

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

**Date of change:** December 2017

**Topic:** Multiple

**Change description:** Documentation changes made in support of PTF UI52856 APAR PI90862 – Improve discard support in External Load

- Topic "What's new" in chapter "Overview"
- Topic "Restrictions and considerations for loading from an external file" in chapter "Loading data from an external file"
- Topic "Example JCL: Loading from an external file" in chapter "Syntax"
- Topic "Syntax diagram: Load from an external file" in chapter "Syntax"
- Topic "Syntax definitions: Load from an external file" in chapter "Syntax"
- Topic "Accelerator Loader messages" in chapter "Troubleshooting"

---

## Chapter "Overview"

**Topic: "What's new"**

Add the following description:

The handling of discarded records when loading from an external file using DB2 Analytics Accelerator has changed, behaving more like the DB2 LOAD utility. Under some conditions, the Accelerator Loader may now complete with a different return code than it would have previously.

- When you load only the accelerator, a discard data set is now supported when running a load with a SYSREC data set. The DISCARDS keyword can also be specified to force Accelerator Loader to fail once a specified number of records are discarded.
- The IGNORE keyword of the DB2 LOAD utility is now supported. The IGNORE keyword controls how different types of discards are handled by Accelerator Loader. Ignored discards are not written to the discard data set and do not count towards the discard limit.
- A data conversion error will now cause Accelerator Loader to fail unless a discard data set is provided or IGNORE(CONV) is specified in the LOAD control cards. Previously, the record that caused the violation would be discarded and processing would continue.
- Accelerator Loader will now end with RC=0 even when records are discarded, provided the associated discard reasons are being IGNORED. Previously, any discard would cause Accelerator Loader to end with RC=4. Since loads from an Accelerator Loader server data source do not support a discard data set, these loads will complete with RC=4 any time there are discards regardless of IGNORE settings.

See "Discard data set restrictions and considerations" for more information.

---

## Chapter "Loading data from an external file"

**Topic: "Restrictions and considerations for loading from an external file"**

Remove the following bulleted item from sections "Considerations for loading only the accelerator" and "Discard data set restrictions and considerations":

- When you load only the accelerator, discard data sets are not supported. If you specify the DISCARD DDN option, the utility terminates with errors. Accelerator Loader does not write constraint-violating records to a DISCARD DDN data set.

Add the following bulleted item to "Considerations for loading only the accelerator" and "Discard data set restrictions and considerations":

- When you load only the accelerator, a discard data set is supported when running a load with a SYSREC data set. The discard data set cannot be a TEMPLATE. It must be specified as a DD statement in the JCL. Use the DISCARDN keyword to communicate the DDNAME to Accelerator Loader. The DISCARDS keyword can also be specified to force Accelerator Loader to fail once a specified number of records are discarded. The DISCARDS keyword is valid only when used with a discard data set. Discard data sets are not supported when loading data from an Accelerator Loader server data source or when loading from multiple partition-level SYSREC data sets.

Add the following additional bulleted items to "Discard data set restrictions and considerations":

- The IGNORE keyword of the DB2 LOAD utility is supported. The IGNORE keyword controls how different types of discards are handled by Accelerator Loader. Accelerator Loader can discard a record for the following reasons, each of which can be ignored via the IGNORE clause:
  - The record does not satisfy a WHEN clause. Specify IGNORE(WHEN) to ignore these discards. If a discard data set is not provided, this type of discard is ignored automatically.
  - The partition key for the record is out of the range of any loaded partition. Specify IGNORE(PART) to ignore these discards. If a discard data set is not provided, this type of discard is ignored automatically.
  - A data conversion error occurs when building the DB2 format internal row. Specify IGNORE(CONV) to ignore these discards.

Multiple reasons can be combined in the IGNORE clause, such as in the following example:

**IGNORE(WHEN, PART, CONV)**

Ignored discards are not written to the discard data set and do not count towards the discard limit. No record-level messages are generated for ignored discards. Record-level messages are written for each non-ignored discard. These messages identify the record number and describe why it was discarded. To avoid flooding the spool with these record-level messages, only the first 1000 non-ignored discards are reported in this way.

**Note:** IGNORE settings VALPROC, IDERROR, and DUPKEY are ignored by Accelerator Loader and passed to the DB2 LOAD utility.

- A data conversion error will cause Accelerator Loader to fail unless a discard data set is provided or IGNORE(CONV) is specified in the LOAD control cards.
- Accelerator Loader will end with RC=0 even when records are discarded, provided the associated discard reasons are being IGNOREd. Since loads from an Accelerator Loader server data source do not support a discard data set, these loads will complete with RC=4 any time there are discards regardless of IGNORE settings.

---

## Chapter "Syntax"

### Topic: "Example JCL: Loading from an external file"

Add the following example:

#### Example: Using a discard data set when loading only the accelerator

To use a discard data set when loading only the accelerator with a SYSREC data set, you must manually update your JCL to include the following items:

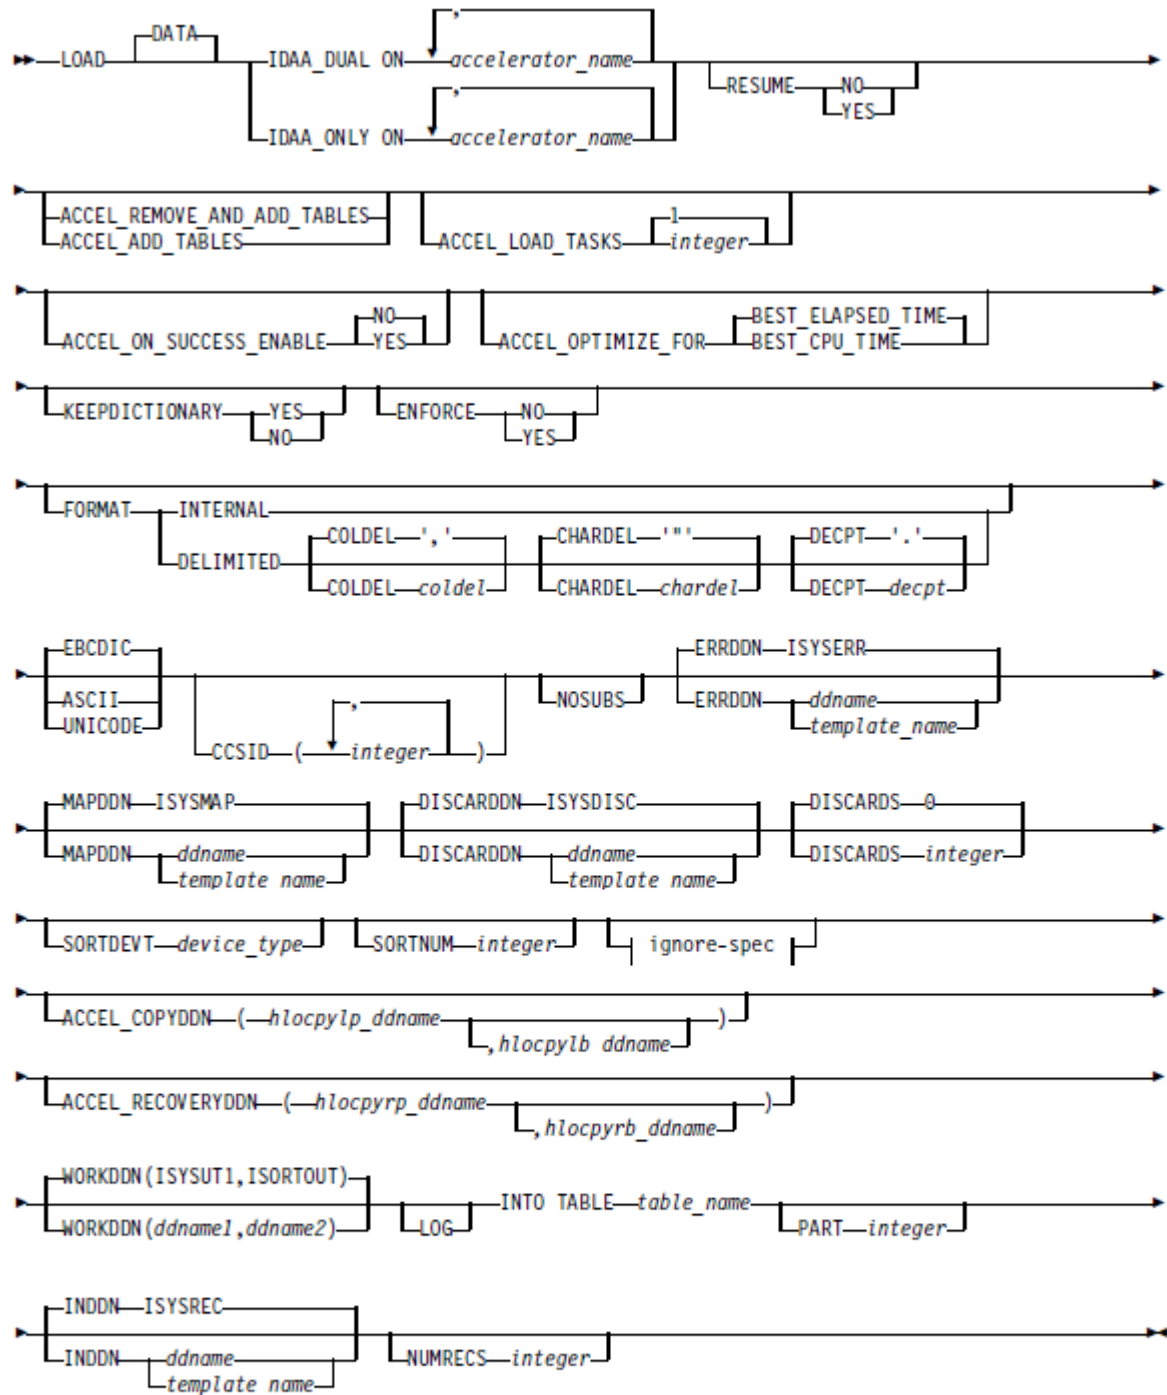
- A DD statement that specifies the discard data set, and the DISCARD DN keyword that specifies the DDNAME
- Optionally, the DISCARDS keyword

The following example shows sample JCL:


```
//HLOD0100 EXEC PGM=DSNUTILB,
//          REGION=0000M,
//          PARM=('DA1A')
//STEPLIB DD DISP=SHR,DSN=RSQA.IBMTAPE.SHLOLOAD
//          DD DISP=SHR,DSN=DA1A.SDSNEXIT
//          DD DISP=SHR,DSN=DSN.VA10.SDSNLOAD
//ISYSREC DD DISP=SHR,DSN=USER.DA1A.SYSREC
//HLODISCD DD DISP=SHR,DSN=USER.DISCARD.DATASET
//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 DA1AACCL
        ACCEL_ON_SUCCESS_ENABLE NO
        INDDN ISYSREC
        ACCEL_LOAD_TASKS 1
        REPLACE
        DISCARDS 10
        UNICODE
        DISCARD DN HLODISCD
        WORKDDN(ISYSUT1,ISORTOUT)
        INTO TABLE
            "AASRC1"."EMPPROJECT"
/*
//*
```

## Topic: "Syntax diagram: Load from an external file"

Replace the diagram with the following updated diagram:



### Ignore-spec:



## Topic "Syntax definitions: Load from an external file"

Add the following syntax elements:

### **DISCARDS** *integer*

Specifies the maximum number of source records that are to be written on the discard data set. *integer* can range from 0 to 2147483647. This keyword is valid only when used with a discard data set. If the discard maximum is reached, the load abnormally terminates, the discard data set is empty, and you cannot see which records were discarded. You can either restart the job with a larger limit, or terminate the utility.

DISCARDS 0 specifies that you do not want to set a maximum value. The entire input data set can be discarded. The default value is 0.

### **IGNORE**

Specifies that the load ignores records that it rejects for the specified reasons. If discarding is specified, no ignored rows are loaded or written to the DISCARD data set. If discarding is not specified, ignored records do not cause the load to terminate.

#### **WHEN**

Specifies that records that do not satisfy the WHEN clause are ignored.

#### **PART**

Specifies that records that do not satisfy any partition being loaded are ignored.

#### **CONV**

Specifies that records that cause a conversion error are ignored.

Multiple reasons can be combined in the IGNORE clause, such as in the following example:

**IGNORE (WHEN, PART, CONV)**

Ignored discards are not written to the discard data set and do not count towards the discard limit. No record-level messages are generated for ignored discards. Record-level messages are written for each non-ignored discard. These messages identify the record number and describe why it was discarded. To avoid flooding the spool with these record-level messages, only the first 1000 non-ignored discards are reported in this way.

**Note:** IGNORE settings VALPROC, IDERROR, and DUPKEY are ignored by Accelerator Loader and passed to the DB2 LOAD utility.

---

## Chapter "Troubleshooting"

### Topic: "Accelerator Loader messages"

Add the following messages:

**HLOP9970E Discard datasets are not supported with IDAA\_ONLY and partition level SYSREC datasets.**

**Explanation:** Discard data sets cannot be specified with the IDAA\_ONLY keyword when loading from multiple partition level SYSREC data sets.

**User response:** Remove the discard data sets from the JCL and LOAD statement and resubmit the job. Alternatively, if each SYSREC data set has the same field specifications, eliminate the INDDN keyword from the INTO TABLE clauses and concatenate all the SYSREC data sets under a single SYSREC DDNAME in the JCL.

**HLOU4033E The discard data set DDNAME <ddname> is missing from the JCL.**

**Explanation:** The DDNAME specified on the DISCARD DD keyword was not specified in the JCL. When running an IDAA\_ONLY load, a template cannot be used for the discard dataset.

**User response:** Specify a valid discard data set in a DD statement in the JCL.

**HLOU4034E Discard record LRECL incompatible with discard DS LRECL.  
Record=<record\_number>.**

**Explanation:** A discarded SYSREC record could not be written to the discard data set because the record length exceeds the discard data set LRECL. This error can occur when the SYSREC is a concatenation of multiple data sets. The product sets the LRECL of the discard data set to the LRECL of the first SYSREC data set in the concatenation. An attempt to discard a record from a subsequent SYSREC data set with a larger LRECL will fail with this error.

**User response:** Make sure the SYSREC data set with the largest LRECL is first in the concatenation.

**HLOU4035E The discard limit has been reached.**

**Explanation:** The limit specified on the DISCARDS keyword has been reached. Accelerator Loader terminates with errors before the input data set is fully processed.

**User response:** Eliminate the DISCARDS keyword or increase the limit value, and resubmit the job.

**HLOU4036I Number of records written to discard dataset=<record\_count>.**

**Explanation:** This informational message indicates the number of SYSREC records that have been written to the discard data set.

**User response:** No action is required.

**HLOU5366I Record=<record\_number> discarded due to partition key value.**

**Explanation:** The record was discarded because its partitioning key did not belong to any partition included in the load.

**User response:** No action is required.

**HLOU5701I Number of records not loaded=<record\_count>**

**Explanation:** SYSREC records were not loaded for any of the following reasons: the record did not satisfy a WHEN clause condition; the partitioning key for the record did not belong to any partition included in the load; a data conversion error occurred. The reported count of records not loaded includes ignored discards.

**User response:** No action is required.