#37 zNibbler (z/OS' zAware)

zTidBits Series

zEC12 RAS - Prologue

RAS is the acronym for Reliability, Availability and Serviceability. The zEC12 system reliability, availability, and serviceability (RAS) strategy is a building-block approach developed to meet the client's stringent requirements of achieving continuous reliable operation. Those building blocks are error prevention, error detection, recovery, problem determination, service structure, change management, and measurement and analysis.

The initial focus is on preventing failures from occurring in the first place. This is accomplished by using *Hi-Rel* (highest reliability) components; using screening, sorting, burn-in, and run-in; and by taking advantage of technology integration. For Licensed Internal Code and hardware design, failures are eliminated through rigorous design rules; design walk-through; peer reviews; element, subsystem, and system simulation; and extensive engineering and manufacturing testing.

The RAS strategy is focused on a recovery design that is necessary to mask errors and make them transparent to customer operations. An extensive hardware recovery design has been implemented to detect and correct memory array faults. In cases where total transparency cannot be achieved, you can restart the server with the maximum possible capacity.

Some of the RAS improvements of the zEC12 are:

- Improved error detection for the L3 / L4 memory cache
- IBM System z Advanced Workload Analysis Reporter in order to detect abnormal behavior of z/OS
- OSA firmware changes to increase concurrent MCL capability
- Digital Temperature Sensor (DTS) and On Chip Temperature Sensor on the Processor Unit (PU) chips.

This issue will target zAware which uses unsupervised machine learning and rules to determine runtime anomaly.

zAware

The IBM System z Advanced Workload Analysis Reporter (IBM zAware) provides a smart solution for detecting and diagnosing anomalies in z/OS 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 z/OS 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.

Systems behaving badly !!!

Several 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 using 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 — 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 a model of normal system behavior and uses it to compare to current message log data from the connected z/OS system. IBM zAware monitors the z/OS operations log (OPERLOG), which contains all messages written to the z/OS console, including messages that are suppressed by message flooding automation. Through *deep analytics*, which is a data-mining process, 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 improves 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 z/OS 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.

Finding the culprit when a problem occurs

Using the **Analysis** page in the IBM zAware GUI, you can answer key questions to diagnose a systemp roblem. Figure 1 illustrates the elements of the **Analysis** page that help you answer this question: “Which z/OS system is behaving abnormally?”

1. To begin your search, check the **Analysis Source** and adjust it as necessary. By default, all monitored systems are displayed in the **Analysis** page but you can use **Change Source** to modify the display to show systems in a specific sysplex or to show specific systems.
2. Review the names of the z/OS systems that are connected to the IBM zAware server, which are listed in the **System** column. You might need to scroll this list depending on the number of systems.
3. View the bar graphs in the **Anomaly Scores** column. Each rectangle in the graph represents a 10-minute interval; the number of unique message IDs that are issued during that interval determines the height of the rectangle.
   IBM zAware uses unsupervised machine learning and IBM rules to determine interval anomaly scores:
   * Through ***unsupervised machine learning***, the IBM zAware server extracts and organizes message scores data to build a model of behavior for each monitored client. This training process is repeated over time, with the frequency determined by the training interval, which enables the server to update and refine each client model.

   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 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.
   * ***z/OS experts*** at IBM know, based on decades of IBM experience with testing and using z/OS systems, which message IDs are likely to indicate potential problems. Message IXC101I, for example, indicates that a system is being removed from a sysplex. For message IXC101I and other messages that are known to indicate potential problems, the IBM zAware server is programmed to assign the highest anomaly score to the intervals in which those messages are issued.
4. Compare the **color of each rectangle** to the **Score key** to determine the relative anomaly score for the interval. The anomaly score indicates unusual patterns of message IDs within that interval, as compared to the model of normal system behavior. Intervals containing relatively normal, common messages receive a low score and lighter blue color, and intervals containing more unusual messages receive a higher score and darker blue or gold color. A high score indicates unusual message IDs or unusual patterns of message IDs compared to the system model.

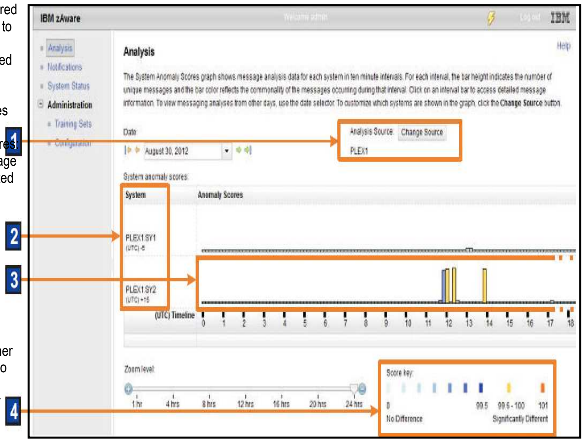

Analysis

The System Anomaly Scores graph shows message analysis data for each system in ten minute intervals. For each interval, the bar height indicates the number of unique messages and the bar color reflects the commonality of the messages occuring during that interval. Click on an interval bar to access detailed message information. To view messaging analyses from other days, use the date selector. To customize which systems are shown in the graph, click the **Change Source** button.

Date: August 30, 2012

System anomaly scores:

Figure 1. *Elements of the* **Analysis** *page that help identify the problematic z/OS system.*

From the interval anomaly scores shown in Figure 1, you can determine that SYSB is a potential source of the problem because its interval anomaly scores show that many unique messages are being issued over many intervals.

Your next diagnostic question to answer is “When did this system start misbehaving?” To determine when the system began behaving abnormally, you can use various controls in the **Analysis** page to look at intervals in the past.

1. The **Date** field displays the currently selected Coordinated Universal Time (UTC) date. You can change the current day to another date.

2. The **Timeline** marks the hours of the day in UTC, using the 24-hour clock. You can move the slider to display hours that do not show in the **Analysis** page.

   For a sysplex that contains systems that operate in different time zones, IBM zAware converts the times to UTC times to align the current times across monitored systems.

3. The **Zoom level** slider controls how many hours of the day are displayed in the **Analysis** page. You can move the slider to display all hours or only a few hours.

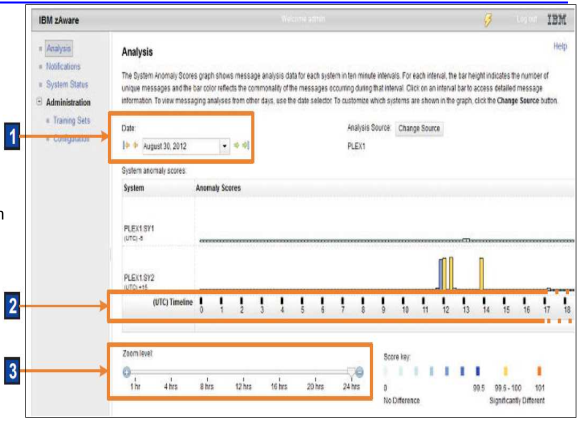

Analysis

The System Anomaly Scores graph shows message analysis data for each system in ten minute intervals. For each interval, the bar height indicates the number of unique messages and the bar color reflects the commonality of the messages occurring during that interval. Click on an interval bar to access detailed message information. To view messaging analyses from other days, use the date selector. To customize which systems are shown in the graph, click the **Change Source** button.

Date: August 30, 2012

Analysis Source: Change Source

PLEX1

System anomaly scores:

System | Anomaly Scores

PLEX1.SY1 (UTC) -5

PLEX1.SY2 (UTC) +11

(UTC) Timeline 0 1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18

Zoom level: 1 hr 4 hrs 8 hrs 12 hrs 16 hrs 20 hrs 24 hrs

Score key: 0 No Difference 99.5 99.6 - 100 101 Significantly Different

Figure 2. *Elements of the* **Analysis** *page that change the time interval display*

Using the **Timeline** and **Zoom level** together, you can change the display to focus on anomaly scores in a specific period.

Through these controls, you can alter the display to determine the time at which SYSB began to issue unusual messages. Hovering your cursor over the time-interval rectangle displays a summary containing the time, the number of unique message IDs issued within that time interval, and the interval anomaly score.

In summary, the **Analysis** page helps you determine which system is behaving abnormally, how many unique messages it is issuing, and when the abnormal behavior began.

To pinpoint and diagnose the problem with SYSB, you might need to ask several questions:

- What message IDs are unusual?
- How often did the unusual message get issued?
- Are messages issued in context within an expected pattern?
- Is a specific z/OS component or application issuing unusual messages?
- Within the 10-minute interval, when did the message ID first appear?

To answer these diagnostic questions, use the **Interval view**, which is illustrated in Figure 3. You can display the **Interval view** for any of the 10-minute time intervals by clicking the colored rectangle in the bar graph of the **Analysis** page.

The **Interval view** provides the following diagnostic details:
1. The system name, the date, and the time interval that you selected from the **Analysis** page.
2. The **Interval anomaly score**, 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. Details about each message issued during the interval. These details include:
   Anomaly Score
   Indicates the difference in expected behavior for this specific message ID within the Selected 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 10-minute interval. This interval score is a function of the rarity score, the number of times that the message appears within this interval, and whether the message appeared in context. Higher scores indicate greater contribution to the interval anomaly score. The interval contribution score ranges from 0 through the largest number that the Java double data type supports.
   Message Context
   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). A message that is issued out of context (without the other messages in the same cluster) might indicate a problem.
   Message ID
   Provides the message identifier. The message ID itself is a link to LookAt, the online z/OS message help facility. For IBM messages that are enabled for LookAt, clicking the link opens a new browser window that displays the description of the message, which often includes a recommended action to correct the issue. If the message is not available through LookAt, you can find an online description by using the Internet search engine of your choice.
   Message Example or Message Summarization
   Provides either the full message text for the first occurrence of this message within the interval, or a summary of the common message text that was issued for each occurrence of the same message. To control the content displayed in this column, click either **View Message Full Text** or **View Message Summary** from the **Actions** list.

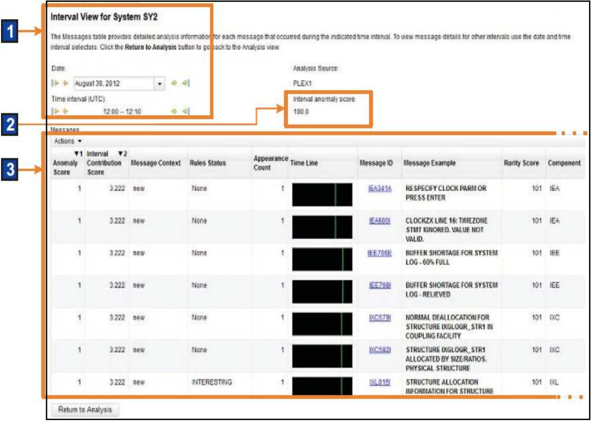

Interval View for System SY2

The Messages table provides detailed analysis information for each message that occurred during the indicated time interval. To view message details for other intervals use the date and time interval selectors. Click the **Return to Analysis** button to go back to the Analysis view.

Date: August 30, 2012

Time interval (UTC): 12:00 – 12:10

Analysis Source: PLEX1

Interval anomaly score: 100.0

Messages

▼1 Anomaly Score	Interval ▼2 Contribution Score	Message Context	Rules Status	Appearance Count	Time Line	Message ID	Message Example	Rarity Score	Component
1	3.222	new	None	1		IEA341A	RESPECIFY CLOCK PARM OR PRESS ENTER	101	IEA
1	3.222	new	None	1		IEA600I	CLOCKZX LINE 16: TIMEZONE STMT IGNORED. VALUE NOT VALID.	101	IEA
1	3.222	new	None	1		IEE766E	BUFFER SHORTAGE FOR SYSTEM LOG - 60% FULL	101	IEE
1	3.222	new	None	1		IEE768I	BUFFER SHORTAGE FOR SYSTEM LOG - RELIEVED	101	IEE
1	3.222	new	None	1		IXC579I	NORMAL DEALLOCATION FOR STRUCTURE IXGLOGR_STR1 IN COUPLING FACILITY	101	IXC
1	3.222	new	None	1		IXC582I	STRUCTURE IXGLOGR_STR1 ALLOCATED BY SIZE/RATIOS. PHYSICAL STRUCTURE	101	IXC
1	3.222	new	INTERESTING	1		IXL015I	STRUCTURE ALLOCATION INFORMATION FOR STRUCTURE	101	IXL

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

To increase your insight into the problem, you can sort these details by column. Select **Sort Multiple** from the **Actions** list and select the columns that you want to sort.

Monitoring behavior after a change

The **Analysis** page and **Interval view** 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

In such cases, you can use the **Analysis** page and **Interval view** 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 and **Interval view** 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) that is dedicated to running the IBM zAware server. This LPAR runs on a IBM zEnterprise EC12 (zEC12) central processor complex (CPC). Figure 4 illustrates the major elements of an IBM zAware configuration.

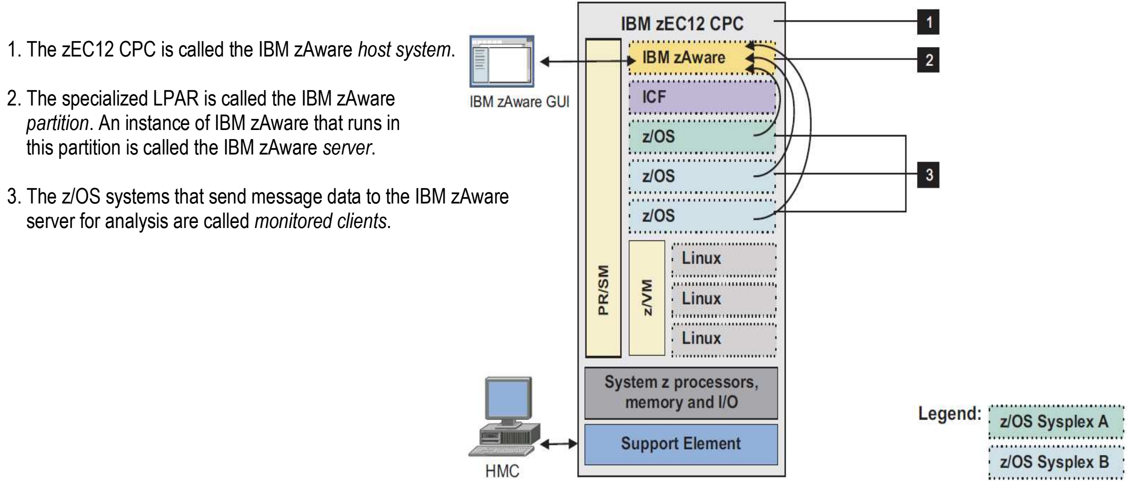

1. The zEC12 CPC is called the IBM zAware *host system*.
2. The specialized LPAR is called the IBM zAware *partition*. An instance of IBM zAware that runs in this partition is called the IBM zAware *server*.
3. The z/OS systems that send message data to the IBM zAware server for analysis are called *monitored clients*.

Figure 4. *An IBM zAware partition running in a zEC12 CPC*

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 5 on next page illustrates another possible configuration with additional z/OS monitored clients running on a IBM zEnterprise 196 (z196) CPC.

- 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 (zEC12). Only one z/OS system resides on the host system; the other z/OS system resides on the z196.
- 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 z196.
- The z/OS system shown at the top of the z196 on the 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 running on the host system.
- 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.

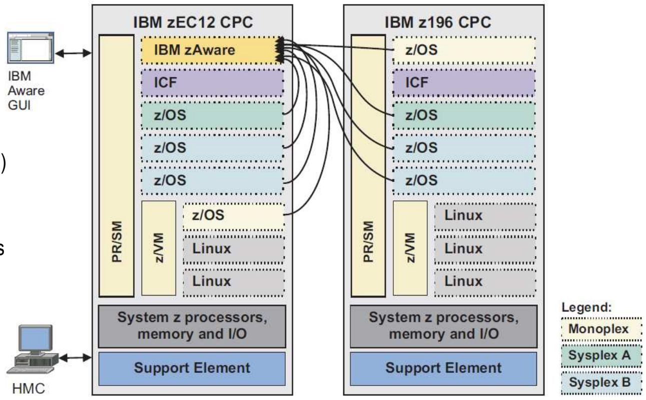

Figure 5. *An IBM zAware partition supporting clients in one zEC12 CPC and one z196*

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 **Notifications**, you can view operational messages that the IBM zAware server issues.
- Through **System Status**, you can view information about the z/OS monitored clients (systems) that are connected to the IBM zAware server.
- Through **Administration** > **Configuration**, you can modify default values that control the analytics engine, add or remove storage devices, manage security mechanisms, view the sysplex topology of monitored clients, and assign priming data that the IBM zAware server uses to build system behavior models. Each model, which is called an IBM zAware model, is a description of normal behavior that is generated for a specific z/OS monitored client.
- Through **Administration** > **Training Sets**, you can view information about the generation of IBM zAware models, which are periodically updated by the server.

Figure 6 illustrates how system management products can use the analytical data that is presented in the IBM zAware GUI.

1. 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). This data is equivalent to the information that is available through the **Analysis** page and **Interval view** in the IBM zAware GUI.
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.

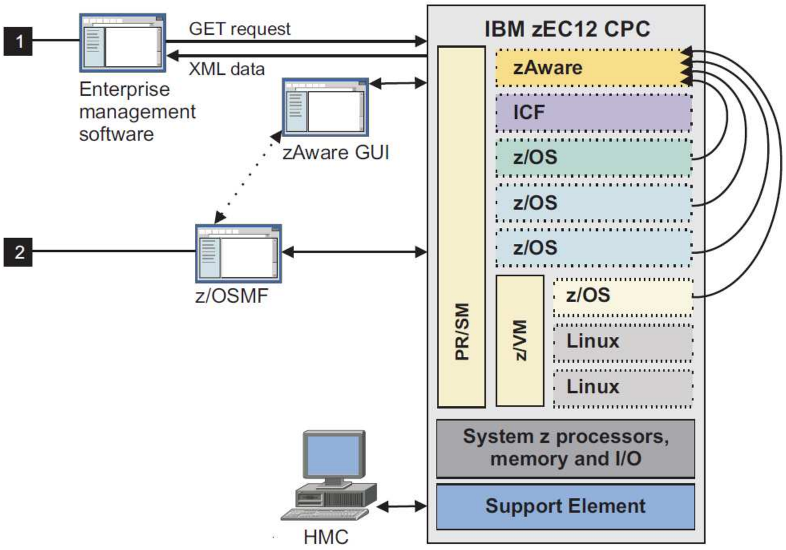

Figure 6. *System management products using IBM zAware analytical data*

Notes:

IBM zAware is available with IBM zEnterprise EC12 (zEC12) models. You can order IBM zAware through the following feature codes.

0011 The IBM zAware enablement feature for a zEC12 that serves as the IBM zAware host system.
0101 The IBM zAware connections feature that represents the quantity of monitored clients to be connected to the IBM zAware server on the host system.
0102 The IBM zAware connections that represents the quantity of monitored clients to be connected to a zEC12 that serves as a disaster recovery (DR) system.

The IBM zAware partition that runs on the host or DR 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).
- 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 z/OS clients that the server monitors.

– – –