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

Date of change: May 2017
Topic: Multiple
Change description: Documentation changes made in support of PTF UI47120 APAR PI79298 – Load accelerator and DB2 from a virtualized data source or remote DBMS

The following topics have been updated or introduced:

In chapter "Overview":
Topic "What's new"

In chapter "Loading data from non-DB2, remote DB2, and remote system sources":
Topic "Accelerator Loader server restrictions and considerations"

In chapter "Syntax", section "Loading data from non-DB2, remote DB2, and remote system sources":
Topic "Example JCL"
Topic "Customizing the JCL to load the accelerator and DB2"
Topic "Syntax diagram: Load from a non-DB2, remote DB2, or remote system source"
Topic "Syntax definitions: Load from a non-DB2, remote DB2, or remote system source"

In chapter "Troubleshooting":
Topic "Accelerator Loader messages"

Chapter "Overview"

Topic: "What's new"
Add the following description:

When loading data from a virtualized data source or remote DBMS, you can now load data to both the accelerator and DB2. Previously, when loading from these sources, you could load to the accelerator only. To use this feature, you must manually edit the JCL generated by the Accelerator Loader studio.

Chapter "Loading data from non-DB2, remote DB2, and remote system sources"

Topic: "Accelerator Loader server restrictions and considerations"
Remove the following restriction:

- You use the server to load the accelerator only. Attempting to load data to DB2 causes IBM DB2 Analytics Accelerator for z/OS to fail.

Chapter "Syntax", section "Loading data from non-DB2, remote DB2, and remote system sources"

Topic: "Example JCL"
Add the following example:

Example 3: Load the accelerator and DB2 with data from a virtualized data source
The following figure contains example JCL to load both the accelerator and DB2 with source data from a virtualized data source using the Accelerator Loader server.

```
//HLOD0100 EXEC PGM=DSNUTILB,
//         REGION=1024M,
//         PARM=('QAA5','USER01.LOAD')
//STEPLIB  DD DISP=SHR,DSN=QDS5.SDSNEXIT
//         DD DISP=SHR,DSN=DSN.VA10.SDSNLOAD
//HLODUMMY DD DUMMY
//SYSPRINT DD SYSOUT=*
//UTPRINT  DD SYSOUT=*
//SYSIN    DD *
    TEMPLATE ISYSREC
        DSN 'USER01.DA1A.HLOLAB2D.HLOTS2.SYSREC'
        DISP(SHR,KEEP,KEEP)
    TEMPLATE ISYSERR
        DSN &US..IDSE.&DB..&TS..&UQ.
        DISP(MOD,CATLG,CATLG)
        SPACE (10,100) CYL
    TEMPLATE ISYSMAP
        DSN &DB..&SN..&US..&JO.
        DISP(MOD,CATLG,CATLG)
        SPACE (10,100) CYL
    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
    EXEC SQL
        DECLARE HLVCSR CURSOR FOR
        SELECT * FROM USER1.SOURCE_DB2_TABLE
    ENDEXEC
    LOAD DATA
        IDAA_DUAL ON RDSBACC1
        REPLACE
        LOG NO NOCOPYPEND
        ENFORCE NO
        ACCEL_CURSOR HLVCSR
        ACCEL_HLV_SSID HLVS
        ACCEL_DATA_SERVER HSZ3
        ACCEL_REMOVE_AND_ADD_TABLES
        ACCEL_ON_SUCCESS_ENABLE YES
        ACCEL_LOAD_TASKS 1
        INTO TABLE "TSADO"."RemoteLoad"
/*
//
```

Topic: "Customizing the JCL to load the accelerator and DB2"
Add as a new topic.

The following steps describe the changes you must make to customize the JCL generated by the Accelerator Loader studio to load the accelerator and DB2.

About this task
You can load data to both the accelerator and DB2 when loading data from a virtualized data source or remote DBMS using the Accelerator Loader server. To use this feature, you must manually edit the JCL generated by the Accelerator Loader studio.

Before making the required changes to the generated JCL as described in the following procedure, review the details of the changes to be made, as follows:

- Update the LOAD control cards to use the `IDAA_DUAL` keyword. The Accelerator Loader studio includes the `IDAA_ONLY` keyword in the generated JCL. You must manually replace `IDAA_ONLY` with `IDAA_DUAL` to load to both the accelerator and DB2.
- Increase the DB2 utility work data set allocations to provide enough work space for the DB2 LOAD utility. The minimum space the Accelerator Loader studio provides for the utility work data sets (SYSUT1, SORTOUT, SYSMAP, SYSERR) is not sufficient for most loads. You can also replace the JCL DD statements for these work data sets with TEMPLATE statements.
- If there are indexes on the DB2 table, add the SORTDEVT (and optionally SORTNUM) DB2 LOAD control cards to the LOAD statement to provide enough SORT work space for the index builds.
- Consider adding the NUMRECS keyword to the INTO TABLE clause. NUMRECS specifies the number of records to be loaded. The DB2 LOAD utility uses the NUMRECS value to size various work data sets. If you omit the NUMRECS keyword, Accelerator Loader passes DB2 LOAD a default value of 100 million.
- If parallelism is used, review the CREATE TABLE DDL generated by the Accelerator Loader studio. Parallelism requires the DB2 table to be range-partitioned by the Accelerator Loader generated column "ACCEL PARTITION KEYCOL". If the Accelerator Loader studio has been used to generate the CREATE TABLE DDL, the table is created with the number of partitions equal to the degree of parallelism. For example, if the degree of parallelism is specified as 10, the table will be created with 10 partitions. All the loaded data must be able to fit in those 10 partitions. You may need to edit the CREATE TABLE DDL generated by the Accelerator Loader studio to ensure the VSAM data sets for the table are large enough to accommodate all the data. Consider adding the DSSIZE, COMPRESS or STOGROUP keywords to the CREATE TABLE statement.

Restrictions
DISCARD data sets are not supported when loading both the accelerator and DB2 from an Accelerator Loader server data set. If DISCARD data sets are provided, Accelerator Loader will fail with the following message:

```
HLOP9953E Discard datasets are not supported when keyword 'ACCEL_CURSOR' is specified.
```

Procedure

1. Generate JCL from the Accelerator Loader studio. For more information, see “Generating JCL”.
2. In the JCL generated by the Accelerator Loader studio, make the following changes:
    a. Replace the `IDAA_ONLY` keyword with `IDAA_DUAL`.
    b. Increase the allocations for the DB2 utility work data sets (SYSUT1, SORTOUT, SYSMAP, SYSERR) to provide enough work space for the DB2 LOAD utility. Optionally, you can replace the JCL DD statements for these work data sets with TEMPLATE statements.
    c. If there are indexes on the DB2 table, add the SORTDEVT (and optionally SORTNUM) DB2 LOAD control cards to the LOAD statement to provide enough SORT work space for the index builds.
    d. Optional: Add the NUMRECS keyword to the INTO TABLE clause to specify the number of records to be loaded.
    e. If parallelism is used, perform the following steps:
        i. Review and update, if necessary, the CREATE TABLE DDL generated by the Accelerator Loader studio to ensure the VSAM data sets for the table are large enough to accommodate all the data.
        ii. Optional: Add the DSSIZE, COMPRESS or STOGROUP keywords to the CREATE TABLE statement.

Topic: "Syntax diagram: Load from a non-DB2, remote DB2, or remote system source"
Replace the syntax diagram with the following updated diagram:

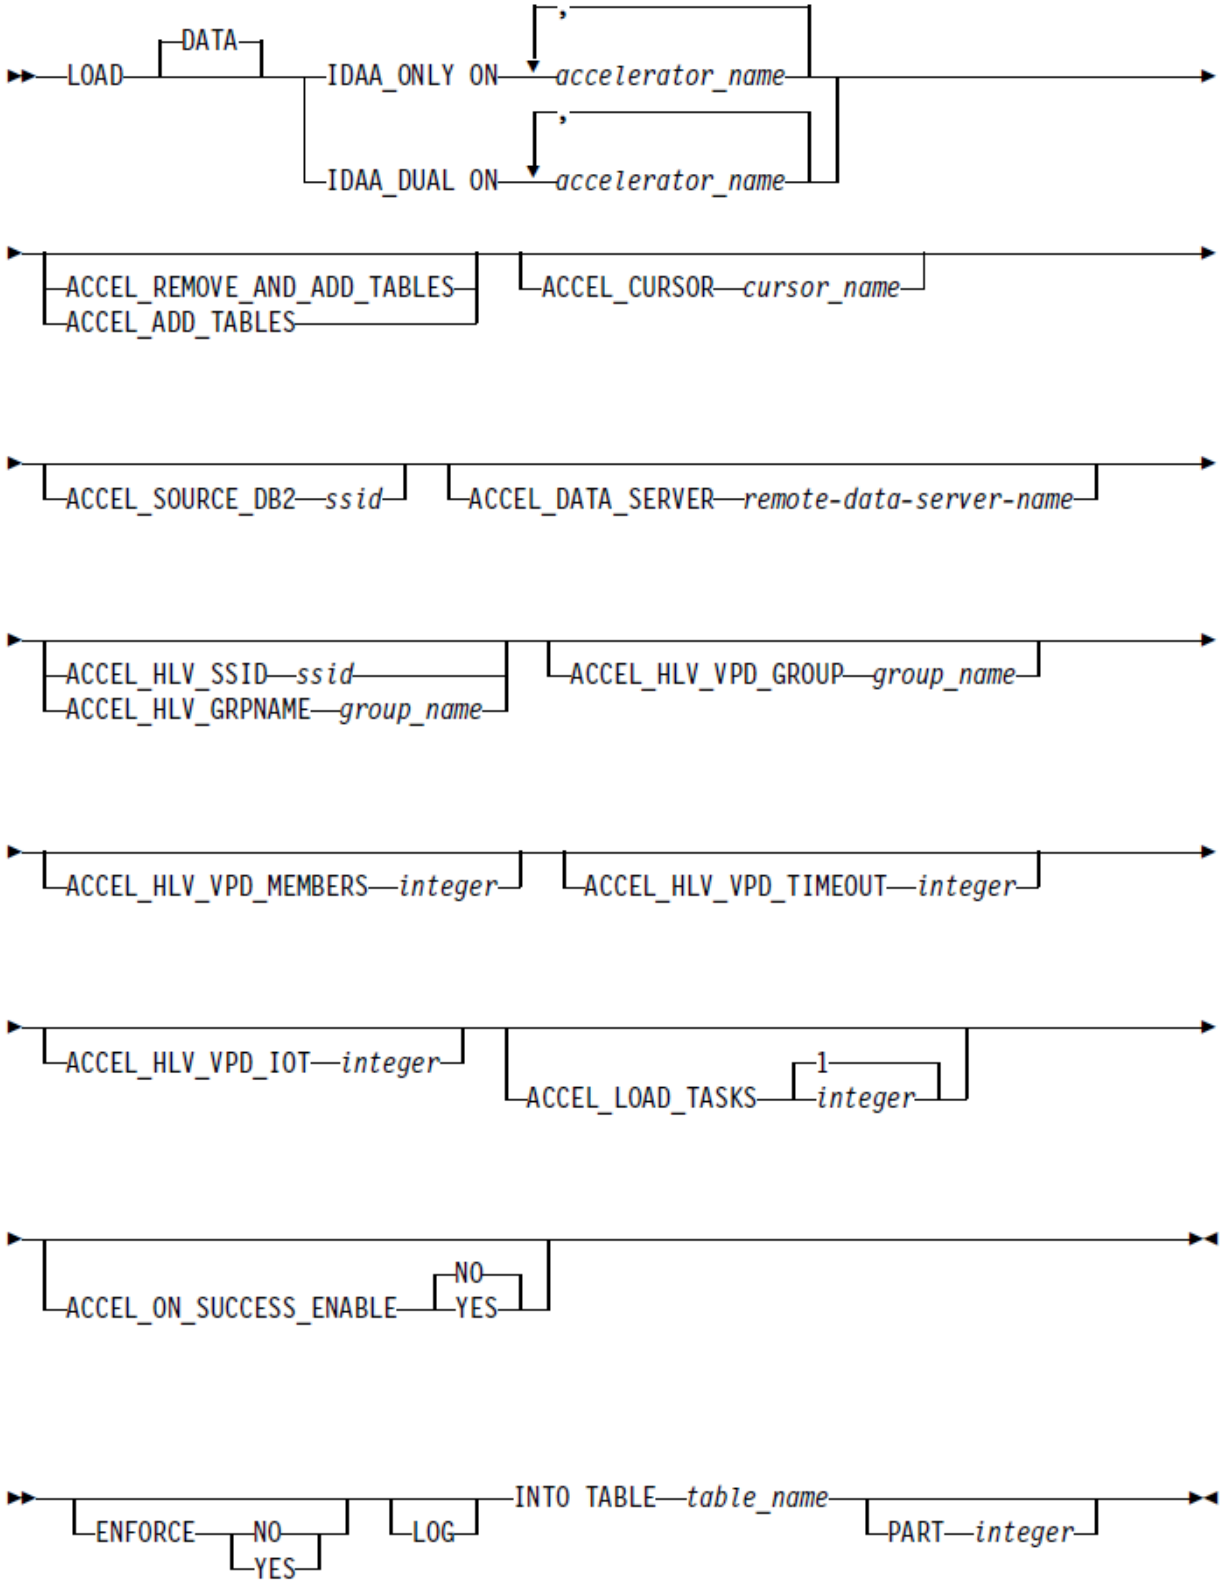

Topic: "Syntax definitions: Load from a non-DB2, remote DB2, or remote system source"
Add or update the following descriptions.

IDAA_DUAL ON *accelerator_name,accelerator_name*

Indicates that you want to load data to up to four accelerators and also to DB2. Replace *accelerator_name* with the name of the accelerators that you want to load, separating multiple accelerator names with a comma. This option is not generated by the Accelerator Loader studio. To use this option, you must manually edit the JCL generated by the Accelerator Loader studio. For more information, see "Customizing the JCL to load the accelerator and DB2".

IDAA_ONLY ON *accelerator_name,accelerator_name*

Indicates that you want to load data to up to four accelerators, and do not want to load to DB2. Replace *accelerator_name* with the names of the accelerators that you want to load. If the load job specifies LOAD REPLACE, existing data in the DB2 table or partition is deleted. This keyword is the default option that is generated by the Accelerator Loader studio.

Chapter "Troubleshooting"

Topic: "Accelerator Loader messages"
Replace the following message:

HLOP9953E Discard datasets are not supported when keyword <*keyword*> is specified.
Explanation: Discard data sets cannot be specified with the named Accelerator Loader keyword. The IDAA_ONLY and ACCEL_CURSOR keywords cannot be used when discard data sets are provided. Remove the SYSDISC ddname from the JCL or remove the 'DISCARDDN' keyword from the LOAD statement and resubmit the job.
User Response: Correct the syntax and resubmit the job.