can add a backup device, but only when its equivalent in-use device is no longer part of the IBM zAware storage configuration. Message AIFP0016E is issued when an administrator uses the IBM zAware GUI to add a backup device that contains the same data as a device that is currently in use. The device IDs supplied in the message text indicate the device number of the backup and its equivalent in-use storage device.

Consider the following example, for which a storage administrator has defined two sets of 3390 DASD for IBM zAware:

- Devices 3001 through 3005 are reserved to contain data that IBM zAware requires for normal operations.
- Devices 9110 through 9115 are reserved to contain backup copies of the data from in-use devices 3001-3005.

1. Through the IBM zAware GUI, the administrator initially configures IBM zAware to use only devices 3002 and 3004.
2. Through other tools or interfaces, the administrator also sets up replication such that the content of device 3002 is periodically copied to device 9112, and the content of device 3004 is copied to device 9114.
3. After successful replication, if an administrator adds backup device 9112 to the storage configuration while device 3002 is in use, IBM zAware detects the duplicate devices and issues message AIFP0016E.

System action: IBM zAware operations cannot continue until one of these devices is removed through the **Configuration** > **Data Storage** tab of the IBM zAware GUI.

Response: With a user ID assigned to the Administrator role, use the **Add and Remove Devices** action on the **Data Storage** tab to remove one of the identified devices from the IBM zAware storage configuration.

AIFP0017I **The removal of storage device** *device_id* **is pending. To complete the removal of this device, select the Apply Pending Removals action on the Data Storage tab.**

Explanation: An IBM zAware administrator used the **Add and Remove Devices** action on the **Configuration** > **Data Storage** tab to remove the device identified by *device_id* from the storage configuration. IBM zAware was unable to immediately remove this device.

System action: In the Data Storage Devices table, IBM zAware displays the status of this device as Pending Removal.

Response: On the **Configuration** > **Data Storage** tab, select the **Apply Pending Removals** action on the Data Storage Devices table.

- If the removal operation completes successfully, IBM zAware changes the device status to Available.
- If the device remains in Pending Removal state after IBM zAware processes an administrator request to apply pending removals, the device status is displayed as Pending Removal (Remove Failed).

AIFP0018W *resource* **shortage.** *number* **percent usage detected. An administrator should determine whether additional capacity can be added, or usage reduced.**

Explanation: IBM zAware detected that the current usage of either storage or memory, as indicated by the variable *resource* in the message text, has reached a level at which corrective action might be required to avoid a critical shortage. The variable *number* indicates the percentage of storage or memory that is currently in use.

A storage shortage might require an administrator to take corrective action. A memory shortage can be a temporary condition that is the result of a relatively sudden increase in activity; for example, when IBM zAware is processing a training request in addition to processing message traffic from monitored clients. In this case, the memory shortage might be resolved without any intervention.

System action: IBM zAware continues operating, periodically reissuing this message until corrective action resolves the storage or memory shortage. If the shortage reaches a critical threshold, IBM zAware begins to issue message AIFP0019E, with greater frequency.

Response: Corrective action depends on the type of resource shortage.

For a storage shortage

- Consider reducing the number of monitored systems that are currently connected.
- Check the Data Storage Devices table on the Data Storage page to determine whether any additional devices are available. Use **Add and Remove Devices** to add devices to the IBM zAware storage configuration.
- Check the retention settings on the **Analytics** tab on the Configuration page. Reducing the retention times for instrumentation data, training models, or analysis results might help to prevent future storage shortages, but the effect is not immediate.

For a memory shortage

- Consider reducing the number of monitored systems that are currently connected.

- Check the Training Sets page to determine whether any training requests are in progress or queued for processing; if so, the increase in memory usage might be temporary. Check the Notifications page for additional occurrences of message AIFP0018W, or for message AIFP0019E. If the percent of memory in use continues to increase, consider cancelling one or more of the queued training requests.

AIFP0019E **Critical *resource* shortage. *number* percent usage detected. An administrator should determine whether additional capacity can be added, or usage reduced.**

Explanation: IBM zAware detected that the current usage of either storage or memory, as indicated by the variable *resource* in the message text, has reached a critical threshold. The variable *number* indicates the percentage of storage or memory that is currently in use.

A critical storage shortage requires an administrator to take immediate corrective action. In contrast, a memory shortage might be a temporary condition that is the result of a relatively sudden increase in activity; for example, when IBM zAware is processing a training request in addition to processing message traffic from monitored clients. In this case, the memory shortage might be resolved without any intervention.

System action: IBM zAware continues operating, frequently reissuing this message until corrective action resolves the storage or memory shortage. If the shortage exceeds this critical threshold, results are unpredictable.

Response: Corrective action depends on the type of resource shortage.

For a critical storage shortage

1. Immediately reduce the number of monitored systems that are currently connected.
2. Use **Add and Remove Devices** on the Data Storage page to add available devices to the IBM zAware storage configuration.
3. Check the retention settings on the **Analytics** tab on the Configuration page. Reducing the retention times for instrumentation data, training models, or analysis results might help to prevent future storage shortages, but the effect is not immediate.

For a critical memory shortage

- Consider reducing the number of monitored systems that are currently connected.
- Check the Training Sets page to determine whether any training requests are in progress or queued for processing; if so, the increase in memory usage might be temporary. Check the Notifications page for additional occurrences of message AIFP0019E. If the percent of memory in use continues to increase, consider cancelling one or more of the queued training requests.

For information about changing the memory resources that are defined for a logical partition, see *z Systems PR/SM Planning Guide*, SB10-7162.

AIFP0020E **Storage device *device_id* in the active configuration has become corrupt. IBM zAware cannot continue until all in-use storage devices have been replaced by a backup set of devices.**

Explanation: IBM zAware has detected that an in-use storage device, identified by *device_id*, has become corrupt and cannot be used. This error condition can occur when an in-use storage device is overwritten by a process running on another partition. To avoid this condition, always use the IBM zAware GUI to remove the device from the in-use storage configuration before allowing other partitions to access the device.

System action: IBM zAware operations stop, and cannot continue until all in-use storage devices have been replaced by a backup set of devices. The in-use devices are unmounted and, on the **Administration** > **Configuration** > **Data Storage** tab, are shown as available.

Response: With a user ID that is mapped to the Administrator role, navigate to the **Data Storage** tab in the IBM zAware GUI. Use the **Add and Remove Devices** function to replace the formerly in-use devices with their equivalent backup set of devices. Make sure that you select the "Preserve data" option when adding the backup devices, so IBM zAware does not overwrite the replicated data when adding the devices.

If your installation does not have backup copies of IBM zAware data, you must reconfigure the IBM zAware environment.

AIFP0021E **IBM zAware could not remove the storage device *device_id* because the remaining devices do not have sufficient space to store the data from the device to be removed. Add more storage devices to the configuration before removing this device.**

Explanation: Through the IBM zAware graphical user interface (GUI), an administrator used the **Add and Remove Devices** function on the Data Storage page to remove one or more devices from the IBM zAware storage configuration. IBM zAware could not remove the device identified by *device_id* because the remaining devices in the storage configuration do not have sufficient space to store the data from the device to be

removed. IBM zAware rejects any attempt to reduce storage below the amount that it is currently using.

System action: IBM zAware does not remove the device from the list of in-use devices, and indicates the device status as In Use (Remove Failed). If the remove request included multiple storage devices, some of which remain to be processed, IBM zAware continues to process the request for those devices.

Response: If you must remove the device from the storage configuration, use **Add and Remove Devices** to add more storage devices, then retry the request to remove the device.

AIFP0022E **The IBM zAware storage configuration does not contain enough space to create a swap file that is used to manage potential memory resource constraints. Add an additional** *storage_amount* **to the storage configuration.**

Explanation: IBM zAware automatically reserves a percentage of in-use storage to create a swap file that it uses only as necessary to manage memory resource constraints. The current amount of in-use storage is not sufficient to create the swap file. The variable *storage_amount* indicates the additional storage amount, in gigabytes (GB), that IBM zAware requires for the swap file.

System action: IBM zAware continues operating.

Response: None required. However, unless an IBM zAware administrator adds additional storage through the **Administration** > **Configuration** > **Data Storage** tab or reduces the number of connected clients, the IBM zAware partition might run out of memory. If the partition is running out of memory, IBM zAware issues message AIFP0018W to indicate when storage usage has reached a level at which corrective action might be required to avoid a critical storage shortage. A critical storage shortage is indicated by message AIFP0019E.

AIFT0001I **A request to build a model for** *name* **started on** *date_time***.**

Explanation: IBM zAware started to process a queued request to build a model for the system or group identified by *name*. This request can be either a training operation that IBM zAware schedules automatically at the end of the training interval, or a training operation that an administrator explicitly requested. The variable *date_time* indicates when IBM zAware began to process the request.

System action: IBM zAware continues to process the training request.

Response: None.

AIFT0002I **A request to build a model for** *name* **completed successfully on** *date_time***.**

Explanation: IBM zAware successfully processed a request to build a model for the system or group identified by *name*. The variable *date_time* indicates when IBM zAware finished processing the request. The resulting model replaces the prior model, if any, for this system or group.

System action: IBM zAware begins using the new model.

Response: None.

AIFT0003I **An IBM zAware administrator cancelled a queued request to build a model for** *name***. The request was cancelled on** *date_time***.**

Explanation: Through the IBM zAware graphical user interface (GUI), an administrator cancelled a queued request to build a model for the system or group identified by *name*. This request can be either a training operation that IBM zAware schedules automatically at the end of the training interval, or a training operation that an administrator explicitly requested.

System action: IBM zAware discards the training request on the date and time indicated by the variable *date_time*. If a model exists already, IBM zAware uses that model for analysis. If a model does not exist, IBM zAware cannot provide analysis results for the system or group of systems indicated by *name*.

Response: None required. If a model does not exist, consider resubmitting a training request. If you do not manually request training, IBM zAware automatically schedules a training request according to the analytics configuration settings that are in effect.

AIFT0004E **A request to build a model for** *name* **failed on** *date_time***. On the Notifications page, look for additional error messages that provide more detail about the cause of the failure and provide possible corrective actions.**

Explanation: IBM zAware started to process a queued request to build a model but failed to complete the training operation. The training request, for the system or group identified by *name*, failed at *date_time*.

System action: IBM zAware continues processing other requests, if any.

Response: On the **Notifications** page, look for additional error messages that provide more detail about the cause of the failure. Follow the corrective actions that are provided for those additional messages.

AIFT0101I **A request to build a model for *name* did not complete successfully because the date range did not contain any days for which data is available for training. If you are building an initial model from priming data, send additional priming data within the training period date range and retry the training request. If you are replacing an existing model, the action required to resolve this error depends on various factors. See the user response for this message.**

Explanation: IBM zAware attempted to create a model of system behavior for the system or group identified by *name*. This training process was not successfully completed because data was not available for the dates in the training period.

System action: IBM zAware issues message AIFT0004E to indicate that the training request failed. If a model already exists, IBM zAware uses it to analyze the current data. If a model does not exist, IBM zAware cannot provide analysis results for the system or group of systems identified by *name*.

Response: One of the following actions, which require a user ID that is mapped to the Administrator role, might correct the cause of the failure.

- If an administrator is priming IBM zAware with prior data and a model does not exist yet, consider sending additional days of log data for this system or systems in the group, and retrying the training request for the system or group. For information about sending more log data, see the appropriate configuration topic for the type of monitored client in Part 4, "Configuring IBM zAware and its monitored clients," on page 93.
- If an administrator modified the default training set such that days that represent normal system activity are excluded from the training period, use **Manage Model Dates** to restore excluded dates to the training set, and retry the training request. If you need more information, see "Excluding dates from a model" on page 226.
- If an administrator modified the default training period but more days are required to build the model, go to the appropriate **Configuration** > **Analytics** tab, increase the "Training period" value, and apply the change. After reconnecting monitored clients as necessary, retry the training request. If you need more information about reconnecting monitored clients, see "Starting and stopping data collection for your monitored systems" on page 203.

AIFT0102I **A request to build a model for *name* did not complete successfully. Analysis of the data did not result in a usable pattern because of insufficient training set data. If you are building an initial model from priming data, send additional priming data and retry the training request. If you are replacing an existing model, the action required to resolve this error depends on various factors. See the user response for this message.**

Explanation: IBM zAware attempted to create a model of system behavior for the system or group identified by *name*. This training process was not successfully completed because the data for the model did not satisfy the requirements for successfully building a model. For information about building models, see the appropriate planning topic for the type of monitored client in Chapter 11, "Planning to create IBM zAware models," on page 87.

System action: IBM zAware issues message AIFT0004E to indicate that the training request failed. If a model for this system or group exists already, IBM zAware uses that model for analysis. If a model does not exist, IBM zAware cannot provide analysis results for the system or group of systems identified by *name*.

Response: One of the following actions, which require a user ID that is mapped to the Administrator role, might correct the cause of the failure.

- If an administrator is priming IBM zAware with prior data and a model does not exist yet, consider sending additional days of log data for this system or systems in the group, and retrying the training request for the system or group. For information about sending more log data, see the appropriate configuration topic for the type of monitored client in Part 4, "Configuring IBM zAware and its monitored clients," on page 93.
- If a model was previously created and an administrator changed the days selected to be part of the training set, use **Manage Model Dates** to restore excluded dates to the training set, and retry the training request. If you need more information, see "Excluding dates from a model" on page 226.
- If a model was previously created and an administrator changed the training period, go to the appropriate **Configuration** > **Analytics** tab to restore the original value, and apply the change. After reconnecting monitored clients, as necessary, retry the training request. If you need more information about reconnecting monitored clients, see "Starting and stopping data collection for your monitored systems" on page 203.
- If a model was created as a result of the immediately previous training request (that is, 30 days ago, if your installation is using the default training interval), you can ignore this message as long as you

are confident that the prior model still represents the current system behavior. If you have changed the workload on the system or group, however, consider increasing the training period or adding more training set data, and retrying the training request.

- If the existing model was created over 60 days ago and one or more training attempts have failed, contact IBM Support.

AIFT0103I **A request to build a model for** *name* **failed. Analysis of the data did not result in a usable model because of insufficient training set data. The training set for this request started on** *date-time* **and ended on** *date-time*, **contained** *number* **unique messages, and included a total of** *number* **analysis intervals, of which** *number* **contained usable data. The length of each analysis interval is** *number* **minutes. Only** *number* **messages appeared in at least** *number* **intervals.**

Explanation: IBM zAware attempted to create a model of system behavior for the system or group identified by *name*. This training process was not successfully completed because the data used for the model did not satisfy the requirements for successfully building a model. The training set used for this request consists of log data for the date range indicated in the message text; note that some individual dates in that range might have been explicitly excluded from training by an IBM zAware administrator. The message text also indicates the number of unique messages and details about the analysis intervals in the training set data.

System action: IBM zAware removes the training request from the queue. If a model exists already, IBM zAware uses that model for analysis. If a model does not exist, IBM zAware cannot provide analysis results for the system or group of systems identified by *name*.

Response: One of the following actions, which require a user ID that is mapped to the Administrator role, might correct the cause of the failure.

- If an administrator is priming IBM zAware with prior data and a model does not exist yet, consider sending additional days of log data for this system or systems in the group, and retrying the training request for the system or group. For information about sending more log data, see the appropriate configuration topic for the type of monitored client in Part 4, "Configuring IBM zAware and its monitored clients," on page 93.
- If a model was previously created, go to the appropriate **Configuration** > **Analytics** tab to increase the length of the training period, and apply the change. After reconnecting monitored clients, as necessary, retry the training request. If you need more information about reconnecting monitored clients, see "Starting and stopping data collection for your monitored systems" on page 203.

If the error is not resolved through one of these actions, the instrumentation data for this system does not provide sufficient information for IBM zAware to construct patterns.

AIFT0104I **Starting on** *date_time*, **IBM zAware successfully received** *number* **lines of data from** *system_name* **in which** *number* **lines contained messages. The received data can be either current** *log_type* **data from a monitored client or priming data for one or more clients.**

Explanation: IBM zAware received a specific type of instrumentation data, indicated by *log_type*, from the monitored client identified by *system_name*. This message indicates when IBM zAware started receiving the data, how many lines of data were successfully received, and how many of these lines contained messages.

System action: IBM zAware receives and stores the data.

Response: None.

AIFT0105E **Starting on** *date_time*, **IBM zAware received** *number* **lines of data from** *system_name* **in which** *number* **lines contained messages. A significant number of lines in the received data do not contain valid message IDs. The received data can be either current** *log_type* **data from a monitored client or priming data for one or more clients. Check the data being sent to ensure that it meets IBM zAware formatting requirements.**

Explanation: IBM zAware received a specific type of instrumentation data, indicated by *log_type*, from the monitored client identified by *system_name*. This message indicates when IBM zAware started receiving the data, how many lines of data were successfully received, and how many of these lines contained messages. A significant amount of this data did not contain valid message identifiers (IDs). A likely cause is that the data does not conform to the required format. To review IBM zAware requirements for data from supported types of monitored clients, see the prerequisites listed in Chapter 2, "Prerequisites for configuring and using IBM zAware," on page 13.

System action: IBM zAware receives and stores the data containing messages.

Response: Corrective action depends on the type of monitored client and the type of received data.

- For z/OS monitored clients:

  - For current OPERLOG data, check the OPERLOG data for systems connecting at *date_time*. Make sure that the OPERLOG configuration conforms to the requirements listed in "Configuring z/OS monitored clients to send data to the IBM zAware server" on page 111.
  - For priming data, check the job used to run the z/OS bulk load client for IBM zAware on the z/OS priming system. Make sure that the input data sets are sequential and contain SYSLOG data in the appropriate format, and rerun the z/OS bulk load client to resend the priming data.
- For Linux monitored clients, check the system log (syslog) data for systems connecting at *date_time*. Syslog data must be formatted according to the Internet Engineering Task Force (IETF) syslog protocol RFC 5424, which includes 4-digit years and time zone information. Additionally, each individual message that is transmitted must be preceded by the length of the message; this convention is known as octet framing. In addition:
  - The Linux system must correctly, consistently, and uniquely identify itself in the host name portion of the syslog message. IBM zAware interprets different but equivalent host name specifications to be different systems.
  - When sending syslog messages, the Linux system must provide a correct time stamp, including the Coordinated Universal Time (UTC) offset.

AIFT0106I **A request to build a model for** *name* **failed. Analysis of the data did not result in a usable model because of insufficient training set data. The training set for this request started on** *date-time* **and ended on** *date-time*, **contained** *number* **unique messages, and included a total of** *number* **analysis intervals, of which** *number* **contained usable data. The length of each analysis interval is** *number* **minutes.**

Explanation: IBM zAware attempted to create a model of system behavior for the system or group identified by *name*. This training process was not successfully completed because the data used for the model did not satisfy the requirements for successfully building a model. The training set used for this request consists of log data for the date range indicated in the message text; note that some individual dates in that range might have been explicitly excluded from training by an IBM zAware administrator. The message text also indicates the number of unique messages and details about the analysis intervals in the training set data.

System action: IBM zAware issues message AIFT0004E to indicate that the training request failed. If a model for this system or group exists already, IBM zAware uses that model for analysis. If a model does not exist, IBM zAware cannot provide analysis results for the system or group of systems identified by *name*.

Response: One of the following actions, which require a user ID that is mapped to the Administrator role, might correct the cause of the failure.

- If an administrator is priming IBM zAware with prior data and a model does not exist yet, consider sending additional days of log data for this system or systems in the group, and retrying the training request for the system or group. For information about sending more log data, see the appropriate configuration topic for the type of monitored client in Part 4, "Configuring IBM zAware and its monitored clients," on page 93.
- If a model was previously created and an administrator changed the days selected to be part of the training set, use **Manage Model Dates** to restore excluded dates to the training set, and retry the training request. If you need more information, see "Excluding dates from a model" on page 226.
- If a model was previously created and an administrator changed the training period, go to the appropriate **Configuration** > **Analytics** tab to restore the original value, and apply the change. After reconnecting monitored clients, as necessary, retry the training request. If you need more information about reconnecting monitored clients, see "Starting and stopping data collection for your monitored systems" on page 203.
- If a model was created as a result of the immediately previous training request (that is, 120 days ago, if your installation is using the default training interval), you can ignore this message as long as you are confident that the prior model still represents the current system behavior. If you have changed the workload on the system or group, however, consider increasing the training period or adding more training set data, and retrying the training request.
- If two training periods have elapsed since the existing model was created, and one or more training attempts have failed, contact IBM Support.

AIFT0107I **Starting on** *date-time*, **IBM zAware successfully processed** *number* **lines of** *log_type* **data from** *system_name* **in which** *number* **lines contained messages.**

Explanation: IBM zAware issues this message when a monitored client, which is identified by *system_name*, stops sending instrumentation data. A client stops sending data when the connection between the client and IBM zAware ends, or when the IBM zAware analytics engine is stopped or recycled.

The message text indicates the date and time when the monitored system most recently started sending data, the specific type of instrumentation data, which is indicated by *log_type*; how many lines of data were successfully processed while the client was connected; and how many of these lines contained messages.

System action: Unless the analytics engine was intentionally stopped through the **System Status** tab in the IBM zAware graphical user interface (GUI), IBM zAware continues normal operations.

Response: None required. Unless the monitored client was intentionally disconnected, you might need to take corrective action.

- If the client was unintentionally disconnected, check the system log of the monitored client for messages that report communication problems between that client and the IBM zAware server.
- If the client was disconnected because the IBM zAware analytics engine was stopped or recycled, you might need to reconnect the monitored client. For more information about reconnecting monitored clients, see "Starting and stopping data collection for your monitored systems" on page 203.

AIFT0108E **Starting on *date-time*, IBM zAware processed *number* lines of *log_type* data from *system_name*. None of these lines contained messages, so this data could not be used to produce analysis results.**

Explanation: IBM zAware issues this message when a monitored client, which is identified by *system_name*, stops sending instrumentation data. A client stops sending data when the connection between the client and IBM zAware ends, or when the IBM zAware analytics engine is stopped or recycled. If IBM zAware could not determine the system name or another unique identifier for the sender, IBM zAware replaces the variable *system_name* with the word "unknown".

This message indicates that, since the date and time when the monitored system most recently started sending data, IBM zAware processed but could not parse any messages from the system data it received. The message text indicates the most recent system connection time (*date-time*), the specific type of instrumentation data (*log_type*) sent by the system, and the total lines of system data that were processed (*number*).

System action: Unless the analytics engine was intentionally stopped through the **System Status** tab in the IBM zAware graphical user interface (GUI), IBM zAware continues normal operations. If a model exists for this monitored system, IBM zAware displays an anomaly score of zero for intervals during which the system was connected, but its data did not yield any messages for analysis.

Response: None required. This message might indicate a problem with the log data that this system sent to IBM zAware. To determine whether any problems exist with the system log data, make sure that the system meets IBM zAware requirements for monitored clients, and check the system log for this time period. For details about system requirements, see Chapter 2, "Prerequisites for configuring and using IBM zAware," on page 13.

Unless the monitored client was intentionally disconnected, you might need to take corrective action to restore the connection.

- If the client was unintentionally disconnected, check the system log of the monitored client for messages that report communication problems between that client and the IBM zAware server.
- If the client was disconnected because the IBM zAware analytics engine was stopped or recycled, you might need to reconnect the monitored client. For more information about reconnecting monitored clients, see "Starting and stopping data collection for your monitored systems" on page 203.

AIFT0109E **IBM zAware received an invalid client identification message from IP address *address*, and has closed the connection. Determine whether the remote system is a monitored client. If so, correct its configuration to send properly formatted data; otherwise, take corrective action to prevent the remote system from accessing IBM zAware.**

Explanation: IBM zAware received an identification message from a remote system with the IP address indicated by the variable *address*. This identification message was not formatted according to IBM zAware requirements for monitored clients.

System action: IBM zAware closes the connection with the remote system.

Response: Determine whether the remote system is a monitored client. If so, correct its configuration to send properly formatted data. For formatting requirements, see Chapter 2, "Prerequisites for configuring and using IBM zAware," on page 13.

Otherwise, take corrective action to prevent the remote system from accessing IBM zAware.

AIFT0110I **A request to build a model for *name* was received. The training set for this request started on *date-time* and ended on *date-time*, contained *number* unique messages, and included a total of *number* analysis intervals, of which *number* contained usable data. The length of each analysis interval is *number* minutes.**

Explanation: IBM zAware received a request to create a model of system behavior for the system or group identified by *name*. The training set used for this request consists of log data for the date range indicated in the message text. The message text also indicates the number of unique messages and analysis intervals in the training set data, and the length of the analysis interval in minutes.

System action: IBM zAware issues message AIFT0002I to indicate successful completion of this request.

Response: None.

AIFT0111W **Starting on *date-time*, IBM zAware processed *number* lines of *log_type* data from *system_name* in which *number* lines contained messages. Although IBM zAware used the messages for analysis, the low number of lines containing messages might indicate a problem. Consider checking the log data and system configuration for potential problems.**

Explanation: IBM zAware issues this message when a monitored client, which is identified by *system_name*, stops sending instrumentation data. A client stops sending data when the connection between the client and IBM zAware ends, or when the IBM zAware analytics engine is stopped or recycled.

This message indicates that, since the date and time when the monitored system most recently started sending data, IBM zAware processed but could not parse messages from a significant amount of the system data it received. The message text indicates the most recent system connection time (*date-time*), the specific type of instrumentation data (*log_type*) sent by the system, the total lines of system data that were processed (*number*), and the number of lines that contained messages (*number*).

System action: Unless the analytics engine was intentionally stopped through the **System Status** tab in the IBM zAware graphical user interface (GUI), IBM zAware continues normal operations. If a model exists for this monitored system, IBM zAware displays an anomaly score of zero for intervals during which the system was connected, but its data did not yield enough messages for analysis.

Response: None required. This message might indicate a problem with the log data that this system sent to IBM zAware. To determine whether any problems exist with the system log data, make sure that the system meets IBM zAware requirements for monitored clients, and check the system log for this time period. For details about system requirements, see Chapter 2, "Prerequisites for configuring and using IBM zAware," on page 13.

Unless the monitored client was intentionally disconnected, you might need to take corrective action to restore the connection.

- If the client was unintentionally disconnected, check the system log of the monitored client for messages that report communication problems between that client and the IBM zAware server.
- If the client was disconnected because the IBM zAware analytics engine was stopped or recycled, you might need to reconnect the monitored client.

  For more information about reconnecting monitored clients, see "Starting and stopping data collection for your monitored systems" on page 203.

AIFT0112I **A request to build a model for *name* was received. The training set for this request started on *date-time* and ended on *date-time*, contained *number* unique messages, and included a total of *number* analysis intervals, of which *number* contained usable data. The length of each analysis interval is *number* minutes. Only *number* messages appeared in at least *number* intervals.**

Explanation: IBM zAware received a request to create a model of system behavior for the system or group identified by *name*. The training set used for this request consists of log data for the date range indicated in the message text. The message text also indicates the number of unique messages and analysis intervals in the training set data, the length of the analysis interval in minutes, and the number of unique messages that appeared in the minimum number of intervals.

System action: IBM zAware issues message AIFT0002I to indicate successful completion of this request.

Response: None.

AIFT0150E **While processing a *type* request, IBM zAware attempted to access storage that it had been using, but the attempt failed because of an I/O read error. Possible causes of this error include an access problem with the physical storage device, a corrupted file, or other read errors. Investigate and correct the problem with the storage device and reactivate the IBM zAware partition.**

Explanation: The IBM zAware server attempted to access storage that it had been using, but the attempt failed because of an I/O read error. Possible causes of this error include an access problem with the physical storage device, a corrupted file, or other read errors. The variable *type* indicates the type of request that IBM zAware was processing when it encountered the error.

System action: The IBM zAware server cannot access its storage.

Response:

1. Ensure that the storage volumes assigned to the IBM zAware partition are physically connected to the IBM zAware host system.
2. Through the Hardware Management Console (HMC), deactivate and activate the IBM zAware partition after access to storage is corrected.
   a. Deactivate the partition:
      1) Stop the IBM zAware analytics engine and the data transmission from monitored clients.

a) Go to the **Systems** > **System Status** tab, and click **Stop** ( ) to stop the analytics engine. This action prevents the IBM zAware server from accepting any data transmission from clients.

b) To prevent monitored systems from experiencing communication errors, stop them from transmitting data.

- For monitored z/OS clients, use the **SETLOGR** command on each z/OS system to prevent the systems from attempting to reestablish the TCP/IP connection to the IBM zAware server.

  ```
  SETLOGR FORCE,ZAIQUIESCE,ALL
  ```

- For monitored Linux systems, stop the syslog daemon, by using the appropriate command for the type of syslog daemon that is in use on the Linux system.

2) Deactivate the IBM zAware partition. Use the **Deactivate** task in the Hardware Management Console (HMC). For authorization requirements and other information about the **Deactivate** task, see HMC/SE topics in IBM Knowledge Center, at http://www.ibm.com/support/knowledgecenter/

b. Activate the partition.

The following steps describe one method of activating the partition through the HMC:

1) Select the image for the IBM zAware partition.

2) From the **Daily** task group, open the **Activate** task. The Activate Task Confirmation window is displayed.

3) Review the confirmation text to decide whether to proceed with the task. If you want to continue this task, click **Yes**. The Activate Progress window opens to indicate the progress of the activation and the outcome.

4) Click **OK** to close the window when the activation completes successfully.

Otherwise, if the activation does not complete successfully, follow the directions on the window to determine the problem and how to correct it.

3. Reconnect the monitored clients that were previously connected to the IBM zAware server.

- To reconnect each z/OS system in your IBM zAware configuration to the server, use the **SETLOGR** command on each z/OS system.

  ```
  SETLOGR FORCE,ZAICONNECT,LSNAME=SYSPLEX.OPERLOG
  ```

- To reconnect a Linux system, restart the syslog daemon, using the appropriate command for the type of syslog daemon and for the type of initialization (init) process that is in use on the Linux system.

AIFT0151E **While processing a *type* request, IBM zAware attempted to access storage that it had been using, but the attempt failed because of an I/O write or lock error. The probable cause of this error is insufficient storage capacity. To correct this problem, an administrator must determine whether additional storage devices can be added to increase capacity.**

Explanation: The IBM zAware server has used all of the available capacity of the storage devices that are currently assigned for its use. The variable *type* indicates the type of request that IBM zAware was processing when it encountered the error.

System action: IBM zAware fails any attempt to write additional information to storage; this information includes analysis of operations log (OPERLOG) data and system models that are created through the training process.

Existing data that is already stored might be available for display through the IBM zAware graphical user interface (GUI).

Response:

1. With a user ID that is assigned to the Administrator role, go to the **Administration** > **Configuration** > **DataStorage** page in the IBM zAware GUI. To determine whether any additional volumes are available for use, sort the Data Storage Devices table by the Status column to search for available devices. If additional volumes are available, follow the procedure in "Adding and removing storage devices" on page 195 to add these devices.
2. If additional storage devices are not available, work with your storage administrator to determine whether any additional physical volumes that are attached to the IBM zAware host system can be used. If so, complete the following steps:

a. Through the Hardware Management Console (HMC), deactivate the IBM zAware partition:

1) Stop the IBM zAware analytics engine and the data transmission from monitored clients.

a) Go to the **Systems** > **System Status** tab, and click **Stop** ( ) to stop the analytics engine. This action prevents the IBM zAware server from accepting any data transmission from clients.

b) To prevent monitored systems from experiencing communication errors, stop them from transmitting data.

- For monitored z/OS clients, use the **SETLOGR** command on each z/OS

system to prevent the systems from attempting to reestablish the TCP/IP connection to the IBM zAware server.

```
SETLOGR FORCE,ZAIQUIESCE,ALL
```

- For monitored Linux systems, stop the syslog daemon, by using the appropriate command for the type of syslog daemon that is in use on the Linux system.

2) Deactivate the IBM zAware partition. Use the **Deactivate** task in the Hardware Management Console (HMC). For authorization requirements and other information about the **Deactivate** task, see HMC/SE topics in IBM Knowledge Center, at http://www.ibm.com/support/knowledgecenter/

b. To add additional volumes to the IO configuration, use the instructions in Chapter 12, "Configuring network connections and storage for the IBM zAware partition," on page 95.

c. Use the HMC to activate the IBM zAware partition.

The following steps describe one method of activating the partition through the HMC:

1) Select the image for the IBM zAware partition.

2) From the **Daily** task group, open the **Activate** task. The Activate Task Confirmation window is displayed.

3) Review the confirmation text to decide whether to proceed with the task. If you want to continue this task, click **Yes**. The Activate Progress window opens to indicate the progress of the activation and the outcome.

4) Click **OK** to close the window when the activation completes successfully.

Otherwise, if the activation does not complete successfully, follow the directions on the window to determine the problem and how to correct it.

d. Through the IBM zAware GUI, return to the **Administration** > **Configuration** > **DataStorage** tab, and add the additional devices by following the procedure in "Adding and removing storage devices" on page 195.

e. Reconnect the monitored clients that were previously connected to the IBM zAware server.

- To reconnect each z/OS system in your IBM zAware configuration to the server, use the **SETLOGR** command on each z/OS system.

```
SETLOGR FORCE,ZAICONNECT,LSNAME=SYSPLEX.OPERLOG
```

- To reconnect a Linux system, restart the syslog daemon, using the appropriate command for the type of syslog daemon and for the type of initialization (init) process that is in use on the Linux system.

3. Consider resending any missing instrumentation data to IBM zAware for analysis and modeling. To send data that was lost while IBM zAware did not have sufficient storage capacity to analyze and save client data, use the instructions in the appropriate configuration topic for the type of monitored client in Part 4, "Configuring IBM zAware and its monitored clients," on page 93.

AIFT0152E **While processing a *type* request, IBM zAware attempted to write to storage but the attempt failed. This condition might be resolved without administrator intervention. Check the Analysis page in the IBM zAware GUI; if analysis of current data has stopped or the GUI seems to hang, deactivate and reactivate the IBM zAware partition.**

Explanation: IBM zAware server attempted to write to an in-use storage device but the attempt failed. Possible causes of this error include an access problem with the physical storage device, a disk failure, contention for locks, or other error conditions. The variable *type* indicates the type of request that IBM zAware was processing when the write attempt failed.

System action: Licensed Internal Code automatically gathers and sends diagnostic information to IBM if the "Call Home" feature is enabled on the IBM zAware host system.

Response: To determine whether corrective action is required, check the **Analysis** page in the IBM zAware graphical user interface (GUI). If analysis of current data from monitored clients is in progress, no corrective action is required. If analysis of current data seems to have stopped and results are not displayed for recent intervals, deactivate and reactivate the IBM zAware partition. Complete the following steps:

1. Through the Hardware Management Console (HMC), deactivate the partition:

a. Stop the IBM zAware analytics engine and the data transmission from monitored clients.

1) Go to the **Systems** > **System Status** tab, and click **Stop** ( ) to stop the analytics engine. This action prevents the IBM zAware server from accepting any data transmission from clients.

2) To prevent monitored systems from experiencing communication errors, stop them from transmitting data.

- For monitored z/OS clients, use the **SETLOGR** command on each z/OS system to prevent the systems from attempting to reestablish the TCP/IP connection to the IBM zAware server.

```
SETLOGR FORCE,ZAIQUIESCE,ALL
```

- For monitored Linux systems, stop the syslog daemon, by using the appropriate command for the type of syslog daemon that is in use on the Linux system.

b. Deactivate the IBM zAware partition. Use the **Deactivate** task in the Hardware Management Console (HMC). For authorization requirements and other information about the **Deactivate** task, see HMC/SE topics in IBM Knowledge Center, at http://www.ibm.com/support/knowledgecenter/

2. Activate the partition.

   The following steps describe one method of activating the partition through the HMC:

   a. Select the image for the IBM zAware partition.

   b. From the **Daily** task group, open the **Activate** task. The Activate Task Confirmation window is displayed.

   c. Review the confirmation text to decide whether to proceed with the task. If you want to continue this task, click **Yes**. The Activate Progress window opens to indicate the progress of the activation and the outcome.

   d. Click **OK** to close the window when the activation completes successfully.

      Otherwise, if the activation does not complete successfully, follow the directions on the window to determine the problem and how to correct it.

3. Through the IBM zAware GUI, navigate to the **Notifications** page to determine whether IBM zAware issued another AIFT0152E message after reactivation. If IBM zAware has not issued another AIFT0152E message, continue to step 4; otherwise, skip to step 5.

4. Repeat the following steps, as necessary, to reconnect the monitored clients that were previously connected to the IBM zAware server.

   a. Use the appropriate command to reconnect each type of monitored client.

      - To reconnect each z/OS system in your IBM zAware configuration to the server, use the **SETLOGR** command on each z/OS system.

        ```
        SETLOGR FORCE,ZAICONNECT,LSNAME=SYSPLEX.OPERLOG
        ```

      - To reconnect a Linux system, restart the syslog daemon, using the appropriate command for the type of syslog daemon and for the type of initialization (init) process that is in use on the Linux system.

   b. After each client has successfully reconnected, navigate to the **Notifications** page to determine whether IBM zAware issued another AIFT0152E message after the reconnection. If IBM zAware has not issued another AIFT0152E message, continue to reconnect clients; otherwise, continue to step 5.

   c. After you have reconnected all monitored clients, continue to step 5.

5. Determine whether the current memory and processor resources allocated to the IBM zAware partition are sufficient for the number of monitored clients, and make any necessary adjustments. Use the System Activity display for the IBM zAware partition to view information about processor resources. To access the System Activity display, use the **Monitors Dashboard** task in the HMC for the IBM zAware host system.

   To check the recommendations for memory and processor resources, see "Estimating processor and memory resources" on page 49.

   For information about changing the processor and memory resources that are defined for a logical partition, see *z Systems PR/SM Planning Guide*, SB10-7162.

6. Regardless of whether you adjusted processor or memory resources in step 5, continue to check the **Notifications** page to determine whether IBM zAware issues another AIFT0152E message. If IBM zAware has not issued another AIFT0152E message after more than 24 hours of operation, no further action is necessary; otherwise, continue to step 7.

7. Contact IBM Support.

   For additional information, see Chapter 29, "Reporting IBM z Advanced Workload Analysis Reporter (IBM zAware) problems to IBM," on page 281. If you do not have a maintenance contract with IBM, use the questions at the end of that topic to collect the information that you might need to report the problem to your hardware maintenance provider.

AIFT0999E **IBM zAware detected an internal error while processing a request for *name*. The request type is *request_type*. IBM zAware collects diagnostic data and retries the request; depending on the outcome of the second attempt, corrective action might not be required. See the message description for further details.**

Explanation: While processing a request for the monitored client identified by *name*, IBM zAware detected an internal error. The variable *request_type* indicates the type of request that IBM zAware was processing:

Train IBM zAware was building the model for the system or the model group to which it belongs. This processing can be initiated by IBM zAware based on the configured training interval value, or by an administrator who used the **Request Training** action through one of the **Administration** > **Training Sets** tabs.

Analyze
IBM zAware was analyzing the current instrumentation data that the system was sending.

Upload IBM zAware was receiving either current instrumentation data from the system, or priming data for one or more systems.

System action: Licensed Internal Code automatically gathers and sends diagnostic information to IBM if the "Call Home" feature is enabled on the IBM zAware host system. IBM zAware terminates the request.

For a Train request
IBM zAware does not create a model when a training request fails. If a model existed prior to the failure, IBM zAware continues to use that model for analysis and attempts to retry the failed training request on the next day (UTC). For an automated build to be scheduled, the client must be connected to the IBM zAware server.

For an Analyze or Upload request
IBM zAware immediately attempts to retry the failed request.

Response: Corrective action depends on the type of request, and on the outcome of the second attempt to process the request.

For a Train request
You can either wait for IBM zAware to retry the training request on the next day, or go to the appropriate **Administration** > **Training Sets** tab and select the **Request Training** action to attempt to rebuild the system model immediately. If the retry attempt fails, deactivate and reactivate the IBM zAware partition and retry the training request.

For an Analyze or Upload request
If the retry attempt fails, use the appropriate commands to disconnect and reconnect the system from IBM zAware, and retry the request.

- For a z/OS monitored client:
  1. Disconnect *system_name* from the IBM zAware server by issuing the **SETLOGR** command:

     ```
     SETLOGR FORCE,ZAIQUIESCE,ALL
     ```

  2. Reconnect the client by issuing the **SETLOGR** command from *system_name*:

     ```
     SETLOGR FORCE,ZAICONNECT,LSN=SYSPLEX.OPERLOG
     ```

  3. For an Upload request only, resend priming data for *name* through the z/OS bulk load client for IBM zAware.
- For a Linux monitored client, restart the syslog daemon, using the appropriate command for the type of syslog daemon and for the type of initialization (init) process that is in use on the Linux system.

If the request fails again, deactivate and reactivate the IBM zAware partition and retry the request.

To deactivate and reactivate the IBM zAware partition, complete the following steps:

1. Through the Hardware Management Console (HMC), deactivate the partition:
   a. Stop the IBM zAware analytics engine and the data transmission from monitored clients.
      1) Go to the **Systems** > **System Status** tab, and click **Stop** ( ) to stop the analytics engine. This action prevents the IBM zAware server from accepting any data transmission from clients.
      2) To prevent monitored systems from experiencing communication errors, stop them from transmitting data.
         - For monitored z/OS clients, use the **SETLOGR** command on each z/OS system to prevent the systems from attempting to reestablish the TCP/IP connection to the IBM zAware server.

           ```
           SETLOGR FORCE,ZAIQUIESCE,ALL
           ```

         - For monitored Linux systems, stop the syslog daemon, by using the appropriate command for the type of syslog daemon that is in use on the Linux system.

   b. Deactivate the IBM zAware partition. Use the **Deactivate** task in the Hardware Management Console (HMC). For authorization requirements and other information about the **Deactivate** task, see HMC/SE topics in IBM Knowledge Center, at http://www.ibm.com/support/knowledgecenter/
2. Activate the partition.

   The following steps describe one method of activating the partition through the HMC:
   a. Select the image for the IBM zAware partition.
   b. From the **Daily** task group, open the **Activate** task. The Activate Task Confirmation window is displayed.
   c. Review the confirmation text to decide whether to proceed with the task. If you want to continue this task, click **Yes**. The Activate Progress window opens to indicate the progress of the activation and the outcome.
   d. Click **OK** to close the window when the activation completes successfully.

      Otherwise, if the activation does not complete successfully, follow the directions on the window to determine the problem and how to correct it.

3. Reconnect the monitored clients that were previously connected to the IBM zAware server.
   - To reconnect each z/OS system in your IBM zAware configuration to the server, use the **SETLOGR** command on each z/OS system.

     ```
     SETLOGR FORCE,ZAICONNECT,LSNAME=SYSPLEX.OPERLOG
     ```
   - To reconnect a Linux system, restart the syslog daemon, using the appropriate command for the type of syslog daemon and for the type of initialization (init) process that is in use on the Linux system.

If the problem persists, request IBM support by generating a Type V Viewable PMH (PMV) record to report this message ID and reason code.

For additional information, see Chapter 29, "Reporting IBM z Advanced Workload Analysis Reporter (IBM zAware) problems to IBM," on page 281. If you do not have a maintenance contract with IBM, use the questions at the end of that topic to collect the information that you might need to report the problem to your hardware maintenance provider.

Index

J

L

M

N

O

P

R

S

T

U

V

X

Z

Printed in USA