

Updates that apply to IBM® DB2® Analytics Accelerator Loader for z/OS® V2R1 User's Guide (SC27-6777-00)

Date of change: August, 2016

Topics: Multiple

Change description: Documentation corrections and updates in support of enhancements provided in the following PTF:

- PTF UI39451

Chapter 1. Overview

Topic: What's new

Add the following descriptions of enhancements.

Use Accelerator Loader to load data to multiple accelerators on the same DB2 subsystem in parallel from a single LOAD utility statement. This feature is referred to as high availability load (HAL).

Topic: Features and benefits

Add the following description:

High Availability Load

Use Accelerator Loader to load data to multiple accelerators in parallel from a single LOAD utility statement. This feature is referred to as high availability load (HAL). To use HAL, two or more accelerators must be configured on the same DB2 subsystem.

- The HALOAD utility loads existing DB2 table data to up to four paired accelerator tables. Each accelerator table must be configured on a different accelerator. The utility runs as a batch job and is not invoked under the control of the DSNUTILB LOAD utility.
 - The Accelerator Loader server can load data from an external sequential data set or from a virtualized data source to up to four accelerator tables. The tables can be standard accelerator tables paired to DB2 or accelerator only tables (AOT). This functionality is implemented as extended syntax to the DSNUTILB LOAD utility.
 - The ISPF interface generates a batch JCL job that can load up to four accelerators.
-

Topic: Scenarios

Add the following description:

Loading data to multiple accelerators

You have existing DB2 table data that you need to load into multiple accelerators in parallel. To accomplish this, use the HALOAD utility.

You need to load data from a SYSREC data set into multiple accelerators and optionally to DB2. In the syntax of your batch job, specify up to four accelerator names in the ON clause of the IDAA_DUAL or IDAA_ONLY keyword. Alternatively, use the ISPF panels to select multiple accelerators and generate JCL.

You need to load non-DB2 from a virtualized data source through the Accelerator Loader server. In the Accelerator Loader studio, you can select multiple accelerators when generating the JCL.

Chapter 3: Customizing DB2 Analytics Accelerator**Topic: Copying the started task PROC (required)**

Add a step to the procedure:

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

Chapter 5: Loading DB2 table data to multiple accelerators

Add information to the chapter description:

You can generate JCL that loads data to as many as four IBM DB2 Analytics Accelerator for z/OS (accelerators) in parallel by using the HALOAD utility. This process is called a high availability load (HALOAD).

Topic: Restrictions and considerations for loading data to multiple accelerators from a DB2 table

The following is a new topic that will be added to the chapter:

Review usage restrictions and considerations before loading data to multiple accelerators from a DB2 table (high availability load).

Note: In addition to the information in this section, make sure that your system meets requirements as listed in Set up your environment prior to customization.

- Two or more accelerators must be configured on the same DB2 subsystem.
- The Multi load profile type and HALOAD utility support up to four accelerators.

Topic: Using the ISPF interface to create or edit a high availability load (Multi) profile

The following is a new topic that will be added to the chapter:

You can specify options to generate JCL to load one to four accelerators from one or more DB2 tables. After specifying the options, save them as a Multi load profile.

Before you begin

Review the information in “Restrictions and considerations for loading data to multiple accelerators from a DB2 table.”

About this task

When you add a DB2 table to the profile, you can filter on tables, views from a single base table, or aliases. The product resolves the view or alias to the base table space and includes the base table space in the generated JCL. A view that was created from a join of more than one table is not supported. The product checks for the existence of the specified DB2 table before generation. However, if you specify partitions, the product cannot validate the partitions, but uses the partitions as specified when generating JCL.

Procedure

1. From the main menu, select Manage Loader Profiles and press Enter.
2. On the Manage Loader Profiles panel, to filter existing profiles by profile name or creator name, specify a wildcard pattern using an asterisk (*) and press Enter.
3. On the Manage Loader Profiles panel, perform one of the following steps:
 - To create a new profile, issue the CREATE command, and then on the Create Profile panel, select the type of profile to create.
 - To create a new profile by copying an existing profile, type C in the Cmd line next to the profile that you want to copy.
 - To edit an existing profile, type E in the Cmd line next to the profile.
4. On the Load Accelerator(s) from DB2 Table(s) panel, specify a name and processing options for the profile.
5. To add a DB2 table to the profile, issue the T panel command.

6. On the Enter Table and Creator Like to Display panel, specify an object creator name and object name pattern and press Enter to display matching objects.

You can specify a table, a view, or an alias.

7. On the DB2 Table Selection panel, use the S line command to select the tables to add to the profile and press Enter.

8. Return to the previous panel by pressing F3.

9. To select the accelerators onto which you want to load data, issue the A panel command, and on the DB2 Analytics Accelerator Selection panel, select accelerators and press Enter.

10. To edit the table column definitions, issue the C command.

11. To define options for a template DD, complete the following steps:

- a. On the Load Accelerator(s) and DB2 from External File panel or the Load Accelerator(s) from External File panel, specify Yes in the Update field for any template DD.
- b. On the Template Specification panel, issue the TEMPLATE command, and edit the template data set name mask.
- c. On the DSN Template panel, specify qualifier codes to create the data set name mask.
- d. To see the resulting DSN mask, issue the SHOW command.
- e. To save and return to the previous panel, press PF3.
- f. Update the template options as needed.
- g. To save and return to the previous panel, press PF3.

Topic: Generating JCL

The following step will be added as step 8:

8. Select up to four accelerators to load from the list of **Available Accelerators**.

Chapter 9: Using and managing load profiles

Add the following bullet to the types of supported profiles:

- Multi specifies options for loading data to up to four accelerators from one or more DB2 tables (high availability load)

Topic: Using the batch interface to build a load job from a profile

Add the following information to the example JCL provided in the subsection.

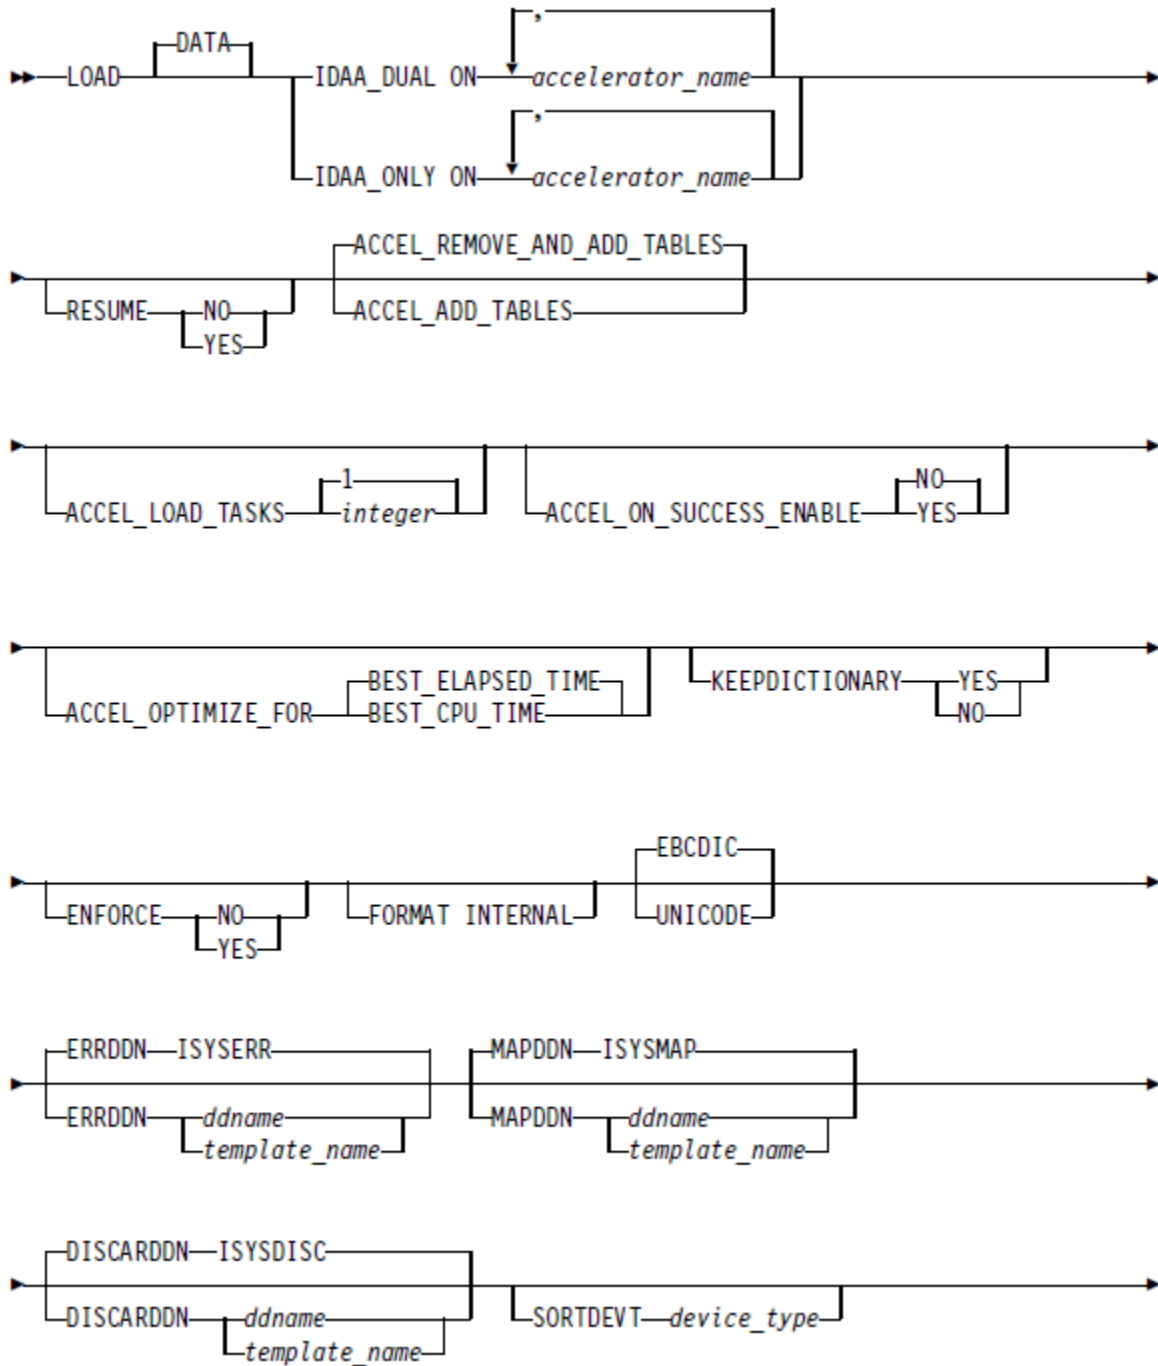
High availability load (Multi) profile

```
//SYSIN DD *
<JOBPREFIX>='HLO'
<TARGET SSID>='DA1A'
<PROFILE SSID>='QA1A'
<PROFILE TYPE>='MULTI'
<PROFILE NAME>='MULTI LOAD PROFILE'
<PROFILE CREATOR>='TSNSB'
<LOAD BY PARTITION>='YES'
<OUTPUT-DSN>='HLO.HLODSN.OUT'
<TABLE>
<TABLE NAME>='&LT;MY TABLE1&GT; TABLE&APOS; &QUOT;NAME&QUOT; '
<TABLE CREATOR>='TABLECREATOR1'
<PARTITION>='ALL'
<SYSREC-TEMPLATE-NAME>='ISYSDISC'
<SYSREC-TEMPLATE-DSN>='&AMP;US..IDSD.&AMP;DB..ABC&AMP;PA.'
<PARALLELISM>='20'
</TABLE>
<TABLE>
<TABLE NAME>='TABLENAME2'
<TABLE CREATOR>='TABLECREATOR2'
<PARTITION>='1-2,4:5'
</TABLE>
<TABLE>
<TABLE NAME>='#VERY LONG TABLE NAME 12345678901234567890123456789012
3456789012345678901234567890123456789012345678901234567890#'
<TABLE CREATOR>='#TABLECREATOR3#'
</TABLE>
/*
```

Chapter 10: Syntax

Topic: Loading from an external file

Subtopic: Syntax diagram: Load from an external file



Topic: Loading data to multiple accelerators:**Subtopic: Loading multiple accelerators with high availability load**

Example: Load DB2 table data to multiple accelerators

The following sample syntax shows the HALOAD command, used to perform a high availability load. The following requirements apply:

- The HLOUHALO program must be specified on the EXEC card.
- The HALOAD ACCEL control card must be present in the utility syntax in lieu of LOAD DATA
- The DB2 SSID must be passed on the PARM card.
- The DB2 load library and the product load library must be included on the STEPLIB or JOBLIB.
- REGION=0M must be specified on the JOB card or the EXEC card.

```
//HLOUHALO JOB , 'SAMPLE HALOAD', CLASS=A, MSGCLASS=X,  
// REGION=0M, NOTIFY=&SYSUID  
//*  
//HLOUHALO EXEC PGM=HLOUHALO, PARM='DSNA'  
//STEPLIB DD DISP=SHR, DSN=HLO.PR0210.SHLOLOAD  
// DD DISP=SHR, DSN=DSN.VA10.SDSNLOAD  
//*  
//HLODUMMY DD DUMMY  
//SYSPRINT DD SYSOUT=*  
//*  
//SYSIN DD *  
HALOAD ACCEL (IDAAS01, IDAAS05)  
FROM TABLE  
HLOTEST.TABLE01  
 , HLOTEST.TABLE02  
PART(1,3,9,12:15)  
ACCEL_REMOVE_AND_ADD_TABLES  
ACCEL_ON_SUCCESS_ENABLE YES  
ACCEL_LOAD_TASKS 5  
/*
```


Subtopic: Syntax diagram: Loading multiple accelerators

Review syntax for a job that loads DB2 table data to up to four accelerators.

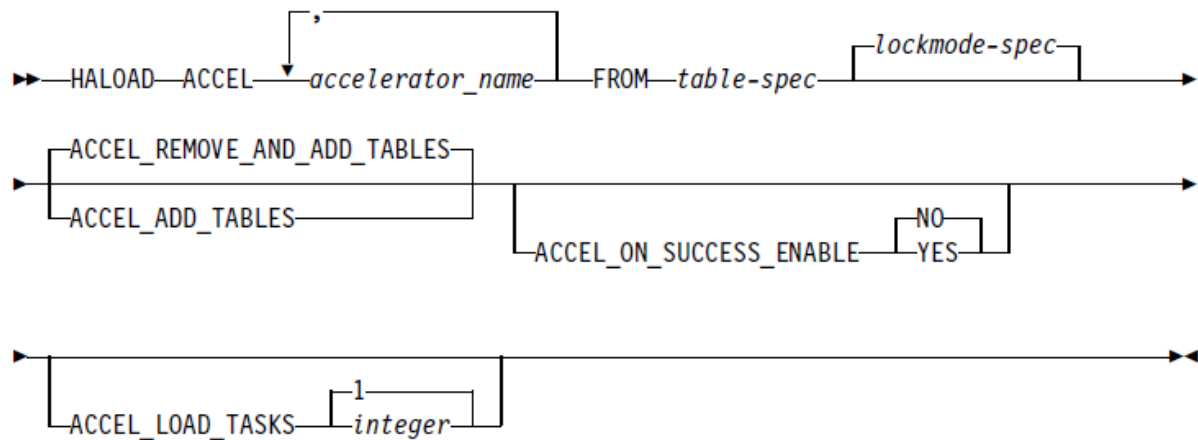





Table spec



Part Spec



Lockmode spec:



Subtopic: Syntax definitions: Loading multiple accelerators

DB2 Analytics Accelerator Loader supports the following syntax elements (presented alphabetically) when you are loading multiple accelerators

ACCEL_ADD_TABLES

Add missing tables to the accelerator before starting the load job. This option can be specified in the ISPF interface by using the Add table to Accelerator field on the Load Accelerator(s) and DB2 from External File panel or the Load Accelerator(s) from External File panel.

ACCEL_REMOVE_AND_ADD_TABLES

Remove and re-add existing tables to the accelerator before starting the load job.

This option can be specified in the ISPF interface by using the Add table to Accelerator field on the Load Accelerator(s) and DB2 from External File panel or the Load Accelerator(s) from External File panel.

ACCEL_LOAD_TASKS integer

Specifies the number of partitions to load into the accelerator and optionally into DB2 in parallel when loading from an external file. Valid values are in the range 1 - 20.

This value cannot exceed the value of the IBM DB2 Analytics Accelerator for z/OS parameter AQT_MAX_UNLOAD_IN_PARALLEL, which indicates the maximum number of partitions that can be loaded in parallel. If AQT_MAX_UNLOAD_IN_PARALLEL is set to 2, the maximum number of partitions that can be written to the accelerator at one time is 2, regardless of the value that you specify for Parallel load tasks ACCEL_LOAD_TASKS.

Specify a value for NUMRECS also. For more information, see the description of the NUMRECS option.

The default value is 1.

This option can be specified in the ISPF interface by using the Load tasks field on the Load Accelerator(s) and DB2 from External File panel or the Load Accelerator(s) from External File panel.

ACCEL_LOCKMODE

Specifies the protection level while tables on an accelerator are being loaded. The protection level ensures a consistent state of the data with respect to the specified entity. The lock mode that you specify is passed to the ACCEL_LOAD_TABLES stored

procedure call. If the target table is enabled for replication, lock mode of ROW is required.

Valid value are:

- **NONE:** (Default) No locking at all. However, only committed data is loaded into the table because the DB2 data is unloaded with isolation level CS and SKIP LOCKED DATA.
- **TABLE:** Protects just the table that is currently being loaded.
- **TABLESET:** Protects all tables to be loaded against changes during the load operation.
- **PARTITIONS:** Protects the table space partition containing that part of the table that is currently being loaded. With this setting, an unpartitioned table is always locked completely.
- **ROW:** Protects just the row or page that is being loaded against updates. DB2 data is unloaded with isolation level CS, but in contrast to lock mode NONE, rows locked by an application are not skipped. It is the recommended choice for loads in connection with continuous incremental updates.

Note: Consider the implications of using this lock mode, especially in connection with incrementally updated tables. You can safely use it if you have enabled continuous incremental updates. For more information, see the information about enabling continuous incremental updates in the documentation for the IBM DB2 Analytics Accelerator for z/OS product.

ACCEL_ON_SUCCESS_ENABLE YES | NO

Controls whether query acceleration is enabled for the table after a successful load. If DB2 discards any rows during the load, query acceleration is not enabled. Valid values are YES and NO.

This option can be specified in the ISPF interface by using the **Acceleration on success** field on the Load Accelerator(s) and DB2 from External File panel or the Load Accelerator(s) from External File panel.

HALOAD

ACCEL

accelerator_name, accelerator_name

Specifies that you want to load one to four accelerators from one or more DB2 tables by using the HALOAD utility. Specify each accelerator name separated with a comma.

PART (partition)

When using the HALOAD command, specify partitions as follows:

- Enclose partition numbers in parentheses.
- Specify each partition number by its one- to four-character physical partition number.
- Separate individual partition numbers by commas (.).
Example, (1, 01, 001)
- Partition ranges:
 - Specify partition ranges in the format lesser value:larger value.

- Separate partition ranges by a colon (:) between the range boundaries.
- Range boundaries are inclusive.

Example: The following partition specification would load table partitions

1,2,3,4,7,12,15,16,17,18,19, and 20:

PART (1:4,7,12,15:20)

TABLE *table_name*

Specifies that you want to load data from one or more DB2 tables. Specify each table name, separating names with a comma.