

Date of change: March 2017

Topic: Multiple

Change description: Documentation changes made in support of PTF UI44867 APAR PI70981 Accelerator data backup and recovery

The following topics have been updated or introduced:

In chapter “Overview”:

Topic: “What's new”

Topic: “Features and benefits”

Topic: “Scenarios”

In chapter “Customizing DB2 Analytics Accelerator Loader”:

Topic: “Copying the started task PROC (required)”

In chapter “Loading data from an external file”

Topic: “Restrictions and considerations for loading from an external file”, subtopic “Considerations for loading only the accelerator”

Topic: “Using the ISPF interface to create or edit a profile to load from an external file”

New chapter “Backing up and recovering accelerator data” with following new topics:

New topic: “About backup copy datasets”

New topic: “Restrictions and considerations for backing up and recovering accelerator data”

New topic: “Using the ISPF interface to back up accelerator data using the BACKUP utility”

New topic: “Using the ISPF interface to back up accelerator data using an inline backup”

New topic: “Using the ISPF interface to recover accelerator data”

In chapter “Using and managing load profiles”:

Topic: “Introduction”

In chapter “Syntax”:

Topic: “Loading from an external file”, subtopics “Example JCL: Loading from an external file”, “Syntax diagram: Load from an external file”, “Syntax definitions: Load from an external file”

New topic: “Creating a backup using the BACKUP utility”

In chapter “Troubleshooting”:

Topic: Messages and codes

In chapter “Reference”:

Topic: “Accelerator Loader terminology”

Topic: “Main menu”

Topic: “Manage Loader Profiles panel”

Topic: “Create Profile panel”

Topic: “Save Profile panel”

Topic: “Load Accelerator from External File panel”

Topic: “Rename Profile panel”

New topic: “Back up Accelerator Table panel”

New topic: “Accelerator Table Selection panel”

New topic: “Copy Data Set Parameters panel”

New topic: “Recover Accelerator Table(s) from a Backup panel”

New topic: “Recovery Table List panel”

New topic: “Recovery Table Selection panel”

New topic: “Backup Copy Selection panel”

New topic: “Use alternate backup panel”

Chapter “Overview”

Topic: “What’s new”

Add the following description:

Use Accelerator Loader to back up and recover data that resides only in the accelerator. This data can be in an accelerator-only table or an accelerator-shadow table that has been loaded to the accelerator only.

Topic: “Features and benefits”

Add the following description:

Accelerator backup and recovery

Use Accelerator Loader to back up and recover data that resides only in the accelerator. This data can be in an accelerator-only table or an accelerator-shadow table that has been loaded to the accelerator only.

Note: Because the data resides only in the accelerator, the standard DB2 COPY and RECOVER utilities cannot be used.

Backup copies can be created using either of the following methods:

- *Backup utility.* The Accelerator Loader backup utility lets you create a full copy by fetching all of the data from the accelerator table and writing out a copy. To use this method, you can generate JCL from a Backup profile.
- *Inline copy.* An accelerator only load can optionally be configured to create a backup copy as the data is loaded to the accelerator. This is the most efficient way to create a backup. This method creates a full copy when running the Accelerator Loader with the syntax LOAD REPLACE and an incremental copy when running the Accelerator Loader with the syntax LOAD RESUME. To use this method, you can generate JCL using an Accelerator only profile.

To generate recovery JCL, you can use the Recovery profile.

The backup and recovery feature supports four copy data sets: a primary and backup copy for the local site, and a primary and backup copy for the remote recovery site. Copy datasets are registered in the backup copy registration table, HLOUCOPY; each DB2 system on which Accelerator Loader is installed has its own copy of this table.

Topic: “Scenarios”

Add the following description:

Backing up or recovering accelerator data

You have data that resides only in the accelerator, either in an accelerator-only table or in an accelerator-shadow table that has been loaded to the accelerator only. You need to make a backup or to recover this data, but because the data resides only in the accelerator, the standard DB2 COPY and RECOVER utilities cannot be used.

To back up or recover this accelerator data, you can use features available in Accelerator Loader:

- To back up accelerator data, you can make a full copy using the Accelerator Loader backup utility, or you can make full or incremental copies inline during an accelerator only load. You can generate backup JCL using the Backup profile or the Accelerator only profile, as appropriate.
- To recover accelerator data, you can generate recovery JCL using the Recovery profile.

For more information, see chapter “*Backing up and recovering accelerator data*”.

Chapter “Customizing DB2 Analytics Accelerator Loader”

Topic: “Copying the started task PROC (required)”

Replace the following step in the procedure:

7. If you plan to use the high availability load utility (HALOAD) or the backup utility, ensure that the product load library is in the JOBLIB or STEPLIB.

Chapter “Loading data from an external file”

Topic: “Restrictions and considerations for loading from an external file”, subtopic “Considerations for loading only the accelerator”

Remove following restriction:

- You need a backup of the data. (DB2 does not provide backup or recovery of the data.)

Add following bullet:

- When you load only to the accelerator, you can create an inline backup copy as the data is loaded to the accelerator.

Topic: “Using the ISPF interface to create or edit a profile to load from an external file”

Add the following step to the end of the procedure.

12. To create an inline backup copy for the target table, specify data set names for the Inline copy data sets options. Inline backup copies can be created for accelerator-only tables or accelerator-shadow tables that have been loaded to the accelerator only.

New chapter “Backing up and recovering accelerator data”

Add as new chapter after the “Loading DB2 table data to multiple accelerators” chapter.

You can use Accelerator Loader to back up and recover data that resides only in the accelerator. This data can be in an accelerator-only table or an accelerator-shadow table that has been loaded to the accelerator only.

Note: Because the data resides only in the accelerator, the standard DB2 COPY and RECOVER utilities cannot be used.

Backup copies can be created using either of the following methods:

- *Backup utility.* The Accelerator Loader backup utility lets you create a full copy by fetching all of the data from the accelerator table and writing out a copy. To use this method, you can generate JCL from a Backup profile.
- *Inline copy.* An *inline copy* is a backup copy of an accelerator table that is created as the data is loaded to the accelerator. An accelerator only load can optionally be configured to create an inline copy, and this is the most efficient way to create a backup. This method creates a full copy when running the Accelerator Loader with the syntax LOAD REPLACE and an incremental copy when running the Accelerator Loader with the syntax LOAD RESUME. To use this method, you can generate JCL using an Accelerator only profile.

To generate recovery JCL, you can use the Recovery profile.

New topic: “About backup copy datasets”

The Accelerator Loader backup and recovery feature supports four copy data sets: a primary and backup copy for the local site, and a primary and backup copy for the remote recovery site. You can create backups for the local site only or the recovery site only. A backup copy for either site may be created only when a primary copy is also being created for that site. You can specify data set allocation parameters within your Backup profile. The backup program determines and sets the RECFM, LRECL, and BLKSIZE.

Copy datasets are registered in the backup copy registration table, HLOUCOPY; each DB2 system on which Accelerator Loader is installed has its own copy of this table.

New topic: “Restrictions and considerations for backing up and recovering accelerator data”

Review the following usage restrictions and considerations before performing a backup or recovery of your accelerator data.

- Accelerator Loader backup and recovery features are for data that resides *only* in the accelerator.
- Because the data resides only in the accelerator, standard DB2 COPY and RECOVER utilities cannot be used for backup and recovery functions.
- Backup copies created by the Accelerator Loader backup feature are not in standard DB2 image copy format. Accelerator tables can be recovered using the Accelerator Loader recovery feature only.
- Removing a table from the accelerator invalidates any backup copies that have been created. If the table is removed from the accelerator and then added back, a new full backup copy should be created.

New topic: “Using the ISPF interface to back up accelerator data using the BACKUP utility”

The Accelerator Loader BACKUP utility lets you create a full copy of your accelerator table by fetching all of the data from the accelerator table and writing out a copy. To use the BACKUP utility, you can create a Backup profile from which you can generate backup JCL.

A Backup profile is a reusable group of options for building jobs to back up data on an accelerator using the BACKUP utility. You can create a profile that saves your selections and reuse the profile to perform future backups of your accelerator data.

About this task

To build the JCL required to run the BACKUP utility, you must specify the accelerator table to back up and the copy data sets to create. You can specify up to four copy data sets to create: a primary and backup copy for the local site, and a primary and backup copy for the remote recovery site. You can create backups for the local site only or the recovery site only. A backup copy for either site can be created only when a primary copy is also being created for that site.

Procedure

1. From the main menu, select **Back up Accelerator table** and press Enter.
2. On the Back up Accelerator Table panel, issue the **TABLE** command.
3. On the Enter Table and Creator Like to Display panel, specify a table creator and table name pattern for the accelerator tables to list, and press Enter. You can use an asterisk (*) in the fields Table creator like and Table name like. Case sensitivity of these fields is controlled by the Case sensitive option.
4. On the Accelerator Table Selection panel, type **S** in the Cmd field next to the table to back up, and press Enter. Only one table can be selected. After you select a table, an asterisk appears in the Cmd field.
5. Press **PF3** to exit the panel.
6. Under Copy data sets options, specify up to four copy data sets to create, and specify **YES** in the Update field to specify data set parameters.
7. If you specified **YES** for Update, on the Copy Data Set Parameters panel for each copy data set, specify allocation parameters for the backup copy data set, and press **PF3**.
8. Optional: To save the Backup profile, specify a name and description for the profile, and issue the **SAVE** command.
9. To build the backup JCL, issue the **BUILD** command.

New topic: “Using the ISPF interface to back up accelerator data using an inline backup”

An *inline copy* is a backup copy of an accelerator table that is created as the data is loaded to the accelerator. An accelerator only load can optionally be configured to create an inline copy. This method creates a full copy when running the Accelerator Loader with the syntax **LOAD REPLACE** and an incremental copy when running the Accelerator Loader with the syntax **LOAD RESUME**.

To create an inline copy when data is loaded to the accelerator, you can use an Accelerator only profile from which you can generate JCL. To create an Accelerator only profile, use the **Load Accelerator(s) from External File** option from the main menu. Use the Inline copy data sets options to specify up to four copy data sets to create: a primary and backup copy for the local site, and a primary and backup copy for the remote recovery site. You can create backups for the local site only or the recovery site only. A backup copy for either site can be created only

when a primary copy is also being created for that site. Specify YES in the Update field to specify data set parameters.

For more information on defining the Accelerator only profile, see chapter “*Loading data from an external file*”.

New topic: “Using the ISPF interface to recover accelerator data”

To recover accelerator data, you can create a Recovery profile from which you can generate recovery JCL.

A recovery profile is a reusable group of options for building jobs to recover data on an accelerator. You can create a profile that saves your selections and reuse the profile to perform any future recoveries of data.

About this task

For Accelerator Loader recovery JCL, you must provide the accelerator table to recover and the backup copy data set to use for recovery. To determine which backup copy data set to use, you must first decide the point in time to which to recover. The point in time option in the ISPF interface controls how the copy data set is selected for each table and applies to all tables selected for recovery. The following options are available for Point in time:

- **CURRENT:** The backup data set for each table will be chosen automatically when the recovery JCL is generated.
- **TIMESTAMP:** The backup data set for each table will be chosen automatically using values specified in Timestamp end point and Time zone of timestamp fields.
- **SELECTED:** You must manually specify a backup data set for each selected table.

Procedure

1. From the main menu, select **Recover Accelerator table(s) from a backup** and press Enter.
2. On the Recover Accelerator Table(s) from a Backup panel, issue the TABLES command.
3. On the Recovery Table List panel, issue the ADD command.
4. On the Enter Table and Creator Like to Display panel, specify a table creator and table name pattern for the accelerator tables to list, and press Enter. You can use an asterisk (*) in the fields Table creator like and Table name like. Case sensitivity of these fields is controlled by the Case sensitive option.
5. On the Recovery Table Selection panel, type S in the Cmd field next to a table to recover or use the ALL command to select all tables, and press Enter. After you select a table, an asterisk appears in the Cmd field.
6. Press PF3 to exit the panel.
7. To manually specify the backup copy data set to use for the recovery:
Note: To use a manually-selected backup copy data set requires the **SELECTED** point in time recovery option, specified in a later step.
 - a. On the Recovery Table List, type B in the Cmd field next to the table, and press Enter.
 - b. On the Backup Copy Selection panel, type S next to a backup copy data set and press Enter.
 - c. Press PF3 to exit the panel.
8. Press PF3 to exit the Recovery Table List panel.
9. On the Recover Accelerator Table(s) from a Backup, specify your target options:
 - a. To add missing tables to the accelerator before starting the recover job, specify YES for **Add table to Accelerator**.
 - b. To enable query acceleration for the table after a successful load, specify YES for **Acceleration on success**.
10. Specify your recovery point options, as follows. These settings apply to all tables that are selected.
 - a. Specify **CURRENT**, **TIMESTAMP** or **SELECTED** for the **Point in time** to which to recover.
 - b. If you specified **TIMESTAMP** for your point in time, specify the **Timestamp end point** and **Time zone of timestamp** values.
11. Optional: To save the Recovery profile, specify a name and description for the profile, and issue the SAVE command.
12. To build the recovery JCL, issue the BUILD command.

Chapter “Using and managing load profiles”

Add the following profile types:

- **Backup** specifies options for backing up a table defined to the accelerator.
- **Recovery** specifies options for recovering a table defined to the accelerator.

Chapter “Syntax”

Topic: “Loading from an external file”

Subtopic: “Example JCL: Loading from an external file”

Add new examples.

Example: Creating a backup using an inline copy

An *inline copy* is a backup copy of an accelerator table that is created as the data is loaded to the accelerator. This method uses an Accelerator Loader accelerator only load, as follows:

- During a LOAD REPLACE, a full copy is created.
- During a LOAD RESUME, an incremental copy is created.

To use this method, you can generate JCL using an Accelerator only profile and include values for Inline copy data sets options.

The following example JCL shows an accelerator only load that includes options for an inline copy to four copy data sets.

```
//HLOD0100 EXEC PGM=DSNUTILB,
//          REGION=0000M,
//          PARM=('UB1A')
//STEPLIB DD DISP=SHR,DSN=HLO.PR00210.LOADLIB
//          DD DISP=SHR,DSN=UB1A.SDSNEXIT
//          DD DISP=SHR,DSN=DSN.VB10.SDSNLOAD
//ISYSREC DD DISP=SHR,DSN=USER1.HLO.SYSREC
//HLOCPYLP DD DSN=USER1.HLO1533.LP,DISP=(NEW,CATLG,DELETE),
//          SPACE=(CYL,(1,1))
//HLOCPYLB DD DSN=USER1.HLO1533.LB,DISP=(NEW,CATLG,DELETE),
//          SPACE=(CYL,(1,1))
//HLOCPYRP DD DSN=USER1.HLO1533.RP,DISP=(NEW,CATLG,DELETE),
//          SPACE=(CYL,(1,1))
//HLOCPYRB DD DSN=USER1.HLO1533.RB,DISP=(NEW,CATLG,DELETE),
//          SPACE=(CYL,(1,1))
//HLODUMMY DD DUMMY
//SYSPRINT DD SYSOUT=*
//UTPRINT DD SYSOUT=*
//SYSIN DD *
    TEMPLATE ISYSUT1
      DSN '&US..IDSU.&DB..&TS..&UQ.'
      DISP(MOD,DELETE,CATLG)
      SPACE (10,100) CYL
    TEMPLATE ISORTOUT
      DSN '&US..IDSO.&DB..&TS..&UQ.'
      DISP(MOD,DELETE,CATLG)
      SPACE (10,100) CYL
    LOAD DATA
      IDAA_ONLY ON 'UB1AACCI'
      ACCEL_ON_SUCCESS_ENABLE NO
      INDDN ISYSREC
      ACCEL_LOAD_TASKS 1
      RESUME YES
      ACCEL_COPYDDN(HLOCPYLP, HLOCPYLB)
      ACCEL_RECOVERYDDN(HLOCPYRP, HLOCPYRB)
      WORKDDN(ISYSUT1,ISORTOUT)
      INTO TABLE
        "USER1"."TSTTBL"
/*
//*
```

Example: Recovering accelerator table data

To recover accelerator table data, use option **Recover Accelerator table(s) from a backup** from the main menu to create a new Recovery profile, from which you can generate recovery JCL. If multiple tables are selected, the generated JCL will include multiple Load steps, one for each table.

Warning! Manually generating recovery JCL is not recommended. The highly recommended procedure for generating recovery JCL is to use the ISPF interface. See “*Using the ISPF interface to recover accelerator data*”.

Because this is a load from FORMAT INTERNAL SYSREC datasets, no field specifications are needed.

```
//HLOD0100 EXEC PGM=DSNUTILB,
//          REGION=0000M,
//          PARM=('UB1A')
//STEPLIB DD DISP=SHR,DSN=HLO.PRD0210.LOADLIB
//          DD DISP=SHR,DSN=UB1A.SDSNEXIT
//          DD DISP=SHR,DSN=DSN.VB10.SDSNLOAD
//ISRECAAA DD DSN=USER1.HLO.LP,DISP=SHR
//HLODUMMY DD DUMMY
//SYSPRINT DD SYSOUT=*
//UTPRINT DD SYSOUT=*
//SYSIN DD *
  LOAD DATA INDDN ISRECAAA
  REPLACE
  IDAA_ONLY ON UB1AACCL
  LOG NO NOCOPYPEND
  ACCEL_REMOVE_AND_ADD_TABLES
  ACCEL_ON_SUCCESS_ENABLE NO
  FORMAT INTERNAL
  INTO TABLE
  USER1.TSTTBL
/*
```

Subtopic: “Syntax diagram: Load from an external file”

Add the following elements to the syntax diagram for the load-from-external job.

```
>----- ACCEL_COPYDDN( hlocpylp_ddname ----- ) ----->
                        |                               |
                        +- ,hlocpylb_ddname +-
>----- ACCEL_RECOVERYDDN( hlocpyrp_ddname ----- ) ->
                        |                               |
                        +- ,hlocpyrb_ddname +-

```

Subtopic “Syntax definitions: Load from an external file”

Add the following syntax elements.

ACCEL_COPYDDN *hlocpylp_ddname,hlocpylb_ddname*

Specifies the DD names for the backup data sets for the local site. *hlocpylp_ddname* is the DD name for local site primary copy data set, and *hlocpylb_ddname* is the DD name for the local site backup copy data set. This option is needed only when local site copies are being created.

The default DD names that are generated in the JCL are HLOCPYLP and HLOCPYLB; however, you can use any value for your DD names as long as they match a DD name in the JCL.

ACCEL_RECOVERYDDN *hlocpyrp_ddname,hlocpyrb_ddname*

Specifies the DD names for the backup data sets for the remote recovery site. *hlocpyrp_ddname* is the DD name for the recovery site primary copy data set, and *hlocpyrb_name* is the DD name for the recovery site backup copy data set. This option is needed only when recovery site copies are being created.

The default DD names that are generated in the JCL are HLOCPYRP and HLOCPYRB; however, you can use any value for your DD names as long as they match a DD name in the JCL.

Syntax definitions:

tbcreator.tbname

Specifies the accelerator table to back up or recover.

ACCEL_COPYDDN *hlocpylp_ddname,hlocpylb_ddname*

Specifies the DD names for the backup data sets for the local site. *hlocpylp_ddname* is the DD name for local site primary copy data set, and *hlocpylb_ddname* is the DD name for the local site backup copy data set. This option is needed only when local site copies are being created.

The default DD names that are generated in the JCL are HLOCPYLP and HLOCPYLB; however, you can use any value for your DD names as long as they match a DD name in the JCL.

ACCEL_RECOVERYDDN *hlocpyrp_ddname,hlocpyrb_ddname*

Specifies the DD names for the backup data sets for the remote recovery site. *hlocpyrp_ddname* is the DD name for the recovery site primary copy data set, and *hlocpyrb_name* is the DD name for the recovery site backup copy data set. This option is needed only when recovery site copies are being created.

The default DD names that are generated in the JCL are HLOCPYRP and HLOCPYRB; however, you can use any value for your DD names as long as they match a DD name in the JCL.

Chapter “Troubleshooting”

Topic: “Messages and codes”

Add the following messages.

HLO594E Invalid profile type. Enter one of the listed values (ALL,1-7).

Explanation: The specified value is not valid. Valid values are 1 for Dual, 2 for Accelerator only, 3 for Consistent, 4 for Image Copy, 5 for Multi, 6 for Backup, 7 for Recovery, and ALL.

User response: Enter a valid value.

HLO621E Invalid value. Enter a number from 1 to 7.

Explanation: The specified value is not valid.

User response: Enter a valid value.

HLO638E Disposition is required. Enter a valid data set disposition.

Explanation: You must specify a data set disposition.

User response: Enter a valid data set disposition value.

HLO639E Invalid disposition. See the documentation for valid syntax.

Explanation: You must specify a valid z/OS data set disposition as documented in the *z/OS MVS JCL Reference* for DD statement DISP parameter.

User response: Enter a valid DD disposition.

HLO640E Space units field is invalid. Specify TRK or CYL.

Explanation: The specified space units value is not valid.

User response: Specify a valid value.

HLO641E Expiration date is invalid. Specify a valid value.

Explanation: The expiration date value must be exactly in YYYYDDD format. The year in expiration date must be in range of 1999 and higher. The day in the expiration date must be in the range of 1 to 366.

User response: Specify a valid value.

HLO642E Retention period date is invalid. Specify a numeric value.

Explanation: The field requires a numeric value.

User response: Specify a numeric value.

HLO643E Data set type is invalid. Specify EXTREQ, EXTPREF, LARGE, BASIC, or blank.

Explanation: The specified data set type value is not valid.

User response: Specify a valid value.

HLO644E At least one of local site primary, local site backup, recovery site primary, or recovery site backup copy data sets must be specified.

Explanation: All copy data set names are empty.

User response: Specify a valid value for a copy data set name.

HLO645E A backup copy for either site may only be created when a primary copy is also being created for that site. Specify a valid primary copy data set value.

Explanation: A backup copy data set name is specified without specifying a primary copy data set name.

User response: Specify a valid primary copy data set value.

HLO646E The retention period and expiration date fields cannot be entered at the same time.

Explanation: You entered a value in both the Expiration date and Retention period fields. This combination is not allowed.

User response: Clear the value from either the Expiration date or Retention period field.

HLO647E: The *member name* DB2 ZPARMs member does not exist.

Explanation: A valid DB2 ZPARMs member value is required.

User response: Specify a valid DB2 ZPARMs member on the DB2 Subsystem Parameters panel for the DB2 subsystem.

HLO648E Load *entry name* entry point from *DD name* DD has failed. RC = *code*, reason = *code*.

Explanation: LOAD macro failed with the specified codes.

User response: Check *entry_name* member existence in data sets of the DD.

HLO649E The *data set name* copy data set specified for *creator.name* DB2 table does not exist.

Explanation: The specified data set was not found in HLOUCOPY table.

User response: Specify a valid copy data set for the table on the Select Copy Data Set panel.

HLO650E A full copy does not exist for the specified point in time for *creator.name* table and *site type* site type.

Explanation: There is no suitable full copy in HLOUCOPY table for the specified DB2 table, point in time and site type.

User response: Specify a valid point in time for the DB2 table.

HLO651E A usable full copy does not exist for *creator.name* table and *site type* site type.

Explanation: Adding columns to the table or altering the definition of any column renders unusable all copies prior to the table change.

User response: Specify another DB2 table for recovery.

HLO652E The most recent full copy for the specified Point in time for *creator.name* table is not usable.

Explanation: Adding columns to the table or altering the definition of any column renders unusable all copies prior to the table change.

User response: Specify a valid point in time or another DB2 table for recovery.

HLO653E A backup copy data set is not specified for *creator.name* recovery table with Point in time value set to SELECTED.

Explanation: You have specified SELECTED for the Point in time field on the Recover Accelerator Table(s) from a Backup panel. It means the recovery process will use the selected backup data set for each specified table.

User response: Specify a valid backup copy data set for the table using the 'B' line command on the Recovery Table List panel.

HLO654E Point in time is **TIMESTAMP**. The Timestamp end point value is required.

Explanation: The Point in time value **TIMESTAMP** directs the product to recover up to the end point that is specified in the Timestamp end point field.

User response: Change the Point in time to **CURRENT** or enter a Timestamp end point.

HLO655E Invalid Point in time value. Valid values are **CURRENT**, **TIMESTAMP**, and **SELECTED**.

Explanation: The specified value is not valid.

User response: Enter a valid value.

HLO656E Point in time is not **TIMESTAMP**, but a Timestamp end point was specified. Change Point in time to **TIMESTAMP** or delete the Timestamp end point value.

Explanation: The Timestamp end point value must be empty if Point in time is CURRENT or SELECTED.

User response: Change the Point in time to TIMESTAMP or remove the Timestamp end point.

HLO657E No usable backup copy data sets found for creator.name recovery table.

Explanation: You can choose backup copies based on full copies that were created after the table was altered only.

User response: Make a new backup copy for the table or choose another table for recovery.

HLO658E The data set name backup copy data set for creator.name table does not exist on MVS.

Explanation: The product detected the backup copy as available for use for the table based on options specified on the Recover Accelerator Table(s) from a Backup panel.

User response: Specify another point in time to make the table recovery.

HLO714E SQL error text (SQL code: code. Program: program name. Statement: line number. Type: type).

Explanation: The product encountered an error with the specified code.

User response: See *DB2 for z/OS Messages* for information about the code. If you are unable to determine the cause of the error, contact IBM Software Support. Have available the listing that contains this message.

HLOP9875E: Identifier_type value length error. The value must be 1 to maximum_name_length characters long.

Explanation: An identifier or name of the specified type has an invalid length. The identifier must be 1 to MAX_LENGTH characters long

User response: Correct the control cards and re-submit the Loader job.

HLOP9956E: Keyword_name is only valid with IDAA_DUAL and IDAA_ONLY type LOADs.

Explanation: The indicated keyword can only be specified with the IDAA_DUAL or IDAA_ONLY keywords.

User response: Correct the LOAD utility syntax and resubmit the job.

HLOU4050E: Accelerator copy DDNAME copy_ddname is missing from the JCL.

Explanation: A copy DDNAME required to create a backup of an accelerator table was not specified in the JCL.

User response: Add the missing DDNAME to the Loader JCL, or remove it from the ACCEL_COPYDDN or ACCEL_RECOVERYDDN keywords, and then re-submit the Loader job.

HLOU4051E: Accelerator copy DDNAME copy_ddname specified for more multiple copy types.

Explanation: An ACCEL_COPYDDN or ACCEL_RECOVERYDDN DD value was specified for multiple types of copies. The DD must be specified only once.

User response: Ensure each DDNAME specified on the ACCEL_COPYDDN or ACCEL_RECOVERYDDN keywords is specified only once. Correct the Loader syntax and re-submit the job.

HLOU4053E: DEVTYPE failed. DDNAME: copy_ddname RC: devtype_rc

Explanation: The DEVTYPE service returned an error. The DDNAME and error RC are included in the message.

User response: Contact IBM Software Support. Provide support with the full output from the failed Accelerator Loader job.

HLOU4054E: Accelerator copy datasets must reside on tape or DASD. DDNAME: copy_ddname

Explanation: An attempt was made to create an Accelerator backup on a medium other than tape or DASD.

User response: Change the Accelerator Loader job to create the Accelerator backup on tape or DASD and resubmit the job.

HLOU4055E: Multiple datasets are concatenated to Accelerator backup DDNAME: copy_ddname

Explanation: An Accelerator backup DDNAME must not refer to a concatenation of datasets.

User response: Correct the Accelerator Loader JCL and resubmit the job.

HLOU4056E: COPY FUNC= copy_function failed. The copy task has terminated.

Explanation: An attempt to create an Accelerator backup copy has failed.

User response: Review the job log messages to determine why the copy subtask failed. If you need additional help please contact IBM Software Support.

**HLOU4057E: System_call failed. RC: *return_code*
Module: *calling_module* DDNAME:
*ddname***

Explanation: The named system service failed. As a result the Loader was unable to create or register the copy datasets.

User response: Review the job log messages for additional error messages. If you need additional help please contact IBM Software Support.

**HLOU4058E: Copy registration failed. RC:
return_code RSN: *reason_code***

Explanation: The server was unable to register the copy datasets. A negative reason code indicates that the server encountered an SQL error. The reason code in this case is the SQLCODE.

User response: Review the server log messages for more information regarding the cause of the error. If you need additional assistance please contact IBM Software Support.

**HLOU4059I: The following copy dataset(s) have
been successfully registered:**

Explanation: The requested backup copies have been created and registered in the product's copy dataset registration table.

User response: None.

**HLOU4060E: Copy must be a physical sequential
dataset. DDNAME: *ddname* DSORG:
*xdataset_organization***

Explanation: The named copy dataset has an unsupported DSORG. Copy datasets must be physical sequential datasets. They cannot be partitioned or indexed sequential datasets

User response: Correct the Loader JCL and resubmit the job.

**HLOU4061E: Copy registration check failed. RC:
return_code RSN: *reason_code***

Explanation: The server was unable to verify that the requested copy data sets are not already registered in SYSIBM.SYSCOPY or HLOUCOPY. This registration check is performed before the copies are created to protect the recoverability of this and other DB2 objects.

User response: Review the server log messages for more information regarding the cause of the error. If you need additional assistance please contact IBM Software Support.

**HLOU4062E: Dataset used for previous *copy_type*
copy. DDNAME: *ddname* TIME:
*registration_time***

Explanation: A utility has been invoked to backup an Accelerator table to a data set which is a duplicate of one already recorded in SYSIBM.SYSCOPY or HLOUCOPY. If the specified dataset is cataloged a matching DSNAMES exists. If the specified dataset is not cataloged a matching DSNAMES, DSVOLSER and FILESEQNO exists.

User response: Change the name of the copy dataset and re-run the job.

**HLOU4063E: No Accelerator copy datasets have
been provided in the JCL.**

Explanation: An Accelerator backup or inline copy has been requested, but no copy datasets were provided in the JCL.

User response: Add 1 or more copy datasets to the JCL. Copy datasets can be specified either with an HLOCOPY DD statement in the JCL, or by specifying the ACCEL_RECOVERYDDN or ACCEL_COPYDDN keywords in the LOAD or BACKUP_ACCELERATOR command syntax. Correct the JCL and re-submit the job.

**HLOU4064E: \$ABPCOPY FUNC=PUTREC
failed. RC: *return_code***

Explanation: The Loader encountered an error while writing to an Accelerator copy dataset.

User response: Check the job log for additional messages that may provide more details on the type of I/O error that occurred. If you need assistance please contact IBM Software Support.

**HLOU4065E: Keyword <IDAA_ONLY> is
required for inline copies.**

Explanation: Inline copies were requested either by specifying the ACCEL_COPYDDN or ACCEL_RECOVERYDDN keywords, or by including the HLOCOPY DDNAME in the JCL. The inline copy feature is available only when IDAA_ONLY is also specified in the LOAD control cards.

User response: Correct the JCL or control cards and re-submit the job.

**HLOU4066E: Inline copies are not supported on
partial loads.**

Explanation: Inline copies were requested either by specifying the ACCEL_COPYDDN or ACCEL_RECOVERYDDN keywords, or by including the HLOCOPY DDNAME in the JCL. The table being loaded is index partitioned or partitioned by range but only a subset of the table

partitions are being loaded. Inline copies are only available when all partitions participate in the load.
User response: Correct the JCL or control cards and re-submit the job.

HLOU4067E: Function GET_ACCEL_GROUP failed. , RC=return_code, RSN=reason_code.

Explanation: An error was encountered when the product tried to determine if the specified accelerator name was a group name. This failure could be caused by an SQL error. Check the started task log for additional error messages.

User response: If you are unable to resolve the problem contact IBM Software Support.

HLOU4068I: Accelerator_name is an accelerator group. The following members will be loaded:

Explanation: The specified accelerator group has been resolved to its member accelerators. All members in the group will be loaded.

User response: None.

HLOU4069I: ...accelerator_name

Explanation: This informational message lists a member of an accelerator group. This message is issued repeatedly in conjunction with HLOU4068I to display all of the members of an accelerator group.

User response: None.

**HLOU4072I: Fetch: fetch_time Waits: wait_count
Wait Time: milliseconds Recs:
record_count Buffs: backup_buffers
MRF: multi_row_fetch_factor**

Explanation: This is an informational message issued when tracing has been activated.

User response: None.

HLOU4073I: Accelerator Backup Utility execution started.

Explanation: The accelerator backup utility has started execution.

User response: None required.

HLOU4074I: Number of rows copied=row_count

Explanation: This informational message indicates the number of rows that have been written to the accelerator copy datasets.

User response: None.

HLOU4075I: copy_type: copy_dsname

Explanation: This informational message is issued for each copy dataset that is created and registered in the product's copy registration table. The message text includes the name of the copy dataset, the site that the copy is for (local or recovery) and whether the copy is the site's primary or backup copy.

User response: None.

HLOU4076E: Error converting data row to DB2 internal format. Copy task terminating due to errors

Explanation: A data conversion error has prevented the backup utility from building a DB2 formatted row for the copy dataset. The backup utility will terminate.

User response: Please contact IBM Software Support. Provide the support representative with the full output of the failed accelerator backup job.

**HLOU4077E: Unsupported data type. Accelerator backup not allowed.
Column=column_name,
Type=column_type**

Explanation: Accelerator backups are only available when the column data types of the target table are limited to: CHAR, VARCHAR, INTEGER, SMALLINT, BIGINT, FLOAT, REAL, DECIMAL, DATE, TIME, TIMESTAMP, GRAPHIC, VARGRAPHIC

User response: Consider dropping and recreating the table to eliminate the unsupported data types, or select a different table to back up.

HLOU4080E: db2_error_msg

Explanation: An error was encountered during an SQL or DB2 instrumentation facility interface (IFI) operation. This message contains the text of the message that the DB2 DSNTIAR message formatting routine issued when the error occurred.

User response: For more information about the error, refer to the IBM DB2 messages documentation.

Chapter “Reference”

Topic: “Accelerator Loader terminology”

Add the following terms:

Accelerator-archived table/partition

A table that has one or more of its partitions moved to the accelerator using the IBM DB2 Analytics Accelerator for z/OS High-performance Storage Saver (HPSS) feature.

Accelerator-only table

A table that exists on the accelerator only.

Accelerator-shadow table

A table that exists both in DB2 for z/OS and IDAA.

Inline copy

A backup copy of an accelerator table that is created as the data is loaded to the accelerator.

Non-accelerator DB2 table

A table that exists in DB2 for z/OS only.

Topic: “Main menu”

Add the following options:

Back up Accelerator table

Opens the Back up Accelerator table panel, on which you can specify options to generate JCL to back up a table defined to the accelerator. These options can be saved to a Backup profile.

Recover Accelerator table(s) from a backup

Opens the Recover Accelerator table(s) from a backup panel, on which you can specify options to generate JCL to recover a table defined to the accelerator. These options can be saved to a Recovery profile.

Topic: “Manage Loader Profiles panel”

Add the following Type/Profile type:

- **Backup** specifies options for backing up a table defined to the accelerator.
- **Recovery/recover** specifies options for recovering a table defined to the accelerator.

Topic: “Create Profile panel”

Add the following types of Accelerator Loader profiles:

Back up Accelerator table

Opens the Back up Accelerator Table panel. Use this panel to specify options for a Backup profile, which is used for backing up a table defined to the accelerator.

Recover Accelerator table(s) from a backup

Opens the Recover Accelerator Table(s) from a Backup panel. Use this panel to specify options for a Recovery profile, which is used for recovering a table defined to the accelerator.

Topic: “Save Profile panel”

Add the following Type/Profile type:

- **Backup** specifies options for backing up a table defined to the accelerator.
- **Recovery** specifies options for recovering a table defined to the accelerator.

Topic: “Load Accelerator from External File panel”

Add the following fields:

Inline copy data sets options:

Local primary

Specifies the primary copy data set for the local site. To update the data set parameters, specify **Yes** in the Update field.

Local backup

Specifies the secondary (backup) copy data set for the local site. You can create a backup copy for the local site only when a primary copy for the local site is also being created. To update the data set parameters, specify **Yes** in the Update field.

Recovery primary

Specifies the primary copy data set for the remote recovery site. To update the data set parameters, specify **Yes** in the Update field.

Recovery backup

Specifies the secondary (backup) copy data set for the remote recovery site. You can create a backup copy for the remote recovery site only when a primary copy for the remote recovery site is also being created. To update the data set parameters, specify **Yes** in the Update field.

Topic: “Rename Profile panel”

Add the following Type/Profile type:

- **Backup** specifies options for backing up a table defined to the accelerator.
- **Recovery** specifies options for recovering a table defined to the accelerator.

New topic: “Back up Accelerator Table panel”

Use this panel to specify options for generating JCL to use the BACKUP utility to back up a table defined to the accelerator. These options can be saved to a Backup profile. All of the following commands are available on the editable version of the panel. On the view-only version of the panel, the **BUILD** command is available.

- **TABLE**: Type this command on the command line to open the Enter Table and Creator Like to Display panel and subsequently the Accelerator Table Selection panel from which you can select the table to back up.
- **SAVE**: Type this command on the command line to save your specifications to a Backup profile.
- **BUILD**: Type this command on the command line to build JCL for the profile to the specified data set.

The following fields are available on this panel.

Creator

The profile creator.

Name

The name of the profile, up to 30 characters. (Using meaningful names for profiles makes them easier to locate and reuse; for example, “Objects for Employee App.”)

Share option

Controls how other users can use a profile:

U (Update)

Other users can update the profile.

V (View only)

Other users can view the profile.

N (No access)

Other users cannot view nor update the profile.

Description

A description of the profile.

Schema

Table Name

Accelerator

These fields display values for the currently selected table. To change the table, issue the **TABLE** command.

Copy data sets options:

Local primary

Specifies the primary copy data set for the local site. To update the data set parameters, specify **Yes** in the Update field.

Local backup

Specifies the secondary (backup) copy data set for the local site. You can create a backup copy for the local site only when a primary copy for the local site is also being created. To update the data set parameters, specify **Yes** in the Update field.

Recovery primary

Specifies the primary copy data set for the remote recovery site. To update the data set parameters, specify **Yes** in the Update field.

Recovery backup

Specifies the secondary (backup) copy data set for the remote recovery site. You can create a backup copy for the remote recovery site only when a primary copy for the remote recovery site is also being created. To update the data set parameters, specify **Yes** in the Update field.

New topic: “Accelerator Table Selection panel”

Use this panel to choose a table from the generated list for which to create a backup. After you select the table to back up, an asterisk appears in the Cmd field. Only one table can be selected. The following commands are available.

- **DEFAULT:** Type this command on the command line to sort the panel contents in default order.
- **S:** Type this command in the Cmd field next to the table that you want to select.

The following fields are available:

Table creator like

The table creator search criteria. Wildcard values are allowed. The wildcard character is an asterisk (*).

Note: Case sensitivity of this field is controlled by the Case sensitive field on the Enter Table and Creator Like to Display panel.

Table name like

The table name search criteria. Wildcard values are allowed. The wildcard character is an asterisk (*).

Note: Case sensitivity of this field is controlled by the Case sensitive field on the Enter Table and Creator Like to Display panel.

Table Name

The table name.

Creator

The user ID of the table space creator.

Accelerator Name

A unique name for the accelerator server. This is the name by which the accelerator server is known to the local DB2 accelerated query tables.

Created Timestamp

The time when the CREATE statement was executed for the table.

Altered Timestamp

The time when the table was last altered.

Refresh Time

The timestamp when the data was last refreshed. If the data was not refreshed, this column contains the default timestamp ('0001-01-01.00.00.00.000000').

New topic: “Copy Data Set Parameters panel”

Use this panel to define the data set allocation parameters for backup copy data sets. The title of this panel reflects the type of backup copy data set being created or updated (Local Primary, Local Backup, Recovery Primary, Recovery Backup).

The following fields are available:

Disposition

A valid z/OS data set disposition as documented in the *z/OS MVS JCL Reference*.

Data set type

A valid z/OS data set type value as documented in the *z/OS MVS JCL Reference*.

Management class**Storage class**

The SMS management and storage classes for the backup copy data set created by Accelerator Loader.

Volume serial

The volume serial number to use for the backup copy data set created by Accelerator Loader. To let SMS select the volume on which to allocate the data set, leave the field blank.

Device type

The device type to use for the backup copy data set created by Accelerator Loader. To let SMS select the device type on which to allocate the data set, leave this field blank.

Data class

The SMS data class (up to 8 alphanumeric characters) to use for the backup copy data set created by Accelerator Loader.

Space units

The allocation unit to be used when allocating the backup copy data set.

Space primary**Space secondary**

The primary and secondary allocation quantities of space to use when allocating the backup copy data set. The unit of measure that you specify in the Space units field is used.

Tape device options:

Expiration date

Specifies the expiration date for a new data set. On and after the expiration date, the data set can be deleted or written over by another data set. This value is valid for tape device only.

Retention period

Specifies the retention period for a new data set to help reduce the chance of later accidental deletion. After the retention period, the data set can be deleted or written over by another data set. This value is valid for tape device only.

New topic: “Recover Accelerator Table(s) from a Backup panel”

Use this panel to specify options for generating JCL to recover a table defined to the accelerator. These options can be saved to a Recovery profile. All of the following commands are available on the editable version of the panel. On the view-only version of the panel, the TABLES and BUILD commands are available.

- **TABLES:** Type this command on the command line to open the Recovery Table List panel from which you can select one or more tables to recover.
- **ACCELERATOR:** Type this command on the command line to open the DB2 Analytics Accelerator Selection panel where you can select the accelerator on which you want to recover data.
- **SAVE:** Type this command on the command line to save your specifications to a Recovery profile.
- **BUILD:** Type this command on the command line to build JCL for the profile to the specified data set.

The following fields are available on this panel:

Creator

The profile creator.

Name

The name of the profile, up to 30 characters. (Using meaningful names for profiles makes them easier to locate and reuse; for example, “Objects for Employee App.”)

Share option

Controls how other users can use a profile:

U (Update)

Other users can update the profile.

V (View only)

Other users can view the profile.

N (No access)

Other users cannot view nor update the profile.

Description

A description of the profile.

Target options:

Accelerator(s)

Specifies the name of the accelerator(s) on which to recover data. To display a list of the existing accelerators, type a question mark (?) in the field and press Enter.

Add table to Accelerator

Indicates whether to add missing tables to the accelerator before starting the recover job.

N (No)

(default) Do not add tables.

A (Add)

Add missing tables. This setting generates the ACCEL_ADD_TABLES keyword in the LOAD command.

R (Refresh)

Add missing tables; remove and re-add existing tables. This setting generates the ACCEL_REMOVE_AND_ADD_TABLES keyword in the LOAD command.

Acceleration on success

Controls whether Accelerator Loader enables query acceleration for the table after a successful load.

Valid values are as follows:

Yes

Enable a table for acceleration after a successful load. The control card ACCEL_ON_SUCCESS_ENABLE YES is generated into the JCL.

No

(default) Do not enable a table for acceleration after a successful load. The control card ACCEL_ON_SUCCESS_ENABLE NO is generated into the JCL.

Recovery point options:

Point in time

Specifies the point in time for recovery.

CURRENT

Applies to all selected tables. Backup data sets for each table will be automatically chosen at time of recovery JCL generation.

TIMESTAMP

Applies to all selected tables. Backup data sets for each table will be automatically chosen using values specified in Timestamp end point and Time zone of timestamp fields.

SELECTED

A backup data set for each selected table will be specified on the Recovery Table List panel, which is accessed by the TABLES command.

Timestamp end point

Indicates the end point at which all selected tables will be recovered. Specify the timestamp in the format (YYYY-MM-DD-hh.mm.ss.nnnnnn). For convenience, you can copy the end point from the HLOUCOPY table and paste it into this field.

Time zone of timestamp

Specifies the timezone of the Timestamp end point value.

New topic: “Recovery Table List panel”

Use this panel to identify the tables for which to recover data and to specify backup copies to use for the recovery. The tables in this list are included in the Recovery profile. All of the following commands are available on the editable version of the panel. On the view-only panel, TABLES and BUILD are available.

- **ADD:** Type this command on the command line to open the Enter Table and Creator Like to Display and subsequently the Recovery Table Selection panel from which you can select one or more tables to recover.
- **D:** Type this command in the Cmd field next to the table name to delete the table from the profile.
- **B:** Type this command in the Cmd field next to the table name to specify the backup copy. This command opens the Backup Copy Selection panel, on which you can select the backup copy.

The following fields are available on this panel.

Creator

The profile creator.

Name

The name of the profile, up to 30 characters. (Using meaningful names for profiles makes them easier to locate and reuse; for example, “Objects for Employee App.”)

Share option

Controls how other users can use a profile:

U (Update)

Other users can update the profile.

V (View only)

Other users can view the profile.

N (No access)

Other users cannot view nor update the profile.

Description

A description of the profile.

Table Name

The table name.

Creator

The user ID of the table space creator.

Copy Data Set Name

The name of the backup copy data set.

Copy Type

The type of copy written to the backup copy data set.

INC Incremental copy

FULL Full copy

Copy Created Timestamp (Local Time)

The backup copy data set creation timestamp in local time.

Copy Created Timestamp (UTC)

The backup copy data set creation timestamp in UTC.

New topic: “Recovery Table Selection panel”

Use this panel to choose the tables from the generated list for which to recover data from backup copies. After you select a table to recover, an asterisk appears in the Cmd field.

The following commands are available.

- **ALL:** Type this command on the command line to select all displayed tables.
- **DEFAULT:** Type this command on the command line to sort the panel contents in default order.
- **S:** Type this command in the Cmd field next to the table that you want to select.

The following fields are available:

Table creator like

The table creator search criteria. Wildcard values are allowed. The wildcard character is an asterisk (*).

Note: Case sensitivity of this field is controlled by the Case sensitive field on the Enter Table and Creator Like to Display panel.

Table name like

The table name search criteria. Wildcard values are allowed. The wildcard character is an asterisk (*).

Note: Case sensitivity of this field is controlled by the Case sensitive field on the Enter Table and Creator Like to Display panel.

Table Name

The table name.

Creator

The user ID of the table space creator.

Database

The database name.

Tablespace

The table space name.

Altered Timestamp (Local Time)

The timestamp of the table creation/alter in Local time.

New topic: “Backup Copy Selection panel”

Use this panel to choose the backup copy for recovery of a table defined to the accelerator. After you select a backup copy, an asterisk appears in the Cmd field.

The following command is available:

- S: Type this command in the Cmd field next to the backup copy that you want to select.

The following fields are available:

Schema

Table Name

These fields display values for the currently selected table.

Copy Data Set Name

The name of the backup copy data set.

Created Timestamp (Local Time)

The backup copy data set creation timestamp in local time

Created Timestamp (UTC)

The backup copy data set creation timestamp in UTC.

Copy Type

The type of copy written to the backup copy data set.

INC Incremental

FULL Full

Copy Site Type

The type of backup copy data set.

LP Local Primary

LB Local Backup

RP Recovery Primary

RB Recovery Backup

Share Level

The share level of the backup copy data set.

Change

Reference

Unit Type

The type of device unit.

DASD Disk device

Tape Tape device

Unit

The name of device unit.

New topic: “Use alternate backup panel”

Use this panel to specify an alternate backup copy data set after attempting to use a backup copy for recovery that is not usable.

Schema**Table name**

These fields display values for the selected table to recover.

Selected backup:

Copy data set

The name of the backup copy data set that is not usable.

Creation time in UTC

The backup copy data set creation timestamp in UTC.

Alternate backup:

Copy data set

The name of an alternate backup copy data set.

Creation time in UTC

The backup copy data set creation timestamp in UTC.