z/OS V2R4 Communications Server IBM zERT Network Analyzer database administration enhancements

Contents

List of Tables

Chapter
1

New Function Summary

Topics:

- IBM zERT Network Analyzer database administration enhancements

IBM zERT Network Analyzer database administration enhancements

z/OS Management Facility (z/OSMF) V2R4 with the IBM zERT Network Analyzer APAR PH16223, provides additional flexibility in IBM zERT Network Analyzer's Db2 for z/OS database schema definitions and reduces the access privileges required by the IBM zERT Network Analyzer's database user ID. The supplied database schema tooling now supports customized values for the database schema name, index names and even table names along with many other operational parameters that were already configurable.

Dependencies: No new dependencies are introduced. The IBM zERT Network Analyzer requires z/OSMF to be installed and a type 4 JDBC connection to Db2 for z/OS 11 or higher.

To enable IBM zERT Network Analyzer database administration enhancements, perform the tasks in Table 1: IBM zERT Network Analyzer database administration enhancements on page 8.

Table 1: IBM zERT Network Analyzer database administration enhancements

Task/Procedure	Reference

Stop z/OSMF if it is running
 	
• If you are installing the IBM zERT Network Analyzer for the first time, your Db2 for z/OS database administrator (DBA) will need to create up your customized database using the provided schema tooling.
• If you already have the IBM zERT Network Analyzer running:
 	
  • Your DBA needs to either upgrade the existing database schema to the new schema version or else drop your existing database and re-create it with customized schema, index and/or table names at the new schema version.
  • Regardless of which approach you decide, your DBA uses the provided database schema tooling to implement the changes. This includes deciding on and defining the desired number of IBM zERT Network Analyzer query result table partitions.
	See Db2 for z/OS customization for the IBM zERT Network Analyzer task in *z/OS Management Facility Configuration Guide*

Once the IBM zERT Network Analyzer database is updated to or created at the new schema version , start z/OSMF and the IBM zERT Network Analyzer plugin. If you are using a newly created database, you need to fill in and save the required JDBC connectivity settings in the IBM zERT Network Analyzer's database settings panel.	See Connect the IBM zERT Network Analyzer task with the Db2 for z/OS database in *z/OS Management Facility Configuration Guide*
Once the correct database settings are stored and the IBM zERT Network Analyzer successfully connects to the database, you can verify the database schema information in the database settings panel.	See *Working with the Database Settings panel* of IBM zERT Network Analyzer online help, *Analysis* category under the IBM z/OS Management Facility online help

To find all related topics about IBM zERT Network Analyzer database administration enhancements, see Table 2: All related topics about IBM zERT Network Analyzer database administration enhancements on page 9.

Table 2: All related topics about IBM zERT Network Analyzer database administration enhancements

Book name	Topics
	• Updating z/OS for the IBM zERT Network Analyzer plug-in • Db2 for z/OS customization for the IBM zERT Network Analyzer task • Recovering unavailable partition
z/OS Management Facility online help	IBM zERT Network Analzyer online help

Chapter

2

z/OS Management Facility Configuration Guide

Topics:

- Updating z/OS for the IBM zERT Network Analyzer Plug-in
- Problems when using IBM zERT Network Analyzer

Updating z/OS for the IBM zERT Network Analyzer Plug-in

If you selected to configure the IBM zERT Network Analyzer plug-in, you might have system customization to perform, as described in this topic.

This topic contains the following information:

- Authorize users to the IBM zERT Network Analyzer task
- Db2 for z/OS customization for the IBM zERT Network Analyzer task
- Install Java Database Connectivity
- Connect IBM zERT Network Analyzer task with the Db2 for z/OS database.

The examples in this topic use RACF commands. If your installation uses an external security manager other than RACF, your security administrator can refer to these examples when creating equivalent commands for your environment.

For a summary of the required profile authorizations for the IBM zERT Network Analyzer plug-in, see *Resource authorizations for the IBM zERT Network Analyzer plug-in*.For a summary of the required profile authorizations for the IBM zERT Network Analyzer plug-in, see *Resource authorizations for the IBM zERT Network Analyzer plug-in*. IBM provides job IZUNASEC in SYS1.SAMPLIB to assist you with performing these updates. The job contains RACF commands for creating the required security authorizations.

Authorize users to the IBM zERT Network Analyzer task

Users of the IBM zERT Network Analyzer task require access to resources that are protected by the profile *<SAF-prefix>*.ZOSMF.ZERT_NETWORK_ANALYZER in class ZMAFPLA. Your z/OS security administrator must perform additional steps to create the necessary authorizations.

1. During the z/OSMF configuration process, edit the IZUNASEC job before you run it.
2. Add the names of the users who should be authorized to access the IBM zERT Network Analyzer task.

```
/*  Connect the users of the zERT Network Analyzer to the       */
/*  zERT Network Analyzer group                                 */
CONNECT USER1 GROUP(IZUZNA)
CONNECT USER2 GROUP(IZUZNA)
/*  End connect the users to zERT Network Analyzer group        */
```

3. Save your changes and run the updated IZUNASEC job.

Db2 for z/OS customization for the IBM zERT Network Analyzer task

The IBM zERT Network Analyzer task stores and queries SMF data in a Db2 for z/OS database. Before you can use the task, this database must be created in a suitable Db2 for z/OS subsystem and the connectivity information for the database must be configured in the IBM zERT Network Analyzer.

Requirement: A given IBM zERT Network Analyzer database must only be accessed by one IBM zERT Network Analyzer plug-in at a time. Concurrent access to a single database by more than one IBM zERT Network Analyzer plug-in generates unpredictable results.

IBM provides tooling and templates in the SYS1.SAMPLIB data set to help your local Db2 for z/OS database administrator (DBA) create the IBM zERT Network Analyzer database in a manner that conforms to local naming and resource conventions. Collectively, this tooling is called the IBM zERT Network Analyzer database schema tooling and it consists of the following members in SYS1.SAMPLIB:

- The IZUZNADT and IZUZNADA members contain Data Definition Language (DDL) templates for creating the required database objects using two different approaches.

  - IZUZNADT contains a template for creating the required database objects using a fixed schema name (SYSIBM_EZB_ZNADB) and fixed table names that the IBM zERT Network Analyzer depends on. If your local naming conventions allow for these fixed names, then you should use this template.
  - IZUZNADA contains a template for creating the required database objects using your own customized schema and/or table names. It also creates aliases for those tables using the fixed schema name and table names that the IBM zERT Network Analyzer depends on. If your local naming conventions require you to change the schema name and/or the table names, you should use this template.

  The templates contain all the appropriate DDL commands using variables for a wide variety of resource names and parameter values that the DBA may want to control. The DBA should set values for some or all these variables that will comply with the local Db2 for z/OS conventions and allocation strategies.

  Note: The IZUZNADT and IZUZNADA prolog commentary explain each of the variables they include as well as the required variable syntax.

- The IZUZNADI member of the SYS1.SAMPLIB data set is a sample variable substitution data set in which the DBA specifies the desired values for each of the template variables that are defined in the templates. The IZUZNADI sample specifies default values for each of the variables, but the DBA can override any of the variables as needed.

- The IZUZNADG DDL generation REXX exec reads the specified template and a customized variable substitution data set and generates a data set that contains a complete set of customized DDL for creating a new IBM zERT Network Analyzer database or to update an existing database to the most current schema level (for applying service).

  Note: Instructions for using the IZUZNADG exec are available by running the exec with the --HELP parameter.

Tip: When an IBM zERT Network Analyzer PTF requires a database schema change, the database schema tooling is updated as part of the PTF. The templates are always updated, and in some cases, the exec and/or variable substitution data set may also be updated. Before generating the DDL updates for such a PTF, make sure you are using the exec and templates that come with the PTF. In addition, if any new variables have been introduced, you must ensure that your customized variable substitution data set contains values for all of the new variables. The PTF's IZUZNADI variable substitution data set contains default values for the new variables.

After running the IZUZNADG exec, your DBA can use the resulting DDL data set in SPUFI or as input to a local JCL stream used to create the new IBM zERT Network Analyzer database in the Db2 for z/OS subsystem.

1. Determine the local resource requirements and location for the IBM zERT Network Analyzer database.

   Your DBA decides which Db2 for z/OS subsystem contains the IBM zERT Network Analyzer database objects as well as the specific Db2 for z/OS resources (storage pools, buffer pools, and so forth) to be allocated to these objects.

   Tip: You should initially deploy the IBM zERT Network Analyzer database and service on a test system. The test system should be a place where you can easily get familiar with the operation of the service and can better understand the Db2 and system resource requirements when running queries against your imported SMF record data. Depending on the number of imported SMF records and the complexity of your queries, you might consider initially limiting query execution to specific times of day or specific systems to minimize system impacts.

   - **Determine the maximum number of query reports that will be open at any given time**

     Your DBA needs to specify the number of partitions for a subset of tables in the IBM zERT Network Analyzer database called "Query Result Tables." These partitioned tables hold intermediate query results that are displayed through the network analyzer's **Report** tab. The IBM zERT Network Analyzer assigns one partition in each query result table to each active IBM zERT Network Analyzer query report for as long as that report is open in the web browser. Because of this, you must work with your DBA to determine an appropriate number

of partitions to ensure that your database has enough partitions to support your community of IBM zERT Network Analyzer users.

The number of partitions is determined by the *<QRTParts>* variable in the DDL templates and is controlled by the DBA in the variable substitution data set used with the IZUZNADG exec (see the IZUZNADI member of SYS1.SAMPLIB).

To calculate the number of partitions you need, consider the maximum number of users that will be using the IBM zERT Network Analyzer at the same time as well as the number of reports each user might have open at any given time (a single user can have multiple reports open at one time, with each one in its own web browser tab). Multiply those two numbers together to determine the maximum possible number of open reports. You should create at least that many partitions for each query result table. You might also want to increase this value by an appropriate percentage to ensure there is some room for growth over time.

The following equation summarizes the above calculation.

```
NumOpenReports = MaxNumberOfUsers * MaxNumberOfReportsPerUser
<QRTParts> = NumOpenReports + ( NumOpenReports * ExtraSpace% )
```

- **Decide whether to use a separate database for the query result tables**

  Two variables in the DDL templates control whether the query result tables are created in the same database as the persistent IBM zERT Network Analyzer tables, or in the same database.

<database>	specifies the name of the database that contains the IBM zERT Network Analyzer's persistent tables (all tables except the query result tables). When you are creating a brand new IBM zERT Network Analyzer database, set this variable to any valid database name. By default, the IZUZNADI sample sets this to 'ZNADB'.
<QRTDatabase>	specifies the name of the database that contains the query result tables. You may choose to store the query result tables in the same database as the persistent tables or in a different database. To use the same database, set the *<QRTDatabase>* and *<database>* variables to the same value. To use different databases, specify different values. By default, the IZUZNADI sample sets this to 'ZNAQRDB', placing the Query Result Tables in their own database.

- **Determine how much table space to allocate in the database**

  The amount of Db2 for z/OS table storage required by IBM zERT Network Analyzer varies, but you can use the following guidelines to estimate the table storage required in your environment:

  - Start with an allocation of 20 MB to hold the core security session data and operational data related to data management and user-built queries. If you are collecting zERT data from an unusually large number of unique security sessions across the z/OS systems, you might eventually need to add to this amount over time.
  - Estimate the space that is required to store the maximum number of SMF Type 119 zERT Summary (subtype 12) records that will be represented in the database at one time and add that amount to the initial

20 MB. To do so, see details in Table 3: How to estimate the space for SMF Type 119 zERT Summary (subtype 12) records on page 15.

Table 3: How to estimate the space for SMF Type 119 zERT Summary (subtype 12) records

Step	Task	Formula	Example
1	Estimate the number of unique security sessions that typically exist across the set of z/OS systems from which you are collecting zERT data.	Use local procedures to estimate this value.	Assume that 10,000 unique security sessions typically exist across all of the zERT-monitored systems. **UniqueSessions = 10,000**
2	Determine the maximum number of SMF intervals to be represented in the IBM zERT Network Analyzer database at a single time.	**MaxIntervals = ((1440 / SMFInterval) * Days)** where: **1440** = the number of minutes in a day; **SMFInterval** = the SMF interval in minutes as defined in your SMFPRMxx parmlib member. Note that if you use different SMF intervals across the zERT-monitored z/OS systems, use the average interval length across the different systems here; **Days** = the maximum number of days' worth of SMF data that is stored in the database at a single time.	Assume the average SMF interval is set to 20 minutes and you plan to store 30 days of SMF data in the zERT Network Analyzer database: MaxIntervals = (1440 / 20) * 30 = 72 * 30 = **2160**
3	Determine how many SMF 119 subtype 12 records are collected over the Days value and imported into the database.	**MaxRecords = UniqueSessions * MaxIntervals**	MaxRecords = 10,000 * 2160 = **21,600,000**
4	Determine how much DASD storage is required to store the data for the maximum number of records. Each SMF 119 subtype 12 record consumes about 500 bytes of DASD storage above and beyond the base space allocation of 20 MB.	**DASDSpaceMB = 20 + ((MaxRecords * 500) / 1,048,576)**	DASDSpaceMB = 20 + ((21,600,000 * 500) / 1,048,576) = 20 + 10,299 = **10,319MB or 10GB**

According to the example shown in the table, you need to allocate a total of 10 GB of space in the storage group that is used for your IBM zERT Network Analyzer database's table spaces.

2. Create the IBM zERT Network Analyzer database.

    Your DBA should use the database schema tooling described above to create the database for your environment.

3. Define the IBM zERT Network Analyzer database user ID.

   This is the z/OS user ID that is permitted to connect to, store data into, and query data in the IBM zERT Network Analyzer database. The IBM zERT Network Analyzer uses this user ID to communicate with the Db2 for z/OS database and to perform all the operations in its database.

   The database user ID must be given the INSERT, SELECT, UPDATE, DELETE privileges for the IBM zERT Network Analyzer database tables to ensure proper operation of the IBM zERT Network Analyzer's various functions as described below.

   If you use the IZUZNADT template or if you use the IZUZNADA aliasing template with SQL GRANT access controls, then grant INSERT, SELECT, UPDATE, DELETE privileges to:

   - SYSIBM_EZB_ZNADB.APPL
   - SYSIBM_EZB_ZNADB.DATAMGMTHISTORY
   - SYSIBM_EZB_ZNADB.DATASET
   - SYSIBM_EZB_ZNADB.SECURITY_SESSION
   - SYSIBM_EZB_ZNADB.SESSION_STATISTICS
   - SYSIBM_EZB_ZNADB.IPSEC_INFO
   - SYSIBM_EZB_ZNADB.SSH_INFO
   - SYSIBM_EZB_ZNADB.TLS_INFO
   - SYSIBM_EZB_ZNADB.TOPOLOGY
   - SYSIBM_EZB_ZNADB.OPENJPA_SEQUENCE_TABLE
   - SYSIBM_EZB_ZNADB.QUERY
   - SYSIBM_EZB_ZNADB.SCOPE_FLTR
   - SYSIBM_EZB_ZNADB.SCOPE_FLTR_ENDPT
   - SYSIBM_EZB_ZNADB.SCOPE_FLTR_SYSSPEC
   - SYSIBM_EZB_ZNADB.SEC_FLTR
   - SYSIBM_EZB_ZNADB.SEC_IPSEC_FLTR
   - SYSIBM_EZB_ZNADB.SEC_SSH_FLTR
   - SYSIBM_EZB_ZNADB.SEC_TLS_FLTR
   - SYSIBM_EZB_ZNADB.FILTEREDSECURITYSESSIONIDS
   - SYSIBM_EZB_ZNADB.TCPSERVER_SUMMARIES
   - SYSIBM_EZB_ZNADB.TCPCLIENT_SUMMARIES
   - SYSIBM_EZB_ZNADB.EEPEER_SUMMARIES
   - SYSIBM_EZB_ZNADB.TCPSERVER_CLIENTDETAILS
   - SYSIBM_EZB_ZNADB.TCPCLIENT_CLIENTDETAILS
   - SYSIBM_EZB_ZNADB.EEPEER_CLIENTDETAILS
   - SYSIBM_EZB_ZNADB.TCPSERVER_CLEARSECURITYSESSIONDETAILS
   - SYSIBM_EZB_ZNADB.TCPSERVER_IPSECSECURITYSESSIONDETAILS
   - SYSIBM_EZB_ZNADB.TCPSERVER_SSHSECURITYSESSIONDETAILS
   - SYSIBM_EZB_ZNADB.TCPSERVER_TLSSECURITYSESSIONDETAILS
   - SYSIBM_EZB_ZNADB.TCPCLIENT_CLEARSECURITYSESSIONDETAILS
   - SYSIBM_EZB_ZNADB.TCPCLIENT_IPSECSECURITYSESSIONDETAILS
   - SYSIBM_EZB_ZNADB.TCPCLIENT_SSHSECURITYSESSIONDETAILS
   - SYSIBM_EZB_ZNADB.TCPCLIENT_TLSSECURITYSESSIONDETAILS
   - SYSIBM_EZB_ZNADB.EEPEER_CLEARSECURITYSESSIONDETAILS
   - SYSIBM_EZB_ZNADB.EEPEER_IPSECSECURITYSESSIONDETAILS

   If you use the IZUZNADA aliasing template with SAF-based access controls on Db2 for z/OS objects, the INSERT, SELECT, UPDATE, DELETE privileges must be granted for the base tables (not the aliases) as specified by the values of the following template variables (these values are your custom schema and table names):

   - <schema>.<appTable>
   - <schema>.<dmhistTable>

- <schema>.<dsTable>
- <schema>.<secsessTable>
- <schema>.<sessstatsTable>
- <schema>.<ipsecTable>
- <schema>.<sshTable>
- <schema>.<tlsTable>
- <schema>.<topoTable>
- <schema>.<queryTable>
- <schema>.<scopeFltrTable>
- <schema>.<scopeFltrEndptTable>
- <schema>.<scopeFltrSysspecTable>
- <schema>.<secFltrTable>
- <schema>.<secIpsecFltrTable>
- <schema>.<secSshFltrTable>
- <schema>.<secTlsFltrTable>
- <schema>.<openjpaTable>
- <schema>.<fssIdsTable>
- <schema>.<tsrvrsTable>
- <schema>.<tclntsTable>
- <schema>.<teepTable>
- <schema>.<tsrvrcTable>
- <schema>.<tclntcTable>
- <schema>.<teepcTable>
- <schema>.<tsrvrCSessTable>
- <schema>.<tsrvrISessTable>
- <schema>.<tsrvrSSessTable>
- <schema>.<tsrvrTSessTable>
- <schema>.<tclntCSessTable>
- <schema>.<tclntISessTable>
- <schema>.<tclntSSessTable>
- <schema>.<tclntTSessTable>
- <schema>.<teepCSessTable>
- <schema>.<teepISessTable>

4. Collect the connectivity information that is required to link the IBM zERT Network Analyzer service with the Db2 for z/OS database to be used by the service.

   The DBA must provide database connectivity information to the person setting up the IBM zERT Network Analyzer service. This information includes:

   - The hostname or IP address on which the Db2 for z/OS subsystem is running
   - The TCP port number of the subsystem
   - The database location name, which is the value of the LOCATION parameter of the DSNJU003 utility
   - The database user ID
   - The password for the database user ID
   - The JDBC classpath for the Db2 for z/OS JDBC driver on the system where the IBM zERT Network Analyzer executes

   See Connect IBM zERT Network Analyzer task with the Db2 for z/OS database for how to use this connectivity information.

5. (**Optional**) Modify the setting of the DSN6SYSP URLGWTH parameter.

   Depending on the size of the SMF dump data sets that you plan to import into IBM zERT Network Analyzer and the setting of the DSN6SYSP URLGWTH parameter, you might see one or more DSNJ031I messages when

importing the SMF dump data sets. You can modify the setting of the DSN6SYSP URLGWTH parameter to reduce the number of DSNJ031I messages. For more information, see *UR LOG WRITE CHECK field (URLGWTH subsystem parameter)* in *Installing and migrating Db2*.

6. **(Optional)** Define additional 4K and 32K work files to be used by IBM zERT Network Analyzer

By default, small 4K and 32K work files are defined for use by the Db2 for z/OS subsystem. You might need to increase the size of the 4K and 32K work files to allow IBM zERT Network Analyzer to operate more efficiently. For more information, see *DSNTIP9: Work file database panel* in *Installing and migrating Db2*.

Restriction:

- Db2 for z/OS packages associated with the NULLID collection will be used when IBM zERT Network Analyzer plug-in connects to the Db2 for z/OS subsystem. The collection-id is an optional parameter when binding Db2 for z/OS packages where NULLID is the default collection-id value for distributed applications such as IBM zERT Network Analyzer. For more information, review the DSNTIJLC and DSNTIJLR Db2 for z/OS jobs.

- APPLCOMPAT bind option for the DB2 for z/OS packages associated with dynamic SQL statements must be set to its respective Db2 for z/OS release level. This means that a Db2 12 for z/OS subsystem must use an APPLCOMPAT value of at least V12R1M500 and a Db2 11 for z/OS subsystem must use an APPLCOMPAT value of V11R1. APPLCOMPAT is an optional parameter when binding Db2 for z/OS packages where the default APPLCOMPAT value is the APPLCOMPAT subsystem parameter. For more information, review the DSNTIJLC and DSNTIJLR Db2 for z/OS jobs.

Install Java Database Connectivity

The IBM zERT Network Analyzer task uses the Java Persistence API (JPA) to access the contents of the Db2 database. JPA, in turn, uses Java Database Connectivity (JDBC). If you have not already installed JDBC, or you have not run the DB2Binder utility as part of that installation, you have additional customization steps to perform.

Follow the instructions in *Installing the IBM Data Server Driver for JDBC and SQLJ as part of a Db2 installation* in *Programming for Db2 for z/OS*.

Connect the IBM zERT Network Analyzer task with the Db2 for z/OS database

You must provide the IBM zERT Network Analyzer task with Db2 for z/OS database connectivity information before using the task for any additional functions.

1. Launch the IBM zERT Network Analyzer task. The first time you launch the task, you are directed immediately to the **Database Settings** panel.

   - To open the IBM zERT Network Analyzer task using the z/OSMF traditional view, expand the **Analysis** category in the navigation area, and select **IBM zERT Network Analyzer**.
   - To open the IBM zERT Network Analyzer task using the z/OSMF desktop view, click the **IBM zERT**

   Network Analyzer icon   .

2. Enter the database connectivity information that is provided to you by the database administrator (DBA) as part of completing the Db2 for z/OS customization for the IBM zERT Network Analyzer task procedure.

   Note: After you save the connectivity information, the IBM zERT Network Analyzer task restarts using the configuration information. You must stop and restart your browser session when the task restarts.

Problems when using IBM zERT Network Analyzer

This section provides a procedure you can use to send troubleshooting documentation to IBM Support.

Steps for sending information to IBM Support

For a failure in the IBM zERT Network Analyzer task, use the following procedure to gather troubleshooting documentation. This information can be used by IBM Support to diagnose the problem.

1. Use the **Applications Settings** sub-panel on the IBM zERT Network Analyzer task to ensure that the proper level of logging is being used. By default, IBM zERT Network Analyzer writes Info level messages to the z/OSMF runtime logs, but writes no logging information to the IBM zERT Network Analyzer logs.

   After the necessary level of logging is in place, re-create the problem that was reported, if needed.

2. Transfer the z/OSMF runtime log files that contain IBM zERT Network Analyzer logging data.

   During normal operations, z/OSMF collects its runtime data (log messages and trace messages) in log files. z/OSMF runtime data is created on the server or sent to the server by the client. Both types of messages are written to the z/OSMF runtime log files.

   In addition, if IBM zERT Network Analyzer logging is enabled, the IBM zERT Network Analyzer task writes some runtime data (log messages and trace messages) to the z/OSMF runtime log files.

   z/OSMF creates the log files in the product logs directory, which is, by default, `/global/zosmf/data/logs`. z/OSMF names the log files IZUG*n*.log, where *n* is a numeral in the range 0 - 9.

   z/OSMF creates log files in a "cascading" manner. The most current log file is always named IZUG0.log. When this log file reaches its predefined limit, z/OSMF saves it as IZUG1.log and begins writing to a new IZUG0.log file. When the IZUG0.log file is again full, z/OSMF saves it as IZUG1.log after it renames the existing IZUG1.log file to IZUG2.log. z/OSMF continues this process, saving each log file under the next available name, up to a maximum of ten log files. Thereafter, z/OSMF discards the oldest log file (IZUG9.log) whenever a new log file must be created.

3. Transfer the IBM zERT Network Analyzer log files that contain IBM zERT Network Analyzer logging data.

   During normal operations, when logging is enabled, IBM zERT Network Analyzer collects runtime data (log messages and trace messages) specific to IBM zERT Network Analyzer processing in its own separate log files. Only server-side information is written to the IBM zERT Network Analyzer runtime log files.

   By default, IBM zERT Network Analyzer creates its log files in the `/global/zosmf/data/app/<release>/debug` folder, where `<release>` is the current level of IBM zERT Network Analyzer. The format of `<release>` is ZNA*xxxx*, where *xxxx* represents the z/OSMF release level. For example, IBM zERT Network Analyzer running on z/OSMF V2R3 would use ZNAV2R3 for `<release>`. The full name of the default folder would be `/global/zosmf/data/app/ZNAV2R3/debug`. IBM zERT Network Analyzer names the files ZnaDbgn.log, where n is a numeric value in the range 0 - 24.

   IBM zERT Network Analyzer creates log files in a "cascading" manner. The most current log file is always named ZnaDbg0.log. When this log file reaches its predefined limit, IBM zERT Network Analyzer saves it as ZnaDbg1.log and begins writing to a new ZnaDbg0.log file. When the ZnaDbg0.txt file is again full, IBM zERT Network Analyzer saves it as ZnaDbg1.log after it renames the existing ZnaDbg1.log file to ZnaDbg2.log. IBM zERT Network Analyzer continues this process, saving each log file under the next available name, up to the current maximum of log files. Thereafter, IBM zERT Network Analyzer discards the oldest log file (ZnaDbgn.log) whenever a new log file must be created.

   By default, the maximum number of log files maintained is 10. You can modify that setting on the IBM zERT Network Analyzer **Applications Settings** sub-panel, up to a maximum of 25 possible log files.

4. Package the files and send them to IBM Support.

Recovering unavailable partitions

In unusual and infrequent situations, an active partition of a query result table might become unavailable. When a partition is in unavailable state, the IBM zERT Network Analyzer is unable to use it to store query results or to free the partition. In order to make the partition available for use, database administrator must recover the partition.

To determine the partitions that are unavailable, the IBM zERT Network Analyzer logs the following message for each unavailable partition when the plugin starts up:

```
'IZUET0049E Error freeing up partition partition-id during IBM zERT Network
Analyzer startup'
```

where *partition-id* is the partition identifier.

This message may be preceded by an IZUET0034 message:

```
IZUET0034E: Exception encountered: stack dump caused by "UNSUCCESSFUL
EXECUTION CAUSED BY AN UNAVAILABLE RESOURCE, REASON reason-code,TYPE OF
RESOURCE resource-type, AND RESOURCE NAME resource-name SQLCODE=-904,
SQLSTATE=57011"
```

For more information on the SQL error code -904, see SQL error codes in the Db2 for *z/OS Codes Book*.

To recover an unavailable partition, see the following steps.

1. Run the Db2 for z/OS RECOVER utility on all query result table spaces associated with the partition-id in the IZUET0049E message. The RECOVER utility recovers all data to the current state or to a previous point in time by restoring a copy and applying log records or by undoing all committed work. In order to use this utility, you must have one the following authorities:

   • RECOVERDB privilege for the database
   • DBADM or DBCTRL authority for the database
   • System DBADM authority
   • DATAACCESS authority
   • SYSCTRL or SYSADM authority

   Use the TABLESPACE option on the RECOVER utility control statement to recover all query result table spaces associated with the partition-id. To recover multiple table spaces, repeat the TABLESPACE option before each specified table space in the following statement:

   ```
   RECOVER TABLESPACE database-name.tablespace-name DSNUM unavailable-
   partition-number
   ```

   Where *database-name* is the value specified for the <QRTSDatabase> variable, *tablespace-name* is the value specified for each <*Space> variable, and *unavailable-partition-number* is the partition-id value in the IZUET0049E message.

   You can find the variables mentioned above in the Query Result Tables database variables section of your customized version of the IZUZNADI dataset used to generate the DDL that created the IBM zERT Network Analyzer database.

   For example:

   ```
   RECOVER TABLESPACE QRTSDB.FSSIDSS DSNUM 1
           TABLESPACE QRTSDB.TCLNCDS DSNUM 1
     TABLESPACE QRTSDB.TSRVSSS DSNUM 1
    TABLESPACE QRTSDB.EPEERSS DSNUM 1
    …
     TABLESPACE QRTSDB.TCLNSS   DSNUM 1
   ```

   For more information on syntax and examples for the recovery untility, see *Recover* and *Recovering a table space* in Db2 for *z/OS Utilities Book*.

2. Run the Db2 online utility. For detailed information on the different ways to invoke and steps to run the utility, see *Invoking Db2 online utilities* in the Db2 for z/OS book.

3. After recovering the partitions, you must recycle z/OSMF for IBM zERT Network Analyzer to detect the availability of the partitions.

Index

I

IBM zERT Network Analyzer
      Analysis 12
      authorize users 12
      Connect Db2 19
      Db2 for z/OS customization 12
      Install Java Database Connectivity 19
IBM zERT Network Analyzer task
      collecting information for troubleshooting 20
      common problems 20
      recovering unavailable partition 21
      troubleshooting 20

T

troubleshooting
      IBM zERT Network Analyzer task 20