

Version 1 Release 1

*IBM Z Table Accelerator
Administration Guide*



Notices

© Copyright International Business Machines Corporation 2020

© Data Kinetics Ltd. 2020

Note to U.S. Government Users — Documentation related to restricted rights — Use, duplication or disclosure is subject to restrictions set forth in GSA ADP Schedule Contract with IBM Corp.

Contents

- Figures..... vii**
- Tables..... ix**
- Preface.....xi**
 - Audience for this Guide..... xi
 - What is Covered in this Guide..... xi
 - Naming protocol..... xi
 - What You Should Know to Use this Guide..... xi
 - Glossary..... xii
 - Conventions used in this guide..... xiii
 - What This Guide Contains..... xiii
 - Additional IBM Z Table Accelerator references..... xiv
- Chapter 1. Installation Modifications..... 1**
- Chapter 2. Master Password Feature.....5**
- Chapter 3. tableSPACE Reporting.....7**
 - Generating the tableSPACE report for batch, IMS and VTS Agent..... 7
 - Generating the tableSPACE report for CICS..... 7
 - tableSPACE report samples and description..... 8
- Chapter 4. Program Call Server..... 13**
 - Initialization of the Program Call Server..... 13
 - Server Messages..... 13
 - Re-initialization of the Program Call Server..... 13
 - Restarting the IBM Z Table Accelerator Program Call Server..... 14
 - Shutting down the Program Call Server..... 14
- Chapter 5. IMS..... 15**
 - JCL..... 15
 - Pre-loading IBM Z Table Accelerator modules..... 15
 - Application/transaction initialization..... 16
 - Transaction abends..... 16
 - VTS-TSR access..... 16
- Chapter 6. CICS..... 17**
 - JCL..... 17
 - Definitions..... 17
 - Journal..... 17
 - Programs..... 17
 - Transactions..... 17
 - CICS initialization and shutdown..... 18
 - Transaction TAST..... 18
 - Restrictions..... 18
 - Transaction TADR..... 18
 - The TAOPT dataset..... 19

Chapter 7. VTS (Virtual Table Share)	21
Overview.....	21
VTS Agent operations.....	21
Startup.....	21
Initialization of a VTS-TSR.....	22
TAOPT parameters.....	22
VTS-TSR table loading sample JCL.....	22
VTS-TSR refresh process.....	23
VTS Agent shutdown.....	23
VTS Agent recovery.....	23
VTS commands.....	24
VTS operational considerations.....	24
Reducing virtual storage when moving tables into VTS.....	24
7/24 VTS.....	24
Capacity.....	24
Loading a VTS-TSR.....	24
VTS messages.....	25
Internal logic errors.....	25
VTS user messages.....	25
Chapter 8. Best Practices	27
Protecting IBM Z Table Accelerator libraries.....	27
Concurrent multiple updates.....	27
Read-Only libraries.....	27
Protecting IBM Z Table Accelerator tables.....	27
Table read and write passwords.....	27
TSR LOCK-LATCH.....	28
Time out while waiting for a lock.....	28
Performance.....	28
Shared tables: Virtual Table Share (VTS).....	28
Linking IBM Z Table Accelerator.....	28
WLM priority level consideration.....	28
Single TAPARM per thread.....	28
Table access optimization.....	29
Multitasking.....	29
IBM Z Table Accelerator library definition considerations.....	29
Library Blocksize.....	29
IBM Z Table Accelerator library space usage.....	29
BDAM/QSAM IBM Z Table Accelerator libraries.....	29
VSAM IBM Z Table Accelerator libraries.....	29
VSAM IBM Z Table Accelerator libraries and LSR pools.....	30
<i>Enhanced Data Integrity</i> option.....	30
Library directory caching.....	30
TSR Space Allocation.....	31
IBM Z Table Accelerator modules that can reside in the Link Pack Area (LPA).....	31
Chapter 9. IBM Z Table Accelerator Diagnostic Information	33
IBM Z Table Accelerator use of MVS Enqueues.....	33
IBM Z Table Accelerator Library Diagnostics.....	33
Appendix A. IBM Z Table Accelerator run-time options	35
CICSJRNL—CICS Journal File ID.....	35
HASH_HI_DEN_LIM—High Density Limit for Hash Indexes.....	35
HASH_LOW_DEN_LIM—Low Density Limit for Hash Indexes.....	35
LIBnn, ML—IBM Z Table Accelerator Library List.....	35
LISTOPTIONS—List Parameter Options.....	36

LOCKTIMER—Lock Timer Wait Value.....	36
LOCKTIMEWTO—Lock Timer Message Wait Value.....	36
MAXNMTAB—Maximum Number of Tables.....	36
MTRETAIN—Retain Rows and Index Areas.....	37
MULTITASKING—Multitasking.....	37
OVERRIDES—Allow Changes to Status Switches.....	37
RACF_LIBACCESS—Perform Checks on Protected Libraries.....	38
STROBE—Strobe Interval.....	39
STROBEMETHOD—Strobe Method.....	39
SUPPRESS_DUMPS—Suppress IBM Z Table Accelerator Messages and Dumps.....	40
SWITCHES—Status Switches.....	40
TABLEWAITRC—Table Open Enqueue Wait Time.....	41
TABLEWAITWTO—Table Open Enqueue Report Time.....	42
TSR_ALGORITHM—Optimize TSR Usage.....	42
TSR_WARNING_FREQ—Frequency of TSR Allocation Warnings.....	42
TSR_WARNING_PCT—Percentage of TSR Allocation for Warning Activation.....	42
TSRSIZE—tableSPACE Region Size.....	43
VTSNAME—Specifying the Name of a VTS-TSR.....	43
ZEROROWS—Zero Data Table Rows on De-allocation.....	43
Appendix B. IBM Z Table Accelerator messages.....	45
IBM Z Table Accelerator return codes.....	45
IBM Z Table Accelerator messages.....	66
DKJTEEXEC error messages.....	82
Appendix C. PC Server implementation in a CICS environment.....	93
Implementation and Restart/Recovery Considerations.....	93
RESTART/RECOVERY Considerations.....	93
NOTES and CAUTIONS.....	94
Notices.....	95
Trademarks.....	96

Figures

- 1. Sample JCL for batch, IMS, VTS Agent tableSPACE report generation..... 7
- 2. Sample JCL for CICS tableSPACE report generation 8
- 3. Sample strobe report: local TSR (STROBE=5) 9
- 4. Sample strobe report: VTS-TSR (STROBEMETHOD=3, STROBE=5) 10
- 5. Sample VTS Agent Startup JCL 21
- 6. Sample JCL for loading tables23
- 7. Sample JCL for VTS refresh23

Tables

1. Default option overrides	1
2. tableSPACE report field definitions	10
3. IBM Z Table Accelerator modules.....	15
4. Overrides default values	38
5. Switches default values	40
6. IBM Z Table Accelerator messages and error codes.....	45
7. Messages and error codes for VTS Agent, CICS, Batch, etc.	67
8. IBM Z Table Accelerator batch utility messages.....	83

Preface

This guide provides the information and tools for the IBM Z® Table Accelerator administrator. It provides general information that relates to most environments, including CICS-specific administration information and Virtual Table Share (VTS), that is available with Release 1.1.0.

Note: Keep this guide in a secure location. The information provided could allow an individual to override normal IBM Z Table Accelerator software security provisions. The guide should remain in the possession of the IBM Z Table Accelerator administrator or the individual who is authorized to perform administrative functions.

Audience for this Guide

This guide is intended for individuals responsible for the administration of the IBM Z Table Accelerator software.

The responsibilities of IBM Z Table Accelerator administrator may include the following tasks:

- Act as the single point of contact for Technical Support.
- Install and configure IBM Z Table Accelerator components.
- Coordinate implementation of IBM Z Table Accelerator upgrades.
- Maintain policies and procedures for IBM Z Table Accelerator administration and use.
- Set up users and user authorizations.
- Define IBM Z Table Accelerator libraries.
- Maintain company-wide table libraries.
- Maintain security rules for IBM Z Table Accelerator libraries.
- Define and control naming standards for IBM Z Table Accelerator tables.
- Obtain the IBM Z Table Accelerator Master Password.
- Assist system support with IBM Z Table Accelerator start-up jobs for all online regions.

What is Covered in this Guide

This guide provides the background information needed by the IBM Z Table Accelerator administrator to set up and maintain a suitable working environment for the IBM Z Table Accelerator product and any additional IBM Z Table Accelerator-related software products.

Naming protocol

All executables begin with DKJ for easy identification, a prefix that has been reserved exclusively by IBM for IBM Z Table Accelerator.

What You Should Know to Use this Guide

A working knowledge of a programming language (for example, COBOL or Assembler), compiler procedures, a linkage editor, and the MVS JCL systems is required to perform the tasks detailed in this guide. The Administrator should be familiar with the IBM Z Table Accelerator concepts and facilities and the MVS environment.

Glossary

The terms defined in this glossary are used frequently in this document.

Alternate Index	An Alternate Index is an Index that may be defined for a Data Table. The Alternate Index has an Alternate Index definition (ALT-DEFINITION) that defines the key, organization, and search order. Alternate Indexes are optional, and there is no limit to the number of Alternate Indexes a Data Table may have.
Data Table	A Data Table is the actual raw data. Each Data Table has a table definition (DT-BLOCK) that is used to generate the Index for the Data Table.
Delivered defaults	The defaults that are delivered with the product. Also known as <i>factory defaults</i> .
Index	An Index is defined for each Data Table. A Data Table Index is generated dynamically when a table is opened or defined based on the information in the table definition (DT-BLOCK).
Installation defaults	The defaults set at implementation time by an administrator, which may or may not be the same as the delivered defaults. Defined using the DKJOPTGN file. (These defaults may be overridden by an individual application using the TAOPT file.)
Multitasking Batch	An MVS region that implements multitasking by attaching multiple Task Control Blocks (TCBs). This can include a batch job that attaches several subtasks or a transaction processing region that implements multitasking through multiple TCBs.
Table Expansion	Dynamic allocation of space for tables in the TSR when the current space allocated becomes insufficient.
Temporary Table	A temporary table exists only within a TSR, and is created by the DT command (or IA). It is never stored in a library. A temporary table can be distinguished from a library table using the GD command output—if found, a temporary table will show no dataset name.
TSR, Local TSR	Table Space Region. A data space of up to 2G is used by IBM Z Table Accelerator to house tables. The data space is owned by an application in the associated address space. The application uses IBM Z Table Accelerator to access data within the tables.
TSR of record	If there is no VTS name in the TA-SUBSYSTEM field of the TA-PARM when a command is being executed, the TSR of record is the local TSR, otherwise it is the VTS-TSR that has its name in the TA-PARM.
VTS	Virtual Table Share. The term “VTS” refers to the VTS product, which permits TSRs to be shared among applications. These shared TSRs are called VTS-TSRs.
VTS-TSR, Shared TSR	A Virtual Table Share (VTS) Table Space Region (TSR) is a shared TSR, and resides in a shared data space. Applications can access tables within a VTS-TSR, and use the information as if it were within their local TSRs.

VTS Agent	The DKJVAGNT program, which initializes VTS-TSRs in IBM Z Table Accelerator, and then sits idle until the VTS-TSR is to be terminated.
-----------	--

Conventions used in this guide

This guide uses conventions to differentiate code and typed commands, to display the names of parameters, and to describe specific version and release levels of the IBM Z Table Accelerator product.

Convention	Description
code examples and commands	Code examples and commands appear in this type of font: <code>this is an example of the font.</code>
MAXNMTAB	Names of parameters appear in upper case simply for ease of reading; actual case used is upper or lower or a mixture.
Version	Following IBM standards, the term <i>version</i> refers to a generation of a software product that has significant new code or new functionality. <i>Version</i> is a more general term than <i>release</i> . For example, <i>Version 1</i> includes <i>Release 1.1</i> and <i>Release 1.2</i> , and is equivalent to <i>Release 1.x</i> .
Release	Following IBM standards, the term <i>release</i> refers to a program or set of programs which represent a specific revision to the base version of a software product. For example, <i>Release 1.1</i> is a term that is used to identify the first release of <i>Version 1</i> . Subsequent releases made available under the Version 1 umbrella, such as <i>Release 1.2</i> , will provide additional revisions to the base product.
Modification Level	Following IBM standards, the term <i>modification level</i> refers to the application of specific program enhancements and error corrections to the release of a software product. For example, <i>Release 1.1.3</i> is at <i>modification level 3</i> , and <i>Release 1.1.0</i> is at <i>modification level 0</i> .
MVS	MVS is a generic term which is used when referring to z/OS and other related IBM operating systems.

What This Guide Contains

Chapter 1 presents the default values for IBM Z Table Accelerator parameters and explains how to change them if desired.

Chapter 2 describes the use of the Master Password for IBM Z Table Accelerator tables and how to obtain it.

Chapter 3 demonstrates how to use the tableSPACE Reporting feature to generate usage statistics for IBM Z Table Accelerator.

Chapter 4 presents the Program Call Server feature of IBM Z Table Accelerator, explains what it does, and shows the messages it will generate in a variety of circumstances.

Chapter 5 provides information on how to tailor the IMS environment to accommodate IBM Z Table Accelerator.

Chapter 6 provides information on how to tailor the CICS environment to accommodate IBM Z Table Accelerator.

Chapter 7 discusses the administrative issues that arise when using VTS (Virtual Table Share) product with IBM Z Table Accelerator.

Chapter 8 contains information on optimizing and organizing the IBM Z Table Accelerator environment that should be of interest to IBM Z Table Accelerator administrators.

Chapter 9 identifies and describes the IBM Z Table Accelerator diagnostics tools for the IBM Z Table Accelerator administrator.

Appendix A lists and explains the IBM Z Table Accelerator run-time option parameters.

Appendix B contains all of the IBM Z Table Accelerator error codes and messages that can be encountered during the normal installation, administration and operation of the product.

Appendix C discusses the implementation and restart/recovery and other considerations of using the IBM Z Table Accelerator PC Server in a CICS environment.

Additional IBM Z Table Accelerator references

This guide is one of a series that describe IBM Z Table Accelerator:

- *IBM Z Table Accelerator Implementation Guide*
- *IBM Z Table Accelerator Concepts and Facilities Guide*
- *IBM Z Table Accelerator Batch Utilities Guide*
- *IBM Z Table Accelerator Administration Guide*
- *IBM Z Table Accelerator Programming Guide*
- *IBM Z Table Accelerator Quick Reference Guide*

Chapter 1. Installation Modifications

IBM Z Table Accelerator has factory defaults; these can be modified for your site at installation using the DKJOPTGN file. Applications can further override options on an individual basis, using the TAOPT file. Table 1 on page 1 lists all of the applicable parameters, provides a brief description, and shows default values. For a complete description of all options, see Appendix A, “IBM Z Table Accelerator run-time options,” on page 35. For further information see *IBM Z Table Accelerator Implementation Guide*.

<i>Table 1. Default option overrides</i>		
Parameter	Description	Default †
CICSJRNL	CICS Journal File ID (CICS only) - Specifies the CICS journal file number if strobe records are written.	99
HASH_HI_DEN_LIM	High Density Limit for Hash Indexes - Upper density limit for hash tables opened.	900
HASH_LOW_DEN_LIM	Low Density Limit for Hash Indexes - Lower density limit for hash tables opened.	600
LIBnn, ML	IBM Z Table Accelerator Library List - Default library names for a IBM Z Table Accelerator Library List (LIB-LIST).	LIB00=MAINLIB
LISTOPTIONS	List Parameter Options - Specifies whether to list execution time parameters.	
LOCKTIMERC	Lock Timer Wait Value - Specifies # of seconds that IBM Z Table Accelerator should wait for a lock (0= wait indefinitely).	
LOCKTIMEWTO	Lock Timer Message Wait Value - Number of seconds between messages indicating that the transaction is waiting for an internal lock (0=no message).	30
MAXNMTAB	Maximum Number of Tables - Specifies the maximum number of open tables allowed in the TSR.	Depends on TSR size - see <i>Programming Guide</i> for details
MTRETAIN	Retain Rows and Index Areas - Determines whether the allocated row and index areas are to be retained when an MT command is issued.	
MULTITASKING	Multitasking Option(Batch only) - Specifies whether this region will allow multiple subtasks to access IBM Z Table Accelerator.	
RACF_LIBACCESS	Perform Checks on Protected Libraries - Specifies whether IBM Z Table Accelerator should check if a IBM Z Table Accelerator library is SAF interface protected.	

Table 1. Default option overrides (continued)

Parameter	Description	Default †
STROBE	Strobe Interval - Specifies frequency of strobe records (in # of IBM Z Table Accelerator calls or minute intervals).	
STROBEMETHOD	Strobe Method - Specifies when IBM Z Table Accelerator strobe reporting should be triggered for the VTS-TSR. <ol style="list-style-type: none"> 0, 1, 2 indicates that strobe should be triggered by number of IBM Z Table Accelerator calls indicated by the STROBE option 3 indicates that strobe should be triggered by minute intervals indicated by the STROBE option 	
SUPPRESS_DUMPS	Suppress IBM Z Table Accelerator messages and dumps - Specifies whether to suppress IBM Z Table Accelerator messages and dumps when a transaction abends outside IBM Z Table Accelerator code. Only for the CICS environment.	NONE
SWITCHES	Status Switches - Specifies status switch defaults (Y or N): <ol style="list-style-type: none"> whether to abend on IBM Z Table Accelerator user errors (1-99;1001-1099) whether to wait for tables that are opened for write whether to return empty rows from hash tables whether to automatically open tables on read accesses whether to trace last 10 IBM Z Table Accelerator commands 	YNYYN for batch/ CICS NNYYN for IMS
TABLEWAITRC	Table Open Enqueue Wait Time - Number of seconds IBM Z Table Accelerator should wait for exclusive table enqueue (applies only if WAIT switch is on).	
TABLEWAITWTO	Table Open Enqueue Report Time - Number of seconds between messages indicating that the transaction is waiting for an enqueue.	
TSR_ALGORITHM	Optimize TSR Usage - Determines whether the Performance (P), or Default (D) attributes are predominant when allocating space in a TSR.	
TSR_WARNING_FREQ	Frequency of TSR Allocation Warnings - Number of seconds between messages indicating the TSR is more than TSR_WARNING_PCT allocated.	0000
TSR_WARNING_PCT	Percentage of TSR Allocation for Warning Activation - Indicates percentage of TSR allocation allowed before a message is generated.	85%

Table 1. Default option overrides (continued)

Parameter	Description	Default †
TSRSIZE	tableSPACE Region Size - Specifies the size of Data Space to be allocated for open tables for a region.	10M
VTSNAME	Specifies the name of a Shared TSR (VTS-TSR).	
ZEROROWS	Zero Data Table Rows on De-allocation - Determines whether the data rows area should be zeroed when table space in a TSR is deallocated.	

†The default settings listed here are delivered defaults of IBM Z Table Accelerator. Your site may have set the installation defaults differently.

Chapter 2. Master Password Feature

The IBM Z Table Accelerator tables can be protected on the library (read and write passwords) and when loaded in memory (LOCK-LATCH). It is important to note that these passwords offer only limited protection against unauthorized use of IBM Z Table Accelerator tables.

In the event that a table password (read, write or LOCK-LATCH) is lost or forgotten, IBM Z Table Accelerator provides a safety mechanism that may be used by the IBM Z Table Accelerator administrator. The table may be accessed by supplying a master password in place of the normal read or write password or in place of the LOCK-LATCH if the table is open in-memory. Closing a table using the master password as the LOCK-LATCH releases the LOCK-LATCH.

Please contact Technical Support for more information on how to obtain the master password.

Chapter 3. tableSPACE Reporting

IBM Z Table Accelerator can produce a report summarizing your use of IBM Z Table Accelerator. This report, known as the tableSPACE Report, includes such information as the space occupied by each table in the tableSPACE Region (TSR), the number of accesses for each table, and much more. These statistics provide information on usage during the life of the region and can be used to optimize performance in all environments.

You control the frequency with which this information is collected. For a local TSR, this is done using the STROBE parameter. For example, if you specify a strobe interval of 10,000, the statistics collection routine is called after every 10,000 calls to IBM Z Table Accelerator. The tableSPACE report can be printed on demand or when IBM Z Table Accelerator terminates.

For a VTS-TSR, the STROBEMETHOD parameter is also required. A value of 0, 1 or 2 for STROBEMETHOD means that the STROBE parameter still represents the number of calls after which to collect IBM Z Table Accelerator usage statistics. A value of 3 for STROBEMETHOD means that STROBE now represents the number of minute intervals after which to collect usage statistics. The default value for STROBEMETHOD is 0.

For STROBEMETHOD of 0, 1 or 2, strobe processing is prevented from waiting for the strobe report. If the number of calls in the STROBE parameter is set to a small value and IBM Z Table Accelerator calls are made more frequently than strobe reports can be produced, the frequency of reports may change due to strobe reports being bypassed.

For IBM Z Table Accelerator performance reasons, IBM Z Table Accelerator strobe is handled by a sub-task. Since the sub-task is running concurrently with IBM Z Table Accelerator calls, the status of the TSR can change while the strobe data is being gathered. This may lead to apparent inconsistencies in the reports, such as the total number of tables open, and the actual number listed.

Generating the tableSPACE report for batch, IMS and VTS Agent

To generate a strobe report for the batch, IMS, or VTS Agent environments, add the following to your JCL. Figure 1 on page 7 is a sample.

```
/** INSERT YOUR JOB CARD HERE
/**
/**TAOPT DD DISP=SHR,DSN=*your.prefix*.taopt_file
/**
/**TATSRT DD SYSOUT=*
/**
```

Figure 1. Sample JCL for batch, IMS, VTS Agent tableSPACE report generation

Figure 3 on page 9 and Figure 4 on page 10 show examples of a local TSR and VTS-TSR report. “tableSPACE report samples and description” on page 8 describe the definitions for the report fields.

Generating the tableSPACE report for CICS

tableSPACE usage statistics are generated and written to the CICS Journal File after every strobe interval if the run time options, STROBE and CICSJRNL are not set to zero. A tableSPACE report is then generated by executing the IBM Z Table Accelerator utility program DKJTSTRB in batch. It reads the journalled data and produces a formatted tableSPACE report.

With this release, we assume that your CICS journal files are written to the MVS system logger. To activate the Strobe feature, you must understand the TAOPT mechanism and, in particular, the STROBE and CICSJRNL TAOPT run-time parameters.

For details on how to set up your CICS Journal File (CICSJRNL) and strobe interval (STROBE), see the *IBM Z Table Accelerator Implementation Guide*.

Member DKJTSTRB of the IZTA.CNTL library from the installation media contains a sample job to produce the tableSPACE report. Before executing the job, customize the job according to your installation standards. [Figure 2 on page 8](#) shows a sample for DKJTSTRB.

```
/* Insert your job card here
/*
/* Note: use this job if you use CICS/TS
/*
/* This job extracts information from the CICS interface journal
/* file and produces the Table Space report
/*
/* Change the parameters below as required
/*
/* Extract Strobe records and produce Strobe reports
/*
//          EXEC PGM=DKJTSTRB
//STEPLIB  DD DISP=SHR,DSN=your.prefix.IZTA.LOAD          <=====
//SYSPRINT DD SYSOUT=*,RECFM=FBA
//SYSUDUMP DD SYSOUT=*
//TADUMP   DD SYSOUT=*
//TALOG    DD DSN=*your.journal.logstream.name*         <=====
//          BLKSIZE=32760,
//          SUBSYS=(LOGR,DFHLGCNV)
//TAREPORT DD SYSOUT=*
/*
```

Figure 2. Sample JCL for CICS tableSPACE report generation

tableSPACE report samples and description

Figure 3 on [page 9](#) is a sample of a tableSPACE report for a local TSR. The report can be generated for any number of strobe cycles; this example shows three strobe cycles with STROBE set to 5). Note that each strobe shows the CALL COUNT value incremented by 5.

The report has a header and then fields that detail each strobe cycle. The information for each strobe cycle has the same format. The last section of the report provides details on each table that is open at the end of the strobe interval, its usage, and status in the TSR.

```

***** TOP OF DATA *****
                IZTA 1.1.0 STROBE REPORT                PAGE: 1
TABLE SPACE REGION NAME LOCAL_DS TSR REGION START 2017/07/07 15:06:26
TSR SIZE 209,715,200 STROBE INTERVAL 5
***** START OF STROBE NUMBER 1 *****
TIME OF STROBE 2017/07/07 15:06:26 CALL COUNT 5
TSR USED 552,960 TSR USED HWM 552,960
LENGTH OF OPEN TABLES 77,824 LENGTH OF TABLES HWM 77,824
NUMBER OF OPEN TABLES 1 NUMBER OF TABLES HWM 1
***** O P E N T A B L E S *****
TABLE NAME ACCESS LENGTH READS UPDATES TOTAL ACCESSES
TADRTUTR R 77,824 4 0 4
***** END OF STROBE NUMBER 1 *****

***** START OF STROBE NUMBER 2 *****
TIME OF STROBE 2017/07/07 15:06:27 CALL COUNT 10
TSR USED 552,960 TSR USED HWM 552,960
LENGTH OF OPEN TABLES 77,824 LENGTH OF TABLES HWM 77,824
NUMBER OF OPEN TABLES 1 NUMBER OF TABLES HWM 1
***** O P E N T A B L E S *****
TABLE NAME ACCESS LENGTH READS UPDATES TOTAL ACCESSES
TADRTUTR R 77,824 9 0 9
***** END OF STROBE NUMBER 2 *****

***** START OF STROBE NUMBER 3 *****
TIME OF STROBE 2017/07/07 15:06:27 CALL COUNT 15
TSR USED 552,960 TSR USED HWM 552,960
LENGTH OF OPEN TABLES 77,824 LENGTH OF TABLES HWM 77,824
NUMBER OF OPEN TABLES 1 NUMBER OF TABLES HWM 1
***** O P E N T A B L E S *****
TABLE NAME ACCESS LENGTH READS UPDATES TOTAL ACCESSES
TADRTUTR R 77,824 14 0 14
***** END OF STROBE NUMBER 3 *****
***** END OF STROBE REPORT *****

```

Figure 3. Sample strobe report: local TSR (STROBE=5)

Figure 4 on page 10 is a sample of a tableSPACE report in minute intervals for a VTS-TSR using STROBEMETHOD=3. This example shows three strobe cycles with STROBE set to 5. Note that each strobe shows the TIME OF STROBE value incremented by 5 minutes.

The rest of the report is similar to that for a local TSR. Table 2 on page 10 describes each field of the report.

```

IZTA 1.1.0 STROBE REPORT                                PAGE: 1
TABLE SPACE REGION NAME  VTSKM355   TSR REGION START 2017/10/31 11:25:32
TSR SIZE                 10,485,760  STROBE INTERVAL 5
***** START OF STROBE NUMBER 1 *****
TIME OF STROBE          2017/10/31 11:30:32  CALL COUNT      26
TSR USED                45,056             TSR USED HWM    45,056
LENGTH OF OPEN TABLES 16,384             LENGTH OF TABLES HWM 16,384
NUMBER OF OPEN TABLES 2                 NUMBER OF TABLES HWM 2
***** O P E N   T A B L E S *****
TABLE NAME ACCESS      LENGTH      READS      UPDATES  TOTAL ACCESSES
TBL01      W           8,192      4         11       15
TBL02      W           8,192      0         11       11
***** END OF STROBE NUMBER 1 *****
***** START OF STROBE NUMBER 2 *****
TIME OF STROBE          2017/10/31 11:35:32  CALL COUNT      26
TSR USED                45,056             TSR USED HWM    45,056
LENGTH OF OPEN TABLES 16,384             LENGTH OF TABLES HWM 16,384
NUMBER OF OPEN TABLES 2                 NUMBER OF TABLES HWM 2
***** O P E N   T A B L E S *****
TABLE NAME ACCESS      LENGTH      READS      UPDATES  TOTAL ACCESSES
TBL01      W           8,192      4         11       15
TBL02      W           8,192      0         11       11
***** END OF STROBE NUMBER 2 *****
***** START OF STROBE NUMBER 3 *****
TIME OF STROBE          2017/10/31 11:40:32  CALL COUNT      30
TSR USED                45,056             TSR USED HWM    45,056
LENGTH OF OPEN TABLES 16,384             LENGTH OF TABLES HWM 16,384
NUMBER OF OPEN TABLES 2                 NUMBER OF TABLES HWM 2
***** O P E N   T A B L E S *****
TABLE NAME ACCESS      LENGTH      READS      UPDATES  TOTAL ACCESSES
TBL01      W           8,192      4         11       15
TBL02      W           8,192      4         11       15
***** END OF STROBE NUMBER 3 *****
***** END OF STROBE REPORT *****

```

Figure 4. Sample strobe report: VTS-TSR (STROBEMETHOD=3, STROBE=5)

Table 2. tableSPACE report field definitions		
Section	Field name	Contains
Header information	tableSPACE Region name	LOCAL_DS if strobe is for a local TSR. VTS-TSR name if the strobe is for a VTS-TSR
	TSR Region start	The date and time that the TSR Region started
	TSR size	The size in bytes of the TSR
	Strobe interval	The strobe interval for the region
Start of strobe	Strobe number	The strobe number for which the following information is being written
	Time of strobe	The date and time at which the statistics were captured for this strobe
	Call count	The total number of IBM Z Table Accelerator calls that have been made at the end of this strobe interval. Only calls that reference the TSR are counted (calls such as BN, LS, ML, CS, and LL are not counted). The call count can accommodate up to 13 digits (maximum value of 9999,999,999,999).

Table 2. tableSPACE report field definitions (continued)

Section	Field name	Contains
	TSR used	The amount of TSR that is in use at the end of the strobe interval
	TSR used HWM	The high water mark is the maximum amount of tableSPACE used up to now
	Length of open tables	The total size of all open tables in the TSR at the end of the strobe interval. The length field can accommodate up to 11 digits (maximum value of 99,999,999,999).
	Length of tables HWM	The highwater mark is the maximum size of all open tables in the TSR at any time since the region has been active. The length field can accommodate up to 11 digits (maximum value of 99,999,999,999).
	Number of open tables	The total number of all open tables in the TSR at the end of the strobe interval (note that this value may not match the number of tables listed, due to the use of sub-tasking in performance data collection)
	Number of tables HWM	The highwater mark is the maximum number of all open tables in the TSR at any time since the region has been active
Open tables	Table name	The name of the table occupying that position in the TSR.
	Access	R for read or W for write
	Length	The number of bytes in the table. The length field can accommodate up to 11 digits (maximum value of 99,999,999,999).
	Reads	The number of reads in this TSR. The number of reads can be up to 13 digits (maximum value of 9999,999,999,999).
	Updates	The number of inserts, deletes, or replaces in this TSR. The number of updates can be up to 13 digits (maximum value of 9999,999,999,999).
	Total accesses	The sum of Reads and Updates. The total accesses can be up to 13 digits (maximum value of 9999,999,999,999).

Note: The page number field at the top of each page of the strobe report can be up to 5 digits long (maximum value of 99999).

For further information on the use of the tableSPACE Report and the STROBE parameter, see the *IBM Z Table Accelerator Programming Guide*.

Chapter 4. Program Call Server

IBM Z Table Accelerator uses Program Call (PC) functionality, provided by IBM, to allow IBM Z Table Accelerator program calls (PCs) to be available for IBM Z Table Accelerator regions. This functionality is provided by the IBM Z Table Accelerator PC Server.

The PC functionality supported by the PC Server provides the functions needed for IBM Z Table Accelerator multi-tasking, for CICS thread-safe support, for accessing a VTS-TSR, as well as other special functions that are required by IBM Z Table Accelerator.

By using program calls, IBM Z Table Accelerator can support those tasks that require special authorization. The PC Server also provides authorized address spaces for each LPAR in which IBM Z Table Accelerator can run the following special functions:

- End-of-task exit clean-up for each region
- Enqueues, dequeues and cross-memory posts entailed in accessing IBM Z Table Accelerator libraries in conjunction with a VTS-TSR
- Monitor the connections and disconnections of each VTS-TSR
- Enqueues and dequeues for accessing IBM Z Table Accelerator in a CICS thread-safe application. For more details on PC Server implementation in a CICS environment, see [Appendix C, “PC Server implementation in a CICS environment,”](#) on page 93.

Note: The PC Server is needed, must be installed, and must be running at all times before IBM Z Table Accelerator initializes in a region.

Initialization of the Program Call Server

The IBM Z Table Accelerator PC Server is required for each LPAR (MVS image) that is running IBM Z Table Accelerator. Each PC Server must be initialized before the first use of IBM Z Table Accelerator on the LPAR.

Initialization of the IBM Z Table Accelerator PC Server is achieved by either submitting a batch job or a started task. It is recommended that it be set up as a started task and initiated at system initialization. Sample JCL is provided in the installation library, **your.prefix.CNTL**. For more information, see the *IBM Z Table Accelerator Implementation Guide*. Your systems programmer can assist you with the initialization of the Server.

The modules for the PC Server can only reside in an LPA, STEPLIB or SYSLIB; they cannot be loaded from LINKLIB.

Server Messages

Messages from the IBM Z Table Accelerator PC Server will indicate successful initialization. If the PC Server is not available, messages will be returned indicating this condition, with explanations where applicable.

For message details, see [Appendix B, “IBM Z Table Accelerator messages,”](#) on page 45.

Re-initialization of the Program Call Server

If the Server is brought down for any reason, all regions that have already completed IBM Z Table Accelerator initialization will continue to operate using the Program Calls that were established when these jobs initialized. Any regions which attempt to initialize while the Server is down will fail. If the Server is brought back up at a different modification level, all new regions will initialize at the modification level of the new Server. All regions that were running before the previous Server came down and after the current Server was brought up will continue to run at the level of the previous Server.

Notes:

1. For CICS, the TAST TERM transaction may be used, followed by the TAST TAINIT transaction, to re-attempt initialization. However, do not attempt to do this without first calling Technical Support.
2. See [Appendix C, “PC Server implementation in a CICS environment,”](#) on page 93 for PC Server implementation in a CICS environment.

Restarting the IBM Z Table Accelerator Program Call Server

If the IBM Z Table Accelerator PC Server was brought down and needs to be restarted at the same level, resubmit the batch job or started task which was used to initialize the PC Server.

If the IBM Z Table Accelerator PC Server was brought down and needs to be restarted at a different level, modify the batch job or started task to point to the new load library and submit it.

Refer to the *IBM Z Table Accelerator Implementation Guide* for details on the job to use to start the IBM Z Table Accelerator PC Server.

Shutting down the Program Call Server

In order to shutdown the IBM Z Table Accelerator PC Server before system shutdown, all VTS regions that are running under the Server must first be shutdown. The MVS STOP (P) command can be used to shutdown the Server.

IBM Z Table Accelerator PC Server shutdown

The IBM Z Table Accelerator PC Server may be shutdown normally by issuing, from an MVS console, a stop command to the PC Server job or started task. The form of the command is:

P jobname

Chapter 5. IMS

This chapter provides information on how to tailor the IMS environment to accommodate IBM Z Table Accelerator.

JCL

Add the following DD statements to the JCL of your BMPs or MPRs that will be using IBM Z Table Accelerator:

1. To the STEPLIB concatenation, add

```
// DD DISP=SHR,DSN=your.prefix.LOAD
```

2. Add, for each IBM Z Table Accelerator library to be referenced

```
//libname DD DISP=SHR,DSN=your.prefix.table_lib
```

Here, libname is the name that will be used in LIB-LIST of applications calling IBM Z Table Accelerator to reference the library specified by the DSNNAME value.

3. Add, if desired, a DD for the TAOPT file (a sequential file or PDS member, of format F or FB, record length 80).

```
//TAOPT DD DISP=SHR,DSN=your.prefix.taopt_file
```

This file is read whenever IBM Z Table Accelerator initializes. See Appendix A, “IBM Z Table Accelerator run-time options,” on page 35 for more information on TAOPT parameters.

4. If you use the Strobe feature, add a DD for the strobe report, for example,

```
//TATS RPT DD SYSOUT=*
```

For IMS programs that are not BMPs or MPRs, see the *IBM Z Table Accelerator Programming Guide* on how to set up your JCL.

Pre-loading IBM Z Table Accelerator modules

It is recommended that all the module names be added to your MPR pre-load list. Not all of the modules included in [Table 3 on page 15](#) are needed in certain IMS MPR configurations. However, determining which modules are required depends on many factors, including whether your applications call IBM Z Table Accelerator dynamically or statically and whether or not your application load modules have been re-linked with the IBM Z Table Accelerator stub. It is simplest to load all of the names in your MPR pre-load.

Pre-Load Module	Alias of	Size (approx) Bytes	Above line
DKJTROTB ^{1,2}		70000	Yes
DKJTNUCL ²		100000	Yes
DKJIBASE		2000	Yes

¹ All modules listed are reentrant.

² If DKJTNUCL or DKJTROTB is in LPA, DKJTVROT and/or DKJTNTPD should also be in the pre-load list.

<i>Table 3. IBM Z Table Accelerator modules (continued)</i>			
Pre-Load Module	Alias of	Size (approx) Bytes	Above line
DKJTVROT		10	Yes
DKJTNTPD		16	Yes
DKJTNAME		1300	Yes
DKJTCALL		6000	Yes

Application/transaction initialization

At the first call to IBM Z Table Accelerator from an MPR application, IBM Z Table Accelerator initializes itself in the region. This consists of reading any TAOPT file, processing the DKJIBASE module, obtaining storage space for control blocks, and acquiring a Data Space for the local TSR. This initialization occurs only once if no transactions abend. If a transaction abends, all resources (memory, TSR) acquired by IBM Z Table Accelerator are freed; any tables in the TSR vanish without being stored. If the MPR control task re-attaches the application control task, IBM Z Table Accelerator initialization will re-occur at the next call to IBM Z Table Accelerator.

This effectively means the contents of the local TSR can be volatile since any updates made to local tables will be lost in the event any transaction abends.

IBM Z Table Accelerator treats the entire sequence of transactions, from first call until normal shutdown of the MPR (or until a transaction abends) as a single 'thread' of execution. A IBM Z Table Accelerator thread has some attributes that persist for the life of the thread, for example, its active ML list, its status switch settings. Those two statements imply that one transaction will inherit the attributes of any earlier transaction that may have set them differently from the default values. Thus writers of MPR transaction modules that call IBM Z Table Accelerator must explicitly set a suitable ML list and status switch values to meet the application's needs. The TB-TURBO field should be set to low values (binary zeros) at the start of each transaction. It should not be changed after that. If a specific VTS-TSR is to be accessed by the application, its name should be put into the TA-SUBSYSTEM field. If no VTS is required, set TA-SUBSYSTEM to low values (binary zeros).

Essentially this states that each transaction module should start off by initializing as if it were its first invocation, no matter how many times it may be used during the MPR's lifetime.

Transaction abends

When a transaction abend occurs in an MPR, the current state of IBM Z Table Accelerator is lost (TSR, tables in it, memory acquired for control information, including thread control information). If it is subsequently called by another transaction, it will re-initialize as if it were starting up for the first time.

VTS-TSR access

If a VTS region is stopped and restarted while an IMS region has access to the VTS-TSR, transactions attempting to access the VTS-TSR will receive error code 1072 until one of the following occurs:

- the transaction is stopped and restarted
- the transaction abends

Chapter 6. CICS

This chapter provides information on how to tailor the CICS environment to accommodate IBM Z Table Accelerator.

JCL

Add the following DD statements to the JCL of your CICS regions that will be using IBM Z Table Accelerator:

1. To the STEPLIB concatenation, add

```
// DD DISP=SHR,DSN=your.prefix.LOAD
```

2. Add, for each IBM Z Table Accelerator library to be referenced

```
//libname DD DISP=SHR,DSN=your.prefix.table_lib
```

Here, libname is the name that will be used in LIB-LIST of applications calling IBM Z Table Accelerator to reference the library specified by the DSNNAME value.

3. If you are using a sequential file or PDS member for TAOPT, instead of a VSAM file, add a DD for the TAOPT file.

```
//TAOPT DD DISP=SHR,DSN=your.prefix.taopt_file
```

A TAOPT sequential file or PDS member must be defined with format F or FB with record length 80. This file is read whenever IBM Z Table Accelerator initializes. See [Appendix A, “IBM Z Table Accelerator run-time options,”](#) on page 35 for more details on the TAOPT parameters.

Note: The TAOPT file can either be a sequential file or a VSAM file. Omit this DD statement if you use a VSAM TAOPT file.

Definitions

Definitions for CICS are covered in the *IBM Z Table Accelerator Implementation Guide*, including:

Journal

Definitions for journal files are covered in the *IBM Z Table Accelerator Implementation Guide*, in the "Create CICS definitions for IBM Z Table Accelerator journal files" section of the "Implementing the CICS Interface" procedure.

Programs

Definitions for programs are covered in the *IBM Z Table Accelerator Implementation Guide*, in the "Create CICS CSD definitions for IBM Z Table Accelerator programs and transactions" section of the "Implementing the CICS Interface" procedure.

Transactions

Definitions for transactions are covered in the *IBM Z Table Accelerator Implementation Guide*, throughout the "Implementing the CICS Interface" procedure.

CICS initialization and shutdown

Under normal circumstances, IBM Z Table Accelerator is started automatically on CICS initialization, via the PLT (PLTSI)—see the *IBM Z Table Accelerator Implementation Guide*, in the "Implementing the CICS Interface" procedure, for further information.

When a CICS region is terminated normally, IBM Z Table Accelerator termination within the region is handled by the program DKJTSHT, via the PLT (PLTSD). If a user issues an emergency shutdown of a CICS region, PLTSD is not executed, and IBM Z Table Accelerator termination is handled by DKJTCRM. Normal region termination, handled by DKJTSHT via the PLTSD, will ensure that all transactions are completed before proceeding with IBM Z Table Accelerator termination for the region.

Transaction TAST

Under CICS, IBM Z Table Accelerator is now initialized when DKJTCIN is invoked through transaction TAST (or the PLT). In earlier releases, it was initialized by the first application call to IBM Z Table Accelerator. TAST can also assist in resolving IBM Z Table Accelerator software problems in a CICS region.

The TAST INIT function initializes the IBM Z Table Accelerator software in the CICS region, essentially in the same manner as during CICS initialization. The TAST INIT function reloads all modules and allocates the TSR. The TAST INIT function does not preload any tables.

Normal completion is indicated by the following messages:

- IZTA 1 - Resource manager initializing
- IZTA Resource manager initialized

The function TAST TERM terminates IBM Z Table Accelerator in the region. (Before using TAST TERM, contact Technical Support.) Ensure that no transactions are using the IBM Z Table Accelerator software before starting this transaction. The TAST TERM function closes all tables in the local TSR (it does not store them). The TSR and all IBM Z Table Accelerator control blocks and modules are deleted from the region.

Normal completion is indicated by the following messages:

- IZTA Deactivating
- IZTA Deactivated

Restrictions

The TAST transaction is supported for stopping and restarting IBM Z Table Accelerator within a CICS region with restrictions. Stopping IBM Z Table Accelerator results in the TSR being released. Restarting IBM Z Table Accelerator results in a new, empty TSR being created. This transaction is only supported if all transactions accessing IBM Z Table Accelerator have completed or have been purged. If any transactions are still accessing IBM Z Table Accelerator when this transaction executes, unpredictable results, including abends and blocked transactions, can result. Recovery could entail recycling CICS and any affected VTS-TSRs.

The IBM Z Table Accelerator CICS interface does not support the loading of DKJTROTC from LPA.

IBM Z Table Accelerator CICS modules must be loaded from DFHRPL.

Transaction TADR

Transaction TADR is covered in detail in the *IBM Z Table Accelerator Programming Guide*.

The TAOPT dataset

If a VSAM TAOPT file is to be used, it has to be defined in CICS. With full support of QSAM TAOPT files, IBM Z Table Accelerator has to determine how to access the TAOPT dataset. If it is defined under CICS, then CICS I/O is used to read it; if not, and the DD is present, MVS I/O will be used. Note that a CICS TAOPT definition is not sent out.

IBM Z Table Accelerator does not initialize if there are apparent user errors in defining TAOPT. If the CICS definition is found, it is assumed that TAOPT is desired, and an error is generated if the VSAM TAOPT file cannot be found.

Chapter 7. VTS (Virtual Table Share)

This chapter discusses the administrative issues that arise when using VTS (Virtual Table Share). This chapter provides background information needed by the IBM Z Table Accelerator administrator to set up and maintain VTS's using VTS Agent.

Note: VTS allows many regions/applications to share in-memory tables.

Overview

A VTS-TSR is a shared tableSPACE Region (TSR), and resides in a shared dataspace. A VTS-TSR allows shared access to IBM Z Table Accelerator tables across multiple batch, TSO, CICS TS and IMS TM regions in a single LPAR or z/OS image. Each VTS-TSR has a VTS Agent program which manages the Data Space.

By loading one copy of data into a VTS-TSR, the I/O, system paging, and overall memory requirements can be reduced. A VTS-TSR can be referenced by any programming language that uses the standard IBM calling protocols, including C and C++.

You may have several differently named VTS Agents active simultaneously, allowing you to access multiple shared VTS-TSRs, each of which has a maximum size of approximately 2 GB. (Your installation may impose limits on the number or maximum size of Data Spaces.)

Applications accessing a VTS-TSR may use any mix of batch, TSO, CICS TS or IMS TM interfaces. If a table is opened for write in a VTS-TSR by any user application, it is available for update by any other users of the VTS-TSR. The table cannot, however, be opened for write in any other TSR. It is the users' responsibility to coordinate updates to the table in the VTS-TSR. The LOCK-LATCH capability can also be used as a means to control update access to specific tables in a VTS-TSR.

VTS Agent operations

This section describes the operation of the VTS Agent. Figure 5 on page 21 is a sample VTS Agent start-up JCL. In the sample, the VTS Agent to be started will be known as DKJA.

```
/* Insert your job card here
/*
/* Starts a VTS Agent named DKJA
/* to stop, enter "P jobname" or "f jobname,stop"
/*
//VTS      EXEC PGM=DKJVAGNT,TIME=1440
//STEPLIB DD DISP=SHR,DSN=your.prefix.IZTA.AUTHLIB      <=====
//MAINLIB DD DISP=SHR,DSN=your.prefix.MAINLIB          <=====
//SYSUDUMP DD SYSOUT=*
//TAOPT    DD *
LISTOPTIONS=Y
VTSNAME=DKJA
TSREGION=50M
/*
```

Figure 5. Sample VTS Agent Startup JCL

Startup

The VTS Agent may be initiated either as a batch job, by submission of the start-up JCL, or as an MVS Started Task (STC). The STC could be included in the IPL or daily start-up procedures for your installation.

The VTS Agent must be run from an authorized library. In the startup JCL shown in [Figure 5 on page 21](#), the TAOPT data set is of special importance.

Note: The IBM Z Table Accelerator Program Call Server must be initialized before starting a VTS-TSR. See [“Initialization of the Program Call Server” on page 13](#).

Initialization of a VTS-TSR

To process commands, check for the following message that indicates that the VTS is ready:

```
DKJ00600I IZTA Vvrm VTS vtsname initialized
```

Once this message is received, you can use DKJTDRV to execute the necessary commands. [Figure 6 on page 23](#) shows an example of the JCL.

TAOPT parameters

The TAOPT DDNAME is used to set execution-time parameters that override the defaults defined in the VTS implementation process. For detailed descriptions of these parameters, see [Appendix A, “IBM Z Table Accelerator run-time options,” on page 35](#).

The parameters that you may wish to change are:

1. VTSNAME=xxxxxxx where xxxxxxx is the one to eight character name used to identify the VTS-TSR created by this VTS Agent. VTS users use this name to make a logical connection to the VTS Agent. It must be unique for each VTS-TSR created within a single MVS image. The default value is DKJA.
2. TSIZE=nnnnnn. nnnnnn specifies the amount of virtual storage (Data Space size) to be allocated for this VTS-TSR. It needs to be large enough to contain all open tables and control information for this VTS-TSR. The delivered default is 10M. For format details, see [“TSIZE—tableSPACE Region Size” on page 43](#).
3. MAXNTAB=nnnnn. This number allows you to limit the number of tables opened in this VTS-TSR in order to reduce the overhead of the open table index constructed in the TSR. Attempts to open more tables will receive an error. The default is calculated to be the maximum number of tables that could possibly be opened in the TSR.
4. STROBE=nnnnnnnnn which specifies the interval after which to collect VTS-TSR usage data. This information will be collected and reported if STROBE is set to a non-zero value. If the value is set too low this may result in a large amount of strobe output which may exceed the installation maximum for job output. The default is 0, which results in no usage data being collected. This option is used in conjunction with STROBEMETHOD (see [STROBEMETHOD](#)).
5. STROBEMETHOD=n. This option is used in conjunction with STROBE (see [STROBE](#)). If n is 0, 1 or 2, STROBE specifies the call count after which to collect VTS-TSR usage data. If n is 3, STROBE specifies the number of minutes after which to collect VTS-TSR usage data. The default is 0.
6. SWITCHES=xxxxxxx which are one-byte codes that are employed to alter the operation of the IBM Z Table Accelerator interface used by your programs. For more detail, see the *IBM Z Table Accelerator Programming Guide* section on "VTS User access" under "Environmental interfaces".

VTS-TSR table loading sample JCL

Since tables can be opened into the VTS-TSR by any region that accesses it, there may be no need for preloading of tables into the VTS. However, if the VTS-TSR must be preloaded with specific tables, a separate job executing DKJTDRV will need to be run. The job executing DKJTDRV can be cued off the message that is returned indicating that the VTS-TSR is now available:

```
DKJ00600I IZTA Vvrm VTS vtsname initialized
```

Sample JCL is shown in [Figure 6 on page 23](#).

```

//VTSLOAD EXEC PGM=DKJTDRV
//STEPLIB DD DISP=SHR,DSN=*YOUR.PREFIX*.IZTA.LOAD <=====
//MAINLIB DD DISP=SHR,DSN=*YOUR.PREFIX*.MAINLIB <=====
//SYSUDUMP DD SYSOUT=*
//TADUMP DD SYSOUT=*
//SYSOUT DD SYSOUT=*
//CMD DD *
SS,DKJA
OR,ABCD* Open for read all tables starting with letters ABCD
OA,TBL* Open any tables
/*

```

Figure 6. Sample JCL for loading tables

VTS-TSR refresh process

Since tables can be updated, closed, and reopened in the VTS-TSR by any region that accesses it, there may be no need for a separate refresh process. However, if the VTS-TSR must be refreshed external to the applications which access it, DKJTDRV can be used to implement this functionality.

The VTS-TSR refresh process is the execution of a series of commands which may or may not include the RF command (refresh table). For more information on the RF command, see “VTS commands” on page 24 and the chapter on “IBM Z Table Accelerator commands” in the *IBM Z Table Accelerator Programming Guide*.

Figure 7 on page 23 shows example JCL for a job that refreshes a VTS-TSR—in this case, closing some tables and refreshing others. Note that all other tables in the VTS-TSR are not affected.

```

//VTSRFRSH EXEC PGM=DKJTDRV
//STEPLIB DD DISP=SHR,DSN=*YOUR.PREFIX*.IZTA.LOAD <=====
//MAINLIB DD DISP=SHR,DSN=*YOUR.PREFIX*.MAINLIB <=====
//SYSUDUMP DD SYSOUT=*
//TADUMP DD SYSOUT=*
//SYSOUT DD SYSOUT=*
//CMD DD *
SS,DKJA
CL,XX*
RF,TBL* Refresh any tables starting with letters TBL
/*

```

Figure 7. Sample JCL for VTS refresh

VTS Agent shutdown

The VTS Agent may be shutdown normally by issuing, from an MVS console, a stop command to the VTS Agent job/STC. The form of the command is:

```
P jobname
```

Any user attempting to access a terminating VTS Agent receives error code 1072.

VTS Agent recovery

If the VTS Agent is terminated or abends during execution for any reason, the VTS Agent job/STC may be resubmitted to MVS. However, the VTS-TSR will be reinitialized and must be re-populated. If the VTS-TSR had been updated by users, these updates will have been lost. While the VTS Agent is reinitializing,

applications will not be able to access tables in the VTS-TSR, and will receive error code 1072 indicating that the VTS-TSR is unavailable.

VTS commands

A VTS-TSR can be accessed using the same commands, and operates in a similar manner, as the local TSR. The chapter entitled "IBM Z Table Accelerator commands" in the *IBM Z Table Accelerator Programming Guide* describes these commands in detail.

The chapter entitled "DKJTDRV command processor for Batch/TSO" in the *IBM Z Table Accelerator Programming Guide* describes the commands implemented in the DKJTDRV utility. In addition, there are several DKJTDRV commands that are unique to the VTS Agent. They may be accessed by anyone using DKJTDRV. As the IBM Z Table Accelerator Administrator, it is your responsibility to control access to these commands, if desired, and, if so, to take measures to restrict their use.

Some commands for DKJTDRV accept the asterisk (*) as a wildcard in the table name to allow operating on a set of tables. Please refer to the chapter entitled "DKJTDRV command processor for Batch/TSO" in the *IBM Z Table Accelerator Programming Guide* for a complete list of these commands.

Note: To minimize the time during which access to the table is delayed, the affected Data Table and any alternates are reloaded while accesses are continuing with the original open copy in the VTS-TSR. Only after the refreshed copy and any alternates are completely loaded will internal pointers be switched to the refreshed copy. Access is delayed for the instant the internal pointers are switched.

VTS operational considerations

This section provides VTS operational considerations.

Reducing virtual storage when moving tables into VTS

Significant reduction of virtual storage utilization is possible when accessing tables from a VTS-TSR. If your installation uses multiple IMS regions or multiple CICS regions that each require access to the same tables, without VTS the tables are replicated in the TSR local to each region. Since each IMS or CICS region can access tables from a common VTS-TSR, the storage required for local TSRs can be greatly reduced. If a local TSR is not required after moving all the tables into a VTS-TSR, the local TSR size can be reduced to the absolute minimum by setting TSRSIZE to 0. See "[TSRSIZE—tableSPACE Region Size](#)" on page 43.

7/24 VTS

If the VTS-TSR is to remain operational for an extended period of time (for example, seven days a week, 24 hours per day) and IBM Z Table Accelerator usage reporting (see Chapter 3, "[tableSPACE Reporting](#)," on page 7) is enabled for the VTS Agent (STROBE is not = 0), ensure that the output file is large enough to capture all the strobe activity that you wish to see. If the file is not large enough, any strobe overflow records will be suppressed to ensure that the VTS can remain running. When this occurs, the following error messages will be displayed in the JESMSG LG of the VTS Agent job:

```
DKJ00280E QSAM ERROR: Function=PUT QSAM,File=TATSRPT
DKJ00281E Strobe output suppressed
```

Capacity

Each VTS-TSR is a shared Data Space that may be defined to be as large as 2 GB. If additional capacity is needed, an application may refer to any number of VTS-TSRs consecutively, by changing the VTS name in the TA-PARM.

Loading a VTS-TSR

Library caching can reduce the time to load a VTS-TSR.

Library caching may be valuable when:

- the VTS-TSR is large
- the VTS-TSR is pre-loaded with a large number of relatively static tables (not frequently updated in the VTS-TSR)
- no other regions are updating tables in the same library during the VTS load operation (this could require a cache refresh, possibly increasing load time)

For more information about library caching, see [“Library directory caching” on page 30](#).

VTS messages

In a VTS environment, specific VTS messages are written to the MVS console and to the JESMSG LG for the VTS Agent. These messages are in addition to the IBM Z Table Accelerator messages displayed when processing commands. For a complete list of IBM Z Table Accelerator and VTS-specific messages, see [Appendix B, “IBM Z Table Accelerator messages,” on page 45](#).

Messages ending with an I are informational, and messages ending with an E or S will not allow the VTS Agent to continue.

Internal logic errors

The following message is displayed in the event of an internal logic error in the VTS Agent:

```
'DKJ00301S LOGIC ERROR; DUMP TAKEN TO TADUMP'
```

Programs that are accessing the VTS agent will not abend. The programs will continue to run in their environment. The system will return an error code of 1072 on every attempt to access the VTS-TSR.

VTS user messages

IBM Z Table Accelerator/VTS user error messages are the same as those issued by IBM Z Table Accelerator when accessing a local TSR. User error codes are also the same, with the addition of error code 1072 (VTS access failed because the specified VTS is not available). Error codes are set in the ERROR and ERROR-SUBCODE fields of the IBM Z Table Accelerator COMMAND-AREA parameter.

Chapter 8. Best Practices

This chapter contains some information on optimization and organization that should be of benefit to IBM Z Table Accelerator administrators.

Protecting IBM Z Table Accelerator libraries

RACF (or equivalent) protection can be placed on IBM Z Table Accelerator libraries. If an unauthorized user attempts to open a protected library, the open attempt will fail with appropriate messages sent to the job's JES log. An error code is returned to the user, which may result in an abend, depending on the abend status switch setting. In some environments, a RACF error will result in an abend of the task.

Concurrent multiple updates

When a table is identified for updating in a region with the use of the Open for Write command (OW), steps are taken to prevent other MVS tasks on the same or other MVS images from updating the same table. This is achieved by issuing an enqueue request. The enqueue minor name includes the Table Name and Table Library Dataset Name. The enqueue major name is DKJIZTA. The scope of the enqueue is SYSTEMS.

An installation's modifications to Global Resource Serialization (GRS) PARMLIB controls must not prevent the propagation of the major enqueue name throughout the complex. Failure to adhere to this can result in the corruption of the IBM Z Table Accelerator libraries on DASD shared between several MVS images running IBM Z Table Accelerator.

Since these requirements are beyond the scope of IBM Z Table Accelerator software, you must ensure that these conditions are met if the same IBM Z Table Accelerator libraries are to be shared between MVS tasks.

Read-Only libraries

IBM Z Table Accelerator supports customers who wish to prevent updates to a IBM Z Table Accelerator library. In batch and IMS, customers can use LABEL=(,,IN) on the DD statement. In CICS, customers can use the RDO statements READ(YES), ADD(NO), and/or UPDATE(NO). If an application attempts to update a read-only library, IBM Z Table Accelerator will return error code 61 subcode 11 "Operation to write to a Library was terminated. Library read-only" or the calling task abends, depending on the abend switch setting.

Protecting IBM Z Table Accelerator tables

IBM Z Table Accelerator tables can be protected on the library (read and write passwords) and when loaded in-memory (LOCK-LATCH). It is important to note that these passwords offer only limited protection against unauthorized use of IBM Z Table Accelerator tables. More formal security can be implemented via security tools such as Resource Access Control Facility (RACF), eTrust CA-ACF2 Security for z/OS and eTrust CA-Top Secret Security for z/OS.

In the event that a table password is lost or forgotten, please see [Chapter 2, "Master Password Feature,"](#) on page 5.

For more information, please see the *IBM Z Table Accelerator Programming Guide*.

Table read and write passwords

Passwords are used to open tables from the library. However once in memory, these tables can be accessed by any application with access to the memory. A read password protects a table from being opened for either read or write. A write password protects a table from being opened for write. For more information on passwords see the *IBM Z Table Accelerator Programming Guide*.

TSR LOCK-LATCH

LOCK-LATCHES are used by IBM Z Table Accelerator to protect a table in memory. A LOCK-LATCH is used when the table is opened for write and is needed for any future update or store access once it is set. LOCK-LATCHES are not checked during fetch access. For more information on LOCK-LATCHES see the *IBM Z Table Accelerator Programming Guide* and the *IBM Z Table Accelerator Concepts and Facilities Guide*.

Time out while waiting for a lock

To prevent processes from waiting indefinitely for a lock, use the TAOPT/DKJOPTGN LOCKTIMER and LOCKTIMEWTO parameters. See [Appendix A, “IBM Z Table Accelerator run-time options,”](#) on page 35.

Performance

Shared tables: Virtual Table Share (VTS)

Significant reduction of virtual storage utilization is possible when accessing tables from a VTS-TSR. If your installation uses multiple IMS regions or multiple CICS regions that each require access to the same tables, without VTS the tables are replicated in the TSR local to each region. Since each IMS or CICS region can access tables from a common VTS-TSR, the storage required for local TSRs can be greatly reduced or eliminated altogether. If a local TSR is not required after moving all the tables into a VTS-TSR, the local TSR size can be reduced to the absolute minimum by setting the TSRSIZE=0. For more information, see the TSRSIZE parameter description in Appendix A of the *IBM Z Table Accelerator Implementation Guide*.

If multiple users could be updating a table—whether online, in a VTS-TSR or other multitasking environment—be aware that a LOCK-LATCH is available to restrict updates to authorized users/applications. It is not required.

In a read/write VTS environment, it is important to ensure all IBM Z Table Accelerator jobs operate at the same WLM priority level. This is to insure that a low priority job does not lock out a high priority job. This can occur when the low priority job is updating a IBM Z Table Accelerator table and is interrupted by MVS.

Linking IBM Z Table Accelerator

In a CICS environment, it is recommended that DKJTCALL be statically linked, and not defined to CICS. By statically linking DKJTCALL, the PPT is not searched for DKJTCALL and the call is very efficient. If it is called dynamically, it must be defined to CICS. In batch and other non-CICS environments, there may be performance impacts if IBM Z Table Accelerator is not statically linked.

WLM priority level consideration

In a read/write VTS environment, it is important to ensure all IBM Z Table Accelerator jobs operate at the same WLM priority level. This is to insure that a low priority job does not lock out a high priority job. This can occur when the low priority job is updating a IBM Z Table Accelerator table and is interrupted by MVS.

Single TAPARM per thread

A single TAPARM should be used per thread (for batch, this means one TAPARM per TCB; for IMS, one per transaction). To be effective, it must be initialized at the first call to IBM Z Table Accelerator and never reinitialized.

When IBM Z Table Accelerator calls are issued, our fast path "handle" is stored in the TAPARM area. If the TAPARM is reinitialized in each thread, IBM Z Table Accelerator attempts to achieve similar performance by updating both DKJTCALL and DKJTROT at initialization of IBM Z Table Accelerator in the region. This will not be possible when these modules are in key 0 storage.

Re-entrant modules will be loaded into key 0 storage when they are loaded from LPA or from an authorized library.

Table access optimization

Table access can be optimized by using a separate command area for each table. Instead of having to search a TSR directory each time to determine where the table resides in memory, IBM Z Table Accelerator uses a reserved field in the command area (table handle) pointing to the position of the table in the TSR. By keeping a separate command area for each table accessed, this table handle is used to quickly locate a previously opened table.

Multitasking

IBM Z Table Accelerator can be tailored to perform optimally in your environment by setting execution-time parameters and switches for both batch and online environments. In a batch environment only, the Multitasking parameter can be set to Y to allow multiple subtasks within a region to access IBM Z Table Accelerator concurrently. The normal batch default is N.

Note: For a full list of execution-time parameters see the *IBM Z Table Accelerator Implementation Guide*.

IBM Z Table Accelerator library definition considerations

You can increase library blocksize to gain efficiencies in DASD utilization and reduce I/O access, in moving directory and table data between DASD and memory. You can also reduce DASD space requirements if the total data for each table in the library is significantly more than the current blocksize.

First, consider space utilization from the consideration of physical blocksize on 3390-style DASD (the current standard).

Library Blocksize

Each library can have an optimal blocksize selected at the time it is created.

Library blocksizes are user-selectable. You can optimize the blocksize for your tables from 800 bytes up to the size of a track (or 32,760 bytes, whichever is smaller). This flexibility allows you to decrease the time to load tables from the library to a TSR, since fewer I/Os are required.

IBM Z Table Accelerator library space usage

IBM Z Table Accelerator attempts to use library space efficiently. If the BLKSIZE (RECLEN) is larger, it fits more directory entries into each block. But every table definition requires 1 full block for the definition.

The actual data for each table is stored in an integral number of full blocks, with the minimum being one. A very small table (e.g., 10 rows with a row length of 12 bytes) still requires a full block, which results in a very low utilization, even if a blocksize of 800 (the minimum allowed) is used. On the other hand, a very large table (e.g., 50000 rows with a row length of 200 bytes) will have over 99% utilization of its allocated space no matter what the blocksize.

BDAM/QSAM IBM Z Table Accelerator libraries

For BDAM/QSAM libraries, the default block size of 3120 results in 15 blocks per track and 82.6% utilization of the space. Increasing block size to 3174 results in 84% utilization (still at 15 blocks per track). Increasing it to 27998 (2 blocks per track) results in 98.8% utilization. Reducing it to 820 (39 blocks per track) results in 56.4% utilization.

BDAM Fixed (F) libraries are supported; Fixed Block (FB) BDAM libraries are not supported.

VSAM IBM Z Table Accelerator libraries

For VSAM libraries, the default blocksize of 3120 (RECSZ in IDCAMS DEFINE) results in a CISZ of 3584, 13 blocks per track and 71.6% utilization. Increasing the RECSZ to 3577 (the maximum in a 3584 CISZ) results in a 82.1% utilization. Increasing it to 4091 (the maximum in a 4096 CISZ) results in 86.6% utilization. Reducing it to 2041 (2048 CISZ) results in 75.6% utilization; reducing it to 1017 (1024 CISZ) yields 59.2% utilization.

VSAM IBM Z Table Accelerator libraries and LSR pools

If a IBM Z Table Accelerator VSAM library is allocated with DISP=OLD there are no restrictions on SHAREOPTIONS, buffering, strings and buffer pools.

However if it is allocated as DISP=SHR, IBM Z Table Accelerator uses MVS Enqueues (SCOPE = SYSTEMS) and I/O techniques to ensure the library's integrity is maintained. Our standard recommendation has been to use SHAREOPTIONS(3,3) in all environments to prevent overlays of the IBM Z Table Accelerator library directory and table data.

In CICS (and optionally batch), standard performance recommendations have been to add VSAM datasets to buffer pools (LSRPOOLS). However this would undermine IBM Z Table Accelerator library integrity, so for libraries that it accesses, IBM Z Table Accelerator will reset any LSRPOOLID value to zero, effectively making them NSR (Non-Shared Resources). Any software that changes buffer pool settings dynamically for VSAM datasets will create integrity problems for IBM Z Table Accelerator libraries.

IBM Z Table Accelerator cannot maintain library integrity with SHAREOPTIONS(3,3) and any VSAM buffering.

However, with CICS/TS 4.1, IBM is enforcing the restriction that CICS Transaction Isolation is not supported for NSR VSAM files. An AFDK abend will be issued when the library is opened. To accommodate this restriction (which does not apply to BDAM IBM Z Table Accelerator libraries or libraries allocated with DISP=OLD), IBM Z Table Accelerator will ensure library integrity with SHAREOPTIONS(4,4) and LSRPOOLS for the CICS environment.

We recommend that the buffer pool used with SHAREOPTIONS(4,4) be set to a minimum size and not be shared with other users. The buffers will not benefit IBM Z Table Accelerator and could interfere with other applications sharing the buffer pool. The library directory caching feature should provide equivalent or better results than LSR pooling.

Enhanced Data Integrity option

IBM Z Table Accelerator Version 1 Release 1 is compatible with z/OS. z/OS runs only on 64-bit processors. But it fully supports 31-bit applications, which includes customers' CICS and IMS applications (and IBM Z Table Accelerator). IBM is also committed to continuing to support interaction between 31-bit and 64-bit applications in the future.

z/OS provides an option called "Enhanced Data Integrity" that can be enabled in SYS1.PARMLIB. This option is designed to protect physical sequential (DSORG=PS) datasets allocated with DISP=SHR from being concurrently opened for update by multiple users.

This can affect IBM Z Table Accelerator users, since IBM Z Table Accelerator libraries can be allocated with DSORG=PS (even though we only document BDAM (DSORG=DA) and VSAM). IBM Z Table Accelerator internally protects IBM Z Table Accelerator libraries when allocated with DISP=SHR, so "Enhanced Data Integrity" should not be enabled for IBM Z Table Accelerator libraries.

The system programmer can bypass this option for specific datasets.

Ensure that IBM Z Table Accelerator libraries allocated with DSORG=PS are listed in the exclude list in the IFGPSEDI member.

Library directory caching

The IBM Z Table Accelerator library directory can be cached in memory. This can dramatically reduce the time needed to open a large number of tables when a TSR is initially loaded.

Note: In an in-house test, load time for 40,000 tables from a IBM Z Table Accelerator library into a local TSR was reduced from 11 hours to 7 minutes. This was an extreme stress test; your results may vary.

Caching for a library is controlled by coding OPTCD=C on the DD statement for a library. This feature is not supported in a CICS region. The initial loading of a TSR may benefit from the use of the library directory caching feature, and subsequent refreshes may also benefit if more than one table from the library is being refreshed. However, if updates are being done to the library while the loading or refreshing is being

performed, these benefits will be lost. It is also important to note that the library directory caching feature is enabled only at region initialization.

Note: Caching has significant benefit when a library is accessed heavily by a region and other regions are not frequently updating the library. However, caching can have negative performance consequences when multiple regions are frequently updating the library directory.

This feature should provide equivalent or better results than LSR pooling. There is limited support for VSAM IBM Z Table Accelerator libraries using LSR pools. For more information on using LSR pooling, refer to [“VSAM IBM Z Table Accelerator libraries and LSR pools”](#) on page 30.

TSR Space Allocation

IBM Z Table Accelerator allocates space for the TSR in 4K chunks and will only allocate space as needed as tables are opened or defined in the TSR. The minimum amount of space that can be allocated for a TSR is 28K and the maximum is 2G.

IBM Z Table Accelerator initializes the first portion of the TSR to contain control information. For smaller TSR sizes of about 1M to 3M, about 20K of the space is used for control information. For larger TSR sizes, between 1 to 2 percent of the TSR space is used for control information.

When IBM Z Table Accelerator initializes, it does a DSPSERV request for a private dataspace of the TSR size requested. For non-VTS regions, the dataspace is allocated as SCOPE=SINGLE. For VTS regions, the dataspace is allocated as SCOPE=ALL. The MVS IEFUSI exit at your site may limit the number and/or sizes of data spaces allocated by a single job step.

If your region defines the TSR size as 1M, this results in the building of the MVS system control blocks to support a virtual storage of 1M. However, if you only open tables to fill up 500K of the TSR, IBM Z Table Accelerator will only allocate about 520K of space for the TSR, 20K being the overhead for the control information.

If your region defines the TSR size as 100M, this results in the building of MVS system control blocks to support a virtual storage of 100M. However, if you only open or define tables that take up 20M of space, IBM Z Table Accelerator will only allocate about 22M of space for the TSR, 2M being the overhead for the control information.

Taking into account the small overhead for the space required to hold control information on the TSR, if you anticipate a large variance in the amount of space that you will require, we recommend that you define your TSR size to be the upper limit of all space that you expect to use during the life of your region.

IBM Z Table Accelerator modules that can reside in the Link Pack Area (LPA)

The following IBM Z Table Accelerator modules are eligible and recommended to be in LPA for the Batch and IMS interfaces:

DKJTCALL, DKJTNAME, DKJTNUCL, DKJTROT B

Note: When running CICS applications, IBM Z Table Accelerator modules cannot reside in LPA.

Chapter 9. IBM Z Table Accelerator Diagnostic Information

This chapter identifies and describes the IBM Z Table Accelerator diagnostics tools for the IBM Z Table Accelerator administrator.

IBM Z Table Accelerator use of MVS Enqueues

In order to protect the integrity of tables stored in IBM Z Table Accelerator libraries, IBM Z Table Accelerator issues two MVS enqueues, a library enqueue and a table enqueue. The IBM Z Table Accelerator software enqueues are issued with a scope of "SYSTEMS".

The enqueues have a major resource name of "DKJIZTA". For a library enqueue, the minor name consists of the letter "L" followed by the dataset name. For a table enqueue, the minor name consists of the letter "T" followed by the table name and the dataset name. These enqueues are independent of "SYSDSN" enqueues issued by the MVS allocation process, based on DISP=SHR or DISP=OLD.

The IBM Z Table Accelerator library enqueue (minor name "L...") protects the library directory and space allocation. It is obtained as a shared enqueue whenever the library has to be read. It is obtained as an exclusive enqueue when the library is being updated (for example, during a store (ST) of a table). The enqueue is held only during the operation of a single IBM Z Table Accelerator command; it is always released before control is returned to the caller. The scope of "SYSTEMS" protects the library from other IBM Z Table Accelerator accesses not only on the local system but also across any other systems in the GRS group.

The IBM Z Table Accelerator table enqueue (minor name "T...") prevents the library copy of a table from being Opened for Write (OW) in more than one TSR. This enqueue is held across multiple IBM Z Table Accelerator calls; it is released only when the table is no longer Open for Write in the TSR (typically released by a CClose or ReRelease command for the table). The scope of "SYSTEMS" protects the table from any other attempts to update the library copy of this table not only on the local system, but also across any other systems in the GRS group. Besides OW, commands CG, DG, RN, XT also obtain this enqueue.

IBM Z Table Accelerator Library Diagnostics

In order to assist in the diagnosis of IBM Z Table Accelerator Library integrity, two command update segments are written to Block 0 of a IBM Z Table Accelerator library for each update made to the library. The first segment is written at the start of the update command and the second segment is written on successful completion of the update command. If a IBM Z Table Accelerator command abends during the update of a library, the first segment will show the current command being processed and the second segment will contain information from the completion of the previous IBM Z Table Accelerator library update.

Block 0 of the library can be viewed by running the job DKJTPRTM which dumps library contents. Technical Support may request you to dump your library contents in order to assist with the diagnosis of the health of your library.

Appendix A. IBM Z Table Accelerator run-time options

This appendix lists the IBM Z Table Accelerator run-time option parameters. Most occur in all interfaces. If you are running applications across multiple environments any changes to these options in one interface may need to be made in each installed interface. Options can be customized for your site at implementation time, using the DKJOPTGN macro in DKJTxx34 modules. Applications can further override options on an individual basis, using the TAOPT dataset.

CICSJRNLCICS Journal File ID

This parameter identifies the CICS Journal File onto which the TABLE SPACE Report records are to be written by the system, if strobe records are to be written. The delivered default value for the CICSJRNLC parameter is 99.

This option can only be specified for the IBM Z Table Accelerator CICS interface.

This option applies to this region's users, regardless of which TSR they access.

HASH_HI_DEN_LIMHigh Density Limit for Hash Indexes

HASH_LOW_DEN_LIM and HASH_HI_DEN_LIM limit the density of the index for a hash table. These values are designed to prevent performance problems which can occur when inappropriately high values are used when defining hash tables (can result in a lot of key collisions in the table). Other problems occur if the difference between low and high density values is too small. A ratio of 2/3 is now enforced: Low density may not be greater than 2/3 of high density.

HASH_HI_DEN_LIM=nnn must be between 100 and 900 (10% and 90%); the default is 900. We recommend the value be lowered to 500.

This option applies to this region's users, regardless of which TSR they access.

This option can be specified for all IBM Z Table Accelerator interfaces except the VTS interface.

HASH_LOW_DEN_LIMLow Density Limit for Hash Indexes

See above (“HASH_HI_DEN_LIMHigh Density Limit for Hash Indexes” on page 35).

HASH_LOW_DEN_LIM=nnn must be between 10 and 600; the default is 600. We recommend the value be lowered to between 200 and 300.

This option applies to this region's users, regardless of which TSR they access.

This option can be specified for all IBM Z Table Accelerator interfaces except the VTS interface.

LIBnn, MLIBM Z Table Accelerator Library List

Specify in sequence, the names of the libraries (DDNAMEs) in the IBM Z Table Accelerator Library List (LIB-LIST). In DKJOPTGN, this is specified as a comma delimited list under ML. In the TAOPT file, the library names are specified as LIBnn=lib_name. Up to 10 LIBnn may be specified, as long as the numbers are contiguous. For example, LIB01=MAINLIB, LIB02=TESTLIB.

LIBnn numbers must be contiguous AND they must start with LIB01 or LIB02, otherwise IBM Z Table Accelerator initialization will fail.

The delivered default is ML=MAINLIB which is equivalent to supplying a single parameter LIB01=MAINLIB. If no ML is set, IBM Z Table Accelerator uses MAINLIB.

This option applies to this region's users, regardless of which TSR they access.

This option can be specified for all IBM Z Table Accelerator interfaces.

LISTOPTIONS—List Parameter Options

Determines whether or not to list execution time parameters. If LISTOPTIONS=Y is specified, all default parameters and parameters overridden by the TAOPT dataset are listed in the JESMSG LG. LISTOPTIONS=N suppresses this list. The delivered default is LISTOPTIONS=N.

LISTOPTIONS=X is a special setting for DKJOPTGN (DKJT1134). It is the equivalent of LISTOPTIONS=N if the TAOPT DD was not present in the jobstream, and the equivalent of LISTOPTIONS=Y if the TAOPT DD is present. LISTOPTIONS=X applies to DKJOPTGN (DKJT1134) only; it does not apply to TAOPT.

This option applies to this region's users, regardless of which TSR they access.

This option can be specified for all IBM Z Table Accelerator interfaces.

LOCKTIMERC—Lock Timer Wait Value

LOCKTIMERC=nnnnnn specifies the number of seconds (default 0) that IBM Z Table Accelerator should wait for a lock. When the LOCKTIMERC interval has passed, RC=71 is returned or the call abends depending on the status switch setting. A value of LOCKTIMERC=0 specifies that the process will never time out.

The lock controlled by LOCKTIMERC is used internally by IBM Z Table Accelerator to maintain table integrity in the TSR. It is unrelated to the table ENQUEUE that occurs when a table is opened for write (OW).

This option applies to this region's users, regardless of which TSR they access.

This option can be specified for all IBM Z Table Accelerator interfaces except the VTS interface.

LOCKTIMEWTO—Lock Timer Message Wait Value

LOCKTIMEWTO=nnnnnn specifies the number of seconds (default 30) to wait before issuing messages that the process is waiting for a lock. A value of LOCKTIMEWTO=0 specifies that no warning message will be issued.

The lock controlled by LOCKTIMEWTO is used internally by IBM Z Table Accelerator to maintain table integrity in the TSR. It is unrelated to the table ENQUEUE that occurs when a table is opened for write (OW).

When applicable, messages are generated to report that some processing has been blocked.

This option applies to this region's users, regardless of which TSR they access.

This option can be specified for all IBM Z Table Accelerator interfaces except the VTS interface.

MAXNMTAB—Maximum Number of Tables

Determines the number of tables that can be opened in a given local TSR or VTS-TSR. (With the possibility of creating a TSR of up to 2G in size, there can be a great variation in how many tables will actually be in a TSR at any time.) If the value of the MAXNMTAB is not set by the user (or is set to zero) then a default value is calculated.

If an attempt is made to open more tables than is indicated by MAXNMTAB, an error code of 20 will be given: The maximum number of tables has been exceeded. Check MAXNMTAB.

This option applies to the TSR created by this region/job and all users accessing it.

This option can be specified for all IBM Z Table Accelerator interfaces.

User sets the value of MAXNMTAB

To optimize the memory required for system overhead, the user may choose to set the value of MAXNMTAB rather than use the default value. MAXNMTAB can be set to any integer, zero or greater. However, if MAXNMTAB is set to a value greater than the default for the given TSR size, the default is used and the user is notified with a warning message.

After the user sets the value of MAXNMTAB, memory is allocated and an index is created that contains enough entries to support the maximum number of tables indicated. Additional memory is allocated as tables are opened. This memory is reused as tables are closed and new tables opened. This approach keeps the system overhead in line with the number of tables that are opened.

The calculation of the default value of MAXNMTAB

If the user does not set the value MAXNMTAB or the value is zero, the IBM Z Table Accelerator software calculates the default value based on the size of the TSR. Each table in the TSR occupies a minimum of 4K pages. The maximum possible number of tables is 51,455, for a 2G TSR; the default, when unspecified is one table for every 4K page in the TSR.

MTRETAIN—Retain Rows and Index Areas

MTRETAIN=Y|N determines whether the allocated rows and index areas are to be retained when an MT command is issued or whether they are to be reduced to the original allocation before table expansion. The default is "N", which is current processing—not retained.

Also see [“ZEROROWS—Zero Data Table Rows on De-allocation”](#) on page 43.

This option applies to the TSR created by this region/job and all users accessing it.

This option can be specified for all IBM Z Table Accelerator interfaces.

MULTITASKING—Multitasking

Specifies whether this region will allow multiple subtasks to access IBM Z Table Accelerator.

This option can only be specified for the IBM Z Table Accelerator batch interface. All other environments are multitasking in their nature, and will not accept this option.

The MULTITASKING option can be set to Y to allow multiple subtasks within a region to access IBM Z Table Accelerator concurrently. Each subtask may have multiple concurrent subtasks of its own, with no limit to the level of multitasking. The delivered default is N (multitasking not enabled).

This option applies to this region's users, regardless of which TSR they access.

Note:

1. If a batch application attempts to issue a IBM Z Table Accelerator command from a second TCB with Multitasking=N, it will abend with an SOC3 and display message: "DKJ00496E Multitasking requires TAOPT Multitasking=Y be set".
2. Please call Technical Support for assistance when writing applications that access IBM Z Table Accelerator in a multitasking environment.

OVERRIDES—Allow Changes to Status Switches

It is recommended that you do not allow this option to be changed using TAOPT. It should only be changed using DKJOPTGN.

The Status Switches may be changed by the application with the use of the Change Status (CS) command. Overriding individual Status Switches may be inhibited by the settings of the override controls. Overrides are specified as a string of Y's and N's; for example: YYYNYNNN. A value in any of the eight positions represents a change for that override. A * or blank in any of the eight positions represents no change for that override.

The defaults for these values will depend upon the interface. See the Implementation Guide for further details.

<i>Table 4. Overrides default values</i>			
Position	Description	Override setting	Meaning
1	Allow Change To Abend On Errors	Y	The application is allowed to set the ABEND status switch.
		N	The application is not allowed to set the ABEND status switch.
2	Allow Change To Wait For Enqueued Tables	Y	The application is allowed to set the WAIT status switch.
		N	The application is not allowed to set the WAIT status switch.
3	Allow Change To Return Empty Rows In Hash Tables	Y	The application is allowed to set the HASH-EMPTYES-RETURNED status switch.
		N	The application is not allowed to set the HASH-EMPTYES-RETURNED status switch.
4	Allow Change To Permit Implicit Open Of Tables	Y	The application is allowed to set the IMPLICIT OPEN status switch.
		N	The application is not allowed to set the IMPLICIT OPEN status switch.
5	Allow Change To Trace IBM Z Table Accelerator Commands	Y	The application is allowed to set the IBM Z Table Accelerator TRACE status switch.
		N	The application is not allowed to set the IBM Z Table Accelerator TRACE status switch.
6	Reserved for future use; default set to Y.		
7	Reserved for future use; default set to Y.		
8	Reserved for future use; default set to Y.		

The values set by this option apply to this region's users, regardless of which TSR they access.

This option can be specified for all IBM Z Table Accelerator interfaces.

RACF_LIBACCESS—Perform Checks on Protected Libraries

Specifies whether IBM Z Table Accelerator should check if a IBM Z Table Accelerator library is SAF interface protected. Examples of products which implement SAF interface protection are RACF, ACF2 and Top Secret.

This option is only available in the IBM Z Table Accelerator Batch and IMS interfaces.

The following is a description of each parameter value:

- Y indicates that IBM Z Table Accelerator will check whether a library is protected. Any update attempts without update access to the library will fail with error code 61-12.

- N indicates that IBM Z Table Accelerator will not check whether a library is protected. Any update attempts without update access to the library will fail with a S913-38 for IBM Z Table Accelerator BDAM libraries and a U301 for IBM Z Table Accelerator VSAM libraries.

The delivered default for this parameter is N.

STROBE—Strobe Interval

Specifies a numeric value, from 0 to 2,147,483,647(2G – 1), to control the interval after which IBM Z Table Accelerator strobe reporting is triggered. The delivered default is 0.

For the VTS environment, this option is used in conjunction with STROBEMETHOD. The numeric value specified either represents a call count (STROBEMETHOD=0, 1 or 2) or a minute interval (STROBEMETHOD=3). See “STROBEMETHOD—Strobe Method” on page 39 for more details.

For all other environments, the numeric value represents a call count.

This option can be specified for all IBM Z Table Accelerator interfaces.

For the batch, IMS and VTS interfaces, a final strobe report will always be produced if the TATSRPT DD statement is present (even if STROBE is set to 0). In order to suppress strobe reporting, the TATSRPT DD statement must be removed from the job. For STROBE=0, no intermediate strobe report is produced.

TATSRPT is not relevant in the CICS environment. CICS strobe statistics are written to a CICS journal file (DFHJ99) from which the strobe report is generated via a IBM Z Table Accelerator utility (DKJTSTRB). No final strobe is written to the CICS journal if STROBE is set to 0 or the STROBE is set at a value which causes no strobos to be written before step end. This means that for these cases, there will be no STROBE report generated.

For more information on CICS strobe reporting, please refer to the *IBM Z Table Accelerator Administration Guide*.

This option applies to the TSR created by this region/job and all users accessing it.

Note: The numeric value must be specified without commas.

STROBEMETHOD—Strobe Method

Specifies when IBM Z Table Accelerator strobe reporting is triggered for the VTS-TSR.

This option applies only to the VTS environment.

Values for this parameter can be 0 and 3. The delivered default is 0.

The following is a description of each parameter value:

- 0 indicates that the traditional strobe reporting method, by number of calls, using the STROBE=n parameter, will occur
- 3 indicates that strobe reporting will be triggered by time interval instead of call count. STROBE=n will indicate the number of minutes, n, after which a strobe is to be taken.

If STROBE=n is not specified with STROBEMETHOD, the default of STROBE=0 is used and only 1 report will be produced at IBM Z Table Accelerator termination.

STROBEMETHOD can be specified in the TAOPT DD statement or modified in DKJOPTGN source, reassembled and relinked into DKJVBASE for the VTS Agent job.

For STROBEMETHOD of 0, IBM Z Table Accelerator processing is prevented from waiting for the strobe report. If the number of calls in the STROBE parameter is set to a small value and IBM Z Table Accelerator calls are made more frequently than strobe reports can be produced, the frequency of reports may not match the STROBE value.

SUPPRESS_DUMPS—Suppress IBM Z Table Accelerator Messages and Dumps

Specifies whether to suppress IBM Z Table Accelerator messages and dumps when an abend occurs.

This option applies only to the CICS environment.

Values for this parameter can be NONE or SYMPATHETIC. The delivered default is SUPPRESS_DUMPS=NONE.

The following is a description of each parameter value:

- NONE indicates that all IBM Z Table Accelerator messages and dumps will be produced.
- SYMPATHETIC indicates IBM Z Table Accelerator error messages and dumps (TADUMP and IBM Z Table Accelerator CICS dumps) will be suppressed if a transaction abends outside IBM Z Table Accelerator code (i.e., in the associated application code, or in some other non-IBM Z Table Accelerator code). Any abends occurring within IBM Z Table Accelerator code will still produce the IBM Z Table Accelerator error messages and dumps.

SWITCHES—Status Switches

The Status Switches are a series of one-byte codes for altering the operation of the Application Programming Interface used by your programs when requesting a IBM Z Table Accelerator service. Switches are specified as a string of Y's and N's; for example: YYNYNNNN. A value in any of the eight positions represents a change for that override. A * or blank in any of the eight positions represents no change for that override. For information about switch settings, see [Table 5 on page 40](#).

The switches are stored in the parameter module for the interface. As each batch job or online task begins, the run-time values will be set according to the values in the parameter module but may be changed using the values in the TAOPT file. For more information, see Parameters in the MVS Batch section of the *IBM Z Table Accelerator Implementation Guide*.

The application also may change the values with the CS, the Change Status command (unless changing a particular switch is suppressed via the override controls).

The defaults for these values depend upon the interface and are described in the *IBM Z Table Accelerator Implementation Guide* in the appropriate chapter:

- Implementing IBM Z Table Accelerator z/OS Batch
- Implementing the CICS Interface
- Implementing the IMS TM Interface

Position	Default	Switch setting	Meaning
1	Abend on Errors	Y	Abend processing is to be performed on IBM Z Table Accelerator errors 0001-0099 or 1001-1099. Errors 100 to 999 and errors greater than 1099 always abend.
		N	Abend processing is not to be performed on IBM Z Table Accelerator errors 0001-0099 or 1001-1099. User programs could handle these return codes.

Table 5. Switches default values (continued)			
2	Wait for Enqueued Table	Y	IBM Z Table Accelerator is to wait for tables that are enqueued.
		N	IBM Z Table Accelerator is not to wait for such enqueued tables. In this case IBM Z Table Accelerator will terminate with a user 0072 abend if abend processing is enabled, or, will return an error code of 0072: Table unavailable; no wait in the command area if abend processing has been disabled.
		Note: Waiting in an online environment is normally discouraged. In CICS and IMS, the default is set to N	
3	Return Empty Rows in Hash Tables	Y	IBM Z Table Accelerator is to return empty rows for hash tables
		N	IBM Z Table Accelerator is to suppress the return of empty rows for hash tables
4	Allow Implicit Open Of Tables	Y	IBM Z Table Accelerator is to automatically open tables for read on first access
		N	IBM Z Table Accelerator is to suppress automatic opens; an explicit open command, OR, OW, or IA must be issued to open a table
		Note: We recommend setting Allow Implicit Open of Tables to N, especially in multi-user environments. This gives better control over which tables are opened.	
5	Trace IBM Z Table Accelerator Commands	Y	IBM Z Table Accelerator automatically records the last ten commands executed per thread. This is done for diagnostic purposes.
		N	IBM Z Table Accelerator does not trace commands
6	Reserved for future use; specify as N.		
7	Reserved for future use; specify as N.		
8	Reserved for future use; specify as N.		

The values set by this option apply to this region's users, regardless of which TSR they access.

This option can be specified for all except the VTS Interface.

TABLEWAITRC—Table Open Enqueue Wait Time

Specify a value to indicate the number of seconds that a user will wait to obtain the MVS enqueue to open a table for read or write before timeout. If the enqueue is not obtained before the timeout, IBM Z Table Accelerator will return code 72, or will abend, if the *Abend on Errors* switch is set to Y.

The delivered default is 0 (wait forever).

The Wait for Enqueued Table switch (see “[SWITCHES—Status Switches](#)” on page 40) must be set to Y for this parameter to have any effect.

This option applies to this region's users, regardless of which TSR they access.

This option can be specified for all except the VTS Interface.

TABLEWAITWTO—Table Open Enqueue Report Time

Specify a value to indicate the elapsed time before a message will be generated to report that the enqueue has not yet been received.

The delivered default is 30 (0=no messages).

The Wait for Enqueued Table switch (see “SWITCHES—Status Switches” on page 40) must be set to Y for this parameter to have any effect.

This option applies to this region's users, regardless of which TSR they access.

This option can be specified for all except the VTS Interface.

TSR_ALGORITHM—Optimize TSR Usage

TSR_Algorithm=P|D allows the TSR Space Manager to determine if TSR space usage should be maximized or performance should be optimized when TSR space is allocated and deallocated for table usage. The TSR space manager is the IBM Z Table Accelerator component responsible for tracking TSR space usage and for allocating and deallocating TSR space when tables are opened, closed, expanded or reduced in size.

The following is a description of each parameter value:

- P (Performance) indicates that the TSR will be optimized for performance. Space will be assigned to tables within the TSR so as to minimize CPU usage, which may result in a less than optimum use of space.
- D (Default) indicates that there will be no optimization, and there will be no messages regarding optimization. However, messages will be provided regarding current percentage of TSR capacity.

The delivered default is P.

This option applies to the TSR created by this region/job and all users accessing it.

This option can be specified for all IBM Z Table Accelerator Interfaces.

TSR_WARNING_FREQ—Frequency of TSR Allocation Warnings

Specify a value to indicate the frequency of reports generated when the TSR allocation percentage (as defined by TSR_WARNING_PCT) is exceeded. 0 results in a message for every allocation over the specified percentage; 30 results in a wait of at least 30 seconds before repeating the message; 999 results in a wait of just over 16 minutes before repeating the message.

The delivered default is one warning every second (001).

This option applies to this region's users, regardless of which TSR they access. (Different regions accessing a VTS-TSR can set their own thresholds.)

This option can be specified for all IBM Z Table Accelerator Interfaces.

TSR_WARNING_PCT—Percentage of TSR Allocation for Warning Activation

Specify a value to indicate the percentage of TSR allocation allowed before a message will be generated to report that the allocation percentage has been exceeded. The message will be repeated on every allocation of this percentage, subject to the TSR_WARNING_FREQ parameter.

The delivered default is 85 (85%).

This option applies to this region's users, regardless of which TSR they access. (Different regions accessing a VTS-TSR can set their own thresholds.)

This option can be specified for all IBM Z Table Accelerator Interfaces.

TSRSIZE—tableSPACE Region Size

The TSRSIZE parameter is an integer representing the amount of storage to be used for the TABLE SPACE REGION (TSR). IBM Z Table Accelerator uses Data Spaces for all TSRs, whether local or VTS. This ensures that tableSpace memory requirements do not affect other memory requirements in the region. The format for TSRSIZE is:

- nnnnnnnn = bytes
- nnnnnnnK = kilobytes
- nnnnM = megabytes
- nG = gigabytes

Value recommendations

The delivered default value for the TSRSIZE is 10M for all interfaces. The current minimum size of a TSR is 28 KB (but this may change in subsequent maintenance releases). The maximum size of a TSR is 2G (or 2048M or 2097152K). TSRSIZE only accepts up to eight digits, thus you cannot specify 2G as TSRSIZE=2147483648.

Note: If you specify a value of 0, the minimum size will be substituted; if you specify a value of between 1 and the minimum, an error code will be returned.

Other considerations

You must ensure you are allocating enough space for all tables open simultaneously in the region. You can determine the TSR size in one of two ways:

1. use the LT command in the batch program DKJTDRV, CICS transaction TADR,
2. browse strobe reports produced by your job runs.

Failure to set your TSR size parameter appropriately can produce the following error:

Error 92	There is insufficient tableSPACE region available. Increase TSR size.
----------	--

Note: There is a possibility that your system programmers have established limits on the maximum size or the maximum number of Data Spaces. All local TSRs are allocated as SCOPE=SINGLE Data Spaces; VTS-TSRs are allocated as SCOPE=ALL Data Spaces.

This option applies to the TSR created by this region/job and all users accessing it.

This option can be specified for all Interfaces.

VTSNAME—Specifying the Name of a VTS-TSR

The VTSNAME option specifies the name to be assigned to the VTS-TSR when it is initialized by the VTS Agent. The name is used by other regions to access the shared TSR.

Notes:

1. VTSNAME can be 1 to 8 characters. The naming restrictions are identical to those for DDNAMEs in IBM JCL.

This option applies to the TSR created by this region/job and all users accessing it.

ZEROROWS—Zero Data Table Rows on De-allocation

Related to the MTRETAIN parameter (see [“MTRETAIN—Retain Rows and Index Areas”](#) on page 37)—applies only when MTRETAIN = N.

ZEROROWS=Y|N determines whether the data rows area should be zeroed when it is deallocated. The default is Y, which is current processing. Note that the index area is never zeroed, except for hash indexes.

This option applies to the TSR created by this region/job and all users accessing it.

This option can be specified for all IBM Z Table Accelerator interfaces.

Appendix B. IBM Z Table Accelerator messages

This appendix contains all of the IBM Z Table Accelerator error codes and messages that can be encountered during the normal installation, administration and operation of the product. The error codes and messages fall into the following categories:

IBM Z Table Accelerator error codes—error conditions that may be encountered while running IBM Z Table Accelerator (see “IBM Z Table Accelerator return codes” on page 45).

IBM Z Table Accelerator messages—conditions that may be encountered with running IBM Z Table Accelerator in conjunction with the VTS agent, CICS, batch, IMS, and other programs/utilities (see “IBM Z Table Accelerator messages” on page 66).

DKJTEXEC error messages—conditions that may be encountered while running IBM Z Table Accelerator in conjunction with the DKJTEXEC batch utility program (see “DKJTEXEC error messages” on page 82).

IBM Z Table Accelerator return codes

The following IBM Z Table Accelerator error codes and subcodes are returned by IBM Z Table Accelerator as a result of events detected in the processing of commands. The error code is returned in the ERROR field of the COMMAND-AREA parameter, and the subcode is returned in the ERROR-SUB-CODE field.

To distinguish IBM Z Table Accelerator abend codes from those of the operating environment when running IBM Z Table Accelerator with CICS, the IBM Z Table Accelerator abend codes are prefixed with a G for abends under 1000 (for example, G009) and prefixed with V for abends over 1000, where V replaces the 1 (for example, 1072 becomes V072). In all other environments, IBM Z Table Accelerator abend codes appear as listed.

In addition, IBM Z Table Accelerator internal logical errors are indicated in CICS as LGIC abends and in other environments as U301 abends. If an abend causes the IBM Z Table Accelerator control blocks to be purposely dumped by IBM Z Table Accelerator, the transaction dump code generated in CICS is DKL1.

If the "Abend on Errors" switch is set to 'N' the return code and subcode are returned to the caller in the command area. If the "Abend on Errors" switch is set to 'Y' the thread will abend, and a message will be written to the JES log. Below is an example:

```
THE SPECIFIED LIBRARY IS NOT SUITABLE          RC = 0060 SC = 0006
```

In this example, you would reference 60-06 in the following table. Depending upon the error condition, there may or may not be a subcode.

In the event that the application thread abends, the IBM Z Table Accelerator error code is used as the user abend value, and any associated error subcode is used as the abend reason code.

Return Code	Text	Meaning / Instructions
00-00	IBM Z Table Accelerator command executed successfully	No action is indicated.

Table 6. IBM Z Table Accelerator messages and error codes (continued)

Return Code	Text	Meaning / Instructions
00-01	Hash High and/or Low Density values for table adjusted based on region limits	IBM Z Table Accelerator command executed successfully. Table density values were substituted with region limits because: <ul style="list-style-type: none"> - table high and low density values from library were greater than HASH_HI_DEN-LIM or less than HASH_LOW_DEN_LIM respectively - low density value may not have been greater than 2/3 of the high density value.
00-03	The table specified on the open command was already open.	IBM Z Table Accelerator command executed successfully. Table specified was already open in the designated TSR.
01-00	The specified command is invalid.	The two-character Command code specified is not a valid IBM Z Table Accelerator command.
02-00	The specified table must be open for this command.	Open the table before processing this command.
02-01	The table is in the process of being opened by another process.	This table is being opened and this command cannot wait for the open to complete. Retry the request.
02-98	Invalid parameter(s) were specified for the CL command.	Correct the parameters and resubmit the command (the CL command does not take any parameters).
03-00	The table is not closed.	The table must not be open in the TSR for this command.
03-01	The command requires that the table be closed	Close the table in the TSR and resubmit the command.
03-02	The table specified by the LD command has an incorrect row size.	Correct the row size for the LD table and resubmit the command. The row size depends on the DIRTYE specified; see the <i>Programming Guide</i> .
03-03	The table specified by the LD command has an incorrect key size.	Correct the key size for the LD table and resubmit the command. The key size depends on the DIRTYE specified; see the <i>Programming Guide</i> .
03-04	The table specified by the LD command has an incorrect key location.	The key location for the table created by the LD command is always 1. Correct the table definition and resubmit the command.

Table 6. IBM Z Table Accelerator messages and error codes (continued)

Return Code	Text	Meaning / Instructions
03-05	The table specified by the LD command is an alternate index.	The table used by an LD command must be defined as a data-table. Correct the table type and resubmit the command.
03-06	The table specified by the LD command is a linked VTS table	Not applicable for your version of IBM Z Table Accelerator.
03-07	The library specified by the LD command cannot be a VTSNAME	Not applicable for your version of IBM Z Table Accelerator.
03-08	The table is already open from a library with a different DSN	The table cannot be opened from the specified DSN. It is already opened in the TSR from a different DSN.
03-09	The table is in the process of being opened by another process.	This table is being opened. Retry the request after closing the table.
03-10	The command requires that the table be closed.	The table is open in the local TSR. Close the table and resubmit the command.
04-00	The FG command is not allowed; the table is not sequential (S or D)	Change the table organization to sequential (S or D) and resubmit the command.
05-00	The indirect open search criterion was not found in the primary table	The indirect open search criterion set by the SI command did not match any keys in the primary table specified on this command. See the <i>Programming Guide</i> for details on using INDIRECT-OPEN.
06-00	The count specified is invalid	Correct the COUNT value in the COMMAND AREA and resubmit the command.
06-01	The COUNT must be a positive number.	The COUNT value in the COMMAND AREA cannot be negative or zero. Correct and resubmit the command.
06-02	The COUNT exceeds the number of rows	The COUNT value in the COMMAND AREA cannot exceed the number of rows in the table (except for hash tables). Correct and resubmit the command.
06-03	The COUNT points to an empty row in a hash table.	The COUNT value in the COMMAND AREA for RC and DC commands for a hash table must point to a valid row.

Table 6. IBM Z Table Accelerator messages and error codes (continued)

Return Code	Text	Meaning / Instructions
07-00	The library DSN or DDNAME has been changed. Use DV or DW before STore.	If the library DSN or DDNAME for an open table has changed, the DV or DW command must be issued for the table before the table can be STored.
08-00	The generation number specified for this table is invalid.	Correct the generation number specified and resubmit the command. Generation numbers can be specified as absolute numbers from 1 to 255 or as negative numbers relative to current generation. The current generation is 0. (See the <i>Programming Guide</i> for details on specific commands)
09-00	The table is not found.	The table is not open in the TSR and cannot be found in the current library list (ML-list).
09-01	The table could not be found in any of the libraries in the ML list.	Ensure the ML-list includes the library where the table can be found and resubmit the command.
09-02	The table specified could not be found in the linked VTS-TSR.	Not applicable for your version of IBM Z Table Accelerator.
09-03	The table specified could not be found in the designated library.	Either ensure the ML-list includes the library where the table resides or ensure the table is resident in the currently designated library. If the command specifies a specific library DDNAME, ensure the table exists on that library.
10-00	The STATUS-SWITCHES parameter provided for the CS command is invalid	Correct the STATUS-SWITCHES parameter and resubmit the command. (See the <i>Programming Guide</i> for details on specific switches)
11-00	The Number of Generations to keep in the DT parameter must be 1 to 9.	IBM Z Table Accelerator supports up to 9 generations of a table on a library. The number of generations to keep in the DT parameter must be 1 to 9. Correct the value and resubmit the command.
12-00	The table name specified is invalid.	A valid table name is a string of 8 bytes that are not all blanks, all low values, all high values, or :TMPNAME. Please refer to the <i>Programming Guide</i> for more details.
13-00	The command is invalid in this environment	The two-letter command code is invalid. It is not supported in this execution environment.

Table 6. IBM Z Table Accelerator messages and error codes (continued)

Return Code	Text	Meaning / Instructions
13-01	This command is only supported for a VTS Agent.	The strobe-request internal command is only supported in the VTS Agent.
13-02	The DL (Define Library) command is not supported in CICS.	The DL command is only supported in the Batch interface.
13-03	Library expansion process may not be initiated from CICS.	The library expansion process is only supported by the DKJTEXEC Batch utility running in a non-multitasking environment.
13-05	Updates not permitted on a linked VTS table.	Not applicable for your version of IBM Z Table Accelerator.
13-06	Update commands are not permitted when a TSR is read only.	Not applicable for your version of IBM Z Table Accelerator.
13-07	Command not supported in this environment.	The command specified is not supported in this environment.
13-08	IBM Z Table Accelerator Library Access is not allowed for this interface.	Not applicable for your version of IBM Z Table Accelerator.
13-11	The command is invalid against a VTS-TSR.	When a VTS is being accessed either directly via the TAPARM or as a linked list, this command is invalid.
14-00	The ROW-SIZE specified is invalid. Must be from 1 to 32767.	The row size specified must be from 1 to 32767. Correct it and resubmit the command.
15-00	The KEY-SIZE specified is invalid. Must be from 1 to 256 and fit in row.	The key size specified must be from 1 to 256, and must fit in the row based on key location. Correct the key-size and resubmit the command. Note that date-sensitive processing can affect this.
16-00	The KEY-LOCATION specified is invalid. Must be from 1 to row-size.	Correct the key-location specified and resubmit the command. Note that the key location must allow the key size to fit within the row size.
17-00	The key will not fit within the row.	The key will not fit within the row because either the key's location in the row or its length is beyond the bounds of the row. Correct the key-size/key-location and resubmit the command. Note that date-sensitive processing can affect this.

Table 6. IBM Z Table Accelerator messages and error codes (continued)

Return Code	Text	Meaning / Instructions
18-00	Insufficient space is available on the IBM Z Table Accelerator library.	The library must be enlarged. A table cannot be stored on the library because there is insufficient free space available. Note that during the processing of a Store command there must be room in the library to hold 2 copies of the table on the library. There are 2 options: 1) create a new larger library and copy the existing library into it, or 2) expand the library. Both options are supported by the Batch Utility DKJTEXEC. (See the Batch Utilities Guide for details.)
19-00	This table already exists on the target library.	A new table cannot be STored on a library when it already exists on that library. Preceding the STore with a DW command will allow the table to overlay the previous table.
20-00	The maximum number of open tables has been exceeded. Check MAXNMTAB.	The MAXNMTAB option for the region must be increased to allow more tables into the TSR. Temporarily, tables can be CLosed. To change this value permanently, override the value in a TAOPT DD or update DKJTxx34. (See the <i>Implementation Guide</i> for the latter option.)
21-00	The Organization specified is invalid, must be blank, R, U, S, D, H. For date-sensitive processing, Organization can only be S.	Correct the organization specified and resubmit the command.
28-00	Paged tables are not supported.	Not applicable for your version of IBM Z Table Accelerator.
29-00	The secondary table is not located on same library as primary table	When the Open Indirect option is invoked for an Open command, the tables to be opened must reside in the same library as the primary table. (See the <i>Programming Guide</i> for details on Open Indirect, or the SI command for more details).
30-00	The password supplied is invalid.	Accessing this table on a IBM Z Table Accelerator library requires a valid password. Supply the correct password and resubmit the command.

Table 6. IBM Z Table Accelerator messages and error codes (continued)

Return Code	Text	Meaning / Instructions
31-00	The write password is either missing or incorrect.	Opening this table for write access from this IBM Z Table Accelerator library or updating it on this IBM Z Table Accelerator library requires the Write Password. Correct the password and resubmit the command.
32-00	The ST command requires that the table be opened for write.	A table can be STored on a IBM Z Table Accelerator library only if it is Open for Write in the TSR. Execute the OW command before attempting to STore the table.
33-00	A different generation of the table is already open.	Two generations of the same table cannot be open in a TSR at the same time. Close the table and open it again in the region in order to access the latest generation on the library.
39-00	The RN command failed. The new name already exists in the library.	The table rename (RN) command failed because a table with the new name already exists in the library.
40-00	The library DDNAME does not exist.	A library name specified in the library list (ML-list) or a command parameter does not correspond with an allocated DDNAME.
40-01	The specified library DDNAME is not allocated.	<p>Correct the library names in the library list (ML-list) or the command parameter and/or allocate the DDNAME in the region (add a DD statement to the JCL) and rerun the job.</p> <p>For CICS, this could be due to a missing entry in the FCT (File Control Table for BDAM libraries) or the CSD (CICS system definition file for VSAM libraries). The CICS region will need to be recycled for newly added FCT entries.</p>
40-02	The ST command failed because an ML entry is required	When a new table is defined in the TSR without any library DDNAMEs in the library list, the library on which the table is to be stored must be provided either by a DV or DW command before the STore command is issued or by specifying a valid library DDNAME in the ML-list in effect when the STore command is issued.

Table 6. IBM Z Table Accelerator messages and error codes (continued)

Return Code	Text	Meaning / Instructions
40-03	A blank or null DDNAME was used as a library list entry	A library DDNAME must be specified in the library list (ML-list) in effect when the command is being invoked.
40-04	The CA command failed because an ML entry is required	A CA command stores the alternate index definition on the first library in the library list (ML-list). A null or blank library list is not valid for this command.
40-05	The library list DDNAME entry is not allocated.	The current library list (ML-list) contains an entry that does not correspond to an allocated DDNAME. Correct the library list or add the DDNAME to the JCL and rerun or ALLOCATE with AL. For CICS, this could be due to a missing entry in the FCT (File Control Table for BDAM libraries) or the CSD (CICS system definition file for VSAM libraries). The CICS region will need to be recycled for newly added FCT entries.
41-00	The STORAGE-MODE-CODE (SMC) specified must be R or blank	Change the SMC value to "R" and resubmit the command.
42-00	The TAPARM sub parameter TA-FORMAT must be 0.	Correct the value in the TA-FORMAT sub-parameter and resubmit the command. (See the Programming Guide for details on TAPARM sub-parameters.)
43-00	The estimated NUMBER-OF-ROWS value in the DT block is out of range	When a table is defined, the estimated number of rows times the row size cannot exceed 2G. It also cannot be negative.
44-00	The EXPANSION-FACTOR specified in the DT block must be from 1 to 999	Correct the value in this field and resubmit the command. Consult the <i>Programming Guide</i> for the usage of this parameter.
49-00	This command requires more parameters than were supplied.	Supply the minimum number of parameters required for this command. Consult the <i>Programming Guide</i> and resubmit the command.
50-00	The DIRTYTYPE specified for the LD command must be T, L, or blank.	Correct the DIRTYTYPE used by the LD command and resubmit the command. (Consult the <i>Programming Guide</i> for the usage of the DIRTYTYPE parameter)

Table 6. IBM Z Table Accelerator messages and error codes (continued)

Return Code	Text	Meaning / Instructions
51-00	HIGH/LOW DENSITY specified in DT block must be from 1 to 999.	Correct the DENSITY value being used and resubmit the command. Consult the <i>Programming Guide</i> for the usage of these parameters. Note that the specified HIGH-DENSITY and LOW-DENSITY parameters can be adjusted dynamically by the HASH-HI-DEN-LIM and HASH_LOW_DEN_LIM run-time options.
55-00	The SEARCH-METHOD must be either S, Q, B, C, or H	Correct the SEARCH-METHOD value being used and resubmit the command. Consult the <i>Programming Guide</i> for the usage of this parameter.
56-00	The SEARCH-METHOD is incompatible with ORGANIZATION specified	Correct the SEARCH-METHOD and/or the ORGANIZATION field ensuring they are compatible and resubmit the command. (Consult the <i>Programming Guide</i> for the relationship of these two parameters).
58-00	The requested module cannot be found in the load libraries searched.	Add the missing module to the load libraries accessed through STEPLIB, LPA or linklist.
60-00	The specified library is not suitable.	The dataset specified is not a valid IBM Z Table Accelerator library.
60-01	The specified library's data definition must be DASD	Ensure that the library DDName is allocated to DASD, and resubmit the command. Note that if you are using VSAM IBM Z Table Accelerator libraries with LSR, they are not supported in the IBM Z Table Accelerator Batch interface. For more information, see the "Best Practices" chapter in the <i>IBM Z Table Accelerator Administration Guide</i> .
60-02	The specified library's BLKSIZE must be at least 600.	Increase the library BLKSIZE to at least 600 and resubmit the job.
60-03	The specified library must contain at least 10 blocks.	Increase the number of blocks in the library; 10 is the minimum. Correct the allocation and resubmit the command.

Table 6. IBM Z Table Accelerator messages and error codes (continued)

Return Code	Text	Meaning / Instructions
60-04	IBM Z Table Accelerator libraries are limited to 8,388,607 blocks.	Reduce the number of blocks in the library. The maximum for a BDAM library dataset is the number of blocks that can be contained in 65,535 tracks (approximately 4369 cyl). The maximum for a un-extended VSAM RRDS library dataset is the number of blocks that does not exceed 4G bytes of data (approx 6145 cyl). Correct the allocation and resubmit the command.
60-05	The BLKSIZE of the source and target libraries differ.	The process being invoked requires that the source and target libraries have the same block size. Correct the block sizes and resubmit the command.
60-06	The size of the designated library is smaller than that of the source.	Reallocate the target library dataset to contain the same as or more blocks than the source dataset and rerun the job.
60-07	The specified library's BLKSIZE is not known	The library is not a valid IBM Z Table Accelerator library. Ensure that the library has been successfully DEFINED with either the DL command or the DKJTEXEC Batch utility. Rerun this job after the library is successfully defined.
60-08	The library expansion is not complete. Recreate the library.	An I/O error was experienced performing the library expansion. Recreate the library and resubmit the job.
60-09	The target library cannot be a V5 Bridge library.	Not applicable for your version of IBM Z Table Accelerator.
61-00	The library status is invalid.	The IBM Z Table Accelerator library DD is not allocated appropriately
61-01	The library DDNAME may not be concatenated with another dataset.	IBM Z Table Accelerator Library DDNAME concatenation is not supported. Ensure each library has a separate DDNAME and resubmit the command. Note that IBM Z Table Accelerator VSAM libraries using LSR (Local Shared Resources) is not supported by the IBM Z Table Accelerator Batch interface.

Table 6. IBM Z Table Accelerator messages and error codes (continued)

Return Code	Text	Meaning / Instructions
61-02	The DSNAME is open with another DDNAME	Only 1 DDNAME is allowed per library DSNAME. Correct the JCL and resubmit the job. Note that IBM Z Table Accelerator retains the reference of a DDName to a DSName for the life of a region; that relationship is established the first time the library is accessed in the region. This applies only to DDName and DSName relationships established using JCL.
61-03	The DL command or library expansion failed, the library was open	The library was already in use. Ensure the process has exclusive use of the library and resubmit the job.
61-04	The DL command or library expansion failed, DISP must be NEW or OLD	Correct the library disposition in the JCL and resubmit the job.
61-05	The DL command failed; a VSAM library was not defined as reusable.	Redefine the VSAM library as reusable and resubmit the job.
61-06	The designated library access method must be BDAM or VSAM.	Correct the DSORG of the library and resubmit the job.
61-07	The library DISP must be OLD or SHR for this command.	Correct the library disposition in the JCL and resubmit the job.
61-08	OPTCD=C and Z mutually exclusive on IBM Z Table Accelerator library DD statement.	Correct the OPTCD DD statement and resubmit the job. See the <i>Administration Guide</i> for documentation on IBM Z Table Accelerator use of the OPTCD subparameters C and Z.
61-09	The BDAM library must be RECFM=F or FB	This error will be returned in the following cases: <ol style="list-style-type: none"> 1. when initializing a BDAM library and RECFM is not F 2. when accessing a BDAM library and LRECL not equal to BLKSIZE. For BDAM library definitions, only RECFM=F is allowed. For BDAM library access, RECFM= F or FB are allowed if LRECL is equal to BLKSIZE.
61-10	The VSAM library definition must be RRDS (Relative Record dataset).	Correct the VSAM library definition and resubmit the job.

Table 6. IBM Z Table Accelerator messages and error codes (continued)

Return Code	Text	Meaning / Instructions
61-11	A write operation to the library was terminated, library read-only.	No updates are allowed to this library because it is defined with read-only access. In Batch jobs, LABEL=(,,IN) on the DD statement defines a library dataset as read-only; in CICS, the library file definition is read-only. If updates to the library are to be allowed, change the JCL or CICS file definition. If no updates are allowed, inform the owners of the application to change their practices.
61-12	The specified library is RACF protected for this type of access.	The RACF/AUTHORITY does not allow the specified access to the library dataset. Determine whether to request RACF changes or inform the owners of the application to change their practices.
62-00	Format of the specified library is not correct.	Ensure that the library has been initialized using DKJTEXEC and resubmit the job.
64-00	The INDEX parameter specified must be P or blank.	Correct the index parameter and resubmit the command. (Consult the <i>Programming Guide</i> for the usage of the INDEX parameter)
68-00	The COUNT value is too small for the DU command to dump all rows	This Error Code may be OK if only a portion of the rows were to be dumped. If all the rows in the table were to be dumped, Increase the COUNT value to the total number of rows and resubmit the command. (Consult the <i>Programming Guide</i> for the parameters of the DUmp command).
70-00	IBM Z Table Accelerator cleanup unable to complete due to CICS shutdown processing.	IBM Z Table Accelerator's attempts to clean up control blocks in the CICS region at termination cannot complete due to CICS termination processing.
71-00	Lock is unavailable after waiting LOCKTIMERC seconds.	A IBM Z Table Accelerator lock was not available for use after waiting LOCKTIMERC seconds. Resubmit the command. If a longer wait time for the resource is desired, update the TAOPT LOCKTIMERC value and resubmit the command.
72-00	The table is unavailable at this time; WAIT switch is off.	The table was in use at the time of the command. Retry the command.

Table 6. IBM Z Table Accelerator messages and error codes (continued)

Return Code	Text	Meaning / Instructions
72-01	The table is unavailable after waiting TABLEWAITRC seconds.	With the WAIT switch on, the table enqueue was unavailable after waiting TABLEWAITRC seconds. If the wait time is to be increased, update the TAOPT TABLEWAITRC value and resubmit the command. The table enqueue is an MVS enqueue with a major name of DKJTCALL and a minor name starting with T followed by the table name and the library DSN. It is held while a table is open for write in a TSR.
73-00	An update command was issued for a table that was not open for write	An update command was issued for a table that has not been opened for write. See descriptions for other sub-codes for detailed explanation. The command being issued requires that the table be OPEN for WRITE (OW). For correct operation, issue an OW command for the table and retry the command.
73-01	An update command was issued for a table that is open for read in a local or VTS-TSR	An update command was issued for a table that has been opened for read in a local or VTS-TSR and this is not a single-tasking application i.e. MULTITASKING=Y or it is a CICS or IMS region. The command being issued requires that the table be OPEN for WRITE (OW). For correct operation, issue an OW command for the table and retry the command.
73-02	An update command was issued for a table that is open for read in a VTS-TSR	An update command was issued for a table that has been opened read in a VTS-TSR and this is a single-tasking application i.e. MULTITASKING=N. The command being issued requires that the table be OPEN for WRITE (OW). For correct operation, issue an OW command for the table and retry the command.
73-03	An update command was issued against the local TSR for a closed table	An update command was issued against the local TSR for a table that is not open in the TSR, implicit open is ON, and this is not a single-tasking application. The table will not be opened after the command completes. The command being issued requires that the table be OPEN for WRITE (OW). For correct operation, issue an OW command for the table and retry the command.

Table 6. IBM Z Table Accelerator messages and error codes (continued)

Return Code	Text	Meaning / Instructions
73-04	An update command was issued against a VTS-TSR for a closed table	An update command was issued against a VTS-TSR for a table is not open in that TSR and implicit open is ON. The table will not be opened after the command completes. The command being issued requires that the table be OPEN for WRITE (OW). For correct operation, issue an OW command for the table and retry the command.
73-05	An update command was issued for a linked table	Not applicable for your version of IBM Z Table Accelerator.
74-00	The LOCK-LATCH specified is invalid	If a lock latch is set on a table in a TSR, any update command must match the lock latch value. Correct the LOCK-LATCH value and resubmit the command. If the lock-latch is not known, the IBM Z Table Accelerator Administrator can perform the required command overriding the lock-latch with the Master Password.
75-00	Dynamic allocation (AL) or un-allocation (UL) failed	IBM Z Table Accelerator command AL (Allocate library) or UL (Un-allocate library) was not successful.
75-01	TBTSLIB not allowed as parameter in AL command.	TBTSLIB is a reserved DDname and cannot be used with the AL command.
75-02	DDNAME specified in AL command already in use.	Ensure the DDNAME is not already allocated and resubmit the command.
75-03	DSNAME specified in AL command not catalogued.	Ensure the DSNAME for the DDNAME is catalogued and resubmit the command. The AL command only supports catalogued library datasets
75-04	DSN enqueue conflict with another region for AL command.	Determine which region has the SYSDSN enqueue for this table and resolve the conflict. Then reissue the AL command with the appropriate SHR parameter. (Consult the <i>Programming Guide</i> for the usage of the AL command)
75-05	Error in DDNAME specified in AL command.	Correct the DDNAME syntax and resubmit the command
75-06	Syntax error in DSNAME specified in AL command.	Correct the DSNAME syntax and resubmit the command

Table 6. IBM Z Table Accelerator messages and error codes (continued)

Return Code	Text	Meaning / Instructions
75-07	SVC99 failure during execution of AL command.	An MVS error was generated. Contact your MVS systems programmer to determine why the AL command failed.
75-08	DDNAME specified by UL command must be allocated by an AL command.	The UL (Unallocate) command failed because IBM Z Table Accelerator was not used to allocate the DDNAME. Deallocate the DDNAME (TSO FREE) and use the IBM Z Table Accelerator AL command to allocate the library dataset.
75-09	DDNAME specified by UL command not known or library not open.	Ensure the DDNAME is allocated.
75-10	SVC99 failure during execution of UL command.	An MVS error was generated. Contact your MVS systems programmer to determine why the UL command failed.
76-00	OI (Open Indirect) failed: primary table not sequential or KLOC not 9.	Ensure the Open Indirect primary table has sequential organization and the key location is 9. Resubmit the command. (See the <i>Programming Guide</i> for details on Open Indirect, or the SI command for more details.)
77-00	An alternate index is defined for a table that is not TYPE=P.	Ensure the data table is defined as a pointer table (TYPE=P) and resubmit the command
78-00	The AL command failed, the SHARE-STATUS parameter must be S or O	Correct the SHARE-STATUS sub-parameter of the LIBRARY-ALLOC parameter. (See the <i>Programming Guide</i> for the usage of the AL command.)
79-01	The request code for the VV command is invalid.	Not applicable for your version of IBM Z Table Accelerator.
79-02	A VTS initialized as RW cannot be switched to RO mode.	Not applicable for your version of IBM Z Table Accelerator.
79-03	The job name does not match the job holding the RW lock for the VTS.	Not applicable for your version of IBM Z Table Accelerator.
79-04	The time stamp of the job does not match that holding the RW lock for the VTS.	Not applicable for your version of IBM Z Table Accelerator.
79-05	Lock reset failed when switching VTS access to RO.	Not applicable for your version of IBM Z Table Accelerator.

Table 6. IBM Z Table Accelerator messages and error codes (continued)

Return Code	Text	Meaning / Instructions
79-06	Attempt to switch VTS to RW unsuccessful; VTS already locked in RW mode.	Not applicable for your version of IBM Z Table Accelerator.
79-07	Lock failed when switching VTS access to RW.	Not applicable for your version of IBM Z Table Accelerator.
79-08	VTS name mismatch	This is an internal error and cannot be fixed by the client. Please contact Technical Support.
79-09	No pointer to TSRCNTL	This is an internal error and cannot be fixed by the client. Please contact Technical Support.
79-10	No pointer to GCE	This is an internal error and cannot be fixed by the client. Please contact Technical Support.
79-11	VTS-TSR not available	Ensure that the VTS is not terminating or starting up when the command is issued.
79-12	Getmain failed	Try running the entire job in a larger region.
79-13	SAF authorization could not be obtained to switch the access mode for this VTS.	Not applicable for your version of IBM Z Table Accelerator.
79-14	VTS in RO mode cannot be switched to RO.	Not applicable for your version of IBM Z Table Accelerator.
79-15	Program not authorized for switching VTS access mode.	Not applicable for your version of IBM Z Table Accelerator.
80-00	Data table is not open; the use of an alternate index not allowed.	The IA command requires the data table be opened first. The OR command will dynamically open a data table.
80-01	Data table must be opened before IA (Invoke Alternate) is issued.	Open the data table before issuing the IA command.
80-02	IA (Invoke Alternate Index) may not reference a linked VTS table.	Not applicable for your version of IBM Z Table Accelerator.
80-03	Open/Invoke Alternate Index must reference a data table.	Ensure the IA command includes the BASE-TABLE-NAME parameter and resubmit the command. The IA or OR command must reference a data table, either in BASE-TABLE-NAME or the stored definition on the library.

Table 6. IBM Z Table Accelerator messages and error codes (continued)

Return Code	Text	Meaning / Instructions
80-04	Alternate Index cannot be opened as a linked VTS table.	Not applicable for your version of IBM Z Table Accelerator.
81-00	The specified Alternate Index definition is not found	If the complete alternate index definition is not specified in the IA command parameters, the alternate index definition must be found in a library in the current library list (ML-list). Correct this error and resubmit the command.
83-00	Data table INDEX sub-parameter of the DT block must be defined as P	If the base table is specified in the IA command, it must be defined as a pointer table (TYPE=P). Redefine the base table.
85-00	The command is invalid for an Alternate Index	The command being invoked is not supported for use with an alternate-index. Correct the error and resubmit the command.
85-01	An alternate index may not have the same name as the data table.	Correct the alternate-index parameters ensuring that the alternate index and base table names are different and resubmit the command.
85-02	DV and DW commands are not supported for Alternate index tables	The DV and DW commands are only supported for base tables since they apply to ensuing STORE commands. An alternate index is never stored, only the base table is.
86-00	Data table of an Alternate Index may not be a linked VTS table.	Not applicable for your version of IBM Z Table Accelerator.
87-00	The KEY-COUNT must be 1 for the definition of the Alternate Index.	Only a value of 1 for the key-count is supported when defining an alternate-index. Correct the key-count value and resubmit the command.
89-00	An invalid parameter was encountered in the TAOPT file.	During initialization of IBM Z Table Accelerator in the region, an error was discovered in the IBM Z Table Accelerator runtime options in the TAOPT dataset. Correct the TAOPT parameter(s) and resubmit the job.

Table 6. IBM Z Table Accelerator messages and error codes (continued)

Return Code	Text	Meaning / Instructions
90-00	Insufficient storage available in region: increase region size.	There is insufficient main storage available for IBM Z Table Accelerator control blocks. Increase the executing region size and resubmit the job (modify the 'REGION=' parameter for the jobstep).
91-00	I/O Error	Determine the reason for the I/O error from the JESMSG LG. Correct the problem and resubmit the command.
91-01	I/O Error while attempting to read/write to a IBM Z Table Accelerator library	Determine the reason for the I/O access error to the IBM Z Table Accelerator library. Correct the problem and resubmit the command.
91-02	I/O Error while attempting to read from the TAOPT file	Determine the reason for the error trying to read from the TAOPT file. Once resolved, resubmit the job.
91-03	Error in I/O subtask in multitasking environment.	The subtask was unable to complete due to an I/O error when TAOPT MULTITASKING=Y. Correct the problem and resubmit the process.
92-00	Insufficient Table Space Region (TSR) size	The TSR is too small for the tables being opened. Either enlarge the TSR or close unnecessary tables in the TSR to free up the required TSR space.
92-01	The space request is bigger than the TSR	The table to be opened is too large for the current TSR. The TSR size will need to be increased to accommodate the table to be opened.
92-02	Insufficient free space in the TSR: close tables or enlarge TSR.	The TSR is full and no longer has enough free space available to allow the table to be opened, defined, expanded or changed. Either enlarge the TSR appropriately or close unnecessary tables to free up the required TSR space.
92-03	There is insufficient free contiguous space in the TSR	The TSR is full. There is no longer any contiguous memory left to accommodate the index for the table being opened, defined, expanded or changed. Either enlarge the TSR or close unnecessary tables to free up required TSR space.

Table 6. IBM Z Table Accelerator messages and error codes (continued)

Return Code	Text	Meaning / Instructions
92-04	There is insufficient space in the TSR to initiate the TCE	The TSR is full. There is no room to accommodate the IBM Z Table Accelerator control block created when opening a table into the TSR. Either enlarge the TSR or close unnecessary tables to free up the required TSR space.
93-00	The RF command did not complete	Correct the error and resubmit the command.
93-01	The table must be open for read in order to be refreshed	Open the table for read (OR) or ReRelease the table before submitting the RF command.
93-02	The data table or alternates were closed since the table was opened	There is a change in either the alternate or the data-table that is preventing successful completion of the RF command. CL the table and re-open it from the desired library to get the most recent version of the table into the TSR.
93-03	The ROWSIZE changed since the table was opened	RF failed because the row size for the table on the library was different compared to the table opened in the TSR. CL and reopen the table to get the most recent version of the table into the TSR.
93-04	The KEYSIZE changed since the table was opened	RF failed because the key size for the table on the library was different compared to the table opened in the TSR. CL and reopen the table to get the most recent version of the table into the TSR.
93-05	The KEYLOC changed since the table was opened	RF failed because the key-location for the table on the library was different compared to the table opened in the TSR. CL and reopen the table to get the most recent version of the table into the TSR.
93-06	The ORGANIZATION changed since the table was opened	RF failed because the organization for the table on the library was different compared to the table opened in the TSR. CL and reopen the table to get the most recent version of the table into the TSR.

Table 6. IBM Z Table Accelerator messages and error codes (continued)

Return Code	Text	Meaning / Instructions
93-07	The SEARCH METHOD changed since the table was opened	RF failed because the Search-Method for the table on the library was different compared to the table opened in the TSR. CL and reopen the table to get the most recent version of the table into the TSR.
95-00	The transaction was terminated during IBM Z Table Accelerator processing.	IBM Z Table Accelerator cleanup was invoked because a transaction was terminated while IBM Z Table Accelerator was processing. The subcodes document what IBM Z Table Accelerator process was executing when the transaction was terminated. This is usually related to a transaction timeout or cancellation.
95-02	IBM Z Table Accelerator found an invalid row (item) area address	See 95-00 description. This can be caused by an error in the input parameters provided.
95-03	IBM Z Table Accelerator found an invalid row address in the Table Space Region (TSR)	See 95-00 description. Send the documentation (JESMSG LG, TADUMP, any other dumps) to IBM Z Table Accelerator support.
95-04	IBM Z Table Accelerator abended during an index resorting process.	See 95-00 description.
95-05	IBM Z Table Accelerator abended during an index expansion process.	See 95-00 description.
95-06	IBM Z Table Accelerator abended during an index entry insertion process.	See 95-00 description.
95-07	IBM Z Table Accelerator abended during an index entry deletion process.	See 95-00 description.
95-08	IBM Z Table Accelerator abended during an open command.	See 95-00 description
95-09	IBM Z Table Accelerator abended during a close command.	See 95-00 description.
97-00	The ML or LL parameter list is invalid	Correct the ML or LL parameter list and resubmit the command.
97-05	An ML list must not contain a reference to TBTSLIB.	TBTSLIB is a reserved DDname and cannot be used in the ML-list.
97-08	The ML list is invalid in VTS. It contains a reference to a VTS-TSR.	Remove the reference to the VTS-TSR in the ML-list and resubmit the command.

Table 6. IBM Z Table Accelerator messages and error codes (continued)

Return Code	Text	Meaning / Instructions
98-00	An internal error occurred in IBM Z Table Accelerator Shutdown processing.	IBM Z Table Accelerator termination (XX command) problems.
98-01	The XX command is being processed on a subordinate task.	When multi-tasking is ON, the XX command must be performed by the persistent (mother) task. This command will fail if it is submitted by sub-tasks invoked by the persistent (mother) task. Send the documentation (JESMSG LG, TADUMP, any other dumps) to IBM Z Table Accelerator support.
98-02	IBM Z Table Accelerator shutdown process did not complete.	If multi-tasking and a sub-task submitted the XX command, other tasks will first issue a warning, then they will abend. This is the warning. Send the documentation (JESMSG LG, TADUMP, any other dumps) to IBM Z Table Accelerator support.
98-03	The subtask that shutdown IBM Z Table Accelerator has tried to invoke IBM Z Table Accelerator.	If multi-tasking and a sub-task submitted the XX command, no tasks are allowed to issue IBM Z Table Accelerator commands. Send the documentation (JESMSG LG, TADUMP, any other dumps) to IBM Z Table Accelerator support.
1072-00	VTS access failed because the specified VTS-TSR is not available.	VTS access failed because the specified VTS-TSR is not available. Ensure that the VTS job which creates the specified VTS-TSR name is running.
1072-01	VTS access failed because the specified PC Server was not running.	Access to a VTS-TSR requires both the VTS and the application to be running under the same PC Server on the same LPAR.
1072-02	VTS close connection failed.	An attempt to close the VTS connection failed. Refer to message DKJ00303I in the JESMSG LG for more details
1072-03	SAF READ access required for this operation.	Not applicable for your version of IBM Z Table Accelerator.
1072-04	SAF UPDATE access required for this operation.	Not applicable for your version of IBM Z Table Accelerator.
1072-05	VTS could not be accessed due to getmain failure for the VTS control block.	Run the job in a larger region.

Table 6. IBM Z Table Accelerator messages and error codes (continued)

Return Code	Text	Meaning / Instructions
1072-06	VTS could not be accessed due to corrupt VTS control block.	This is an internal error and cannot be fixed by the client. Please contact Technical Support.
1072-07	VTS has stopped running.	The VTS is terminating or is no longer running.
1072-08	VTS could not be found in this version.	The VTS name does not match any VTS in this version.
1072-09	VTS could not be found in the region list of VTS names.	The VTS does not exist in the region list of VTS names for this version.
1072-10	VTS no longer running.	The VTS was running but has been quiesced.
1072-11	VTS re-initialization failed during VTS close.	This is an internal error and cannot be fixed by the client. Please contact Technical Support.
1072-12	VTS re-initialization failed during VTS open.	This is an internal error and cannot be fixed by the client. Please contact Technical Support.
1072-13	VTS not running.	The VTS is not running.
1072-14	VTS name in region list has been zeroed.	This is an internal error and cannot be fixed by the client. Please contact Technical Support.
1072-15	The VTS generation number or time stamp does not match the current generation or time stamp.	Not applicable for your version of IBM Z Table Accelerator.
1072-16	The VTS is not available for transaction processing.	The VTS is terminating or is initializing.
1072-17	SAF CONTROL access required for this operation.	Not applicable for your version of IBM Z Table Accelerator.
1072-18	The VTS is currently locked in RW mode.	Not applicable for your version of IBM Z Table Accelerator.
1072-19	SAF ALTER access required for this operation.	Not applicable for your version of IBM Z Table Accelerator.

IBM Z Table Accelerator messages

Messages issued from the IBM Z Table Accelerator components, such as the VTS Agent, CICS, Batch and others, vary from function to function, and are listed in the table below. These messages can usually be viewed on the JESMSGLOG. Any Severe Errors (S) messages should be referred to IBM Z Table Accelerator Customer Support. Messages are in the format DKJnnnnnA where DKJ is the prefix, followed

by a five digit numeric code and an error code letter. The error code letter indicates the type and severity of the error message:

E—Error (usually indicates user error)

I—Information

W—Warning

A—Action (user action is required)

S—Severe (no further processing is possible)

Msg Code	Text	Meaning / Instructions
DKJ00049S	NO CMDAREA IN CALL TO IZTA. ABENDING.	There was no command area supplied. Ensure that a command area is supplied for the IBM Z Table Accelerator call.
DKJ00100W	Unexpected non-zero return code from user exit	
DKJ00107W	A message was truncated at 207 characters	Information only.
DKJ00200S	Initialization failed	
DKJ00200S	IZTA Initialization failed	
DKJ00201S	IZTA PC Server unavailable	The IBM Z Table Accelerator PC Server is not running. It must be running in order for IBM Z Table Accelerator to initialize in batch jobs, CICS regions, IMS regions and VTS-TSRs.
DKJ00202I	IZTA V<nnn> is initializing for <customername>	Information only.
DKJ00203I	IZTA <interface> license expires on <yyyy/mm/dd>	Information only.
DKJ00204I	Executing in Step=<stepname>	Information only.
DKJ00205I	Executing in JOB Step=<stepname> PROC Step=<stepname>	Information only.
DKJ00206S	IZTA PC Server does not support LOAD	Ensure the latest PC Server is running.
DKJ00207E	DKJTX071 not loaded; System Exits disabled	Not applicable for your version of IBM Z Table Accelerator.
DKJ00207E	DKJTX072 not loaded; User Exits disabled	Not applicable for your version of IBM Z Table Accelerator.
DKJ00208S	Unable to locate IZTA GCA	Ensure the latest PC Server is running.
DKJ00209S	IZTA PC Server unavailable	Ensure the latest PC Server is running.

Table 7. Messages and error codes for VTS Agent, CICS, Batch, etc. (continued)

Msg Code	Text	Meaning / Instructions
DKJ00210E	Module DKJTNUCL not loaded	Contact Technical Support for assistance.
DKJ00210I	Unable to locate IZTA GCA	Information only.
DKJ00211I	Unable to locate IZTA GCA.	Information only.
DKJ00213E	Error in IZTA PC server	Contact Technical Support for assistance.
DKJ00214E	Create CMA Named Token failed	Contact Technical Support for assistance.
DKJ00214E	CMA GETMAIN failed	Contact Technical Support for assistance.
DKJ00216S	\$DOPC GRANDE Failed	Ensure the latest PC Server is running.
DKJ00217S	\$DOPC Failed to Load DKJTROTB.	Check that DKJTROTB is available for an MVS LOAD.
DKJ00218S	\$DOPC Free LX Index Failed	Contact Technical Support for assistance.
DKJ00221E	I/O ERROR:<text>	Contact Technical Support for assistance.
DKJ00222E	QSAM ERROR: FUNCTION=<xxxxxxxx>,FILE=<fffffff>	Contact Technical Support for assistance.
DKJ00224E	I/O PROBLEM: FILE=<xxxxxxxx> FUNC=<fffffff> RESP=<dddd> RESP2=<dddd>	Contact Technical Support for assistance.
DKJ00225W	JOB:<jobname>,ID:<userid>,TSR:<TSRname>, CMD:<cc>,TBL:<tablename>, - WAITING <nnnnn> SECS FOR MAPLOCK HELD BY <jobname2> <task-id>	TSR space allocation is blocked by another process accessing this TSR. <jobname2> identifies the job holding the lock and <task-id> identifies the task id of the transaction holding the lock. Cancelling the job or transaction holding the lock may free up the lock. Warning: The operator must determine if it is safe to cancel the job or transaction. Note: This message is usually intermittent until the IBM Z Table Accelerator command that is blocking the access ends. If the problem persists and is happening in a VTS-TSR, contact Technical Support for instructions on how to alleviate the problem.

Table 7. Messages and error codes for VTS Agent, CICS, Batch, etc. (continued)

Msg Code	Text	Meaning / Instructions
DKJ00226W	Please change DD TAOPTV to TAOPT	Not applicable for your version of IBM Z Table Accelerator.
DKJ00226i	Please change DD TAOPTV to TAOPT	Not applicable for your version of IBM Z Table Accelerator.
DKJ00227W	JOB:<jobname>,ID:<userid>,TSR:<TSRname>,CMD:<cc>,TBL:<tablename>, - WAITING <nnnnnn> SECS FOR GLOBAL TCELOCK HELD BY <jobname2> <task-id>	<p>Access to the table in the TSR is blocked by another process using the table. <jobname2> identifies the job holding the lock and <task-id> identifies the task id of the transaction holding the lock. Cancelling the job or transaction holding the lock may free up the lock.</p> <p>Warning: The operator must determine if it is safe to cancel the job or transaction.</p> <p>Note: This message is usually intermittent until the IBM Z Table Accelerator command that is blocking the access ends. If the problem persists and is happening in a VTS-TSR, contact Technical Support for instructions on how to alleviate the problem.</p>
DKJ00227W	JOB:<jobname>,ID:<userid>,TSR:<TSRname>,CMD:<cc>,TBL:<tablename>, - WAITING <nnnnnn> SECS FOR TCELOCK HELD BY<jobname2> <task-id>	<p>Access to the table in the TSR is blocked by another process using the table. <jobname2> identifies the job holding the lock and <task-id> identifies the task id of the transaction holding the lock. Cancelling the job or transaction holding the lock may free up the lock.</p> <p>Warning: The operator must determine if it is safe to cancel the job or transaction.</p> <p>Note: This message is usually intermittent until the IBM Z Table Accelerator command that is blocking the access ends. If the problem persists and is happening in a VTS-TSR, contact Technical Support for instructions on how to alleviate the problem.</p>
DKJ00228W	JOB <jobname> IN <VTSname> WAITING <nnnnnnnn> SECONDS TABLE ENQ <tablename>	Job is waiting for VTS-TSR table enqueue--not received yet.

Table 7. Messages and error codes for VTS Agent, CICS, Batch, etc. (continued)

Msg Code	Text	Meaning / Instructions
DKJ00229W	JOB:<jobname>,ID:<userid>,TSR:<TSRname>, CMD:<cc>,TBL:<tablename>, - WAITING <nnnnnn> SECS FOR TLBLOCK HELD BY <jobname2> <task-id> - LIB DD:<DDname>,DSN:<DSNname>	Access to the IBM Z Table Accelerator library is blocked by another process in the same region using the library. <jobname2> identifies the job holding the lock and <task-id> identifies the task id of the transaction holding the lock. Cancelling the job or transaction holding the lock may free up the lock. Warning: The operator must determine if it is safe to cancel the job or transaction. Note: This message is usually intermittent until the IBM Z Table Accelerator command that is blocking the access ends.
DKJ00230I	TAOPT Processing	Information only.
DKJ00231I	<keyword>=<value>	Information only.
DKJ00232E	Invalid value	The value entered is invalid.
DKJ00233E	Missing keyword	A keyword is missing.
DKJ00234E	Invalid keyword	The keyword entered is invalid.
DKJ00235E	Value exceeds maximum allowed	The value provided exceeds the maximum allowed value.
DKJ00236E	Keyword not allowed in this environment	The parameter provided is not allowed for this interface.
DKJ00237E	Missing "="	A "=" is missing in the parameter.
DKJ00238E	Invalid quoted string	The string is invalid.
DKJ00239E	Value field too wide	The value entered is too large. Correct and try again.
DKJ00240E	Value less than minimum allowed	The value entered is smaller than the minimum required.
DKJ00241E	Missing value	There is a missing value in the parameter.
DKJ00242E	Keyword (or alias) previously specified	The keyword has already been specified by using another alias name.

Table 7. Messages and error codes for VTS Agent, CICS, Batch, etc. (continued)

Msg Code	Text	Meaning / Instructions
DKJ00243E	Option supported only in DKJOPTGN	This parameter must be entered in the distributed source for DKJV1134 or DKJTxx34 and rebuilt. It cannot be entered in the input for the TAOPT DD card.
DKJ00244I	MAXNMTAB set to <nnnnnn>	Information only.
DKJ00245W	VTSONLY not valid with LIB entries	Not applicable for your version of IBM Z Table Accelerator.
DKJ00246W	Invalid VTS prefix for LIB<nn>=<xxxxxxxx>	Not applicable for your version of IBM Z Table Accelerator.
DKJ00247W	Empty value: LIB<xx> - LIB<yy>	
DKJ00248W	LIB<xx> exceeds max allowed <mm>	
DKJ00249W	MULTOPNX=N is not valid	Not applicable for your version of IBM Z Table Accelerator.
DKJ00250I	TAOPT Parameters: (*= Parameters not taken from TAOPT input file)	Information only.
DKJ00251I	<keyword>=<value>.	Information only.
DKJ00252I	<<<<<<< TAOPT Done >>>>>>	Information only.
DKJ00253I	<message>	Information only.
DKJ00255E	HASH_LOW_DEN_LIM must be <= 2/3 HASH_HI_DEN_LIM	Reset the density limits.
DKJ00256I	TSRSIZE set to LDS size of <nnnnnnn>K	Not applicable for your version of IBM Z Table Accelerator.
DKJ00258I	MAXNMTAB set by preformatted LDS TSR	Not applicable for your version of IBM Z Table Accelerator.
DKJ00260I	DKJTX066 Exit Manager Initializing	Not applicable for your version of IBM Z Table Accelerator.
DKJ00262W	Exit <xxxxxxxx> not loaded and disabled	Not applicable for your version of IBM Z Table Accelerator.
DKJ00263I	Exit <xxxxxxxx> is loaded and enabled	Not applicable for your version of IBM Z Table Accelerator.
DKJ00264I	DKJTX066 Exit Manager Shutdown complete	Not applicable for your version of IBM Z Table Accelerator.

Table 7. Messages and error codes for VTS Agent, CICS, Batch, etc. (continued)

Msg Code	Text	Meaning / Instructions
DKJ00265W	Exit name <xxxxxxxxxxxxxxxx> not found in current release	Not applicable for your version of IBM Z Table Accelerator.
DKJ00266E	PC Service unavailable; required for VTS_TSR use	Ensure that the correct version of the IBM Z Table Accelerator PC Server is running.
DKJ00270E	Insufficient storage to process strobe data; strobe suppressed.	The strobe data file has run out of space so IBM Z Table Accelerator strobe data will be suppressed.
DKJ00271E	Error on CICS WRITE JOURNALNUM; Strobe logging suppressed	
DKJ00272E	WRITE JOURNALNUM RESP=<1234>, RESP2=<1234>	
DKJ00273E	Dynalloc of strobe report failed. Strobe suppressed.	
DKJ00274I	STROBE=0 forced for Stored Procedure	Not applicable for your version of IBM Z Table Accelerator.
DKJ00275E	Module DKJTRSTA not found.	Contact Technical Support for assistance.
DKJ00280E	QSAM ERROR: FUNCTION=<xxxxxxxx>,FILE=<fffffff>	Contact Technical Support for assistance.
DKJ00281E	STROBE OUTPUT SUPPRESSED	IBM Z Table Accelerator strobe report output (for file TATSRPT) has been suppressed due to some error in the file, possibly a file overflow. Try increasing the disk space for the strobe report output file, TATSRPT.
DKJ00282E	I/O ERROR:	Contact Technical Support for assistance.
DKJ00283E	TATSRPT spin failed. See JESlog for messages.	
DKJ00284E	TATSRPT test failed. See JESlog for messages.	
DKJ00286E	SYSPRINT DD MISSING. Run terminated.	A SYSPRINT DD card is required in the job.
DKJ00292W	DEBUG TRACE requested but not started	Contact Technical Support for assistance.
DKJ00292E	DEBUG TRACE requested but not started	Contact Technical Support for assistance.

Table 7. Messages and error codes for VTS Agent, CICS, Batch, etc. (continued)

Msg Code	Text	Meaning / Instructions
DKJ00293S	IZTA must be initialized on QR TCB. Use tran TADR TAINIT or PLTPI.	
DKJ00295A	Specify character for PC server Named Token	
DKJ00300S	SYSTEM FAILURE: <xxxxxxx> Code=<cccc>/<hhhhhhh>, Reason=<rrrr>/<hhhhhhh>	Contact Technical Support for assistance.
DKJ00300S	Call from <xxxxxxx> Offset <xxxx>	Contact Technical Support for assistance.
DKJ00300S	R0-R7 <xxxxxxx xxxxxxx xxxxxxx xxxxxxx>.	Contact Technical Support for assistance.
DKJ00300S	R8-R15 <xxxxxxx xxxxxxx xxxxxxx xxxxxxx>	Contact Technical Support for assistance.
DKJ00301S	Logic error; Dumpid <xxx> taken to TADUMP.	Contact Technical Support for assistance.
DKJ00301S	Logic error dump taken. Dumpcode is LGIC. Dumpid is <nnnnnnn>.	Contact Technical Support for assistance.
DKJ00302S	System ENQ limit exceeded. Task abended.	Contact Technical Support for assistance.
DKJ00303I	CLOS-VTS <VTSname> Failed. Return Code = <nnnn>	IBM Z Table Accelerator internal processing to end its access to a VTS-TSR failed because the VTS-TSR was canceled before the task terminated. This message is informational only.
DKJ00310W	IZTA TSR <nnn>% Allocated. TSR is local TSR. Owned by Job <jobname>	Setting the TSR_WARNING_FREQ = 0 will cause all the messages to be displayed once the TSR space usage threshold is crossed from below the threshold to above it or exactly equal to it. The default FREQ setting is 1 (i.e. only display one message per second, at the most). (The Default TSR_WARNING_PCT is 85.) For more details refer to the TAOPT parameter settings in the IBM Z Table Accelerator Implementation Guide
DKJ00311W	IZTA TSR <nnn>% Allocated. VTS-TSR name is <vtsname>. Owned by Job <jobname>.	Same as for message DKJ00310W, but applies to VTS-TSR.

Table 7. Messages and error codes for VTS Agent, CICS, Batch, etc. (continued)

Msg Code	Text	Meaning / Instructions
DKJ00330I	DDNAME <ddname> is not a IZTA Release 6 format library.	Not applicable for your version of IBM Z Table Accelerator.
DKJ00331W	LIBRARY FREE SPACE MAP INCONSISTENT WITH LEFT BLOCK COUNT	
DKJ00332I	Caching requested for DDname <ddname>. Caching not permitted on this library. Processing with no cache.	
DKJ00340W	Old form ML mixed with use of extended ML	
DKJ00362E	ABEND DURING UPDATE PROCESS RC=95 SC=<xx>	
DKJ00400E	IZTA User Error Abend <xxxx>, Reason=<dddd/hhhh>	For more information, call Technical Support.
DKJ00400E	IZTA User Error Abend <xxxx>, Command=<xx>,<tablename>	For more information, call Technical Support.
DKJ00400E	IZTA User Error Abend <xxxx>, IZTA Initializing	For more information, call Technical Support.
DKJ00400E	IZTA User Error Abend <xxxx>, Reason=<dddd/hhhh>, Command=<xx>,<tablename>	For more information, call Technical Support.
DKJ00400E	IZTA User Error Abend <xxxx>, Reason=<dddd/hhhh>, IZTA Initializing	Abend code G200 indicates a CICS environment; 0200 indicates otherwise. For more information, call Technical Support.
DKJ00411E	TSR Create failed: Failed by site IEFUSI	Contact Technical Support for assistance.
DKJ00412E	TSR Create failed: MVS resource shortage	Contact Technical Support for assistance.
DKJ00413W	SNAP SUPPRESSED; GETMEM SHORTAGE	Contact Technical Support for assistance.
DKJ00414E	TSR Create failed: LDS Smaller than TSR SIZE	Not applicable for your version of IBM Z Table Accelerator.
DKJ00415E	TSRACCESS=RO using empty LDS	Not applicable for your version of IBM Z Table Accelerator.
DKJ00416I	TSR successfully mapped to LDS	Not applicable for your version of IBM Z Table Accelerator.

Table 7. Messages and error codes for VTS Agent, CICS, Batch, etc. (continued)

Msg Code	Text	Meaning / Instructions
DKJ00417I	TSR successfully saved on LDS	Not applicable for your version of IBM Z Table Accelerator.
DKJ00471E	BDAM ERROR: File=<fffffff>, Function=<xxxxxxxx>	Contact Technical Support for assistance.
DKJ00472E	VSAM ERROR: File=<fffffff>, Function=<xxxxxxxx>, Reg15=<rrrrr>, Error Code=<xxxxx>	Contact Technical Support for assistance.
DKJ00472E	I/O ERROR: <text>	Contact Technical Support for assistance.
DKJ00472E	<xxxx> ERROR: File=<fffffff>, Function=<xxxxxxxx>, Resp=<rrrrr>, Resp2=<ppppp> (Resp2=VSAM Response)	Contact Technical Support for assistance.
DKJ00473E	VSAM ERROR: File=<fffffff>, Request=ATTR2, Retcode=<xxx>, Reason=<yyy>, R15=<zzz> DSNAME=<dsname>	Contact Technical Support for assistance.
DKJ00476E	xxAM ERROR: FUNCTION=<xxxxxxxx>, FILE=<fffffff>	Contact Technical Support for assistance.
DKJ00477E	I/O ERROR: <text>	Contact Technical Support for assistance.
DKJ00477E	VSAM ERROR: FUNCTION=<xxxxxxxx>, REG15=<rrrrr>, ERROR CODE=<xxxxx>, FILE=<fffffff>	Contact Technical Support for assistance.
DKJ00490E	ERROR ON TADEBUG FILE, REASON=<xxxx>, DEBUG DISABLED	Contact Technical Support for assistance.
DKJ00494E	Persistent Task ended before end of step	Contact Technical Support for assistance.
DKJ00495I	IN ROTBVTSD	Information only.
DKJ00496E	Multitasking requires TAOPT Multitasking=Y be set	Reset MULTITASKING parameter in TAOPT
DKJ00497E	failed to delete ESTAE	Contact Technical Support for assistance.
DKJ00501S	EXEC CICS FAILURE: <xxxxxxxx> Resp=<cccc/hhhhhhh>, Resp2=<cccc/hhhhhhh>	Contact Technical Support for assistance.
DKJ00511E	TSR Create failed: Failed by site IEFUSI	Contact Technical Support for assistance.

Table 7. Messages and error codes for VTS Agent, CICS, Batch, etc. (continued)

Msg Code	Text	Meaning / Instructions
DKJ00512E	TSR Create failed: MVS resource shortage	Contact Technical Support for assistance.
DKJ00513W	SNAP suppressed; GETMEM shortage	Contact Technical Support for assistance.
DKJ00530I	No storage for QTCWA; QCT function not used.	Information only.
DKJ00531I	Module DKJTRSTA not found; QCT function not used.	Information only.
DKJ00532I	Attaching QCT x	Information only.
DKJ00533I	QCT X Detached	Information only.
DKJ00550I	IZTA <version> - Resource manager initializing	Information only.
DKJ00551I	IZTA Resource manager initialized	Information only.
DKJ00552E	IZTA Initialization Error	Contact Technical Support for assistance.
DKJ00553I	IZTA Initialization Failed	Information only.
DKJ00554I	IZTA Starting	Information only.
DKJ00555I	IZTA started	Information only.
DKJ00556I	IZTA Deactivating	Information only.
DKJ00557I	IZTA Deactivated	Information only.
DKJ00558E	IZTA Deactivation Error	Contact Technical Support for assistance.
DKJ00559I	IZTA Deactivation Failed	Information only.
DKJ00560I	IZTA Restarting	Information only.
DKJ00561I	IZTA Verifying	Information only.
DKJ00562I	IZTA Verify completed	Information only.
DKJ00563I	IZTA - Error Info: Func(<functionname>/<subfunctionname>) Resp(<responsecode>) Resp2(<responsecode>)	This is error information received by IBM Z Table Accelerator when requesting a CICS service function. Refer to the CICS manuals as required.
DKJ00564I	DKJTCIN used in PLT takes one of TAINIT (default) or INIT.	Information only.

Table 7. Messages and error codes for VTS Agent, CICS, Batch, etc. (continued)

Msg Code	Text	Meaning / Instructions
DKJ00565I	Successful <xxxx> INIT must be done before <xxxx> DKJTCALL	Information only.
DKJ00566I	Verification failure	Information only.
DKJ00567E	The IZTA PC Server must be running for IZTA to initialize.	The IBM Z Table Accelerator PC Server is not running. It must be running in order for IBM Z Table Accelerator to initialize in CICS regions.
DKJ00568E	Program DKJTCIN must be defined as CONCURRENCY(QUASIRENT)	Define DKJTCIN as per message and try again.
DKJ00569E	Module DKJCBASE must be current to determine version.	Contact Technical Support for assistance.
DKJ00570I	<xxxx> takes TAINIT or TERM as a parameter	Information only.
DKJ00571I	Phase in of <xxxxxxxx> not allowed by CICS; current in-storage copy used.	Information only.
DKJ00572I	<xxxxxxxx> LP=<xxxxxxxx> EP=<xxxxxxxx> AM=<xx> RC=<xxxxxxxx> UC=<xxxxxxxx>	Information only.
DKJ00573I	Module <xxxxxxxx> not loaded	Information only.
DKJ00574I	Module <xxxxxxxx> not found.	Information only.
DKJ00575I	EX=<xxxxxxxx> EN=<xxxxxxxx> STST=<xxxx> SHST=<xxxx> TAST=<xxxx> TALN=<xxxx> GALN=<xxxx>	Information only.
DKJ00576I	<xxxxxxxxxxxxxxxxxxx/xxxxxxxx> OK	Information only.
DKJ00577I	<xxxxxxxx> LP=<xxxxxxxx> EP=<xxxxxxxx> LEN=<xxxxxxxx> AM=<xxx> KEY=<x>	Information only.
DKJ00578I	IZTA - Shutdown Completed.	Information only.
DKJ00579I	IZTA - Shutdown Aborted: Internal Resource Anchor Block (RCABLK) not found.	Information only. Contact Technical Support for further information about this message.
DKJ00580I	IZTA - Shutdown Aborted: TRUE DKJTRUE not enabled.	Information only. Contact Technical Support for further information about this message.
DKJ00590E	IZTA ERROR DEBUG WTD, REASON=<xxxx>, DEBUG DISABLED	Contact Technical Support for assistance.

Table 7. Messages and error codes for VTS Agent, CICS, Batch, etc. (continued)

Msg Code	Text	Meaning / Instructions
DKJ00595I	Trans error dump taken. Dumpcode is DKL1. Dumpid is <xxxxxxx>	Information only.
DKJ00596E	CICS Assign failed (Resp/Resp2=<nnnnn/ nnnnn>); dump suppressed	Contact Technical Support for assistance.
DKJ00597E	CICS DUMP TRANSACTION failed (Resp/ Resp2=<nnnnn/nnnnn>). Dump for userid <xxxxxxx> not taken.	Contact Technical Support for assistance.
DKJ00600I	IZTA V<vrm> VTS <xxxxxxx> initialized	Information only.
DKJ00601E	GCA cannot be located; VTS stopping	The IBM Z Table Accelerator PC Server is not running. It must be running before any VTS-TSRs can be brought up.
DKJ00602E	VTS Agent incompatible with GCA structure	Contact Technical Support for assistance.
DKJ00603E	Duplicate VTS attempting to start; auto-stopping	Another VTS with the same name is already running.
DKJ00604E	Program <xxxxxxx> is not APF authorized	The program must be APF authorized to run this task.
DKJ00605E	V<vrm> Environment not ready; VTS stopping	The IBM Z Table Accelerator PC Server is not running. It must be running before any VTS-TSRs can be brought up.
DKJ00606E	Incompatible DKJVBASE; VTS stopping	The version of DKJVBASE is not compatible with this release.
DKJ00608E	Not authorized to start VTS resource: <vtsresourcename>	Not applicable for your version of IBM Z Table Accelerator.
DKJ00610E	Error in <xxxxxxx> macro; RC/REASON=<xxx/ yyy>	Contact Technical Support for assistance.
DKJ00630I	Console interface is available	Information only.
DKJ00631E	Invalid command; command ignored	The command supplied is invalid.
DKJ00632W	'REFRESH not supported in IZTA <vv> VTS	The RF command is not supported in VTS.
DKJ00633I	SHUTDOWN	Information only.
DKJ00634I	STOP command received, V<vrm> VTS <xxxxxxx> terminating	Information only.

Table 7. Messages and error codes for VTS Agent, CICS, Batch, etc. (continued)

Msg Code	Text	Meaning / Instructions
DKJ00636W	switching TSRACCESS to RW, VTS <xxxxxxxx>: <nnnn> users may be affected	Not applicable for your version of IBM Z Table Accelerator.
DKJ00637W	TSRACCESS=RW forced RO for locking- <xxxxxxxx>	Not applicable for your version of IBM Z Table Accelerator.
DKJ00638W	switching TSRACCESS to RO, VTS <xxxxxxxx>: <nnnn> users may be affected	Not applicable for your version of IBM Z Table Accelerator.
DKJ00639E	Not authorized to stop VTS resource: <vtsresourcename>	Not applicable for your version of IBM Z Table Accelerator.
DKJ00701S	IZTA PC Server unavailable, required for license options	Not applicable for your version of IBM Z Table Accelerator.
DKJ00702I	TSR size reset to LDS size	Not applicable for your version of IBM Z Table Accelerator.
DKJ00703S	Invalid LDSTSR dataset	Not applicable for your version of IBM Z Table Accelerator.
DKJ00704W	TSRACCESS forced to RO: LDSTSR DISP=SHR	Not applicable for your version of IBM Z Table Accelerator.
DKJ00705I	MAXNMTAB ignored for TSRACCESS=RO	Not applicable for your version of IBM Z Table Accelerator.
DKJ00706W	SYSPLEX=Y incompatible with TSRACCESS=RW; TSRACCESS set to RO	Not applicable for your version of IBM Z Table Accelerator.
DKJ00707E	TPVM incompatible with COMPAT VTS Agent	Not applicable for your version of IBM Z Table Accelerator.
DKJ00708I	StrobeMethod 3 allowed only for VTS	Information only.
DKJ00709E	StrobeMethod 3 required for VTS with non- zero TSR_KEY-PROTECT	Not applicable for your version of IBM Z Table Accelerator.
DKJ00710W	TSRACCESS forced to RO: TSR_KEY_PROTECT set	Not applicable for your version of IBM Z Table Accelerator.
DKJ00800I	IZTA PC server V<vrm> available	The IBM Z Table Accelerator PC Server is now running. It must be running in order for IBM Z Table Accelerator to initialize in batch jobs, CICS regions, IMS regions and VTS-TSRs.
DKJ00802E	PC server <xxxxxxxx> terminating; see message DKJ00803I for reason.	Contact Technical Support for assistance.
DKJ00803I	<variable length message text>	Information only.

Table 7. Messages and error codes for VTS Agent, CICS, Batch, etc. (continued)

Msg Code	Text	Meaning / Instructions
DKJ00804S	System must support 64-bit architecture	Starting with Version 1 Release 1, IBM Z Table Accelerator requires z/OS running on 64-bit hardware which supports the General Instructions Extension Facility (GIEF), the Long Displacement Facility, the Extended-immediate Facility and the Execute-extensions Facility. These are provided with the Z10-EC processor or higher.
DKJ00805W	License checking error <nnn> for <pppppppp>	Not applicable for your version of IBM Z Table Accelerator.
DKJ00806W	Error <nnn>-<rrr> LOADing DKJTLICK; <pppppppp> not licensed	Not applicable for your version of IBM Z Table Accelerator.
DKJ00807W	Error <nnn> ATTACHing DKJTLICK for <pppppppp>	Not applicable for your version of IBM Z Table Accelerator.
DKJ00811I	Invalid command; ignored: <cccccccc>	Information only.
DKJ00812I	STOP command received	Information only.
DKJ00820W	DKJTPC01/DKJTLICK interface mismatch <nnn/nnn>	Not applicable for your version of IBM Z Table Accelerator.
DKJ00821W	PC Server <vvvv> does not support <pppppppp>	Contact Technical Support for assistance.
DKJ00822W	<pppppppp> not supported by licensing routine	Not applicable for your version of IBM Z Table Accelerator.
DKJ00823W	DKJTLICK not authorized to enter key 0	Not applicable for your version of IBM Z Table Accelerator.
DKJ00824W	PC Server <vvvv> not found	Ensure that the PC Server for this release is running.
DKJ00825W	Licensing subroutine returned RC=<rrrrr>	Not applicable for your version of IBM Z Table Accelerator.
DKJ00826W	<pppppppp> license expired <yyyy-mm-dd>; in grace period	Not applicable for your version of IBM Z Table Accelerator.
DKJ00827W	<pppppppp> license expired <yyyy-mm-dd>	Not applicable for your version of IBM Z Table Accelerator.
DKJ00828W	<pppppppp> license will expire <yyyy-mm-dd>	Not applicable for your version of IBM Z Table Accelerator.
DKJ00829W	<pppppppp> license expires today <yyyy-mm-dd>	Not applicable for your version of IBM Z Table Accelerator.

Table 7. Messages and error codes for VTS Agent, CICS, Batch, etc. (continued)

Msg Code	Text	Meaning / Instructions
DKJ00830W	<pppppppp> license module DKJxBASE missing	Not applicable for your version of IBM Z Table Accelerator.
DKJ00831W	<pppppppp> DKJxBASE <vvvv> mismatch <vvvv>	The license module specified for the named product is for a different version.
DKJ00832W	<pppppppp> DKJxBASE return code <rrrr>	Contact Technical Support for assistance.
DKJ00833I	<pppppppp> licensed to <customer>	Not applicable for your version of IBM Z Table Accelerator.
DKJ00844S	Customer Anchor Table slot invalid: RC=<r>.	Contact Technical Support for assistance.
DKJ00847S	cannot start, too many DKJ products	Contact Technical Support for assistance.
DKJ00848S	TPM may be running, run DKJPTERM	Not applicable for your version of IBM Z Table Accelerator.
DKJ00849S	Replacing corrupted GCA	Contact Technical Support for assistance.
DKJ00850I	deleting NT	Information only.
DKJ00851I	NT deleted	Information only.
DKJ00852I	freeing PC code	Information only.
DKJ00853I	PC code freed	Information only.
DKJ00854I	Freeing GCA	Information only.
DKJ00855I	Freeing GCAX	Information only.
DKJ00856I	GCA purged	Information only.
DKJ00857I	Replacing corrupted Dummy TPM block	Not applicable for your version of IBM Z Table Accelerator.
DKJ00858S	PC Server unable to obtain storage	Contact Technical Support for assistance.
DKJ00859I	Completing Dummy TPM/TPV blocks	Not applicable for your version of IBM Z Table Accelerator.
DKJ00870W	LXFRE Bad Return Code	Contact Technical Support for assistance.

Table 7. Messages and error codes for VTS Agent, CICS, Batch, etc. (continued)

Msg Code	Text	Meaning / Instructions
DKJ00871S	\$DOPC GETMAIN Failed	Contact Technical Support for assistance.
DKJ00872E	LXFREE Getmain Failed	Contact Technical Support for assistance.
DKJ00873W	LXFREE Freemain Failed	Contact Technical Support for assistance.
DKJ00874S	\$DOPC SCHEDIRB Macro Failed	Contact Technical Support for assistance.
DKJ00875S	\$DOPC DKJTPC11 Routine Failed	Contact Technical Support for assistance.
DKJ00876S	\$DOPC LOADROTB Service Failed	Contact Technical Support for assistance.
DKJ00877I	DKJTVROT Named Token Failed	Information only.
DKJ00900S	Not licensed for VTS Manager options	Not applicable for your version of IBM Z Table Accelerator.
DKJ00901S	Not licensed for VTS Gate options	Not applicable for your version of IBM Z Table Accelerator.
DKJ00902S	Not licensed for DEFAULT_TSR options	Not applicable for your version of IBM Z Table Accelerator.
DKJ00990S	Incompatible ROOT and NUCLEUS versions	Contact Technical Support for assistance.
DKJ00991E	Parms module not loaded	Contact Technical Support for assistance.
DKJ00993E	Parms module corrupted	Contact Technical Support for assistance.

Note: Abnormal terminations between 99 and 1000 or over 1099, are IBM Z Table Accelerator errors and should be brought to the immediate attention of your IBM Z Table Accelerator Administrator.

DKJTEXEC error messages

Error and audit messages generated by the batch utility program DKJTEXEC are identified and described in alphabetical order in the following table.

Table 8. IBM Z Table Accelerator batch utility messages

Message	Text	Meaning/Instructions
Change definition failed	The attempt to change the definition of the table has failed for the reasons noted above the message.	
Change is not for Alternate Index	The CHANGE ALT= command specified a table that is not an Alternate Index. Use the CHANGE TBL= version of the command.	
Change key size/location failed	The attempt to change the key size and/or key location has failed for reasons noted in preceding messages or to the right of this message.	
Change maxgen failed	The attempt to change the number of generations to be kept has failed for reasons noted in preceding messages or to the right of this message.	
Change not made for reasons above	None of the changes in the above change were performed. The reason for the rejection of this command is noted above the message.	
Change successful	The table definition has been changed as requested.	Information only.
CMD requires specified table to be open	This command requires the specified table to be open. This error should not occur when using DKJTEXEC.	
Command is invalid	The specified command is not a DKJTEXEC command.	
Command is invalid for an Alternate Index	The command cannot be performed on an Alternate Index.	
Command rejected for reason above	The preceding DKJTEXEC Command was rejected for the reasons noted above the message.	
Command requires more parms than given	DKJTEXEC called IBM Z Table Accelerator incorrectly.	

Table 8. IBM Z Table Accelerator batch utility messages (continued)

Message	Text	Meaning/Instructions
Copy complete	The requested operation completed successfully.	
Copy complete except where indicated	The requested copy operation was successfully completed except as noted above the message.	
Copy failed	The requested operation failed. Diagnostic messages are issued to give the reason(s) for the failure.	
Copy of table (all generations) complete	The requested copy operation of all generations of a table was successfully completed.	
Create Alternate Index definition failed	The attempt to create an Alternate Index definition failed for the reasons noted in preceding messages or to the right of this message.	
Create Alternate Index definition successful	The Alternate Index definition has been created successfully.	Information only.
Data table index must be defined as P	DKJTEXEC attempted to invoke an Alternate Index of a table that is not defined to be type P (indexed or pointer).	
Defaults have been set	The defaults entered have been set.	
DEFINE TEMPORARY TABLE FAILED	Internal error defining temporary table for <filename>.	Please increase TSR size and rerun.
Delete failed	The requested delete operation has failed for reasons noted in preceding messages or to the right of this message.	
Delimiter is invalid, = was expected	The next symbol after a keyword must be an equal sign.	
Density parameter must be from 1-999	The density must be numeric and in the range 1-999.	
Destination library too small	The destination library is too small to hold the table(s) to be copied.	

Table 8. IBM Z Table Accelerator batch utility messages (continued)

Message	Text	Meaning/Instructions
Directory is empty	The directory on the specified library is empty.	
DT command (maxgen parm) must be 1-9	The maximum generations to be kept must be numeric in the range 1-9 inclusive.	
Duplicate keyword for this command	The same keyword was specified twice for a command. Fix the command sequence and try again.	Fix the command sequence and try again.
End of data - IBM Z Table Accelerator utility ended	End of data has been reached on CNTLCARD, the DKJTEXEC input file.	Information only.
Error In table definition	An error has been detected in the table definition created from the parameters supplied. Refer to the Error Code for the reason.	
Estimated number-of-rows is out of range	The estimated number of rows is non-numeric, or is too large.	
Expand library failed	The requested operation failed. Diagnostic messages are issued to give the reason(s) for the failure.	
Expand library successful	The requested operation completed successfully.	Information only.
Expansion factor must be from 1 To 999	The expansion factor must be numeric and in the range 1-999.	
Field is greater than 8 characters	All of the keywords and keyword values have a maximum length of eight characters.	Ensure field is 8 characters or less.
Generation number specified is invalid	The generation must be numeric and in the range 0-255.	
Generation requested has been copied	The COPY operation has been completed successfully. A new table (generation number 1) has been added to the target library.	Information only.

Table 8. IBM Z Table Accelerator batch utility messages (continued)

Message	Text	Meaning/Instructions
Index parameter must be P, T or blank	The Index field must be T for true tables, P for indexed (Pointer) tables or blank. However, all tables are treated as Pointer tables internally, so if T is used, it will be equivalent to P.	
Initialization successful	The new IBM Z Table Accelerator library has been initialized successfully.	Information only.
Insufficient local tableSPACE region	The TSR is not large enough to contain the table and IBM Z Table Accelerator internal tables. IBM Z Table Accelerator uses a small part of the TSR for its own internal tables.	
Insufficient space on new library	There is not enough space on the target library for the tables being copied.	
INV storage-mode-code: Must be R, Blank	The storage mode code (SMC) must be R or blank.	
Key size invalid: must be 1-256	The key size must be a number from 1-256, inclusive.	
Keyword and value are incompatible	If the keyword requires a numeric value, the value specified with it must be numeric. Alternatively, the ALLGEN keyword must be followed by = YES.	
Keyword invalid for this command	The keyword specified was not recognized.	
Keyword is incomplete	The keyword has not been followed by an equal sign (=) and a value.	
Keyword is invalid	The keyword specified is not used with this command.	
KLOC invalid: must be 1 to row-size	The key location must be numeric in the range 1 to the row size, inclusive.	

Table 8. IBM Z Table Accelerator batch utility messages (continued)

Message	Text	Meaning/Instructions								
Library is not empty	The target or destination library of an EXPAND LIBRARY request contains at least one table. The target library must be empty for an EXPAND operation.									
Library not initialized	For the reason cited to the right of this message, the library was not initialized.									
Load failed	The requested operation failed. Diagnostic messages are issued to give the reason(s) for the failure.									
Mthd/org parameters are incompatible	This combination of search method and organization will not work. Valid combinations are: <div style="border: 1px solid black; padding: 5px; margin: 10px 0;"> <table style="width: 100%; border-collapse: collapse;"> <thead> <tr> <th style="text-align: left; border-bottom: 1px dashed black;">Organization</th> <th style="text-align: left; border-bottom: 1px dashed black;">Method</th> </tr> </thead> <tbody> <tr> <td>R, U</td> <td>S</td> </tr> <tr> <td>S, D</td> <td>S, B, C</td> </tr> <tr> <td>H</td> <td>H</td> </tr> </tbody> </table> </div>	Organization	Method	R, U	S	S, D	S, B, C	H	H	
Organization	Method									
R, U	S									
S, D	S, B, C									
H	H									
New generation loaded successfully	A new generation of this table has been created from the contents of the FROM dataset.	Information only.								
Newlib DDNAME not Assigned -- check JCL	The requested initiation of the DDNAME specified by the NEWLIB parameter could not be performed since the DDNAME is not defined in the JCL.	Alter JCL as required.								
nnnnn/mmmmm tables exported successfully	nnnnn tables were exported from the IBM Z Table Accelerator library successfully. mmmmm tables were requested to be exported.	Information only.								
nnnnn/mmmmm tables imported successfully	nnnnn tables were imported to the IBM Z Table Accelerator library successfully. mmmmm tables were requested to be imported.	Information only.								
No changes specified	A change command did not identify any fields to be changed.	Identify fields as required.								

Table 8. IBM Z Table Accelerator batch utility messages (continued)

Message	Text	Meaning/Instructions
Not copied, new name same as old name	A COPY table request specified the same value for NEWNAME as for TBL.	
Not enough space on library	There is not enough space on the IBM Z Table Accelerator library for the table being defined, copied, loaded, or expanded.	
Old Alternate Index definition deleted	During a copy of an Alternate Index definition, the old definition was deleted from the target library, but the new definition could not be copied from the source library to the target library.	
On dest lib - insufficient space for copy	The copy operation has been requested to a destination library with insufficient space on it to receive all generations of the tables on the FROM library.	
Org parm invalid: must be R, U, S, D or H	The Organization must be R, U, S, D, or H.	
Paged tables are not supported	Paged tables are no longer supported.	Not applicable for your version of IBM Z Table Accelerator.
Print request completed successfully	The requested print operation has finished with no errors. The listing can be found in TARPT.	Information only.
Rename failed	The requested operation failed. Diagnostic messages are issued to give the reason(s) for the failure.	
Rename successful	The requested rename operation was successfully performed.	Information only.
Requested generation has been cleared	A new generation of this table has been created using the definition of the generation specified. This new generation contains no items.	Information only.
Requested generation has been deleted	The generation of the table specified has been deleted.	Information only.
Requested generation(s) deleted	The requested operation completed successfully.	Information only.

Table 8. IBM Z Table Accelerator batch utility messages (continued)

Message	Text	Meaning/Instructions								
Required keyword missing for this command	One of the required keywords for this command was not contained in the command sequence.	Correct command sequence as required.								
RN fails. new name Already Exists	The table was not renamed because a table with the new name already exists on the library.									
Row size invalid: must be 1-32767	The row size must be numeric in the range 1-32767 inclusive.									
Search-method incompatible with organization	<p>The search method must be S, Q, B, C, or H and must match the organization. Valid combinations are:</p> <table border="1" data-bbox="646 793 1052 968"> <thead> <tr> <th>Organization</th> <th>Method</th> </tr> </thead> <tbody> <tr> <td>R, U</td> <td>S</td> </tr> <tr> <td>S, D</td> <td>S, B, C</td> </tr> <tr> <td>H</td> <td>H</td> </tr> </tbody> </table>	Organization	Method	R, U	S	S, D	S, B, C	H	H	
Organization	Method									
R, U	S									
S, D	S, B, C									
H	H									
Semi-colon is missing from command	The command and its keywords must be followed by a semicolon (;) to indicate the end of the keyword list for the command sequence.									
Source library is empty	The source library in a copy operation contains no tables.									
Table already exists on library	The table could not be defined on, or copied to, the new library because a table of the same name already exists on the target library.									
IZTA error detected	The requested operation failed. Diagnostic messages are issued to give the reason(s) for the failure.									
IZTA internal program error - xxxx	The requested operation failed. Diagnostic messages are issued to give the reason(s) for the failure.									
Table definition successful	The table was defined successfully and generation 1 was stored on the library.									

Table 8. IBM Z Table Accelerator batch utility messages (continued)

Message	Text	Meaning/Instructions
Table is not closed	This error should not occur when using DKJTEEXEC.	Contact Technical Support.
Table is not found	The table was not found on the given library.	
Table name specified is invalid	A valid IBM Z Table Accelerator table name is a string of 8 bytes that are not all blanks, all low values, all high values, or :TMPNAME.	
Table opened for read cannot be stored	A table has been opened for read-only access and a store or write command has been issued.	
Table type is invalid	The table type must be F, V, A, or X.	
Table unavailable. No wait in effect	DKJTEEXEC could not store the table because it is locked by another application. The NO WAIT parameter is in effect.	
Table unloaded successfully	The specified generation of the specified table has been unloaded to the TO dataset.	Information only.
Table updated as requested	The requested update operations on the specified table were performed successfully.	Information only.
Table updated successfully	The requested operation completed successfully.	Information only.
Table xxxxxxxx could not be opened	For reasons noted in preceding messages or to the right of this message, the table xxxxxxxx could not be opened. Since the table must be opened to perform this DKJTEEXEC command, the DKJTEEXEC command could not be successfully completed.	
Table xxxxxxxx export failed	The export of table xxxxxxxx failed for the reasons given in other messages.	
Table xxxxxxxx import failed	The import of table xxxxxxxx failed for the reasons given in other messages.	

Table 8. IBM Z Table Accelerator batch utility messages (continued)

Message	Text	Meaning/Instructions
Table xxxxxxxx import replace failed	The import of table xxxxxxxx failed for the reasons given in other messages.	
Table xxxxxxxx sort error	The sort required for this operation failed for the reason noted.	
Table xxxxxxxx was not stored	The table xxxxxxxx was not stored for the reasons noted in preceding messages or to the right of this message.	
IZTA error detected	A DKJTCALL error was detected in the processing of this command. The explanation of the error is found to the right of this message.	
The count specified is invalid	The count is less than one, or greater than the number of rows in the table.	
The key will not fit within the row	The end of the key cannot exceed the end of the row.	
The library DDNAME does not exist	The LIB=DDNAME is not in the JCL.	Alter JCL as required.
The library status is invalid	IBM Z Table Accelerator library disposition must be NEW for new libraries, and SHR or OLD for existing libraries.	
The password supplied is invalid	An incorrect password has been specified.	
The specified command is invalid	The specified command is not a valid DKJTEXEC command.	
The specified library is not suitable	There are a number of conditions that can create this error. See IBM Z Table Accelerator error codes.	
Write password is missing or incorrect	The operation requires access to a write-protected table, but the write password has not been specified or specified incorrectly.	

Table 8. IBM Z Table Accelerator batch utility messages (continued)

Message	Text	Meaning/Instructions
xxxxxxx field is greater than 8 characters	All keywords and keyword values must be eight characters or fewer.	
<tablename> NOT FOUND IN SOURCE LIBRARY	Specified table is not in the source library.	Please specify correct table associated with EXCLUDE parameter and/or EXCLFILE parameter file.

Appendix C. PC Server implementation in a CICS environment

Implementation and Restart/Recovery Considerations

The IBM Z Table Accelerator PC Server provides support for running CICS threadsafe transactions which access IBM Z Table Accelerator.

In order to prevent IBM Z Table Accelerator application problems, IBM Z Table Accelerator will not initialize in a CICS region if a PC Server at the same maintenance level is not running when IBM Z Table Accelerator is first initialized in a CICS region.

It is intended that the IBM Z Table Accelerator PC Server be started at IPL of an LPAR and never stopped. If it is stopped, it should be immediately restarted. Since the PC Server does not execute any code except during its initialization or termination, it will not abend. However, operations or other personnel authorized to stop or cancel jobs on a z/OS system could stop it.

Once the IBM Z Table Accelerator PC Server is successfully started, the PC-routine module is resident in memory. This module, which is used by IBM Z Table Accelerator transactions, remains in memory and is not released when the PC Server goes down. Therefore, IBM Z Table Accelerator transactions running in regions on that LPAR will continue to process IBM Z Table Accelerator transactions without problems. (They do not access the PC Server directly; they use the PC-routine module in memory.)

However, if these CICS regions are recycled without the PC Server running, the in-memory PC-routine module will not exist, and therefore IBM Z Table Accelerator will not initialize in these regions and no IBM Z Table Accelerator transactions will work. The message displayed in these regions will be:

```
DKJ00567E The IZTA PC Server must be running for IZTA to initialize.
```

If the IBM Z Table Accelerator PC Server is successfully restarted before the CICS regions are recycled, the in-memory PC-routine will be generated, and any newly recycled CICS regions will access the new copy of the in-memory PC-routine. IBM Z Table Accelerator will initialize successfully and IBM Z Table Accelerator transactions will work. The existing CICS regions will continue to access the old copy of the in-memory PC-routine and will continue to work.

RESTART/RECOVERY Considerations

A successful start of the PC Server will display the following message on the system console:

```
DKJ00800I IZTA PC Server Vvrm available
```

If the PC Server is stopped or cancelled by the operator while jobsteps, including CICS regions, are accessing IBM Z Table Accelerator, the jobs will continue to run and access IBM Z Table Accelerator. However, new jobsteps will not be able to start accessing IBM Z Table Accelerator until the PC Server is restarted.

To prevent operational problems related to the PC Server not running, whether after an IPL or due to the PC Server step being stopped or cancelled, the following procedures should be considered:

1. The PC Server should be started automatically at IPL before any batch jobs, CICS regions or IMS MPRs that access IBM Z Table Accelerator are started.
2. The appearance of message DKJ00800I could be used to trigger the start of any jobs that access IBM Z Table Accelerator.
3. Any current facility which monitors availability of critical subsystems should include monitoring the PC Server jobname.
4. If the PC Server is not running at any given time, a restart of the job should be automatically triggered.

5. If the monitor does not trigger it, the appearance of message DKJ00201 or DKJ00567 should be set up to trigger it.
6. There is no danger in attempting to start the PC Server multiple times; subsequent PC Server tasks will end with a message indicating that the PC Server is running (for message details, see [Appendix B, “IBM Z Table Accelerator messages,”](#) on page 45).
7. If CICS regions have started while the PC Server is not running, IBM Z Table Accelerator will fail to initialize; a message will be returned indicating this condition (for message details, see [Appendix B, “IBM Z Table Accelerator messages,”](#) on page 45).

IBM Z Table Accelerator will not be available in these regions even if the PC Server is subsequently started. After the PC Server is started, this message could be used to trigger recycling these regions or, alternatively invoking the TAST TAINIT transaction in the affected regions.

NOTES and CAUTIONS

- The PC Server PC routines are loaded into memory when it is started. They do not overlay any previous IBM Z Table Accelerator PC routines. Although the routines are quite small, if the PC Server is stopped and restarted thousands of time, they could theoretically finally occupy all memory available.
- Only one copy of the PC Server for any IBM Z Table Accelerator release can be running on an LPAR at any time (whether it is initiated as a started task OR as a batch job).

Notices

This information was developed for products and services offered in the U.S.A.

IBM® may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, NY 10504-1785
U.S.A.

For license inquiries regarding double-byte (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

Intellectual Property Licensing
Legal and Intellectual Property Law
IBM Japan, Ltd.
19-21, Nihonbashi-Hakozakicho, Chuo-ku
Tokyo 103-8510, Japan

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law:

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement might not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Licenseses of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

IBM Corporation
224A/101
11400 Burnet Road

Austin, TX 78758
U.S.A.

Such information may be available, subject to appropriate terms and conditions, including in some cases payment of a fee.

The licensed program described in this document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement or any equivalent agreement between us.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. You may copy, modify, and distribute these sample programs in any form without payment to IBM for the purposes of developing, using, marketing, or distributing application programs conforming to IBM's application programming interfaces.

Trademarks

IBM, the IBM logo, and [ibm.com](http://www.ibm.com)[®] are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the Web at "Copyright and trademark information" at <http://www.ibm.com/legal/copytrade.shtml>.

Microsoft and Windows are trademarks of Microsoft Corporation in the United States, other countries, or both.

Other product and service names might be trademarks of IBM or other companies.



SC28-3229-00

