

CICS® Transaction Server for VSE/ESA™



Application Programming Guide

Release 1

CICS® Transaction Server for VSE/ESA™



Application Programming Guide

Release 1

Note!

Before using this information and the product it supports, be sure to read the general information under "Notices" on page 475.

First Edition (June 1999)

This edition applies to Release 1 of CICS Transaction Server for VSE/ESA, program number 5648-054, and to all subsequent versions, releases, and modifications until otherwise indicated in new editions. Make sure you are using the correct edition for the level of the product.

The CICS for VSE/ESA Version 2.3 edition remains applicable and current for users of CICS for VSE/ESA Version 2.3.

Order publications through your IBM representative or the IBM branch office serving your locality.

At the back of this publication is a page entitled "Sending your comments to IBM". If you want to make any comments, please use one of the methods described there.

© **Copyright International Business Machines Corporation 1989, 1999. All rights reserved.**

US Government Users Restricted Rights – Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Contents

Preface	ix
What this book is about	ix
Who should read this book	ix
What you need to know to understand this book	ix
How to use this book	ix
Notes on terminology	ix
What is not covered in this book	ix
Programming interface information	x
Notes on terminology	xi
Road map	xiii

Part 1. Getting started 1

Chapter 1. Preparing your application to run	3
Writing CICS programs	3
Preparing your program	3
The translation process	5
Specifying translator options	6
Translator options	7
Preparing BMS maps	16
Using VSE macros from within an application	18
Chapter 2. Language considerations	19
COBOL considerations	19
C considerations	38
PL/I considerations	44
Assembler-language and HLASM considerations	46

Part 2. Application design 49

Chapter 3. Designing efficient applications	51
Program structure	51
General programming techniques	55
Other suggestions for improving performance	62
Storing data within a transaction	64
Lengths of areas passed to CICS commands	67
EXEC interface stubs	68
Data operations	71
Terminal operations	75
Chapter 4. Sharing data across transactions	83
Common work area (CWA)	83
TCTTE user area (TCTUA)	87
COMMAREA in EXEC CICS RETURN commands	88
Display screen	89
Chapter 5. Affinity among transactions	91
What is transaction affinity?	91
Programming techniques used by CICS application programs	92

Safe programming techniques	94
Unsafe programming techniques	97
Suspect programming techniques	103
Detecting intertransaction affinities	112
Duration and scope of intertransaction affinities	113
Recommendations	117
Chapter 6. Intercommunication considerations	119
Design considerations	119
Transaction routing	120
Function shipping	120
Distributed program link (DPL)	121
Asynchronous processing	132
Distributed transaction processing (DTP)	132
Common Programming Interface (CPI) Communications	132
External CICS interface (EXCI)	133
Chapter 7. Recovery considerations	135
Journaling	135
Syncpointing	137
Chapter 8. Minimizing errors	139
Protecting CICS from application errors	139
Testing applications	139
Chapter 9. Dealing with exception conditions	141
Default CICS exception handling	141
Handling exception conditions by in-line code	142
Modifying the default CICS exception handling	146
Chapter 10. Access to system information	153
System programming commands	153
EXEC interface block (EIB)	153
Chapter 11. Abnormal termination recovery	155
The EXEC CICS HANDLE ABEND command	155
Restrictions on retrying operations	157
Exception handling	158
Trace	161
Monitoring	162
Dumping	163

Part 3. Files and databases 165

Chapter 12. An overview of file control	167
VSAM data sets	167
DAM data sets	169
CICS shared data tables	170
Accessing data sets from CICS application programs	172
Reading records	172
Updating records with the READ command	176
Deleting records	177
Adding records	178

Review of file control command options	179
Avoiding transaction deadlocks	181
KEYLENGTH option for remote data sets	183
Chapter 13. File control—VSAM considerations	185
Record identification	185
CICS locking of VSAM records in recoverable files	186
Chapter 14. File control—DAM considerations	189
Record identification	189
Updating records from DAM data sets	190
Browsing records from DAM data sets	190
Adding records to DAM data sets	191
DAM exclusive control	192
Chapter 15. Database control	193
DL/I databases	193
DATABASE 2 (DB2) databases	193

Part 4. Data communication 195

Chapter 16. Introduction to data communication	199
Basic CICS terms	199
How tasks get started	200
Which transaction?	201
CICS APIs for terminals	203
Topics elsewhere in this book	204
Where to find more information	205
Chapter 17. The 3270 family of terminals	207
Introduction	207
Background	208
The 3270 buffer	210
3270 display data: defining 3270 fields	212
Input from a 3270 terminal	221
Unformatted mode	224
Chapter 18. Basic mapping support	225
Introduction to BMS	225
BMS support levels	226
A BMS output example	228
Creating the map	231
Rules for writing BMS macros	235
Assembling the map	237
Complex fields	239
Sending mapped output: basics	241
Initializing the output map	244
Moving the variable data to the map	244
Setting the display characteristics	244
Control options on the EXEC CICS SEND MAP command	247
Options for merging the symbolic and physical maps	248
Summary: what appears on the screen	249
Using GDDM® and BMS	252

Positioning the cursor	252
Sending invalid data and other errors	253
Receiving data from a display	253
An input-output example	254
The symbolic input map	256
Programming simple mapped input	257
The EXEC CICS RECEIVE MAP command	257
Getting storage for mapped input: INTO and SET	258
Reading from a formatted screen: what comes in	259
Other information from EXEC CICS RECEIVE MAP	260
Support for non-3270 terminals	267
The MAPPINGDEV facility.	272
Sending mapped output: additional facilities	276
Terminal operator paging: the CSPG transaction	280
Page formation: the ACCUM option	282
Formatting text output	288
Message routing: the ROUTE command	294
Using SET	302
Partition support	304
Logical device components	312
BMS support for other special hardware	314
Chapter 19. Terminal control	319
Access method support	319
Terminal control commands	320
VTAM considerations	333
Sequential terminal support	336
Batch data interchange	338
Chapter 20. CICS support for printing	341
Formatting for CICS printers	341
CICS printers: getting the data to the printer	349
Non-CICS printers	352
Printing display screens	354
Chapter 21. CICS interface to POWER	357
CICS interface to VSE/POWER	357
Programming note for spool commands	362

Chapter 22. Report Controller	363
Reports	364
<hr/>	
Part 5. CICS management functions	367
Chapter 23. Interval control	369
Expiration times	369
Request identifiers	371
Chapter 24. Task control	373
Controlling sequence of access to resources	374
Chapter 25. Program control	377
Application program logical levels	378
Link to another program expecting return	378
Passing data to other programs	379
Chapter 26. Storage control and storage protection	391
CICS storage control	391
CICS storage protection	392
CICS-key and user-key storage	392
Deciding what execution and storage key to specify	397
C programs and CICS storage	401
Chapter 27. Transient data control	403
Introduction	403
Intrapartition queues	403
Extrapartition destinations	404
Indirect destinations	404
Automatic transaction initiation (ATI)	405
Chapter 28. Temporary storage control	407
Temporary storage queues	407
Typical uses of temporary storage control	408
Chapter 29. Security control	409
QUERY SECURITY command	409
Non-Terminal transaction security	410
<hr/>	
Part 6. Testing applications	413
Chapter 30. Testing applications: the process	415
Preparing the application and system table entries	415
Preparing the system for debugging	415
Single-thread testing	416
Multithread testing	416
Regression testing	416
Year 2000 testing	417
Chapter 31. CICS Execution diagnostic facility (EDF)	419
Getting started	419
Where does EDF intercept the program?	419

What does EDF display?	420
How to use EDF	433
Security considerations	440
Chapter 32. Temporary storage browse (CEBR)	441
How to use the CEBR transaction	441
What does the CEBR transaction display?	442
The CEBR commands	444
Using the CEBR transaction with transient data	446
Security considerations	447
Chapter 33. Command-level interpreter (CECI)	449
How to use CECI	449
What does CECI display?	450
Additional displays	455
Making changes	459
How CECI runs	460
Security considerations	462

Part 7. Appendix	463
Appendix. CICS commands and their equivalent obsolete macros	465
Bibliography	469
Books from VSE/ESA 2.4 base program libraries	470
Books from VSE/ESA 2.4 optional program libraries	472
Notices	475
Trademarks and service marks	476
Index	477

Preface

What this book is about

This book gives **guidance** about the CICS for VSE application programming interface; it complements the **reference** information in the *CICS Application Programming Reference* manual. For guidance information on debugging CICS applications, see the *CICS Problem Determination Guide*.

Who should read this book

This book is mainly for experienced application programmers. Those who are relatively new to CICS should be able to understand it. If you are a system programmer or system analyst, you will still find it useful.

What you need to know to understand this book

You must be able to program in COBOL, C, PL/I, or assembler language. You need to know about CICS application programming at the *CICS Application Programming Primer (VS COBOL II)* level.

How to use this book

Read the parts covering what you need to know. (Each part has a full table of contents to help you find what you want.) The book is a guide, not a reference manual. On your first reading, it will probably help to work through any one part of it from start to finish.

Notes on terminology

API refers to the CICS command-level application programming interface unless otherwise stated.

ASM is sometimes used as the abbreviation for assembler language.

What is not covered in this book

Guidance for usage of the CICS for VSE Front End Programming Interface is not discussed in this book. See the *CICS Front End Programming Interface User's Guide* for background information about FEPI design considerations and programming information about its API.

Programming interface information

This book documents General-use Programming Interface and Associated Guidance Information and Product-sensitive Programming Interface and Associated Guidance Information provided by CICS.

General-use programming interfaces allow the customer to write programs that obtain the services of CICS.

Product-sensitive programming interfaces allow the customer installation to perform tasks such as diagnosing, modifying, monitoring, repairing, tailoring, or tuning of CICS. Use of such interfaces creates dependencies on the detailed design or implementation of the IBM software product. Product-sensitive programming interfaces should be used only for these specialized purposes. Because of their dependencies on detailed design and implementation, it is to be expected that programs written to such interfaces may need to be changed in order to run with new product releases or versions, or as a result of service.

Product-sensitive Programming Interface and Associated Guidance Information is identified where it occurs by an introductory statement to a chapter or section.

Notes on terminology

The terms listed in Table 1 are commonly used in the CICS Transaction Server for VSE/ESA Release 1 library. See the *CICS Glossary* for a comprehensive definition of terminology.

Term	Definition (and abbreviation if appropriate)
\$(the dollar symbol)	In the character sets and programming examples given in this book, the dollar symbol (\$) is used as a national currency symbol and is assumed to be assigned the EBCDIC code point X'5B'. In some countries a different currency symbol, for example the pound symbol (£), or the yen symbol (¥), is assigned the same EBCDIC code point. In these countries, the appropriate currency symbol should be used instead of the dollar symbol.
BSM	BSM is used to indicate the basic security management supplied as part of the VSE/ESA product. It is RACROUTE-compliant, and provides the following functions: <ul style="list-style-type: none">• Signon security• Transaction attach security
C	The C programming language
CICSplex	A CICSplex consists of two or more regions that are linked using CICS intercommunication facilities. Typically, a CICSplex has at least one terminal-owning region (TOR), more than one application-owning region (AOR), and may have one or more regions that own the resources accessed by the AORs
CICS Data Management Facility	The new CICS Transaction Server for VSE/ESA Release 1 facility to which all statistics and monitoring data is written, generally referred to as "DMF"
CICS/VSE	The CICS product running under the VSE/ESA operating system, frequently referred to as simply "CICS"
COBOL	The COBOL programming language
DB2 for VSE/ESA	Database 2 for VSE/ESA which was previously known as "SQL/DS".

Table 1 (Page 2 of 2). Commonly used words and abbreviations in CICS Transaction Server for VSE/ESA Release 1

Term	Definition (and abbreviation if appropriate)
ESM	<p>ESM is used to indicate a RACROUTE-compliant external security manager that supports some or all of the following functions:</p> <ul style="list-style-type: none"> • Signon security • Transaction attach security • Resource security • Command security • Non-terminal security • Surrogate user security • MRO/ISC security (MRO, LU6.1 or LU6.2) • FEPI security.
FOR (file-owning region)—also known as a DOR (data-owning region)	A CICS region whose primary purpose is to manage VSAM and DAM files, and VSAM data tables, through function provided by the CICS file control program.
IBM C for VSE/ESA	The Language Environment version of the C programming language compiler. Generally referred to as “C/VSE”.
IBM COBOL for VSE/ESA	The Language Environment version of the COBOL programming language compiler. Generally referred to as “COBOL/VSE”.
IBM PL/I for VSE/ESA	The Language Environment version of the PL/I programming language compiler. Generally referred to as “PL/I VSE”.
IBM Language Environment for VSE/ESA	The common runtime interface for all LE-conforming languages. Generally referred to as “LE/VSE”.
PL/I	The PL/I programming language
VSE/POWER	Priority Output Writers Execution processors and input Readers. The VSE/ESA spooling subsystem which is exploited by the report controller.
VSE/ESA System Authorization Facility	The new VSE facility which enables the new security mechanisms in CICS TS for VSE/ESA R1, generally referred to as “SAF”
VSE/ESA Central Functions component	The new name for the VSE Advanced Function (AF) component
VSE/VTAM	“VTAM”

Road map

<i>Table 2. Getting started road map</i>	
If you want to...	Refer to...
Prepare your application to run, and find out the considerations and restrictions that apply when using various programming languages	Part 1, "Getting started" on page 1
Design efficient applications, and consider topics such as: <ul style="list-style-type: none"> • Sharing data across transactions • Transaction affinity • Writing applications that communicate with other CICS systems • Using journaling or syncpointing to help recovery or reconstruction • Making your applications error-free • Dealing with exception conditions • Accessing information that is outside the local environment of your application program • Using the program-level abend exit facility for abnormal termination recovery 	Part 2, "Application design" on page 49
Gain an overview of CICS file control, review VSAM and DAM considerations, and obtain a brief insight into database control	Part 3, "Files and databases" on page 165
Review the CICS facilities for communicating with the terminals through which users access CICS applications (data communication)	Part 4, "Data communication" on page 195
Review the groups of CICS operations that control the processing of tasks within a CICS system: interval control, task control, program control, storage control, transient data control, and temporary storage control	Part 5, "CICS management functions" on page 367
Obtain an insight into testing applications, using the execution diagnostic facility (EDF), temporary storage browse (CEBR), and command-level interpreter (CECI) transactions	Part 6, "Testing applications" on page 413
Identify CICS commands and their equivalent obsolete macros	Appendix, "CICS commands and their equivalent obsolete macros" on page 465

Part 1. Getting started

Chapter 1. Preparing your application to run	3	Global variables	36
Writing CICS programs	3	Comma and semicolon as delimiters	36
Preparing your program	3	Symbolic character definition	36
The translation process	5	Summary of restrictions	37
Specifying translator options	6	COBOL2 translator option	37
Translator options	7	COBOL3 translator option	38
Preparing BMS maps	16	C considerations	38
Using VSE macros from within an application	18	Data declarations needed for C	39
Chapter 2. Language considerations	19	Naming EIB fields	39
COBOL considerations	19	Restrictions	40
COBOL/VSE with the ANSI85 COBOL standards	30	PL/I considerations	44
Literals intervening in blank lines	30	Compilers supported	44
Sequence numbers containing any character	30	Restrictions	44
Lowercase characters supported in all COBOL words	31	OPTIONS(MAIN) specification	45
REPLACE statement	31	Program segments	46
Batch compilation	31	PL/I and dynamic storage	46
Nested programs	32	Assembler-language and HLASM considerations	46
Reference modification	36	Commands contained within macros and COPY code	47
		Invoking assembler-language application programs with a call	47

Chapter 1. Preparing your application to run

This chapter describes what you need to do to make your application run in a CICS environment. It contains the following topics:

- “Writing CICS programs”
- “Preparing your program”
- “The translation process” on page 5
- “Specifying translator options” on page 6
- “Translator options” on page 7
- “Preparing BMS maps” on page 16
- “Using VSE macros from within an application” on page 18

Writing CICS programs

You write a CICS program in much the same way as you write any other program. You can use COBOL, C, PL/I, or assembler language to write CICS application programs. Most of the processing logic is expressed in standard language statements, but you use CICS commands for some functions.

Other chapters in this book tell you the functions that CICS provides and indicate by example how CICS commands are written. However, you need to consult the *CICS Application Programming Reference* manual for programming information, including the full list of options associated with each command and the exact syntax. There are a few restrictions on normal language use under CICS, and these are described in Chapter 2, “Language considerations” on page 19.

CICS allows you to use SQL statements, DLI requests, CPI statements, and CICS Front End Programming Interface (FEPI) commands in your program as well as CICS commands. You need to consult additional manuals for information:

- *SQL Reference* manual and the *Application Programming and SQL Guide* (for SQL)
- *DL/I Application Programming: High Level Programming Interface* and the *DL/I Application Programming: CALL and RQDLI Interfaces* (for DL/I)
- *IBM SAA: CPI Reference* manual and the *SAA Common Programming Interface for Resource Recovery Reference* manual (for CPI)
- *CICS Front End Programming Interface User's Guide* (for programming information about FEPI commands)

Preparing your program

Because the compilers (and assemblers) cannot process CICS commands directly, an additional step is needed to convert your program into executable code. This step is called **translation**, and consists of converting CICS commands into the language in which the rest of the program is coded, so that the compiler (or assembler) can understand them.

CICS provides a translator program for each of the languages in which you may write, to handle both EXEC CICS and EXEC DLI statements. There are three

steps: translation, compilation (assembly), and link-edit. Figure 1 on page 4 shows the process.

Note: If you use EXEC SQL, you need additional steps to translate the SQL statements and bind; see the *Application Programming and SQL Guide* for information about these extra steps.




Figure 1. Preparing an application program

The *CICS System Definition Guide* manual provides more information about translation.

There are a few points to note, however:

- You can specify a number of options for the translation process, and you may need to do this for certain types of programs. If you are using EXEC DLI, for example, you need to tell the translator this fact. “Specifying translator options” on page 6 explains how to specify options, and “Translator options” on page 7 defines the options available.
- The translator may produce error messages, and it is as important to check these messages as it is to check the messages produced by the compiler and

link-editor. See “The translation process” on page 5 for the location of these messages.

- The compiler (or assembler) reads the translated version of your program as input, rather than your original source. This affects what you see on your compiler (assembler) listing. It also means that COPY statements in your source code must not bring in untranslated CICS commands, because it will be too late for the translator to convert them. (Copying in pretranslated code is possible, but not recommended; it may produce unpredictable results.)
- EXEC commands are translated into CALL statements that invoke CICS interface modules. These modules get incorporated into your object module in the link-edit step, and you see them in your link-edit output listing. You can read more about these modules in the *CICS System Definition Guide*.

The translation process

A language translator reads your source program and creates a new one; normal language statements remain unchanged, but CICS commands are translated into CALL statements of the form required by the language in which you are coding.

It requires a partition with a size of at least 250KB and a GETVIS of at least 100KB. This may need to be increased for larger programs.

The calls invoke CICS-provided “EXEC” interface modules, which later get link-edited into your load module. These, in turn, invoke the requested services at execution time.

The translators for all of the languages use one input and two output files:

SYSIPT (translator input) is the file that contains your source program.

SYSPCH (translated source) is the translated version of your source code, which becomes the input to the compile (assemble) step. In this file, your source has been changed as follows:

- The EXEC interface block (EIB) structure has been inserted.
- EXEC CICS and EXEC DLI commands have been turned into CALL statements.
- CICS DFHRESP and DFHVALUE built-in functions have been processed.
- A data interchange block (DIB) structure and initialization call have been inserted if the program contains EXEC DLI statements.

The CICS commands that get translated still appear in the source, but as comments only. Generally the non-CICS statements are unchanged. The output from the translator always goes to an 80-byte fixed-record-length data set.

SYSLST (translator listing) shows the number of messages produced by the translator, and the highest severity code associated with any message. The options used in translating your program also appear, unless these have been suppressed with the NOOPTIONS option.

For COBOL, C, and PL/I programs, SYSLST also contains the messages themselves. In addition, if you specify the SOURCE option of the translator, you also get an annotated listing of the source in

SYSLST. This listing contains almost the same information as the subsequent compilation listing and, therefore, many installations choose to omit it (the NOSOURCE option). One item you may need from this listing which is not present in the compile listing, however, is the line numbers, if the translator is assigning them. Line numbers are one way to indicate points in the code when you debug with the execution diagnostic facility (EDF). If you specify the VBREF option, you also get a list of the commands in your program, cross-referenced by line number, and you can use this as an alternative to the source listing for EDF purposes.

For assembler language programs, SYSLST contains only the translator options, the message count, and maximum severity code. The messages themselves are inserted into the SYSPCH file as comments after the related statement. This causes the assembler to copy them through to the assembler listing, where you can check them. You may also see MNOTEs that are generated by the assembler as the result of problems encountered by the translator.

Specifying translator options

The translator options you can choose are listed in “Translator options” on page 7. You can specify your choices in either of two ways:

1. List them as suboptions of the XOPTS option on the statement that the compiler (assembler) provides for specifying options. These statements are:

Compiler	Statement
COBOL/VSE	CBL, PROCESS
C/VSE	#pragma
PL/I VSE	* PROCESS
High Level Assembler	*ASM

2. List your options in the PARM operand of the EXEC job control statement for the translate step. Most installations use catalogued procedures to translate, compile (assemble) and link CICS programs, and therefore you specify this PARM field in the EXEC job control statement that invokes the procedure.

For example, if the name of the procedure for COBOL programs is COBCOMP, you set translator options for a COBOL program with a statement such as this one:

```
// COBCOMP, PARM=VBREF,QUOTE,SPACE(2),NOBLCARD
```

If you specify an option by one method and the same option or an option that conflicts by the other method, the specifications in the language statement override those in the EXEC statement. Similarly, if you specify multiple values for a single option or options that conflict on either type of statement, the last setting takes precedence. Except for COBOL/VSE programs translated with the ANS185 or COBOL3 option, these statements must precede each source program; there is no way to batch the processing of multiple programs in other languages.

Translator options may appear in any order, separated by one or more blanks or by a comma. If you specify them on the language statement for options, they must

appear in parentheses following the XOPTS parameter, because other options are ignored by the translator and passed through to the compiler.

The following COBOL example shows both translator and compiler options being passed together:

```
CBL LIB XOPTS(QUOTE SPACE(2))
```

These examples show translator options being passed alone:

```
#pragma XOPTS(FLAG(W) SOURCE);  
* PROCESS XOPTS(FLAG(W) SOURCE);  
*ASM XOPTS(NOPROLOG NOEPILOG)
```

If you use the PARM operand of the EXEC job control statement to specify options, the XOPTS keyword is unnecessary, because the only options permitted here are translator options. However, you may use XOPTS, with or without its associated parentheses. If you use XOPTS with parentheses, be sure to enclose all of the translator options. For example, the following forms are valid:

```
PARM='op1 op2 .. opn'
```

```
PARM='XOPTS op1 op2 .. opn'
```

```
PARM=XOPTS(op1 op2 .. opn)
```

but the following is **not** valid:

```
PARM=(XOPTS(op1 op2) opn)
```

(For compatibility with previous releases, the keyword CICS can be used as an alternative to XOPTS, except when you are translating batch EXEC DLI programs.)

Remember, if you alter the default margins for C #pragma statement processing using the PARM operand, the sequence margins should be altered too. You can do this using the NOSEQUENCE option.

For assembler programs, an “*ASM XOPTS” statement must be on the first line of input, and nothing must come before it. This includes comments and assembler statements such as “PRINT ON” or “EJECT.”

Translator options

You can specify the translator options that apply to all languages except where stated otherwise. Table 3 lists all the translator options, the program languages that apply, and any valid abbreviations.

Translator option	COBOL	C	PL/I	Assembler
ANSI85	X	X		
APOST or QUOTE	X			

Table 3 (Page 2 of 2). Translator options applicable to programming language

Translator option	COBOL	C	PL/I	Assembler
CBLCARD or NOCBLCARD	X			
CICS	X	X	X	X
COBOL2	X			
COBOL3	X			
DBCS	X			
DEBUG or NODEBUG	X	X	X	
DLI	X	X	X	X
EDF or NOEDF	X	X	X	X
EPILOG or NOEPILOG				X
EXCI	X	X	X	X
FE or NOFE	X			
FEPI or NOFEPI	X	X	X	X
FLAG(I or W or E or S)	X	X	X	
GDS		X		X
GRAPHIC			X	
LENGTH or NOLENGTH	X		X	X
LINECOUNT(n)	X	X	X	X
LINKAGE or NOLINKAGE	X			
MARGINS(m,n)		X	X	
NATLANG	X	X	X	X
NUM or NONUM	X			
OPMARGINS(m,n[,c])		X	X	
OPSEQUENCE (m,n) or NOOPSEQUENCE		X	X	
OPT or NOOPT				
OPTIONS or NOOPTIONS	X	X	X	X
PROLOG or NOPROLOG				X
QUOTE or APOST	X			
SEQ or NOSEQ	X			
SEQUENCE (m,n) or NOSEQUENCE		X	X	
SOURCE or NOSOURCE		X	X	
SP	X	X	X	X
SPACE(1 or 2 or 3)	X			
SPIE or NOSPIE	X	X	X	X
SYSEIB	X	X	X	X
VBREF or NOVBREF	X	X	X	

The default options are those in bold type in Table 3 on page 7. Many installations change the defaults, however. You can tell what options get used by default at your installation by looking at the SYSLST translator listing output from the translate step (see “The translation process” on page 5). If you want an option that is not the default, you must specify it, as described in “Specifying translator options” on page 6.

ANSI85 (COBOL only)

specifies that the translator is to translate COBOL programs that implement the ANSI85 standards.

Note: This option causes the COBOL2 and NOSEQ options to be used, even if you have specified SEQ.

APOST (COBOL only)

indicates that literals are delineated by the apostrophe ('). This is the default—QUOTE is the alternative. The same value must be specified for the translator step and the following compiler step.

The CICS-supplied COBOL copybooks use APOST.

CBLCARD (COBOL only)

specifies that the translator is to generate a CBL statement. This is the default—the alternative is NOCBLCARD.

CICS

specifies that the translator is to process EXEC CICS commands. It is the default specification in the translator. CICS is also an old name for the XOPTS keyword for specifying translator options, which means that you can specify the CICS option explicitly either by including it in your XOPTS list or by using it in place of XOPTS to name the list. The only way to indicate that there are no CICS commands is to use the XOPTS keyword without the option CICS. You must do this in a batch DL/I program using EXEC DLI commands. For example, to translate a batch DL/I program written in assembler language, specify:

```
*ASM XOPTS(DLI)
```

To translate a batch program written in COBOL and containing EXEC CICS LINK commands, specify:

```
CBL XOPTS(EXCI)
```

COBOL2 (COBOL only)

specifies that the translator is to translate programs compiled by the VS COBOL II (or later) compilers.

COBOL3 (COBOL only)

specifies that the translator is to translate programs compiled by the COBOL/VSE compiler. This option implies the ANSI85 and COBOL2 options.

DBCS (COBOL only)

specifies that the source program may contain double-byte characters. It causes the translator to treat hexadecimal codes X'0E' and X'0F' as shift-out (SO) and shift-in (SI) codes, respectively, wherever they appear in the program.

For more detailed information about how to program in COBOL using DBCS, see the section on DBCS character strings in the *COBOL for VSE/ESA Programming Guide* manual.

If you specify this option, the COBOL2 option is assumed.

DEBUG (COBOL, C and PL/I only)

instructs the translator to produce code that passes the line number through to CICS for use by the execution diagnostic facility (EDF). DEBUG is the default—NODEBUG is the alternative.

DLI

specifies that the translator is to process EXEC DLI commands. If the program is batch, you must specify it with the XOPTS option, that is, XOPTS(DLI), as explained in connection with the CICS option. You must specify it with the XOPTS option, that is, XOPTS(DLI).

EDF

specifies that the execution diagnostic facility is to apply to the program. EDF is the default—the alternative is NOEDF.

EPILOG (Assembler language only)

specifies that the translator is to insert the DFHEIRET macro at the end of the program being translated. DFHEIRET returns control from the issuing program to the program that invoked it. If you want to use any of the options of the EXEC CICS RETURN command, use EXEC CICS RETURN and specify NOEPILOG.

EPILOG is the default—the alternative, NOEPILOG, prevents the translator inserting the macro DFHEIRET. (See the *CICS Application Programming Reference* manual for programming information about the DFHEIRET macro.)

EXCI

specifies that the translator is to process EXEC CICS LINK commands for the External CICS Interface (EXCI). These commands must be used only in batch programs, and so the EXCI translator option is mutually exclusive to the CICS translator option, or any translator option that implies the CICS option. An error message is produced if both CICS and EXCI are specified, or if EXCI and a translator option that implies CICS are specified.

The EXCI option is also mutually exclusive to the DLI option. An error message is produced if both EXCI and DLI translator commands are specified.

The EXCI translator option is specified via XOPTS, that is, XOPTS(EXCI).

FE (COBOL only)

produces translator information messages that print (in hexadecimal notation) the bit pattern corresponding to the first argument of the translated call. This bit pattern has the encoded information that CICS uses to determine which function is required and which options are specified. All diagnostic messages are listed, whatever the FLAG option specifies.

FEPI

allows access to the FEPI API commands of the CICS Front End Programming Interface (FEPI). FEPI is available only if you have installed the CICS Front End Programming Interface. The alternative is NOFEPI. FEPI commands and design are not discussed in this book, but are discussed in the *CICS Front End Programming Interface User's Guide*.

FLAG(I|W|E|S) (COBOL, C, and PL/I only) Abbreviation: F

specifies the minimum severity of error in the translation which requires a message to be listed.

FLAG (I)

All messages.

FLAG (W)

(Default) All except information messages.

FLAG (E)

All except warning and information messages.

FLAG (S)

Only severe and unrecoverable error messages.

Note: The FE option overrides the FLAG option.

GDS (C and assembler language only)

specifies that the translator is to process CICS GDS (generalized data stream) commands. For programming information about these GDS commands, see the *CICS Application Programming Reference* manual.

GRAPHIC (PL/I only)

specifies that the source program may contain double-byte characters. It causes the translator to treat hexadecimal codes X'0E' and X'0F' as shift-out (SO) and shift-in (SI) codes, respectively, wherever they appear in the program.

It also prevents the translator from generating parameter lists that contain the shift-out and shift-in values in hexadecimal form. Wherever these values would ordinarily appear, the translator expresses them in binary form, so that there are no unintended DBCS delimiters in the data stream that the compiler receives.

If the compiler you are using supports DBCS, you need to prevent unintended shift-out and shift-in codes, even if you are not using double-byte characters. Do this by specifying the GRAPHIC option for the translator, so that it does not create them, or by specifying NOGRAPHIC on the compile step, so that the compiler does not interpret them as DBCS delimiters.

For more detailed information about how to program in PL/I using DBCS, see the *PL/I VSE Language Reference* manual.

LENGTH**(COBOL, Assembler and PL/I only)**

instructs the translator to generate a default length if the LENGTH option is omitted from a CICS command in the application program. The alternative is NOLENGTH.

LINECOUNT(n) Abbreviation: LC

specifies the number of lines to be included in each page of translator listing, including heading and blank lines. The value of "n" must be an integer in the range 1 through 255. If "n" is less than 5, only the heading and one line of listing are included on each page. The default is 60.

LINKAGE (COBOL only)

requests the translator to modify the LINKAGE SECTION and PROCEDURE DIVISION statements in top-level programs according to the existing rules.

This means that the translator will insert a USING DFHEIBLK DFHCOMMAREA statement in the PROCEDURE DIVISION, if one does not already exist, and will ensure that the LINKAGE SECTION (creating one if necessary) contains definitions for DFHEIBLK and DFHCOMMAREA.

LINKAGE is the default - the alternative is NOLINKAGE.

MARGINS(m,n[,c]) (C and PL/I only) Abbreviation: MAR

specifies the columns of each line or record of input that contain language or CICS statements. The translator does not process data that is outside these limits, though it does include it in the source listings.

The option can also specify the position of an American National Standard printer control character to format the listing produced when the SOURCE option is specified; otherwise, the input records are listed without any intervening blank lines. The margin parameters are:

m Column number of left-hand margin.

n Column number of right-hand margin. It must be greater than *m*.

Note: When used as a C compiler option, the asterisk (*) is allowable for the second argument on the MARGIN option. For the translator, however, a numeric value between 1 and 100 inclusive must be specified. When the input data set has fixed-length records, the maximum value allowable for the right hand margin is 80. When the input data set has variable-length records, the maximum value allowable is 100.

c Column number of the American National Standard printer control character. It must be outside the values specified for *m* and *n*. A zero value for *c* means no printer control character. If *c* is nonzero, only the following printer control characters can appear in the source:

b (blank)	Skip 1 line before printing.
0	Skip 2 lines before printing.
-	Skip 3 lines before printing.
+	No skip before printing.
1	New page.

The default for C is MARGINS(1,72,0) for fixed-length records, and for variable-length records it is the same as the record length (1,record length,0).

The default for PL/I is MARGINS(2,72,0) for fixed-length records, and MARGINS(10,100,0) for variable-length records.

NATLANG(EN|CH|GE|KA)

specifies what language is to be used for the translator message output by the translator:

EN	(Default) English
CH	Chinese
GE	German
KA	Kanji

(Take care not to confuse this option with the NATLANG API option.)

NOCBLCARD (COBOL only)

specifies that the translator is not to generate a CBL statement. The parameter that the CICS translator normally inserts must be the COBOL/VSE IGYCOPT macro. These parameters are RENT, NODYNAM, and LIB.

NODEBUG (COBOL, C, and PL/I only)

instructs the translator not to produce code that passes the line number through to CICS for use by the execution diagnostic facility (EDF).

NOEDF

specifies that the execution diagnostic facility is not to apply to the program. There is no performance advantage in specifying NOEDF, but the option can be useful to prevent commands in well-debugged subprograms appearing on EDF displays.

NOEPILOG (Assembler language only)

instructs the translator not to insert the DFHEIRET macro at the end of the program being translated. DFHEIRET returns control from the issuing program to the program that invoked it. If you want to use any of the options of the EXEC CICS RETURN command, use EXEC CICS RETURN and specify NOEPILOG. NOEPILOG prevents the translator inserting the macro DFHEIRET. (See the *CICS Application Programming Reference* manual for programming information about the DFHEIRET macro.)

NOFE (COBOL only)

does not produce translator information messages that print the bit pattern corresponding to the first argument of the translated call. NOFE is the default—the alternative is FE.

NOFEPI

disallows access to the FEPI API commands of the CICS Front End Programming Interface (FEPI). NOFEPI is the default—the alternative is FEPI.

NOLENGTH**(COBOL, Assembler and PL/I only)**

instructs the translator not to generate a default length if the LENGTH option is omitted from a CICS command in the application program. The default is LENGTH.

NOLINKAGE (COBOL only)

requests the translator not to modify the LINKAGE SECTION and PROCEDURE DIVISION statements to supply missing DFHEIBLK and DFHCOMMAREA statements.

This means that you can provide COBOL copybooks to define a COMMAREA and use the EXEC CICS ADDRESS command.

LINKAGE is the default—the alternative is LINKAGE.

NONUM (COBOL only)

instructs the translator not to use the line numbers appearing in columns one through six of each line of the program as the line number in its diagnostic messages and cross-reference listing, but to generate its own line numbers. NONUM is the default—the alternative is NUM.

NOOPSEQUENCE (C and PL/I only) Abbreviation: NOS

specifies the position of the sequence field in the translator output records. The default for C is OPSEQUENCE(73,80) for fixed-length records and NOOPSEQUENCE for variable-length records. For PL/I, the default is OPSEQUENCE(73,80) for both types of records.

NOOPTIONS Abbreviation: NOP

instructs the translator not to include a list of the options used during this translation in its output listing.

NOPROLOG (Assembler language only)

instructs the translator not to insert the macros DFHEISTG, DFHEIEND, and DFHEIENT into the program being assembled. These macros define local program storage and initialize at program entry. (See the *CICS Application Programming Reference* manual for programming information about these “prolog” macros.)

NOSEQ (COBOL only)

instructs the translator not to check the sequence field of the source statements, in columns 1-6. The alternative, SEQ, is the default. If SEQ is specified and a statement is not in sequence, it is flagged.

If you specify the ANSI85 option for COBOL, the translator does no sequence checking and the SEQ or NOSEQ option is ignored.

NOSEQUENCE (C and PL/I only) Abbreviation: NSEQ

specifies that statements in the translator input are not sequence numbered and that the translator must assign its own line numbers.

The default for fixed-length records is SEQUENCE(73,80). For variable-length records in C, the default is NOSEQUENCE. For variable-length records in PL/I, the default is SEQUENCE(1,8).

NOSOURCE (C and PL/I only)

instructs the translator not to include a listing of the translated source program in the translator listing.

NOSPIE

prevents the translator from trapping irrecoverable errors; instead, a dump is produced. Use NOSPIE only when requested to do so by the IBM support center.

NOVBREF (COBOL, C, and PL/I only)

instructs the translator not to include a cross-reference of commands with line numbers in the translator listing. (NOVBREF used to be called NOXREF; for compatibility, NOXREF is still accepted.) NOVBREF is the default—the alternative is VBREF.

NUM (COBOL only)

instructs the translator to use the line numbers appearing in columns one through six of each line of the program as the line number in its diagnostic messages and cross-reference listing.

OPMARGINS(m,n[,c]) (C and PL/I only) Abbreviation: OM

specifies the translator output margins, that is, the margins of the input to the following compiler. Normally these are the same as the input margins for the translator. If they are not, the number of columns within the translator output margins must be at least as great as the number of columns in the input margins.

For a definition of input margins and the meaning of *m*, *n*, and *c*, see MARGINS. The default for C is OPMARGINS(1,72,0) and, for PL/I, the default is OPMARGINS(2,72,0).

The maximum *n* value allowable for the OPMARGINS option is 80. The output from the translator is always of a fixed-length record format.

If the OPMARGINS option is used to set the output from the translator to a certain format, it may be necessary to change the input margins for the compiler being used. If the OPMARGINS value is allowed to default, this is not necessary.

OPSEQUENCE(m,n) (C and PL/I only) Abbreviation: OS

specifies the position of the sequence field in the translator output records. For the meaning of *m* and *n*, see SEQUENCE. The default for C is OPSEQUENCE(73,80) for fixed-length records and NOOPSEQUENCE for variable-length records. For PL/I, the default is OPSEQUENCE(73,80) for both types of records.

OPTIONS Abbreviation: OP

instructs the translator to include a list of the options used during this translation in its output listing.

PROLOG (Assembler language only)

instructs the translator to insert the macros DFHEISTG, DFHEIEND, and DFHEIENT into the program being assembled. These macros define local program storage and initialize at program entry. (See the *CICS Application Programming Reference* manual for programming information about these “prolog” macros.) PROLOG is the default—the alternative is NOPROLOG.

QUOTE (COBOL only)

indicates that literals are delineated by the double quotation mark (“”). The same value must be specified for the translator step and the following compiler step.

The CICS-supplied COBOL copybooks use APOST, the default, instead of QUOTE.

SEQ (COBOL only)

instructs the translator to check the sequence field of the source statements, in columns 1-6. SEQ is the default—the alternative is NOSEQ. If SEQ is specified and a statement is not in sequence, it is flagged.

If you specify the ANS185 option for COBOL, the translator does no sequence checking and the SEQ option is ignored.

SEQUENCE(m,n) (C and PL/I only) Abbreviation: SEQ

specifies that statements in the translator input are to be sequence numbered and the columns in each line or record that contain the sequence field. The translator uses this number as the line number in error messages and cross-reference listings. No attempt is made to sort the input lines or records into sequence. If no sequence field is specified, the translator assigns its own line numbers.

The SEQUENCE parameters are:

m Leftmost sequence number column.

n Rightmost sequence number column.

The sequence number field must not exceed eight characters and must not overlap the source program (as specified in the MARGINS option).

The default for fixed-length records is SEQUENCE(73,80).

For variable-length records in C the default is NOSEQUENCE. For variable-length records in PL/I, the default is SEQUENCE(1,8).

SOURCE (C and PL/I only)

instructs the translator to include a listing of the translated source program in the translator listing. SOURCE is the default—the alternative is NOSOURCE.

SP

must be specified for application programs that contain special (SP) CICS commands or they will be rejected at translate time. These commands are ACQUIRE, COLLECT, DISABLE, DISCARD, ENABLE, EXTRACT, INQUIRE, PERFORM, RESYNC, and SET. They are generally used by system programmers. For programming information about these commands, see the *CICS System Programming Reference* manual.

SPACE(1 or 2 or 3) (COBOL only)

indicates the type of spacing to be used in the output listing: SPACE(1) specifies single spacing, SPACE(2) double spacing, and SPACE(3) triple spacing. SPACE(3) is the default.

SPIE

specifies that the translator is to trap irrecoverable errors. SPIE is the default—the alternative is NOSPIE.

SYSEIB

indicates that the program is to use the system EIB instead of the application EIB. The SYSEIB option allows programs to execute CICS commands without updating the application EIB, making that aspect of execution transparent to the application. However, this option imposes restrictions on programs using it, and should be used only in special situations.

A program translated with the SYSEIB option must:

- Execute in AMODE(31), because the system EIB is assumed to be located in "TASKDATALOC(ANY)" storage.
- Obtain the address of the system EIB using the ADDRESS EIB command. (If the program is translated with the SYSEIB option, this command automatically returns the address of the system EIB).

Be aware that the use of the SYSEIB option implies the use of the NOHANDLE option on all CICS commands issued by the program. (Commands should use the RESP option as required.)

VBREF (COBOL, C, and PL/I only)

specifies whether the translator is to include a cross-reference of commands with line numbers in the translator listing. (VBREF used to be called XREF, and is still accepted.)

Preparing BMS maps

If your program uses BMS maps, you need to create the maps. The traditional method for doing this is to code the map in BMS macros and assemble these macros. You actually do the assembly twice, with different output options.

One assembly creates a set of definitions. You copy these definitions into your program using the appropriate language statement, and they allow you to refer to the fields in the map by name.

The second assembly creates an object module that is used when your program actually executes. The process is illustrated in Figure 2 on page 17.




Figure 2. Preparing a map

See the *CICS System Definition Guide* for more information about assembling maps.

Several products can produce the same two outputs from input that you express interactively. These include Screen Definition Facility II for VSE. For more information about SDF II, see *Screen Definition Facility II for VSE/ Primer for CICS/BMS programs*, and the *Screen Definition Facility II for VSE/ General Information* manual.

Whatever way you produce maps, you need to create a map before you compile (assemble) any program that uses it, unless you are using the MAPONLY option. In addition, if you change the map, you usually need to recompile (reassemble) all programs that use it. Some changes affect only the physical map and are not reflected in the corresponding symbolic map used by the program. One of these is a change in field position which does not alter the order of the fields. However, changes in data type, field length, field sequence, and other changes do affect the symbolic map, and it is always safest to recompile (reassemble).

Using VSE macros from within an application

For information about the use of specific VSE macros in CICS application programs, see the *VSE/ESA System Macros Reference* manual.

An application program is not precluded from direct communication with the operating system, but the results of such action can be unpredictable.

Chapter 2. Language considerations

This chapter describes various considerations and restrictions that apply when using different languages:

- “Map sets and partition sets” on page 400
- “C considerations” on page 38
- “PL/I considerations” on page 44
- “Assembler-language and HLASM considerations” on page 46
- “COBOL considerations”

COBOL considerations

You can use only the following COBOL compilers to process your COBOL application programs:

- IBM COBOL for VSE/ESA 1.1

With COBOL, you must use only EXEC CICS commands to invoke operating system services. Some COBOL statements must not be used. Some features of COBOL that are of interest to the CICS programmer are:

- Simplified based addressing, with the replacement of BLL cell manipulation by pointer variables and the ADDRESS special register.
- The ability to use COBOL CALL statements to call assembler language and other COBOL programs.
- The LENGTH special register, which CICS uses to deduce the length of data items. You no longer need to specify a length in a CICS command that specifies a data name in the FROM option, or, if the data is fixed length, in the INTO option.
- The ability to use the RETURN-CODE special register in a CICS application program. This register allows you to set and access return codes in COBOL programs.
- There is no need for the SERVICE RELOAD statement. If included, it is ignored. The COBOL/VSE compiler recognizes references to based addresses and generates code to manipulate them (even when the OPTIMIZE compiler option is used).

Based addressing

CICS application programs need to access data dynamically when the data is in a CICS internal area, and only the address is passed to the program. Examples are:

- CICS areas such as the CWA, TWA, and TCTTE user area (TCTUA), accessed using the EXEC CICS ADDRESS command.
- Input data, obtained by CICS commands such as EXEC CICS READ and EXEC CICS RECEIVE with the SET option.

COBOL/VSE provides a simple method of obtaining addressability to the data areas defined in the LINKAGE SECTION using pointer variables and the ADDRESS special register. Figure 3 on page 20 gives an example of this.

The ADDRESS special register holds the address of a record defined in the LINKAGE SECTION with level 01 or 77. This register can be used in the SET option of the EXEC CICS GETMAIN, EXEC CICS LOAD, EXEC CICS READ, and EXEC CICS READQ commands.

```

WORKING-STORAGE SECTION.
  77 LRECL-REC1          PIC S9(4) COMP.
*
LINKAGE SECTION.
  01 REC-1.
    02 FLAG1            PIC X(1).
    02 MAIN-DATA        PIC X(5000).
    02 OPTL-DATA        PIC X(1000).
  01 REC-2.
    02 ....
*
PROCEDURE DIVISION.

    EXEC CICS READ UPDATE...
      SET(ADDRESS OF REC-1)
      LENGTH(LRECL-REC1)
    END-EXEC.

    IF FLAG1 EQUAL 'Y'
      MOVE OPTL-DATA TO...

    EXEC CICS REWRITE...
      FROM(REC-1)
    END-EXEC.

```

Figure 3. Addressing CICS data areas in locate mode

Figure 3 shows the use of ADDRESS special registers in COBOL/VSE. If the records in the EXEC CICS READ or EXEC CICS REWRITE commands are fixed length, COBOL/VSE does not require a LENGTH option. This example assumes variable-length records. After the read, you can get the length of the record from the field named in the LENGTH option (here, LRECL-REC1). In the EXEC CICS REWRITE command, you must code a LENGTH option if you want to replace the updated record with a record of a different length.

USE FOR DEBUGGING facility under COBOL/VSE

If a “D” is placed in column seven of a COBOL EXEC CICS command, that “D” is also found in the translated CALL statements. This translated command is executed only if the USE FOR DEBUGGING facility is activated. A “D” placed on any line other than the first line of the EXEC CICS statement is not required and is ignored by the translator.

Restrictions

This section describes COBOL language elements that you cannot use under CICS, or whose use is restricted or can cause problems under CICS.

In general, neither the CICS translator nor the COBOL/VSE compiler detects the use of COBOL words affected by the following restrictions. The use of a restricted word in a CICS environment may cause a failure at execution time. However, a

table of CICS-reserved words, IGYCCICS, is supplied with COBOL/VSE. This allows the COBOL compiler to flag any occurrences of COBOL reserved words that conflict with CICS restrictions. How to use and create installation-specific COBOL reserved word tables is documented in the *COBOL/VSE Installation and Customization* guide.

The following restrictions apply to a COBOL program that is to be used as a CICS application program. (See the appropriate COBOL programmer's guide for more information about these functions.)

- If no IDENTIFICATION DIVISION is present, only the CICS commands are expanded.
- If the IDENTIFICATION DIVISION only is present, only DFHEIVAR, DFHEIBLK, and DFHCOMMAREA are produced.
- Statements that produce variable-length areas, such as OCCURS DEPENDING ON, should be used with caution within the WORKING-STORAGE SECTION.
- If you have any CICS applications written in COBOL, you may need to review the COBOL/VSE run-time options in use at your installation.
- You cannot use entry points in COBOL in CICS.
- When a debugging line is to be used as a comment, it must not contain any unmatched quotation marks.
- Do not use COBOL statements that invoke operating system functions. Instead, use CICS commands.

- Do not use the following statements:

ACCEPT	READ
CALL 'literal' with DYNAM	RERUN
CLOSE	REWRITE
DELETE	SORT
DISPLAY	START
MERGE	STOP 'literal'
OPEN	WRITE
READ	

- There are restrictions on the use of the SORT statement. See the *COBOL/VSE Programming Guide* for information on this.
- Do not use:
 - USE declaratives (except USE FOR DEBUGGING)
 - ENVIRONMENT DIVISION and FILE SECTION entries associated with data management, because CICS handles data management
 - User-specified parameters to the main program
- Do not use the following compiler options:
 - DYNAM
 - NOLIB (if program is to be translated)
 - NORENT
 - NUMPROC (PFD)
 - TRUNC
- Use TRUNC(OPT) for handling binary data items if they conform to the PICTURE and USAGE definitions, otherwise use TRUNC(BIN). For example, use TRUNC(BIN) for a data item defined as PIC S9(4) BINARY if it might

receive a value greater than 4 digits. (TRUNC(STD) is the default and TRUNC(BIN) is slower.)

- The length of working storage, plus 80 bytes for storage accounting and save areas, must not exceed 64KB when the COBOL/VSE compiler option DATA(24) is used. If, however, the compiler option DATA(31) is used, up to 128MB are available.
- If the DLI option is specified and an ENTRY statement immediately follows the PROCEDURE DIVISION header, it is recommended that the ENTRY statement be terminated with a period (.).
- COBOL and PL/I application programs cannot be link-edited together. For further information about using COBOL/VSE with other languages, see the *LE/VSE Programming Guide*.
- The following compiler options have no effect in a CICS environment:
ADV
FASTSRT
OUTDD
- If you use EXEC CICS HANDLE CONDITION or EXEC CICS HANDLE AID, you can avoid addressing problems by using SET(ADDRESS of A-DATA) or SET(A-POINTER) where A-DATA is a structure in the LINKAGE SECTION and A-POINTER is defined with USAGE IS POINTER clause.

Restrictions for 31-bit addressing

The following restrictions apply to a COBOL/VSE program running above the 16MB line:

- If the receiving program is link-edited with AMODE(24), addresses passed to it must be 24 bits long.
- If the receiving program is link-edited with AMODE(31), addresses passed to it must be 31 bits long. If the address has previously been accessed as a 24-bit address, you must ensure that the leftmost byte contains zeros.
- The EXEC CICS WAIT EVENT interval control command is not supported when the associated event control block (ECB) resides above the 16MB line. Instead you can use the task control EXEC CICS WAIT EXTERNAL ECBLIST or EXEC CICS WAITCICS ECBLIST commands.
- The COMMAREA is restricted in a mixed addressing mode transaction environment. For a description of the restriction, see “Mixed addressing mode transactions” on page 383.
- The length of working storage, plus the length of the task global table (TGT), plus the length of the register save area (RSA, 60 bytes), plus 24 bytes for SCP accounting, must not exceed 64KB unless the COBOL/VSE DATA(31) compiler option is specified. For programs with DATA(31), working storage is above the 16MB line and the maximum size is 128MB.
- The program must not violate 31-bit addressing rules described in the *VSE/ESA Extended Addressability* manual.

DL/I CALL interface

You should make the following changes to programs that use CALL DL/I:

- Retain the DLIUIB declaration and at least one PCB declaration in the LINKAGE SECTION.
- Change the PCB call to specify the UIB directly, as follows:

```
CALL 'CBLTDLI' USING PCB-CALL  
                    PSB-NAME  
                    ADDRESS OF DLIUIB.
```

- Obtain the address of the required PCB from the address list in the UIB.

There are five restricted calls (ctdli(), ceetdli, asmtkli, cbltdli, and plitdli) and these are documented in the *LE/VSE Programming Guide*.

Figure 4 on page 24 illustrates the whole of the above process. The example in the figure assumes that you have three PCBs defined in the PSB and want to use the second PCB in the database CALL. Therefore, when setting up the ADDRESS special register of the LINKAGE SECTION group item PCB, the program uses 2 to index the working-storage table, PCB-ADDRESS-LIST. To use the *n*th PCB, you use the number *n* to index PCB-ADDRESS-LIST.

```

WORKING-STORAGE SECTION.
  77 PCB-CALL          PIC X(4) VALUE 'PCB '.
  77 GET-HOLD-UNIQUE  PIC X(4) VALUE 'GHU '.
  77 PSB-NAME         PIC X(8) VALUE 'CBLPSB'.
  77 SSA1             PIC X(40) VALUE SPACES.
  01 DLI-IO-AREA.
    02 DLI-IO-AREA1  PIC X(99).
*
LINKAGE SECTION.
  COPY DLIUIB.
  01 OVERLAY-DLIUIB REDEFINES DLIUIB.
    02 PCBADDR       USAGE IS POINTER.
    02 FILLER        PIC XX.
  01 PCB-ADDR-LIST.
    02 PCB-ADDRESS-LIST USAGE IS POINTER
                          OCCURS 10 TIMES.

  01 PCB.
    02 PCB-DBD-NAME  PIC X(8).
    02 PCB-SEG-LEVEL PIC XX.
    02 PCB-STATUS-CODE PIC XX.
*
PROCEDURE DIVISION.
*SCHEDULE THE PSB AND ADDRESS THE UIB
  CALL 'CBLTDLI' USING PCB-CALL PSB-NAME ADDRESS OF DLIUIB.
*
*MOVE VALUE OF UIBPCBAL, ADDRESS OF PCB ADDRESS LIST (HELD IN UIB)
*(REDEFINED AS PCBADDR, A POINTER VARIABLE), TO
*ADDRESS SPECIAL REGISTER OF PCB-ADDR-LIST TO PCBADDR.
  SET ADDRESS OF PCB-ADDR-LIST TO PCBADDR.
*MOVE VALUE OF SECOND ITEM IN PCB-ADDRESS-LIST TO ADDRESS SPECIAL
*REGISTER OF PCB, DEFINED IN LINKAGE SECTION.
  SET ADDRESS OF PCB TO PCB-ADDRESS-LIST(2).
*PERFORM DATABASE CALLS .....
  .....
  MOVE ..... TO SSA1.
  CALL 'CBLTDLI' USING GET-HOLD-UNIQUE PCB DLI-IO-AREA SSA1.
*CHECK SUCCESS OF CALLS .....
  IF UIBFCTR IS NOT EQUAL LOW-VALUES THEN
      ..... error diagnostic code
  .....
  IF PCB-STATUS-CODE IS NOT EQUAL SPACES THEN
      ..... error diagnostic code
  .....

```

Figure 4. Using the DLI CALL interface

Mixing languages

A run unit is a running set of one or more programs that communicate with each other by COBOL static or dynamic CALL statements. In a CICS environment, a run unit is entered at the start of a CICS task, or invoked by an EXEC CICS LINK or an EXEC CICS XCTL command. A run unit can be defined as the execution of a single entry in the processing program table (PPT) even though for dynamic CALL, the subsequent PPT entry is needed for the called program.

Because CICS does not support a mixture of language levels in a run unit, a COBOL run unit can contain only:

- COBOL programs compiled with the COBOL/VSE compiler or VSE/COBOL II compilers.
- Assembler routines

CICS supports only COBOL-to-COBOL and COBOL-to-assembler calls.

However, a CICS transaction can consist of many run units, each of which can be at a different language level. This means that a single transaction can consist of programs compiled by different compilers (including non-COBOL compilers), as long as programs compiled by different compilers communicate with each other by using EXEC CICS LINK or EXEC CICS XCTL commands only.

Calling subprograms from COBOL/VSE

In a CICS system, when control is transferred from the active program to an external program but the transferring program remains active and control can be returned to it, the program to which control is transferred is called a **subprogram**.

Control can be transferred to a subprogram by an EXEC CICS LINK command, a static COBOL call, or a dynamic COBOL call. EXEC CICS LINK can be used to call a DOS/VS COBOL, VS COBOL II or COBOL/VSE compiled program, but a static COBOL call or dynamic COBOL call is allowed to call a VS COBOL II program or a COBOL/VSE program only.

There are three ways of transferring control to a subprogram:

EXEC CICS LINK

The calling program contains a CICS command in one of the following forms:

```
EXEC CICS LINK PROGRAM('subpname')  
EXEC CICS LINK PROGRAM(name)
```

In the first form, the called subprogram is explicitly named as a nonnumeric literal within quotation marks. In the second form, name refers to the COBOL data area with length equal to that required for the name of the subprogram.

Static COBOL call

The calling program contains a COBOL statement of the form:

```
CALL 'subpname'
```

The called subprogram is explicitly named as a literal string.

Dynamic COBOL call

The calling program contains a COBOL statement of the form:

```
CALL identifier
```

The identifier is the name of a COBOL data area that must contain the name of the called subprogram.

Table 4 on page 27 gives the rules governing the use of the three ways to call a subprogram. This table refers to CICS application logical levels. Each LINK command creates a new logical level, the called program being at a level one lower than the level of the calling program (CICS is taken to be at level 0). Figure 5 on page 29 shows logical levels and the effect of RETURN commands and CALL statements in linked and called programs.

The term **run unit**, used in Figure 5 on page 29, is defined under the heading “Mixing languages” on page 24.

When control is passed by an EXEC CICS XCTL command, the program receiving control cannot return control to the calling program by an EXEC CICS RETURN command or a GOBACK statement, and is therefore not a subprogram.

In an ANSI85 unit of compilation, a called nested program is internal to the calling program, and is therefore not a subprogram. See “Nesting—what the application programmer must do” on page 33.

The CALL has the following form:

```
CALL {'PROG'|identifier} USING DFHEIBLK COMMAREA PARM1 PARM2 ...
```

In the called program PROG, the CICS translator inserts DFHEIBLK and DFHCOMMAREA into the LINKAGE SECTION, and into the USING list of the PROCEDURE DIVISION statement. You code the PROCEDURE DIVISION statement normally, as follows:

```
PROCEDURE DIVISION USING PARM1 PARM2...
```

and the translator inserts DFHEIBLK and DFHCOMMAREA into this statement before PARM1.

You must always use the NODYNAM compiler option (the default) when you compile a COBOL program that is to run with CICS, even if the program issues dynamic calls.

<i>Table 4 (Page 1 of 3). Rules governing methods of calling subprograms</i>			
	EXEC CICS LINK	Static COBOL CALL	Dynamic COBOL CALL
Translation	The linked-to subprogram must be translated.	The called subprogram must be translated if it contains CICS commands or references to the EXEC interface block (DFHEIBLK) or to the CICS communication area (DFHCOMMAREA).	
Link-editing	The linked-to subprogram must be compiled and link-edited as a separate program.	The called subprogram must be link-edited with the calling program to form a single phase (but the programs can be compiled separately). This can produce large program modules, and it also stops two programs that call the same program from sharing a copy of that program.	The called subprogram must be compiled and link-edited as a separate phase. It can reside in the SVA or in a library that is shared with other CICS and non-CICS partitions at the same time.
CICS system definition data set (CSD) entries without program autoinstall. (If you use program autoinstall, you do not need an entry in the CSD.)	The linked subprogram must be defined using RDO, unless you issue EXEC CICS LINK with SYSID, and SYSID is not local, in which case, you do not need a definition for the linked-to program. If the linked subprogram is unknown or unavailable, even though autoinstall is active, the LINK fails with the PGMIDERR condition.	The calling program must be defined in the CSD. If program A calls program B, and then program B attempts to call program A, the result will be a COBOL message and abend (1015). See the <i>COBOL/VSE Application Programming and Debugging Guide</i> for information about this abend code.	
		The subprogram is part of the calling program so no CSD entry is required.	The called subprogram must be defined in the CSD. If the called subprogram cannot be loaded or is unavailable even though autoinstall is active, COBOL issues a message and abends (1029).
Return from called subprogram	The linked-to subprogram must return using EXEC CICS RETURN command in the COBOL statement GOBACK.	The called subprogram must return using the COBOL statement GOBACK or EXIT PROGRAM. A called assembler-language subprogram must return using the macro DFHEIRET. The use of EXEC CICS RETURN in the called subprogram terminates the calling program.	
Language of called subprogram	Any language supported by CICS.	VS COBOL II, or COBOL/VSE or assembler language.	
Contents of called or linked-to subprogram	Any function supported by CICS for the language (including calls to external databases, for example, DB2® for VSE/ESA and DL/I) with the exception that an assembler-language subprogram cannot call a lower level subprogram.		

Table 4 (Page 2 of 3). Rules governing methods of calling subprograms

	EXEC CICS LINK	Static COBOL CALL	Dynamic COBOL CALL
Passing parameters to the subprogram	Data can be passed by any of the standard CICS methods (COMMAREA, INPUTMSG, TWA, TCTUA, TS queues) if the called or linked-to subprogram is processed by the CICS translator.		
	If the COMMAREA is used, its address must be passed in the EXEC CICS LINK command. If the linked-to subprogram uses 24-bit addressing, and the COMMAREA is above the 16MB line, CICS copies it to below the 16MB line, and recopies it on return.	The CALL statement may pass DFHEIBLK and DFHCOMMAREA as the first two parameters, if the called program is to issue EXEC CICS requests, or the called program can issue EXEC CICS ADDRESS commands. The COMMAREA is optional but if other parameters are passed, a dummy commarea must also be passed.	In an ANSI85 unit of compilation, a nested program is not a subprogram, and the above rule can be varied. See "Nesting—what the application programmer must do" on page 33.
Storage	On each entry to the linked-to subprogram, a new initialized copy of its working storage is provided, and the run unit is reinitialized (in some circumstances, this can cause a performance degradation).	On the first entry to the called subprogram within a CICS logical level, a new initialized copy of its working storage is provided. On subsequent entries to the called subprogram at the same logical level, the same WORKING STORAGE is provided in its last-used state, that is, no storage is freed, acquired, or initialized. If performance is unsatisfactory with EXEC CICS LINK commands, COBOL calls may give improved results.	
CICS condition/AID and abend handling	On entry to the called subprogram, no abend or condition handling is active. Within the subprogram, the normal CICS rules apply. In order to establish an abend or condition handling environment, that exists for the duration of the subprogram, a new HANDLE command should be issued on entry to the subprogram. The environment so created remains in effect until either a further HANDLE command is issued, or the subprogram returns control to the caller.	If the dynamic COBOL CALL fails, CICS abend handling is not invoked, and you may get a COBOL abend code (1013). See the <i>COBOL/VSE Application Programming and Debugging Guide</i> for information about this abend code. On entry to the called subprogram, COBOL issues a PUSH HANDLE to stack the calling program's condition or abend handlers. In order to establish an abend or condition handling environment that exists for the duration of the subprogram, a new HANDLE command should be issued on entry to the subprogram. The environment that this creates remains in effect until either a further HANDLE command is issued or the subprogram returns control to the caller. When control is returned to the calling program from the subprogram, COBOL unstacks the condition and abend handlers using a POP HANDLE.	

Table 4 (Page 3 of 3). Rules governing methods of calling subprograms			
	EXEC CICS LINK	Static COBOL CALL	Dynamic COBOL CALL
Location of called or linked-to program	Can be remote	Must be local	Must be local




Figure 5. Flow of control between COBOL programs and run units in a CICS environment

COBOL/VSE with the ANSI85 COBOL standards

The COBOL/VSE compiler supports the ANSI85 COBOL standards. The CICS translator option ANSI85 also supports most of these standards. If invoked with the ANSI85 option, the translator uses the COBOL2 option. The COBOL3 option implies both ANSI85 and COBOL2.

CICS support for these standards takes the form of changes to the translator. Because the translator is not a compiler, it is not affected by all the ANSI85 standards.

The standards that affect the translator are:

- Literals intervening in blank lines
- Sequence numbers containing any character
- Lowercase characters supported in all COBOL words
- REPLACE statement
- Batch compilation
- Nested programs
- Reference modification
- Global variables
- Interchangeability of comma, semicolon, and space
- Symbolic character definition

If a standard is not fully supported by the translator, a programming restriction applies. These standards are described under their appropriate headings, and each description is followed by a “Translator action.” This information is summarized in “Summary of restrictions” on page 37. The translator actions assume that you have specified the ANSI85 translator option.

The term **unit of compilation** means a section of source input from which the compiler produces a single object program. A unit of compilation can consist of a containing program and other programs nested within it.

Literals intervening in blank lines

Blank lines can appear anywhere in a COBOL source program. A blank line is a line that contains only blanks between margin C (the continuation column) and margin R (the last character in the line) inclusive.

Translator action: If blank lines occur within literals in a COBOL source program, the translator eliminates them from the translated output, but includes them in the translated listing.

(If the ANSI85 option is not specified, a blank line in a literal causes a translator error.)

Sequence numbers containing any character

In a COBOL source program, the sequence number field can contain any character in the computer’s character set. The sequence number fields can be in any order and need not be unique.

Translator action: The translator does not check on the contents or the sequence of the number fields.

If you specify the SEQ translator option, the translator issues a message saying that the SEQ option has no effect when you specify the ANSI85 option.

Lowercase characters supported in all COBOL words

Lowercase characters can occur anywhere in any COBOL word, including user-defined names, system names, and reserved words. A lowercase character can be used wherever an uppercase character is required by a COBOL compiler that does not conform to ANSI85 standards.

Translator action: The translator listing and output preserve the case of COBOL text as entered.

In addition, the translator accepts:

- Mixed case in translator options
- Mixed case in EXEC CICS commands, both for keywords and for arguments to keywords

(If the ANSI85 option is not specified, the translator expects COBOL words to consist entirely of uppercase characters.)

Notes:

1. The translator does not translate lowercase text into uppercase. Some names in COBOL text, for example, file names and transaction IDs, must match with externally defined names. Such names should always be entered in the same case as the external definition.
2. CBL and PROCESS cards must be in uppercase.

REPLACE statement

COBOL programs can include the REPLACE statement, which allows the replacement of identified text by defined substitution text. The text to be replaced and inserted is specified as in the REPLACING option of the COPY statement, and can be pseudo-text, an identifier, a literal, or a COBOL word. REPLACE statements are processed after COPY statements.

Translator action: The translator accepts REPLACE statements but does not translate text between pseudo-text delimiters, with one exception: CICS built-in functions (DFHRESP and DFHVALUE) are translated wherever they occur. CICS commands should not be placed between pseudo-text delimiters.

Batch compilation

Separate COBOL programs can be compiled together as one input file. An END PROGRAM header statement terminates each program. The END PROGRAM header statement is optional for the last program in the batch.

Translator action: The translator accepts separate COBOL programs in a single input file, and interprets END PROGRAM header statements according to the ANSI85 standards.

Translator options specified as parameters when invoking the translator are in effect for the whole batch, but can be overridden for a unit of compilation by options specified in the CBL or PROCESS card that initiates the unit.

The options for a unit of compilation are determined according to the following order of priority:

1. Options specified in the CBL or PROCESS card that initiates the unit
2. Options specified when the translator is invoked
3. Default options

For more information about compilation, see the *CICS System Definition Guide*.

Nested programs

Under the ANSI85 standard:

- COBOL programs can contain COBOL programs.
- Contained programs are included immediately before the END PROGRAM statement of the containing program.
- A contained program can also be a containing program, that is, it can itself contain other programs.
- Each contained or containing program is terminated by an END PROGRAM statement.

For an explanation of valid calls to nested programs and of the COMMON attribute of a nested program, see the *COBOL/VSE Programming Guide*. An example of a nested program is given in “An example of a nested program” on page 34.

Translator action: The translator treats top-level and nested programs differently.

Top-level programs: The translator translates a top-level program (a program that is not contained by any other program) in the normal way, with one addition. The translator uses the GLOBAL storage class for all translator-generated variables in the WORKING-STORAGE SECTION.

Nested programs: The translator translates nested or contained programs in a special way as follows:

- A DATA DIVISION and LINKAGE SECTION are added if they do not already exist.
- Declarations for DFHEIBLK (EXEC interface block) and DFHCOMMAREA (communication area) are inserted into the LINKAGE SECTION.
- EXEC CICS commands and CICS built-in functions are translated.
- The PROCEDURE DIVISION statement is not modified.
- No translator-generated temporary variables, used for pre-call assignments, are inserted in the WORKING-STORAGE SECTION.

Recognition of nested programs: If the ANSI85 option is specified, the translator assumes that the input source starts with a top-level program if the first noncomment record is any of the following:

IDENTIFICATION DIVISION statement
CBL card
PROCESS card

If the first record is none of these, the translator treats the input as part of the PROCEDURE DIVISION of a nested program. The first CBL or PROCESS card indicates the start of a top-level program and of a new unit of compilation. Any IDENTIFICATION DIVISION statements that are found after the first top-level program indicate the start of a new nested program.

The practical effect of these rules is that nested programs cannot be held in separate files and translated separately. A top-level program and all its directly (and indirectly) contained programs constitute a single unit of compilation, and should be submitted together to the translator.

Positioning of comments: The translator treats comments that follow an END PROGRAM statement as belonging to the next program in the input source. Comments that precede an IDENTIFICATION DIVISION statement appear in the listing after the IDENTIFICATION DIVISION statement.

To avoid confusion, always place comments:

- After the IDENTIFICATION DIVISION statement that initiates the program to which they refer
- Before the END PROGRAM statement that terminates the program to which they refer.

Nesting—what the application programmer must do

1. Submit a top-level containing program and all its directly and indirectly contained programs as a single unit of compilation.
2. In each nested program that contains EXEC commands, CICS built-in functions, or references to the EIB or COMMAREA, code DFHEIBLK and DFHCOMMAREA as the first two parameters of the PROCEDURE DIVISION statement as follows:

```
PROCEDURE DIVISION USING DFHEIBLK
                        DFHCOMMAREA PARM1 PARM2 ...
```

3. In every call to a nested program that contains EXEC commands, CICS built-in functions, or references to the EIB or COMMAREA, code DFHEIBLK and DFHCOMMAREA as the first two parameters of the CALL statement as follows:

```
CALL 'PROGA' USING DFHEIBLK
                        DFHCOMMAREA PARM1 PARM2 ...
```

4. For every call that forms part of the control hierarchy between the top-level program and a nested program that contains EXEC commands, CICS built-in functions, or references to the EIB or COMMAREA, code DFHEIBLK and DFHCOMMAREA as the first two parameters of the PROCEDURE DIVISION statement in the calling program and code DFHEIBLK and COMMAREA as the first two parameters of the CALL DIVISION statement in the called program. This is necessary to allow addressability to the EIB and COMMAREA to be passed to programs not directly contained by the top-level program.
5. If it is not necessary to insert DFHEIBLK and DFHCOMMAREA in the PROCEDURE DIVISION of a nested program for any of the reasons given in points 2, 3, and 4, calls to that program should not include DFHEIBLK and COMMAREA in the parameter list of the CALL statement.

An example of a nested program

A unit of compilation (see Figure 6) consists of a top-level program W and three nested programs, X, Y, and Z, all directly contained by W.




Figure 6. Nested program example—nesting structure

Program W During initialization and termination, calls Y and Z to do initial CICS processing and non-CICS file access. Calls X to do main processing.

Program X Calls Z for non-CICS file access and Y for CICS processing.

Program Y Issues CICS commands. Calls Z for non-CICS file access.

Program Z Accesses files in batch mode.

Applying the rules:

- Y must be COMMON to enable a call from X.
- Z must be COMMON to enable calls from X and Y.
- Y issues CICS commands; therefore:
 - All calls to Y must have DFHEIBLK and a COMMAREA as the first two parameters.
 - Y's PROCEDURE DIVISION statement must have DFHEIBLK and DFHCOMMAREA as the first two parameters.
- Though X does not access the EIB or the communication area, it calls Y, which issues CICS commands. Therefore the call to X must have DFHEIBLK and a COMMAREA as the first two parameters and X's PROCEDURE DIVISION statement must have DFHEIBLK and DFHCOMMAREA as its first two parameters.

Figure 7 on page 35 illustrates the points in "Nesting—what the application programmer must do" on page 33.

```

IDENTIFICATION DIVISION.
PROGRAM-ID.W.
.
PROCEDURE DIVISION.
.
CALL Z.
.
CALL Y USING DFHEIBLK
COMMAREA.
.
CALL X USING DFHEIBLK
COMMAREA.
.

IDENTIFICATION DIVISION.
PROGRAM-ID.X.
.
PROCEDURE DIVISION USING
DFHEIBLK DFHCOMMAREA.
.
CALL Z.
.
CALL Y USING DFHEIBLK
COMMAREA.
.
END PROGRAM X.

IDENTIFICATION DIVISION.
PROGRAM-ID.Y IS COMMON
.
PROCEDURE DIVISION USING
DFHEIBLK DFHCOMMAREA.
.
CALL Z.
.
EXEC CICS ....
.
END PROGRAM Y.

IDENTIFICATION DIVISION.
PROGRAM-ID.Z IS COMMON
.
PROCEDURE DIVISION.
.
END PROGRAM Z.

END PROGRAM W.

```

Figure 7. Nested program example—coding

Reference modification

Reference modification is supported, and is a method of referencing a substring of a character data item by specifying the starting (leftmost) position of the substring in the data item and, optionally, the length of the substring. The acceptable formats are:

```
data-name (starting-position:)  
data-name (starting-position: length)
```

Data-name can be subscripted, or qualified, or both. Both *starting-position* and *length* can be arithmetic expressions.

Translator action: The translator accepts reference modification wherever the name of a character variable is permitted in a COBOL program or in an EXEC CICS command.

Note: If a CICS command uses reference modification in defining a data value, it **must** include a LENGTH option to specify the data length. Otherwise the translator generates a COBOL call with a LENGTH register reference in the form:

```
LENGTH OF (reference modification)
```

This is rejected by the compiler.

Global variables

The GLOBAL variable storage class is supported. A variable defined with the GLOBAL variable storage class in a top-level program (see “Top-level programs” on page 32) can be referred to in any of its nested programs, whether directly or indirectly contained.

Translator action: The translator accepts the GLOBAL keyword.

Comma and semicolon as delimiters

A separator comma is a comma followed by a space. A separator semicolon is a semicolon followed by a space. A separator comma or a separator semicolon can be used as a separator wherever a space alone can be used.

Translator action: The translator accepts the use of a separator comma or a separator semicolon in COBOL statements wherever a space can be used. For example, the translator accepts the statement:

```
IDENTIFICATION; DIVISION
```

The translator does **not** support the use of the separator comma and separator semicolon as delimiters in EXEC CICS commands. The only acceptable word delimiter in an EXEC CICS command is a space.

Symbolic character definition

Symbolic characters can be defined in the SPECIAL-NAMES paragraph after the ALPHABET clause. A symbolic character is a program-defined word that represents a 1-character figurative constant.

Translator action: The translator accepts the use of symbolic characters as specified in the ANSI85 standard.

Note: In general, the compiler does not accept the use of figurative constants and symbolic characters as arguments in CALL statements. For this reason, do

not use figurative constants or symbolic characters in CICS commands, which are converted into CALL statements by the translator. There is one exception to this restriction: a figurative constant is acceptable in a CICS command as an argument to **pass** a value if it is of the correct data type. For example, a numeric figurative constant can be used in the LENGTH option.

Summary of restrictions

The following is a summary of the programming restrictions associated with CICS translator support for the ANSI85 COBOL standards:

- With the ANSI85 option, the translator varies the rule for parameters to be passed by a static COBOL call. For details, see “Nesting—what the application programmer must do” on page 33.
- A REPLACE statement must not contain an EXEC CICS command in pseudo-text.
- Programs cannot use a comma or semicolon as a word separator in a CICS command.
- Programs cannot use a symbolic character as an argument in a CICS command.
- Comments should not precede the IDENTIFICATION DIVISION statement or follow the END PROGRAM statement.
- CICS commands that use reference modification to define a character data value must include a LENGTH option to define the data length.
- A name that must match an external definition, for example a file name or a transaction ID, must be coded in the same case as the external definition.

COBOL2 translator option

If you are using the COBOL2 option you must use the VS COBOL II compiler to process your COBOL application programs.

This compiler is a licensed program that conforms to the standard set by American National Standard COBOL, X3.23-1974.

For information about translating your program and preparing it for execution, see Chapter 1, “Preparing your application to run” on page 3.

Lower-case characters can occur anywhere in any COBOL word, including user-defined names, system names, and reserved words. A lower-case character can be used wherever an upper-case character is required by a COBOL compiler that does not conform to the ANSI85 standards.

Translator action: The translator listing and output preserve the case of COBOL text as entered.

In addition, the translator accepts mixed case in:

- Translator options
- EXEC CICS commands, both for keywords and for arguments to keywords
(If the ANSI85 option is not specified, the translator expects COBOL words to consist entirely of upper-case characters.)

Notes:

1. The translator does not translate lower-case text into upper-case. Some names in COBOL text, for example, file names and transaction IDs, must match with externally defined names. Such names should always be entered in the same case as the external definition.
2. CBL and PROCESS statements must be in upper case.

COBOL3 translator option

If you are using the COBOL3 option you must use the COBOL/VSE compiler to process your application programs.

Lower-case characters can occur anywhere in any COBOL word, including user-defined names, system names, and reserved words.

Translator action: The translator listing and output preserve the case of COBOL text as entered.

In addition, the translator accepts mixed case in:

- Translator options
- EXEC CICS commands, both for keywords and for arguments to keywords
- CBL and PROCESS statements
- Compiler directives such as EJECT and SKIP1

The translator does not translate lower-case text into upper-case. Some names in COBOL text, for example file names and transaction IDs, must match with externally defined names. Such names must always be entered in the same case as the external definition.

C considerations

You can use the following compilers:

- IBM C for VSE/ESA 1.1

All the EXEC CICS commands available in COBOL, PL/I, and assembler applications are also supported in C applications, with the exception of those commands related to nonstructured exception handling:

- EXEC CICS HANDLE CONDITION option (with or without a label)
- EXEC CICS HANDLE AID option (with or without a label)
- EXEC CICS IGNORE CONDITION option
- EXEC CICS PUSH HANDLE
- EXEC CICS POP HANDLE
- EXEC CICS HANDLE ABEND (label)

Use of these commands is diagnosed by the translator.

In C, working storage consists of the stack and the heap. The location of the stack and heap, with respect to the 16MB line, are controlled by the ANYWHERE and BELOW options on the stack and heap run time options. The default is that both the stack and heap are located above the 16MB line.

On return from a C application, any value passed by C via the exit function or the return statement is saved in EIBRESP2.

A set of sample application programs is provided to show how CICS commands can be used in a program written in the C language.

<i>Table 5. Sample programs</i>			
Sample program	Map set	Map source	Transaction ID
DFH\$DMNU Operator instruction (3270)	DFH\$DGA	DFH\$DMA	DMNU
DFH\$DALL Update (3270)	DFH\$DGB	DFH\$DMB	DINQ, DADD, DUPD
DFH\$DBRW Browse (3270)	DFH\$DGC	DFH\$DMC	DBRW
DFH\$DREN Order entry (3270)	DFH\$DGK	DFH\$DMK	DORD
DFH\$DCOM Order entry queue print (3270)	DFH\$DGL	DFH\$DML	DORQ
DFH\$DREP Report (3270)	DFH\$DGD	DFH\$DMD	DREP

The transaction and program definitions are provided in group DFH\$DFLA in the CSD and should be installed using the command:

```
CEDA INSTALL GROUP(DFH$DFLA)
```

The following record description files are provided as C language header files:

- DFH\$DFIL — FILEA record descriptor
- DFH\$DL86 — L860 record descriptor

Data declarations needed for C

The following data declarations are provided by CICS for C:

- Execution interface block definitions (EIB).
- BMS screen attributes definitions—C versions of the DFHBMSCA, DFHMSRCA, and DFHAID files are supplied by CICS, and may be included by the application programmer when using BMS.
- DL/I support—a C language version of DFHDIB is included by the DLI translator if the translator option has been specified. (You have to include DLIUIB if the CALL DLI interface is used.)

The EIB declarations are enclosed in #ifndef and #endif lines, and are included in all translated files. The C compiler ignores duplicated declarations. The inserted code contains definitions of all the fields in the EIB, coded in C..

Naming EIB fields

Within a C application program, fields in the EIB are referred to in lowercase and fully qualified as, for example, dfheiptr->eibtrnid, in contrast to EIBTRNID as used in other CICS applications.

Data types in EIB fields: The following mapping of data types is used:

- Halfword binary integers are defined as “short int”
- Fullword binary integers are defined as “long int”
- Single-character fields are defined as “unsigned char”
- Character strings are defined as “unsigned char” arrays

Restrictions

The following restrictions apply to a C program that is to be used as a CICS application program:

- In a C application, every EXEC CICS command is treated as if it had the NOHANDLE or RESP option specified. This means that the set of “system action” transaction abends that result from a condition occurring, but not being handled, is not possible in a C application. Control always flows to the next instruction, and it is up to the application to test for a normal response.
- EXEC CICS HANDLE ABEND PROGRAM commands are allowed, but you cannot use EXEC CICS PUSH HANDLE or EXEC CICS POP HANDLE.
- On return from an EXEC CICS SEND MAP command with the ACCUM option, if an EXEC CICS HANDLE CONDITION OVERFLOW command is not active, the OVERFLOW condition is not raised, and the RESP field indicates a normal response. If, however, you specify the NOFLUSH option of the EXEC CICS SEND MAP command, the OVERFLOW condition is indicated in the RESP field.
- C programs may call, or be called by, or address application programs written in other languages supported by CICS. This may be done by means of EXEC CICS LINK, EXEC CICS XCTL, and EXEC CICS RETURN commands. Application data can be passed between programs using the COMMAREA option on each command.
- C language programs **must** be link-edited with the attributes AMODE(31), and may reside above the 16MB line. Storage allocation is performed in accordance with the LE run-time options you have established. See the *LE/VSE Programming Guide* for more information about running C applications.
- Simulation of extended precision floating point is not supported.
- C language does not support packed decimal data. The application has access to packed decimal data using the C character-string data type. No native C functions are available to perform arithmetic on this data, but you can write your own functions to manipulate this data.
- You can easily use HOURS, MINUTES, and SECONDS options. You may define expiration times using TIME or INTERVAL options if you provide functions to handle them in your application.
- You can enter all CICS keywords in mixed case, except for CICS keywords on #pragma directives, which must be in uppercase only.
- Do not omit the length operand from commands that require a length operand (for example, EXEC CICS READ, EXEC CICS READNEXT, EXEC CICS READPREV, and EXEC CICS WRITE).
- All native C functions are allowed in the source program, but the following functions are not executable and result in return codes or pointers indicating that the function has failed.

CDUMP	FETCH
CLOCK	RELEASE
CSNAP	SETLOCALE
CTDLI	SVC99
CTEST	SYSTEM
CTRACE	

For further information, see the *LE/VSE C Run-Time Library Reference*.

- Native C functions are implemented in the C run-time library.
- Native C file operations operate only on files opened with *type=memory* specified. I/O to CICS-supported access methods must use the CICS API.
- C uses a null character as an end-of-string marker. CICS, however, does not recognize a null as an end-of-string marker. Therefore you must take care when using the C function, for example (*strcmp*), to operate CICS data areas.
- Two arguments, *argc* and *argv*, are normally passed to a C main function. *argc* denotes the count of the number of variables passed; *argv* is the actual variable string. In CICS, the value of *argc* is 1, *argv[0]* is the transaction ID, and *argv[1]* is NULL.
- Where CICS expects a fixed-length character string such as a program name, map name, or queue name, the translator pads it with blanks up to the required length.

For EXEC DLI commands, the SEGMENT name is padded by the translator if a literal is passed.

- Take care not to use field names that cause the C compiler to abend, even though they are acceptable to the assembler compiler. These include \$, #, and @. See the *C/VSE User's Guide* for more information.
- The COMMAREA is restricted in a mixed addressing mode transaction environment. For a description of the restriction, see "Mixed addressing mode transactions" on page 383.

Passing values as arguments

Arguments in C are copied to the program stack at run time, where they are read by the function. These arguments can either be values in their own right, or they can be pointers to areas of memory that contain the data being passed. Passing a pointer is also known as passing a value by reference.

Other languages, such as COBOL and PL/I, pass their arguments **by reference**, which means that the compiler passes a list of addresses pointing to the arguments to be passed. This is the call interface supported by CICS. To pass an argument by reference, you prefix the variable name with **&**, unless it is already a pointer, as in the case when an array is being passed.

As part of the build process, the compiler may convert arguments from one data type to another. For example, an argument of type **char** may be converted to type **short** or type **long**.

When you send values from a C program to CICS, the translator takes the necessary action to generate code that results in an argument list of the correct format being passed to CICS. The translator does not always have enough information to enable it to do this, but in general, if the argument is a single-character or halfword variable, the translator makes a precall assignment to

a variable of the correct data type and passes the address of this temporary variable in the call.

When you receive data from CICS, the translator prefixes the receiving variable name with **&**, which causes the C compiler to pass it values **by reference** rather than **by value** (with the exception of a character string name, which is left unchanged). Without the addition of **&**, the compiler would copy the receiving variable and then pass the address of the copy to CICS. Any promotion occurring during this copying could result in data returned by CICS being lost.

Table 6 on page 43 shows the rules that apply when passing values as arguments in EXEC CICS commands.

<i>Table 6. Rules for passing values as arguments in EXEC CICS commands</i>		
Data type	Usage	Coding the argument
Character literal	Data-value (S)	The user must specify the character literal directly. The translator takes care of any required indirection.
Character variable (char)	Data-area (R)	The user must specify a pointer to the variable, possibly by prefixing the variable name with & .
	Data-value (S)	The user must specify the character variable directly. The translator takes care of any required indirection.
Character string literal	Name (S)	The user can either code the string directly as a literal string or use a pointer that points to the first character of the string.
Character string variable	Data-area (R) Name (S)	Whether receiving or sending, the argument should be the name of the character array containing the string—the address of the first element of the array.
Integer variable (short, long, or int)	Data-area (R)	The user must specify a pointer to the variable, possibly by prefixing the variable name with & .
	Data-value (S)	The user must specify the name of the variable. The translator looks after any indirection that is required.
Integer constant (short, long, or int)	Data-value (S)	The user must specify the integer constant directly. The translator takes care of any required indirection.
Structure or union	Data-area (S) Data-area (R)	The user must code the address of the start of the structure or union, possibly by prefixing its name with & .
Array (of anything)	Data-area (R) Data-value (S)	The translator does nothing. You must code the address of the first member of the array. This is normally done simply by coding the name of the array, which the compiler interprets as the address of the first member.
Pointer (to anything)	Ptr-ref (R) Data-area (S)	Whether receiving or sending, the argument should be the name of the variable which denotes the address of interest. The translator takes care of the extra level of indirection that is necessary to allow CICS to update the pointer.
Note: (R) indicates “Receiver”, where data is being received from CICS; (S) indicates “Sender”, where data is being passed to CICS.		

EXEC CICS ADDRESS EIB command

The address of the exec interface block (EIB) is not passed as an argument to a C main function. This means that C functions must use the EXEC CICS ADDRESS EIB command to obtain the address of the EIB.

Addressability is achieved by using the command:

```
EXEC CICS ADDRESS EIB(dfheiptr);
```

or by passing the EIB address or particular fields therein as arguments to the CALL statement that invokes the external procedure.

If access to the EIB is required, an EXEC CICS ADDRESS EIB command is required at the beginning of each application. This applies to any commands that include EIBRESP or EIBRESP2 options.

EXEC CICS ADDRESS COMMAREA

The address of the communication area is not passed as an argument to a C main function. This means that C functions must use EXEC CICS ADDRESS COMMAREA to obtain the address of the communication area.

PL/I considerations

Compilers supported

You can use the following PL/I compiler to process your PL/I application programs:

- LE for VSE/ESA PL/I 1.1

Restrictions

The following restrictions apply to a PL/I program that is to be used as a CICS application program. See the *LE/VSE Programming Guide* for more guidance information about these facilities.

- You cannot use the multitasking built-in functions:

COMPLETION
PRIORITY
STATUS

- You cannot use the multitasking options:

EVENT
PRIORITY
TASK

- Do not use the following PL/I statements:

CLOSE	OPEN
DELAY	PUT
DELETE	READ
DISPLAY	REWRITE
EXIT	STOP
GET	UNLOCK
HALT	WRITE
LOCATE	

You are provided with CICS commands for the storage and retrieval of data, and for communication with terminals. (However, you can use CLOSE, PUT, and OPEN for SYSPRT.)

- The COMMAREA is restricted in a mixed addressing mode transaction environment. For a description of the restriction, see “Mixed addressing mode transactions” on page 383.
- You cannot use PL/I Sort/Merge.
- You cannot use static storage (except for read-only data).
- If you declare a variable with the STATIC attribute and EXTERNAL attribute, you must also include the INITIAL attribute. If you do not, such a declaration generates a common CSECT that cannot be handled by CICS.
- You cannot use the PL/I 48-character-set option in EXEC CICS statements.
- You cannot use lowercase keywords in EXEC CICS statements or variable names in your PL/I programs.
- Do not define variables or structures with variable names that are the same as variable names generated by the translator. Care must be taken with the LIKE keyword to avoid implicitly generating such variable names.
- All procedure statement parameters must be in uppercase.
- If a CICS command uses the SUBSTR built-in function in defining a data value, it must include a LENGTH option to specify the data length. If it does not, the translator generates a PL/I call including an invocation of the CSTG built-in function in the form:

```
CSTG(SUBSTR(...,...))
```

This is rejected by the compiler.

OPTIONS(MAIN) specification

If OPTIONS(MAIN) is specified in an application program, that program can be the first program of a transaction, or control can be passed to it by means of an EXEC CICS LINK or an EXEC CICS XCTL command.

An application program in which OPTIONS(MAIN) is not specified cannot be the first program in a transaction, nor can it have control passed to it by an EXEC CICS LINK or EXEC CICS XCTL command, but it can be link-edited to a main program.

The definition of the EIB is generated in each program. However, in programs other than that declared with OPTIONS(MAIN), addressability to the EIB is the user's responsibility.

Addressability is achieved by using the command:

```
EXEC CICS ADDRESS EIB(DFHEIPTR)
```

or by passing the EIB address, or particular fields in it, as arguments to the CALL statement that invokes the external procedure.

If access to the EIB is required, an EXEC CICS ADDRESS EIB command is required at the beginning of each application. This applies to any commands that include EIBRESP or EIBRESP2 options.

Program segments

Segments of programs can be translated by the command language translator, stored in their translated form, and later included in the program to be compiled.

Subsequent copying or manipulating of statements originally inserted by the CICS translator in an application program may produce unpredictable results.

The external procedure must always be passed through the CICS translator, even when all its commands are in included segments.

PL/I and dynamic storage

Storage allocation will be performed in accordance with the LE-run-time options you have established. For more information about running PL/I applications see the LE/VSE Programming Guide.

Assembler-language and HLASM considerations

All the considerations described in this section apply to both assembler language and HLASM.

The following instructions cannot be used in an assembler language program that is to be used as a CICS application program:

- COM (identify blank common control section)
- ICTL (input format control)
- OPSYN (equate operation code)

Working storage is allocated above or below the 16MB line, according to the value of the DATALOCATION parameter on the RDO PROGRAM resource definition.

When using BAKR instructions (branch and stack) to provide linkage between assembler programs, care should be taken that the linked-to program does not issue EXEC CICS requests. If CICS receives control and performs a task switch before the linked-to program returns via a PR instruction (program return), then other tasks might be dispatched and issue further BAKR / PR calls. These will modify the linkage-stack and result in the wrong environment being restored when the original task issues its PR instruction.

If the access register support of the High-Level Assembler (HLASM) is used, you must be in primary addressing mode when invoking any CICS service. The primary address space must be the home address space. All parameters passed to CICS must reside in the primary address space.

The following restrictions apply to an assembler language application program executing in 31-bit mode:

- The COMMAREA is restricted in a mixed addressing mode transaction environment. For a description of the restriction, see "Mixed addressing mode transactions" on page 383.
- All addresses must be handled as 31-bit addresses.
- Instructions that are used differently in 24-bit and 31-bit addressing mode (for example, LA, BAL, BALR, CALL, and AL(3)) must be used correctly.

- The application program must not violate the 31-bit addressing rules described in the *VSE/ESA Extended Addressability* manual.

Note: Some VSE services require the caller to be in AMODE(24) and can lead to unpredictable results if called in AMODE(31). Refer to the VSE/ESA System Macros Reference for details of the specific addressing mode requirements of a VSE service.

Commands contained within macros and COPY code

Macros that generate commands, and COPY code that contains commands, must be translated and stored in the source library in translated form for later inclusion by the assembler.

Invoking assembler-language application programs with a call

Assembler-language application programs that contain commands can have their own RDO program definition. Such programs can be invoked by COBOL, C, PL/I, or assembler-language application programs using EXEC CICS LINK or EXEC CICS XCTL commands (see Chapter 9, “Dealing with exception conditions” on page 141). However, because programs that contain commands are invoked by a system standard call, they can also be invoked by a COBOL, C, or PL/I CALL statement or by an assembler-language CALL macro.

A single CICS application program, as defined in an RDO program definition, may consist of separate CSECTs assembled or compiled separately, but linked together.

An assembler application program that contains commands can be linked with other assembler-language programs, or with programs in one, and only one, of the high-level languages (COBOL, C, or PL/I). When you do this, the high-level language program must be linked first.

If an assembler language program contains command-level calls, and is called from a high-level language program, it requires its own CICS interface stub. The linkage editor message 2143I is issued, but this can be ignored.

Because assembler-language application programs containing commands are always passed the parameters EIB and COMMAREA when invoked, the CALL statement or macro must pass these two parameters followed, optionally, by other parameters.

An assembler-language application program that is called by another begins with the DFHEIENT macro and ends with the DFHEIRET macro. The CICS translator inserts these for you, so if the program contains EXEC CICS commands and is to be passed to the translator you do not need to code these macros.

Part 2. Application design

Chapter 3. Designing efficient applications	51	Additional terminal control considerations	81
Program structure	51	Chapter 4. Sharing data across transactions	83
Program size	52	Common work area (CWA)	83
Choosing between pseudoconversational and conversational design	53	Protecting the CWA	84
General programming techniques	55	TCTTE user area (TCTUA)	87
Virtual storage	55	COMMAREA in EXEC CICS RETURN commands	88
Reducing paging effects	56	Display screen	89
Exclusive control of resources	59	Chapter 5. Affinity among transactions	91
Processor usage	60	What is transaction affinity?	91
Recovery design implications	60	Intertransaction affinity	91
Terminal interruptibility	61	Transaction-system affinity	92
Operational control	61	Programming techniques used by CICS application programs	92
Operating system waits	62	Safe techniques	92
Writing a message to the system console	62	Unsafe techniques	93
Other suggestions for improving performance	62	Suspect techniques	93
Runaway tasks	62	Safe programming techniques	94
Auxiliary trace	63	The COMMAREA	94
The NOSUSPEND option	63	The TCTUA	95
Multithreading	64	Unsafe programming techniques	97
Storing data within a transaction	64	Using the common work area	97
Transaction work area (TWA)	65	Using GETMAIN SHARED storage	98
User storage	65	Using the LOAD PROGRAM HOLD command	98
COMMAREA in EXEC CICS LINK and EXEC CICS XCTL commands	66	Sharing task-lifetime storage	100
Program storage	67	Using the WAIT EVENT command	101
Lengths of areas passed to CICS commands	67	Using ENQ and DEQ commands	102
LENGTH options	67	Suspect programming techniques	103
Journal records	68	Using temporary storage	103
Data set definitions	68	Using transient data	105
Recommendation	68	Using the RETRIEVE WAIT and START commands	106
EXEC interface stubs	68	Using the START and CANCEL REQID commands	108
Assembler language	68	Using the EXEC CICS DELAY and EXEC CICS CANCEL REQID commands	109
Temporary storage	69	Using the EXEC CICS POST and EXEC CICS CANCEL REQID commands	111
Intrapartition transient data	70	Detecting intertransaction affinities	112
GETMAIN SHARED command	71	Intertransaction affinities caused by application generators	113
Your own data sets	71	Duration and scope of intertransaction affinities	113
Data operations	71	Affinity transaction groups	113
DL/I database operations	71	Relations and lifetimes	114
VSAM and DAM data set operations	72	Recommendations	117
Browsing data sets	74		
Logging data sets	74		
Sequential data set access	74		
Terminal operations	75		
Length of the data stream sent to the terminal	75		
Basic mapping support considerations	75		
Page-building and routing operations	78		
Requests for printed output	80		

Chapter 6. Intercommunication considerations	119	An example of exception handling in C	143
Design considerations	119	An example of exception handling in COBOL	145
Programming language	120	Modifying the default CICS exception handling	146
Transaction routing	120	Use of the EXEC CICS HANDLE CONDITION command	147
Function shipping	120	Use of the EXEC CICS HANDLE CONDITION ERROR command	148
Distributed program link (DPL)	121	Use of the EXEC CICS IGNORE CONDITION command	149
Using the distributed program link function	122	Use of the EXEC CICS HANDLE ABEND command	150
Examples of distributed program link	124	EIBRESP and NOHANDLE options	150
Programming considerations for distributed program link	128	How CICS keeps track of what to do	151
Asynchronous processing	132	EXEC CICS PUSH HANDLE and EXEC CICS POP HANDLE commands	152
Distributed transaction processing (DTP)	132	Chapter 10. Access to system information	153
Common Programming Interface (CPI) Communications	132	System programming commands	153
External CICS interface (EXCI)	133	EXEC interface block (EIB)	153
Chapter 7. Recovery considerations	135	Chapter 11. Abnormal termination recovery	155
Journaling	135	The EXEC CICS HANDLE ABEND command	155
Journal records	135	Restrictions on retrying operations	157
Journal output synchronization	136	Exception handling	158
Syncpointing	137	Trace	161
Chapter 8. Minimizing errors	139	Trace entry points	161
Protecting CICS from application errors	139	Monitoring	162
Testing applications	139	Dumping	163
Chapter 9. Dealing with exception conditions	141		
Default CICS exception handling	141		
Handling exception conditions by in-line code	142		
How to use the RESP and RESP2 options	143		

Chapter 3. Designing efficient applications

In this chapter, design changes are suggested that can improve performance and efficiency without much change to the application program itself. Other aspects of application design (such as productivity, readability of the code, usability, standards, and the effort involved) are not discussed. Always try to consider these, however, when deciding what to improve and how to improve it.

In order of priority, think about:

1. Application and system design
2. Task design
3. Program design
4. Program coding

If you have a performance problem that applies in a particular situation, try to isolate the changes you make so that their effects apply only in that situation. After fixing the problem and testing the changes, use them in your most commonly-used programs and transactions, where the effects on performance are most noticeable.

Program structure

Two main aspects of design to consider are:

- Program size
- Choosing between pseudoconversational and conversational design

The original versions of CICS ran on machines without virtual storage. Storage sizes were generally very small by current standards and consequently, storage was almost always the critical resource. As a result, CICS programmers were encouraged to keep everything as small as possible: programs, data areas, EXEC CICS GETMAIN commands, and so on. Programs were loaded on demand, and one copy of the program was shared among all concurrent users.

With the virtual storage concept and the trend toward much larger and less expensive storage, this constraint on storage eased somewhat. Virtual storage seemed almost unlimited at first, although there was still often a shortage of the underlying real storage. Eventually, CICS systems grew so much in volume and complexity that even virtual storage became a constraint, but the trade-off of storage resources against others definitely shifted during this time. Whereas extensive programming efforts and long instruction sequences would once have been invested to save even a modest amount of storage, now both CICS as a system, and the application programmers using it, are willing to trade some storage for processor savings and additional function. Some indications of this are:

- Many programs that are now kept resident in virtual storage, rather than being loaded on demand.
- There is much more extensive use of high-level programming languages.
- The EXEC CICS interface, which requires larger control blocks than the macro interface, but saves enormously in programming effort.

Program size

The early emphasis on small programs led CICS programmers to break up programs into units that were as small as possible, and to transfer control using the EXEC CICS XCTL command, or link using the EXEC CICS LINK command, between them. In current systems, however, it is not always better to break up programs into such small units, because there is CICS processing overhead for each transfer and, for EXEC CICS LINK commands, there is also storage overhead for the register save areas (RSAs).

In fact, for modestly sized blocks of code that are processed sequentially, inline code is most efficient. The exceptions to this rule are blocks of code that are:

- Fairly long and used independently at several different points in the application
- Subject to frequent change (in which case, you balance the overhead of EXEC CICS LINK or EXEC CICS XCTL commands with ease of maintenance)
- Infrequently used, such as error recovery logic and code to handle uncommon data combinations

If you have a block of code that, for one of these reasons, has to be written as a subroutine, the best way of dealing with this from a performance viewpoint is to use a closed subroutine within the invoking program (for example, code that is dealt with by a PERFORM command in COBOL). If it is needed by other programs, it should be a separate program. A separate program can be called, with a CALL statement (macro), or it can be kept separate and processed using an EXEC CICS XCTL or a EXEC CICS LINK command. Execution overhead is least for a CALL, because no CICS services are invoked; for example, the working storage of the program being called is **not** copied.

A called program, however, must be linked into the calling one and so cannot be shared by other programs that need it unless you use special COBOL, C, or PL/I facilities. A called subroutine is loaded as part of each program that CALLs it and hence uses more storage. Thus, subsequent transactions using the program may or may not have the changes in the working storage made to the called program. This depends entirely on whether CICS has loaded a new copy of the program into storage. Furthermore, if the called program is written in COBOL, it cannot issue CICS commands unless you compile it using the COBOL/VSE compiler.

Overhead (but also flexibility) is higher with EXEC CICS XCTL and EXEC CICS LINK commands. Both processor and storage requirements are much greater for EXEC CICS LINK than for EXEC CICS XCTL. Therefore, if the invoking program does not need to have control returned to it after the invoked program is processed, it should use the EXEC CICS XCTL command.

The load module resulting from any application program can occupy up to 16MB of main storage, although this is not recommended. You may get an abend code of APCG if your program size exceeds the available space in the appropriate CICS dynamic storage area (DSA).

Choosing between pseudoconversational and conversational design

In a **conversational** transaction, the length of time spent in processing each of a user's responses is extremely short when compared with the amount of time waiting for the input. A conversational transaction is one that involves more than one input from the terminal, so that the transaction and the user enter into a conversation. A **nonconversational** transaction has only one input (the one that causes the transaction to be invoked). It processes that input, responds to the terminal, and terminates.

Processor speeds, even allowing for accessing data sets, are considerably faster than terminal transmission times, which are considerably faster than user response times. This is especially true if users have to think about the entry or have to enter many characters of input. Consequently, conversational transactions tie up storage and other resources for much longer than nonconversational transactions.

A **pseudoconversational** transaction sequence contains a series of nonconversational transactions that look to the user like a single conversational transaction involving several screens of input. Each transaction in the sequence handles one input, sends back the response, and terminates.

Before a pseudoconversational transaction terminates, it can pass data forward to be used by the next transaction initiated from the same terminal, whenever that transaction arrives. A pseudoconversational transaction can specify what the next transaction is to be, and it does this by setting the transaction identifier of the transaction that handles the next input. However, be aware that, if another transaction is started for that device, it may interrupt the pseudoconversational chain you have designed.

No transaction exists for the terminal from the time a response is written until the user sends the next input and CICS starts the next transaction to respond to it. Information that would normally be stored in the program between inputs is passed from one transaction in the sequence to the next using the COMMAREA or one of the other facilities that CICS provides for this purpose. (See Chapter 4, "Sharing data across transactions" on page 83 for details.)

There are two major issues to consider in choosing between conversational and pseudoconversational programming.

- The effect of the transaction on **contention** resources, such as storage and processor usage. Storage is required for control blocks, data areas, and programs that make up a transaction, and the processor is required to start, process, and terminate tasks. Conversational programs have a **very** high impact on storage, because they last so long, relative to the sum of the transactions that make up an equivalent pseudoconversational sequence. However, there is less processor overhead, because only one transaction is initiated instead of one for every input.
- The effect on **exclusive-use** resources, such as records in recoverable data sets, recoverable transient data queues, enqueue items, and so on. Again, a conversational transaction holds on to these resources for much longer than the corresponding sequence of nonconversational transactions. From this point of view, pseudoconversational transactions are better for quick responses, but recovery and integrity implications may mean that you prefer to use conversational transactions.

CICS ensures that changes to recoverable resources (such as data sets, transient data, and temporary storage) made by a unit of work (UOW) are made completely or not at all. A UOW is equivalent to a transaction, unless that transaction issues EXEC CICS SYNCPOINT commands, in which case a UOW lasts between syncpoints. For a more detailed description of syncpoints and UOWs, see the *CICS Recovery and Restart Guide*.

When a transaction makes a change to a recoverable resource, CICS makes that resource unavailable to any other transaction that wants to change it until the original transaction has completed. In the case of a conversational transaction, the resources in question may be unavailable to other terminals for relatively long periods.

For example, if one user tries to update a particular record in a recoverable data set, and another user tries to do so before the first one finishes, the second user's transaction is suspended.

This has advantages and disadvantages. You would not want the second user to begin updating the record while the first user is changing it, because one of them is working from what is about to become an obsolete version of the record, and these changes erase the other user's changes. On the other hand, you also do not want the second user to experience the long, unexplained wait that occurs when that transaction attempts to READ for UPDATE the record that is being changed.

If you use pseudoconversational transactions, however, the resources are only very briefly unavailable (that is, during the short component transactions). However, unless all recoverable resources can be updated in just one of these transactions, recovery is impossible because UOWs cannot extend across transactions. So, if you cannot isolate updates to recoverable resources in this way, you must use conversational transactions.

The previous example poses a further problem for pseudoconversational transactions. Although you could confine all updating to the final transaction of the sequence, there is nothing to prevent a second user from beginning an update transaction against the same record while the first user is still entering changes. This means that you need additional application logic to ensure integrity. You can use some form of enqueueing, or you can have the transaction compare the original version of the record with the current version before actually applying the update.

There is one further difference between conversational and pseudoconversational transactions. After a user begins a conversational transaction, no messages can be delivered to that user's terminal from any source except the transaction being processed. Broadcast messages, and messages sent by other transactions, are not displayed until the conversational transaction has ended. In a pseudoconversational sequence, however, such messages are delivered as soon as no transaction is running for that user's terminal, which may be immediately after any screen in the sequence.

This has advantages and disadvantages. You may find it a problem if you are not able to send messages to users immediately; for example, if you have to shut down the system at an unscheduled time. However, users of pseudoconversational transactions may find it annoying to have a data entry screen overlaid by a message. The unexpected change in the screen may even cause the next

transaction in the sequence to fail. If this seems likely, design your transaction screens so that messages can occur only in a specific area.

There are factors other than performance overhead to consider when choosing between pseudoconversational and conversational design for CICS applications. The method you choose can affect how you write the application programs. You may need extra CICS requests for pseudoconversations, particularly if you are updating recoverable files. After you have done this, however, operational control (performance monitoring, capacity planning, recovery, system shutdown, and distributing system messages) may be much easier.

General programming techniques

To know how programming techniques can affect the performance and efficiency of the CICS system, it is necessary to understand a little about the environment in which CICS operates. Here, the following factors are considered:

- “Virtual storage”
- “Reducing paging effects” on page 56
- “Exclusive control of resources” on page 59
- “Processor usage” on page 60
- “Recovery design implications” on page 60
- “Terminal interruptibility” on page 61
- “Operational control” on page 61
- “Writing a message to the system console” on page 62
- “Runaway tasks” on page 62
- “Auxiliary trace” on page 63
- “The NOSUSPEND option” on page 63
- “Multithreading” on page 64

Virtual storage

A truly conversational CICS task is one that converses with the terminal user for several or many interactions, by issuing a terminal read request after each write (for example, using either an EXEC CICS SEND command followed by an EXEC CICS RECEIVE command, or an EXEC CICS CONVERSE command). This means that the task spends most of its extended life waiting for the next input from the terminal user.

Any CICS task requires some virtual storage throughout its life and, in a conversational task, some of this virtual storage is carried over the periods when the task is waiting for terminal I/O. The storage areas involved include the TCA and associated task control blocks (including EIS or EIB, and JCA) and the storage required for all programs that are in use when any terminal read request is issued. Also included are the work areas (such as copies of COBOL working storage) associated with this task’s use of those programs.

With careful design, you can sometimes arrange for only one very small program to be retained during the period of the conversation. The storage needed could be shared by other users. You must multiply the rest of the virtual storage requirement by the number of concurrent conversational sessions using that code.

By contrast, a pseudoconversational sequence of tasks requires almost all of its virtual storage only for the period actually spent processing message pairs. Typically, this takes a period of 1—3 seconds in each minute (the rest being time

waiting for operator input). The overall requirement for multiple concurrent users is thus perhaps five percent of that needed for conversational tasks. However, allow for data areas that are passed from each task to the next. This may be a COMMAREA of a few bytes or a large area of temporary storage. If it is the latter, you are normally recommended to use temporary storage on disk rather than in main storage, but that means adding extra temporary storage I/O overhead in a pseudoconversational setup, which you do not need with conversational processing.

The extra virtual storage you need for conversational applications usually means that you need a correspondingly greater amount of real storage. The paging you need to control storage involves additional overhead and virtual storage. The adverse effects of paging increase as transaction rates go up, so minimize its use as much as possible. See “Reducing paging effects” for information about doing so.

Reducing paging effects

Reducing paging effects is a technique used by CICS in a virtual-storage environment. The key objective of programming in this environment is the reduction of page faults. A page fault occurs when a program refers to instructions or data that do not reside in real storage, in which case the page in virtual storage that contains the instructions or data referred to must be paged into real storage. The more paging required, the lower the overall system performance.

Although an application program may be able to communicate directly with the operating system, the results of such action are unpredictable and can degrade performance.

An understanding of the following terms is necessary for writing application programs to be run in a virtual-storage environment:

Locality of reference	The consistent reference, during the execution of the application program, to instructions and data within a relatively small number of pages (compared to the total number of pages in a program) for relatively long periods.
Working set	The number and combination of pages of a program needed during a given period.
Reference set	Direct reference to the required pages, without intermediate storage references that retrieve useless data.

The virtual machine environment affects CICS in the same way as batch programs, and the same coding recommendations apply.

Locality of reference

Keep the instructions processed and data used in a program within a relatively small number of pages (4096-byte segments). This quality in a program is known as **locality of reference**. You can do this by:

- Making the execution of the program as linear as possible.
- Keeping any subroutines you use in the normal execution sequence as close as possible to the code that invokes them.
- Placing code inline, even if you have to repeat it, if you have a short subroutine that is called from only a small number of places.

- Separating error-handling and other infrequently processed code from the main flow of the program.
- Separating data used by such code from data used in normal execution.
- Defining data items (especially arrays and other large structures) in the order in which they are referred to.
- Defining the elements within a data structure in the approximate order in which they are referred to. For example, in PL/I, all the elements of one row are stored, then the next row, and so on. Define an array so that you can process it by row rather than by column.
- Initializing data as close as possible to where it is first used.
- Avoiding COBOL variable MOVE operations and the EXAMINE verb, because these expand into subroutine calls.
- Issuing as few GETMAIN commands as possible. It is generally better for the program to add up its requirements and do one GETMAIN command than to do several smaller ones, unless the durations of these requirements vary greatly.
- Avoiding use of the INITIMG option on an EXEC CICS GETMAIN command, if possible. It causes an immediate page reference to the storage that is obtained, which might otherwise not occur until much later in the program, when there are other references to the same area.

Note: Some of the suggestions above may conflict with your installation's programming standards if these are aimed at the readability and maintainability of the code, rather than speed of execution in a virtual-storage environment. Some structured programming methods, in particular modular programming techniques, make extensive use of the PERFORM command in COBOL (and the equivalent programming techniques in C, PL/I, and assembler language) to make the structure of the program clear. This may also result in more exceptions to sequential processing than are found in a nonstructured program. Nevertheless, the much greater productivity associated with structured code may be worth the possible loss of locality of reference.

Working set

The working set is the number and combination of pages of a program needed during a given period. To minimize the size of the working set, the amount of storage that a program refers to in a given period should be as small as possible. You can do this by:

- Writing modular programs and structuring the modules according to frequency and anticipated time of reference. Do not modularize merely for the sake of size; consider duplicate code inline as opposed to subroutines or separate modules.
- Using separate subprograms whenever the flow of the program suggests that execution is not to be sequential.
- Not tying up main storage awaiting a reply from a terminal user.
- Using command-level file control locate-mode input/output rather than move-mode. Use of multiple temporary storage queues is restricted. For programming information about temporary storage restrictions when using locate-mode input/output with the SET option, see the *CICS Application Programming Reference* manual.

- In COBOL programs, specifying constants as literals in the PROCEDURE DIVISION, rather than as data variables in the WORKING STORAGE section.
- In C and PL/I programs, using static storage for constant data.
- Avoiding the use of EXEC CICS LINK commands where possible, because they generate requests for main storage.

Reference set

Try to keep the overall number of pages that a program uses during normal operation as small as possible. These pages are termed the **reference set**, and they give an indication of the real storage requirement of the program. You can reduce the reference set by:

- Specifying constants in COBOL programs as literals in the PROCEDURE DIVISION, rather than as data variables in the WORKING STORAGE SECTION. The reason for this is that there is a separate copy of working storage for every task executing the program, whereas literals are considered part of the program itself, of which only one copy is used in CICS.
- Using static storage in C and PL/I for data that is genuinely constant, for the same reason as in the previous point.
- Reusing data areas in the program as much as possible. You can do this with the REDEFINES clause in COBOL, the union clause in C, based storage in PL/I, and ORG or equivalents in assembler language. In particular, if you have a map set that uses only one map at a time, code the DFHMSD map set definition without specifying either the STORAGE=AUTO or the BASE operand. This allows the maps in the map set to redefine one another.

Refer to data directly by:

- Avoiding long searches for data in tables
- Using data structures that can be addressed directly, such as arrays, rather than structures that must be searched, such as chains
- Avoiding methods that simulate indirect addressing

No attempt should be made to use overlays (paging techniques) in an application program. System paging is provided automatically and has superior performance. The design of an application program for a virtual-storage environment is similar to that for a real environment. The system should have all modules resident so that code on pages not referred to need not be paged in.

If the program is dynamic, the entire program must be loaded across adjacent pages before execution begins. Dynamic programs can be purged from storage if they are not being used and an unsatisfied storage request exists. Allowing sufficient dynamic area to prevent purging is more expensive than making them resident, because a dynamic program does not share unused space on a page with another program.

Exclusive control of resources

The very fundamental and powerful recovery facilities that CICS provides have performance implications. CICS serializes updates to recoverable resources so that, if a transaction fails, its changes to those resources can be backed out independently of those made by any other transaction. Consequently, a transaction updating a recoverable resource gets control of that resource until it terminates or indicates that it wants to commit those changes with a SYNCPOINT command. Other transactions requiring the same resource must wait until the first transaction finishes with it.

The primary resources that produce these locking delays are data sets, DL/I databases, temporary storage, and transient data queues. The unit on which protection is based is the individual record (key) for files, the program specification block (PSB) for DL/I databases, and the queue name for temporary storage. For transient data, the “read” end of the queue is considered a separate resource from the “write” end (that is, one transaction can read from a queue while another is writing to it).

To reduce transaction delays from contention for resource ownership, the length of time between the claiming of the resource and its release (the end of the UOW) should be minimized. In particular, conversational transactions should not own a critical resource across a terminal read.

Note: Even for nonrecoverable data sets, VSAM prevents two transactions from reading the same record for update at the same time. This enqueue ends as soon as the update is complete, however, rather than at the end of the UOW. Even this protection for a DAM data set, can be relinquished by defining them with “no exclusive control” (SERVREQ=NOEXCTL) in the DFHFCT TYPE=FILE macro for the file.

This protection scheme can also produce deadlocks as well as delays, unless specific conventions are observed. If two transactions update more than one recoverable resource, they should always update the resources in the same order. If they each update two data sets, for example, data set “A” should be updated before data set “B” in all transactions. Similarly, if transactions update several records in a single data set, they should always do so in some predictable order (low key to high, or conversely). You might also consider including the TOKEN keyword with each READ UPDATE command. See “The TOKEN option” on page 181 for information about the TOKEN keyword. Transient data, temporary storage, and user journals must be included among such resources. The *CICS Recovery and Restart Guide* contains further information on the extent of resource protection.

It may be appropriate here to note the difference between CICS enqueues on a VSAM file, and VSAM internal enqueues on a file. VSAM enqueues on the control interval (CI). Because CICS has no information about VSAM enqueue, a SHAREOPTIONS(4) data set that is updated simultaneously from batch and CICS can result in, at best, reduced performance and, at worst, an undetectable deadlock situation between batch and CICS. Avoid such simultaneous updates between batch and CICS. In any case, if a data set is updated by both batch and CICS, CICS is unable to ensure data integrity.

Processor usage

Pseudoconversational tasks require a new task to be created to process each message-pair, and to be deleted when that task has finished. The additional processor usage that this requires is also known as the ATTACH/DETACH overhead. These may include the cost of initializing a new work area for the program that is first entered. (In a conversational task, this area is retained permanently, as already mentioned.)

There may also be extra processor overhead because of extra requests needed to retrieve data passed from the previous task of the pseudoconversation, and possibly to pass data to the next task.

Recovery design implications

Many applications require a succession of interactions with the user to get all the data needed to create a file record. For example, creating an order involves header information such as customer number, date created, and date required.

Some installations may require that only complete orders are entered on the file. A conversational application might create a partial order record and then update it in stages, as the terminal operator enters items. If all the updates are to be committed and backed out together, this means retaining the protective enqueues on records throughout the conversation until the order is complete. You may need to protect both the current order being entered and the stock records that have been decreased by the number of items ordered. Thus, a whole series of enqueues could be carried forward through the conversation for several minutes, and any other user making a conflicting request might wait without warning until the end of the order. This also means that ENQ areas are held in virtual storage for this time.

If you are also using DL/I, you must keep the PSB in question scheduled from just before the first insert or get update request until the end of the order to keep the ENQs. This is a fairly large control block, and it is associated with others that manage a thread into DL/I. To allow multiple conversational users to do a long series of updates would mean a very large allocation of threads into DL/I (the maximum is 255), and a lot of virtual storage for the control blocks.

Other areas that could be involved are the dynamic log that holds all the updates. This too has to last from the first GET UPDATE request to the end of the order. It can overflow into main storage above the 16MB line.

Lastly, if the conversation went on to another order, presumably a syncpoint would be taken to commit the previous one. This could affect the ability of the program to restart after an DL/I VSE deadlock abend. (The DFHREST module would need to be modified to get over this possibility, so that it can be restarted even though an intervening syncpoint has occurred.)

In a pseudoconversational implementation, the above approach is quite impossible because updates on one task are committed independently of any other. Therefore, an order that must be complete “in one piece” must be created by just one task. However many interactions it takes to get all the necessary input, the final task has to be the one that creates the order. Data supplied earlier in the conversation must be saved somewhere between transactions—usually in

temporary storage on disk. That is, you must incur extra overhead in input/output to temporary storage while the order is built up.

If the operator is taking orders over the telephone, with no written backup material on paper, the TS data itself should be made recoverable to avoid the remote client having to dictate the order over again.

To summarize the issue: recovery places separate design constraints on **both** implementations, but the performance cost of the pseudoconversational approach is usually more acceptable.

Terminal interruptibility

When a conversational task is running, CICS allows nothing else to send messages to that task's terminal. This has advantages and disadvantages. The advantage is that unexpected messages (for example, broadcasts) cannot interrupt the user-machine dialogue and, worse, corrupt the formatted screen. The disadvantage is that the end user cannot then be informed of important information, such as the intention of the control operator to shut down CICS after 10 minutes. More importantly, the unwitting failure of the end user to terminate the conversation may in fact prevent or delay a normal CICS shutdown.

Pseudoconversational applications can allow messages to come through between message pairs of a conversation. This means that notices like shutdown warnings can be delivered. This might disturb the display screen contents, and can sometimes interfere with transaction sequences controlled by the RETURN command with the TRANSID option. However, this can be prevented by forcing the terminal into NOATI status during the middle of a linked sequence of interactions (like building one order in the example above), or by judiciously allowing space at the top or bottom of the screen for use by any message being sent to the screen. The ERRATT option is useful here, but does not control **all** messages generated by CICS.

The main problem is that CICS shutdown could occur in mid sequence—in our example, when an order is only partly built. This is because CICS cannot distinguish between the last CICS task of a user transaction and any other. You can guard against this by ensuring that users are warned of any intended shutdown sufficiently far in advance, so they do not start work that they might not complete in time.

Operational control

The CICS system initialization parameter MXT specifies the maximum number of user tasks that can exist in a CICS system at the same time. MXT is invaluable for avoiding short-on-storage (SOS) conditions and for controlling contention for resources in CICS systems. It works by delaying the creation of user tasks to process input messages, if there are already too many activities in the CICS system. In particular, the virtual storage occupied by a message awaiting processing is usually much less than that needed for the task to process it, so you save virtual storage by delaying the processing of the message until you can do so quickly.

Transaction classes are useful in limiting the number of tasks of a particular user-defined type, or class, if these are heavy resource users. See the *CICS Performance Guide* for further information on transaction classes.

To summarize, although conversational tasks may be easier to write, they have serious disadvantages—both in performance (especially the need for virtual storage) and in their effect on the overall operability of the CICS systems containing them. Processors are now larger, with more real storage and more power than in the past, and this makes conversational tasks less painful in small amounts; but if you use conversational applications, you may rapidly run into virtual storage constraint. If you run application programs “above the line”, you will probably encounter ENQ problems before running into virtual storage constraints.

Operating system waits

Avoid using facilities that cause operating system waits. All CICS activity stops when one of these waits occurs, and all transactions suffer response delays. The chief sources of such waits are:

- Extrapartition transient data sets. (See “Sequential data set access” on page 74.)
- Those COBOL, C, and PL/I language facilities that you must not use in CICS programs and for which CICS generally provides alternative facilities. For guidance information about the language restrictions, see Chapter 2, “Language considerations” on page 19.
- SVCs and assembler language macros that invoke operating system services, such as write-to-operator (WTO).

Writing a message to the system console

You can use the EXEC CICS WRITE OPERATOR command to write to the system console and, if necessary, wait for a reply. Currently, many application programs use the unsupported macro DFHWTO to do this.

If you are using unsupported methods for the write-to-operator function (such as the PL/I DISPLAY or COBOL DISPLAY UPON CONSOLE statement), modify your programs to use the supported methods only. For example, replace a PL/I DISPLAY statement with an EXEC CICS WRITE OPERATOR command. If you have programs that contain DFHWTO, WTO, WTOR, PUTR, or PUT macros, or SVC 30 calls, change them to use the EXEC CICS WRITE OPERATOR command. This minimizes the possibility of a partition wait.

Other suggestions for improving performance

Here are some other suggestions that can improve the performance of application code.

Runaway tasks

CICS resets a task’s runaway time (ICVR) only when a task is suspended. An EXEC CICS command cannot be guaranteed to cause a task to suspend during processing because of the unique nature of each CICS implementation. The runaway time may be exceeded causing a task to abend AICA. This abend can be prevented by coding an EXEC CICS SUSPEND command in the application. This causes the dispatcher to suspend the task that issued the request and allow any task of higher priority to run. If there is no task ready to run, the program that issued the suspend is resumed. For further information about abend AICA, see the *CICS Problem Determination Guide*.

Auxiliary trace

Use auxiliary trace to review your application programs. For example, it can show up any obviously unnecessary code, such as:

- Data set EXEC CICS BROWSE from the beginning of a data set instead of using the RIDFLD option to avoid unnecessary EXEC CICS GETNEXT requests
- Too many or too large EXEC CICS GETMAIN commands
- Failure to release storage when it is no longer needed
- Unintentional logic loops
- Failure to unlock records held for exclusive control which are no longer needed

The NOSUSPEND option

The default action for the ENQBUSY, NOJBUFSP, NOSPACE, NOSTG, QBUSY, SESSBUSY, and SYSBUSY conditions is to suspend the execution of the application until the required resource (for example, storage) becomes available, and then resume processing the command. The commands that can give rise to these conditions are EXEC CICS ALLOCATE, ENQ, GETMAIN, WRITE JOURNALNUM, READQ TD, and WRITEQ TS.

On these commands, you can use the NOSUSPEND option (also known as the NOQUEUE option in the case of the EXEC CICS ALLOCATE command) to inhibit this waiting and cause an immediate return to the instruction in the application program following the command.

CICS maintains a table of conditions referred to by the HANDLE CONDITION and IGNORE CONDITION commands in a COBOL application program¹. Execution of these commands either updates the existing entry, or causes a new entry to be made if the condition has not yet been the subject of such a command. Each entry indicates one of the three states described below:

- A label is currently specified, as follows:
HANDLE CONDITION condition(label)
- The condition is to be ignored, as follows:
IGNORE CONDITION
- No label is currently specified, as follows:
HANDLE CONDITION

When the condition occurs, the following tests are made:

1. If the command has the NOHANDLE or RESP option, control returns to the next instruction in the application program. Otherwise, the condition table is scanned to see what to do.
2. If an entry for the condition exists, this determines the action.
3. If no entry exists and the default action for this condition is to suspend execution:

¹ HANDLE CONDITION and IGNORE CONDITION commands are not supported for C programs

- If the command has the NOSUSPEND or NOQUEUE option, control returns to the next instruction.
 - If the command does not have one of these options, the task is suspended.
4. If no entry exists and the default action for this condition is to abend, a second search is made looking for the ERROR condition:
- If found, this entry determines the action.
 - If ERROR is searched for and not found, the task is abended.

For programming information about the use of the NOSUSPEND option, see the *CICS Application Programming Reference* manual.

Multithreading

Multithreading is a technique that allows a single copy of an application program to process several transactions concurrently. For example, the first section of an application program may be processing one transaction. When that section is completed (in general, signaled by the execution of an EXEC CICS command that causes a wait), processing of another transaction using a different section of the application program may take place. (Compare this with single-threading, which is the execution of a program to completion: processing of one transaction is completed before another transaction is started.)

Multithreading requires that all CICS application programs be reentrant; that is, they must be serially reusable between entry and exit points. CICS application programs using the EXEC CICS interface obey this rule automatically.

For COBOL and C programs, reentrancy is ensured by a fresh copy of working storage being obtained each time the program is invoked. Always use the RENT option on the compile or pre-link utility, even for C programs that do not have writable static variables and are naturally reentrant.

Temporary variables and DFHEPTR fields inserted by the CICS translator are usually defined as writable static variables and require the RENT option. For these programs to stay reentrant, variable data should not appear as static storage in PL/I, or as a DC in the program CSECT in assembler language.

Take care if a program involves lengthy calculations. Because an application program retains control from one EXEC CICS command to the next, processing of other transactions is completely excluded. However, the task-control SUSPEND command can be used to allow other transaction processing to proceed; see Chapter 24, "Task control" on page 373 for details. Note that the runaway task time interval is controlled by the RUNAWAY option on an RDO TRANSACTION resource definition and the system initialization parameter ICVR. CICS purges a task that does not return control before expiry of the IVCR-specified interval.

Storing data within a transaction

CICS provides a variety of facilities for storing data within and between transactions. Each one differs according to how available it leaves data to other programs within a transaction and to other transactions; in the way it is implemented; and in its overhead, recovery, and enqueueing characteristics.

Storage facilities that exist for the lifetime of a transaction include:

- “Transaction work area (TWA)” on page 65
- “User storage”
- “COMMAREA in EXEC CICS LINK and EXEC CICS XCTL commands” on page 66
- “Program storage” on page 67

All of these areas are main storage facilities and come from the same basic source—the dynamic storage areas (DSAs) and extended dynamic storage areas (EDSAs). None of them is recoverable, and none can be protected by resource security keys. They differ, however, in accessibility and duration, and therefore each meets a different set of storage needs.

Transaction work area (TWA)

The transaction work area (TWA) is allocated when a transaction is initiated. It lasts for the entire duration of the transaction, and is accessible to all local programs in the transaction. Any remote programs that are linked via a distributed program link command do not have access to the TWA of the client transaction. The size of the TWA is determined by the TWASIZE option on the RDO TRANSACTION resource definition. If this size is nonzero, the TWA is always allocated. See the *CICS Resource Definition Guide* for more information about the TWASIZE option.

Processor overhead associated with using the TWA is minimal. You do not need a GETMAIN command to access it, and you address it using a single ADDRESS command. The TASKDATAKEY option on an RDO TRANSACTION resource definition governs whether the TWA is obtained in CICS-key or user-key storage. (See Chapter 26, “Storage control and storage protection” on page 391 for a full explanation of CICS-key and user-key storage.) The TASKDATALOC option of the RDO TRANSACTION governs whether the acquired storage can be above the 16MB line or not.

The TWA is suitable for quite small data storage requirements and for larger requirements that are both relatively fixed in size and are used more or less for the duration of the transaction. Because the TWA exists for the entire transaction, a large TWA size has much greater effect for conversational than for nonconversational transactions.

User storage

User storage is available to all the programs in a transaction, but some effort is required to pass it between programs using EXEC CICS LINK or EXEC CICS XCTL commands. Its size is not fixed, and it can be obtained (using EXEC CICS GETMAIN commands) just when the transaction requires it and returned as soon as it is not needed. Therefore, user storage is useful for large storage requirements that are variable in size or are shorter-lived than the transaction.

See Chapter 26, “Storage control and storage protection” on page 391 for information about how the USERDATAKEY and CICSDATAKEY options on the EXEC CICS GETMAIN command override the TASKDATAKEY option specified on the associated RDO TRANSACTION resource definition.

The SHARED option of the EXEC CICS GETMAIN command causes the acquired storage to be retained after the end of the task. The storage can be passed in the

communication area from one task to the next at the same terminal. The first task returns the address of the communication area in the COMMAREA option of the EXEC CICS RETURN command. The second task accesses the address in the COMMAREA option of the EXEC CICS ADDRESS command. You must use the SHARED option of the EXEC CICS GETMAIN command to ensure that your storage is in common storage.

The amount of processor overhead involved in an EXEC CICS GETMAIN command means that you should not use it for a small amount of storage. Use the TWA for the smaller amounts, or group them together into a larger request. Although the storage acquired by an EXEC CICS GETMAIN command may be held somewhat longer when using combined requests, the processor overhead and the reference set size are both reduced.

COMMAREA in EXEC CICS LINK and EXEC CICS XCTL commands

A communication area (COMMAREA) is a command-level facility used to transfer information between two programs within a transaction or between two transactions from the same terminal. For information about using COMMAREA between transactions, see “COMMAREA in EXEC CICS RETURN commands” on page 88.

Information in COMMAREA is available only to the two participating programs, unless those programs take explicit steps to make the data available to other programs that may be invoked later in the transaction. When one program links to another, the COMMAREA may be any data area to which the linking program has access. It is often in the working storage or LINKAGE SECTION of that program. In this area, the linking program can both pass data to the program it is invoking and receive results from that program. When a program transfers control (an EXEC CICS XCTL command) to another, CICS may copy the specified COMMAREA into a new area of storage, because the invoking program and its control blocks may no longer be available after it transfers control. In either case, the address of the area is passed to the program that is receiving control, and the CICS command-level interface sets up addressability. See Chapter 25, “Program control” on page 377 for further information.

CICS ensures that any COMMAREA is addressable by the program that receives it, by copying below the 16MB line and/or to the user-key storage where necessary, depending on the addressing mode and the value of the EXECKEY options on the RDO PROGRAM resource definition for the receiving program. See Chapter 26, “Storage control and storage protection” on page 391 for more information about the EXECKEY option.

CICS contains algorithms designed to reduce the number of bytes to be transmitted. The algorithms remove some trailing binary zeros from the COMMAREA before transmission and restore them after transmission. The operation of these algorithms is transparent to the application programs, which always see the full-size COMMAREA.

The overhead for using COMMAREA in an EXEC CICS LINK command is minimal; it is slightly more with the EXEC CICS XCTL and EXEC CICS RETURN commands, when CICS creates the COMMAREA from a larger area of storage used by the program.

The maximum size of a dynamic log record is 32KB. If a COMMAREA is passed to a transaction defined with the RDO DEFINE TRANSACTION option

RESTART(YES), CICS journal control adds a 4-byte header before writing the COMMAREA contents to the dynamic log. Therefore, in this case, the practical limit to the COMMAREA size is 32KB minus 4 bytes, or 32763 bytes.

For a COMMAREA passed between successive transactions in a pseudoconversational sequence in a distributed environment, 32KB is the limit imposed on the total data length by VTAM. This limit applies to the entire transmitted package, which includes control data added by VTAM. The amount of control data increases if the transmission uses intermediate links.

You can also use INPUTMSG on EXEC CICS LINK, EXEC CICS XCTL, and EXEC CICS RETURN commands to pass data to the program being invoked. For more information about the INPUTMSG option, see "INPUTMSG" on page 381.

Program storage

CICS creates a separate copy of the variable area of a CICS program for each transaction using the program. This area is known as **program storage**. This area is called the WORKING-STORAGE SECTION in COBOL, automatic storage in C and PL/I, and the DFHEISTG section in assembler language. Like the TWA, this area is of fixed size and is allocated by CICS without you having to issue an EXEC CICS GETMAIN command. The EXEC CICS interface sets up addressability automatically. Unlike the TWA, however, this storage lasts only while the program is being run, not for the duration of the transaction. This makes it useful for data areas that are not required outside the program and that are either small or, if large, are fixed in size and are required for all or most of the execution time of the program.

Lengths of areas passed to CICS commands

When a CICS command includes a LENGTH option, it usually accepts the length as a signed halfword binary value. This places a theoretical upper limit of 32KB on the length. In practice, the limits are less than this and vary for each command. The limits depend on data set definitions, recoverability requirements, buffer sizes, and local networking characteristics.

LENGTH options

In COBOL, PL/I, and assembler language, the translator defaults certain lengths; this means that they are optional in programs that specify data areas (except in the case of the EXEC CICS ENQ and EXEC CICS DEQ commands).

In C, all LENGTH options must be specified.

See the *CICS Application Programming Reference* manual for programming information, including details of when you need to specify the LENGTH option. Whatever the CICS command, do not let the LENGTH you code exceed 24KB, if possible.

Many commands involve the transfer of data between the application program and CICS. In all cases, the length of the data to be transferred must be provided by the application program.

In most cases, the LENGTH option must be specified if the SET option is used; the syntax of each command and its associated options show whether this rule applies.

For journal commands, the restrictions apply to the sum of the LENGTH and PFXLENG values. (See “Journaling” on page 135.)

Journal records

For journal records, the journal buffer size may impose a limit lower than 24KB. Note that the limit applies to the sum of the LENGTH and PFXLENG values.

Data set definitions

For temporary storage, transient data, and file control, the data set definitions can impose limits lower than 24KB. For details, see the *CICS System Definition Guide* (for information about defining data sets) and the *CICS Resource Definition Guide* (for information about RDO for files).

Recommendation

For any command in any system, 24KB is a good working limit for LENGTH specifications. Subject to user-specified record and buffer sizes, this limit is unlikely either to cause an error or to place a constraint on applications.

You will probably not find a 24KB limit too much of a hindrance; online programs do not often handle such large amounts of data, for the sake of efficiency and response time.

Note: The value in the LENGTH option should never exceed the length of the data area addressed by the command.

EXEC interface stubs

Each application program you write must contain an interface to CICS. This takes the form of an EXEC interface **stub**, which is a function-dependent piece of code used by the CICS high-level programming interface. The stub must be link-edited with your application program to provide communication between your code and the EXEC interface program (DFHEIP). These stubs are invoked during execution of EXEC CICS and EXEC DLI commands.

There are stubs for each programming language. For information about the EXEC interface stubs, see the *CICS System Definition Guide*.

Each EXEC command is translated into a COBOL CALL statement or PL/I CALL statement (as appropriate) by the command translator. The external entry point invoked by the CALL statement is resolved to an entry in the stub.

For C, each EXEC CICS command is translated by the command translator into a C function invocation. The external entry point is invoked by the function and is resolved by an entry in the stub.

Assembler language

Each EXEC command is translated into an invocation of the DFHECALL macro by the command translator and the external entry point invoked by DFHECALL is resolved to an entry in the stub.

Temporary storage

Temporary storage is the primary CICS facility for storing data that must be available to multiple transactions.

Data items in temporary storage are kept in queues whose names are assigned dynamically by the program storing the data. A temporary storage queue containing multiple items can be thought of as a small data set whose records can be addressed either sequentially or directly, by item number. If a queue contains only a single item, it can be thought of as a named scratch-pad area.

Temporary storage is implemented in two different ways. The one used for a particular queue is determined by what is specified on the command that creates the first item. Specifying the MAIN option means that the queue is kept in main storage, in space taken from the dynamic storage area. The AUXILIARY option means that the queue is written to an entry-sequenced VSAM data set. Whichever method you use, CICS maintains an index of items in main storage.

Both these methods have characteristics that you should bear in mind:

- Main temporary storage requires much more virtual storage than auxiliary. In general, use it only for small queues that have short lifetimes or are accessed frequently. Auxiliary temporary storage is specifically designed for relatively large amounts of data that have a relatively long lifetime or are accessed infrequently. You may find it useful to establish a cutoff point of a lifetime of one second to decide which queues should be in main storage and which should be in auxiliary.
- You can make queues in auxiliary storage recoverable, but not queues in main storage:
 - Only one transaction at a time can update a recoverable temporary storage queue. So, if you choose to make queues recoverable, bear in mind the probability of enqueues.
 - Ensure that there are enough buffers and VSAM strings to eliminate as much contention as possible.
- If a task tries to write to temporary storage and there is no space available, CICS normally suspends it, although the task can regain control in this situation by using either a HANDLE CONDITION NOSPACE command, or the RESP or NOHANDLE option on the WRITEQ TS command. If suspended, the task is not resumed until some other task frees the necessary space in main storage or the VSAM data set. This can produce unexplained response delays, especially if the waiting task owns exclusive-use resources, in which case all other tasks needing those resources must also wait.

With auxiliary temporary storage, running out of temporary storage means that all of the control intervals (CIs) in the VSAM data set have been used. To assist in controlling this type of wait, the temporary storage statistics show the maximum number of VSAM CIs used. As this value approaches the total available, enlarge the temporary storage data set. With main temporary storage, running out of temporary storage has the same effects as being short on storage, and may cause the system to fail. Take great care to control overall usage of temporary storage.

- It can be more efficient to use main temporary storage exclusively in very low-volume systems that have no need for recovery. You need to balance the

needs for additional main storage requirement for the VSAM access method and a larger temporary storage program with the need for main storage for the temporary storage records.

The following points apply to temporary storage in general:

- You must use an EXEC CICS command every time data is written to or read from a temporary storage queue, and CICS must find or insert the data using its internal index. This means that the overhead for using main temporary storage is greater than for the CWA or TCTUA. With auxiliary storage, (often the most frequently used), there is usually data set I/O as well, which increases overhead even more.
- You need not allocate temporary storage until it is required; you need keep it only as long as it is required, and the item size is not fixed until you issue the command that creates it. This makes it a good choice for relatively high-volume data and data that varies in length or duration.
- The fact that temporary storage queues can be named as they are created provides a very powerful form of direct access to saved data. You can access scratch-pad areas for terminals, data set records, and so on, simply by including the terminal name or record key in the queue name.
- Resource protection is available for temporary storage.

Intrapartition transient data

Intrapartition transient data has some characteristics in common with auxiliary temporary storage. (See “Sequential data set access” on page 74 for information about **extrapartition** transient data.) Like temporary storage, intrapartition transient data consists of queues of data, kept together in a single data set, with an index that CICS maintains in main storage.

You can use transient data for many of the purposes for which you would use auxiliary temporary storage, but there are some important differences:

- Transient data queue names must be defined in the destination control table (DCT) before CICS is started. You cannot define them arbitrarily at the time the data is created. Thus, transient data does not have the same dynamic characteristics as temporary storage.
- Transient data queues must be read sequentially, and each item can be read only once. After a transaction reads an item, that item is removed from the queue and is not available to any other transaction. In contrast, items in temporary storage queues may be read either sequentially or directly (by item number). They can be read any number of times and are not removed from the queue until the entire queue is purged.

These two characteristics make transient data inappropriate for scratch-pad data but suitable for queued data such as audit trails and output to be printed. In fact, for data that is read sequentially once, transient data is preferable to temporary storage.

- Items in a temporary storage queue can be changed; items in transient data queues cannot.
- Transient data queues are always written to a data set. (There is no form of transient data that corresponds to main temporary storage.)

- You can define transient data queues so that writing items to the queue causes a specific transaction to be initiated (for example, to process the queue). Temporary storage has nothing that corresponds to this “trigger” mechanism, although you may be able to use an EXEC CICS START command to perform a similar function.
- Transient data has more varied recovery options than temporary storage. Transient data queues can be physically or logically recoverable, and you can specify whether you want areas of the entry-sequenced data sets (ESDS) that have been written and read to be reused for new data.
- Because the commands for intrapartition and extrapartition transient data are identical, you can switch easily between the internal CICS facility (intrapartition) and an external data set, described in “Sequential data set access” on page 74. To do this, change only the DCT, not your application programs. Temporary storage has no corresponding function of this kind.

GETMAIN SHARED command

Storage acquired using the SHARED option of the GETMAIN command is not released when the acquiring task ends. This enables one task to leave data in storage for use by another task. The storage is not released until a FREEMAIN command is issued, either by the acquiring task or by another task.

Your own data sets

You can also use your own data sets to save data between transactions. This method probably has the largest overhead in terms of instructions processed, buffers, control blocks, and user programming requirements, but does provide extra functions and flexibility. Not only can you define data sets as recoverable resources, but you can log changes to them for forward recovery. You can specify the number of strings for the data set, (as well as on the temporary storage and transient data sets), to ensure against access contention, and you can use resource security.

Data operations

CICS supports the following:

- “DL/I database operations”
- “VSAM and DAM data set operations” on page 72
- “Browsing data sets” on page 74
- “Logging data sets” on page 74
- “Sequential data set access” on page 74

DL/I database operations

The following recommendations apply to using DL/I with CICS:

- Use command codes with CALL level and keywords with command level to reduce the number of requests whenever appropriate. For example, a DL/I path call is more efficient than a number of individual DL/I calls. With individual DL/I calls, the GN call gives the best performance. Although several DL/I calls may get their information from the DL/I or VSAM buffers, some of the instructions have to be processed within a DL/I call. Consider, therefore, the number of DL/I calls needed for the processing of a transaction.

- It is more efficient to use qualified segment-search areas (SSAs) than to check on “segment found” in the application program.
- Scheduling calls should be issued at the latest possible time, so as to minimize the time that the transaction has exclusive control of the PSB. (This control is released at the end of the UOW, which occurs at the next TERM call, explicit SYNCPOINT command, or the syncpoint implicit in task termination.)
- Be aware of the effects of explicit syncpointing on performance and recovery. See the *CICS Performance Guide* for more details of the performance implications of explicit syncpointing.

VSAM and DAM data set operations

The efficiency of database and data set operations is an important factor in the performance of any CICS system.

In VSAM, the main impact on efficiency, and thus on response time, comes from contention for serial-use resources (record keys, control intervals, and strings), and for storage use and processor overhead. As is usual in these situations, any improvements you make in one area may be at the expense of other areas.

VSAM data sets

To minimize contention delays using VSAM data sets:

- Minimize the time that VSAM resources are reserved for exclusive use. The **exclusive use** enqueue is the way CICS and VSAM prevent concurrent updates. CICS serializes update requests by base cluster record key. VSAM serializes by enqueueing on the control interval (CI), so that no transaction can update a record in the same control interval as another record being updated. The VSAM hold for exclusive use ends when the request is complete in VSAM terms.

For example, in an update operation, exclusive use that starts with the EXEC CICS READ command with the UPDATE option and ends when VSAM has completed the EXEC CICS REWRITE command. For **nonrecoverable** data sets, the CICS exclusive use ends at the same time. For **recoverable data sets**, however, it does not end until the task ends or issues an EXEC CICS SYNCPOINT command. Recoverability is specified by the RECOVERY option of an RDO FILE resource definition (LOG operand of DFHFCT TYPY=FILE macro). See the *CICS Resource Definition Guide* for more information about defining data set recoverability.

Table 7 shows which requests require exclusive use and when that reservation terminates. Remember that, for recoverable data sets, the CICS enqueue on the record key does not end until an EXEC CICS SYNCPOINT command is issued.

Command	Released by VSAM at
READ.. UPDATE	REWRITE/DELETE/UNLOCK
WRITE.. MASSINSERT	UNLOCK
WRITE	Completion of WRITE
DELETE.. RIDFLD	Completion of DELETE

- Hold position in a VSAM data set for as short a time as possible. Table 8 on page 73 shows which commands hold position and when the hold is released.

<i>Table 8. Commands that hold position and when hold is released</i>	
Command	Released by VSAM at
READ.. UPDATE	REWRITE/DELETE/UNLOCK
WRITE.. MASSINSERT	UNLOCK
STARTBR	ENDBR

Each request in progress against a VSAM data set requires at least one string. Requests that hold position tie up a string until a command is issued to release the hold position. Requests that do not hold position release the string as soon as that request is complete.

To minimize processor overhead when using VSAM data sets:

- Use the MASSINSERT option if you are adding many records in sequence. This improves performance by minimizing processor overheads and therefore improves the response times. For ESDSs and KSDSs, adding records with MASSINSERT causes CICS to use sequential VSAM processing. This changes the way VSAM places records within control intervals when a split is required, resulting in fewer splits and less unused space within the affected CIs.
- Use **skip-sequential** processing if you are reading many records in sequence whose keys are relatively close together but not necessarily adjacent. (Skip sequential processing begins with an EXEC CICS STARTBR (start browse) command.) Each record is retrieved with an EXEC CICS READNEXT command, but the key feedback area pointed to by RIDFLD is supplied with the key of the next requested record before the EXEC CICS READNEXT command is issued.
- Use the GENERIC option on the EXEC CICS DELETE command when deleting a group of records whose keys start with a common character string. CICS internally optimizes a generic EXEC CICS DELETE command.

DAM data sets

DAM data sets are less efficient than VSAM because CICS has to do some single-thread processing and issue some operating system waits to handle DAM data set requests. Therefore, if possible, use a relative record VSAM data set or an entry-sequenced data set addressed by relative byte address (RBA) in place of a DAM data set.

If you are using DAM data sets in update mode, be aware that performance is affected dramatically by the means of data set integrity you choose. If you specify **exclusive control** in DFHFCT TYPE=FILE macro SERVREQ operands for a DAM data set, CICS requests the operating system to prevent concurrent updates. However, this involves significant overhead. The sequence of events for a DAM read for update and release, using exclusive control, is:

1. Read record (block)
2. Enqueue
3. Read record (block)
4. CICS wait

5. Record processing
6. CICS file control release
7. DASD seek
8. Operating system wait (synchronous)
9. Dequeue
10. Return to CICS

Another method is to use EXEC CICS ENQ and DEQ commands to ensure integrity. This requires more programming, but avoids the synchronous wait used in exclusive control. You can include the enqueue-dequeue logic in your application program or in a file control exit. Note that this method applies only if the DAM data set is updated in one CICS partition only. If other partitions also update the data set, you must use the operating system mechanism to ensure integrity.

Browsing data sets

A data set browse is often the source of the output in transactions that produce a large number of output screens, which can monopolize system resources. A long browse can put a severe load on the system, locking out other transactions and increasing overall response time, in addition to the overhead needed for BMS, task control, and terminals. This is because CICS control philosophy is based on the assumption that the terminal operator initiates a transaction that accesses a few data records, processes the information, and returns the results to the operator. This process involves numerous waits that enable CICS to do multitasking. However, CICS is not an interrupt-driven multitasking system, so tasks that involve small amounts of I/O relative to processing can monopolize the system regardless of priority. A browse of a data set with many records in a control interval is just such a transaction.

You can prevent this by issuing EXEC CICS DELAY or EXEC CICS SUSPEND commands periodically, so that other tasks can get control. If the browse produces paged output, consider breaking the transaction up in one of the ways suggested in "Page-building and routing operations" on page 78.

Logging data sets

CICS provides options to log some or all types of activity against a data set. Logging updates enables you to reconstruct data sets from backup copies, if necessary. You may also want to log reads for security reasons. Again, you have to balance the need for data integrity and security against the performance effects of logging. These are the actual operations needed to do the logging and the possible delays caused because of the exclusive control that logging implies.

Sequential data set access

CICS provides a number of different sequential processing options. Temporary storage and intrapartition transient data queues (already discussed in "Temporary storage" on page 69 and in "Intrapartition transient data" on page 70) are the most efficient to use, but they must be created and processed entirely within CICS.

The following methods apply to sequential data sets that must be processed **externally** to CICS:

- Extrapartition transient data
- VSAM entry-sequenced data sets

Extrapartition transient data is the CICS way of handling standard sequential (SAM) data sets. It is the least efficient of the two forms of sequential support listed, because CICS has to issue operating system waits to process the data sets, as it does when handling DAM. Moreover, extrapartition transient data sets are not recoverable. VSAM ESDS datasets, on the other hand, are recoverable within limitations, and processing is more efficient. The recovery limitation is that records added to an ESDS during an uncompleted UOW cannot be removed physically during the backout process, because of VSAM restrictions. They can, however, be flagged as deleted by a user exit routine.

Terminal operations

Some design factors related to communicating with terminals may affect performance. The following sections discuss these:

- “Length of the data stream sent to the terminal”
- “Basic mapping support considerations”
- “Page-building and routing operations” on page 78
- “Requests for printed output” on page 80
- “Additional terminal control considerations” on page 81

Length of the data stream sent to the terminal

Good screen design and effective use of 3270 hardware features can significantly affect the number of bytes transmitted on a teleprocessing link. It is particularly important to keep the number of bytes as small as possible because, in most cases, this is the slowest part of the path a transaction takes. The efficiency of the data stream therefore affects both response time and line usage.

Basic mapping support considerations

When building a formatted data stream with basic mapping support (BMS), bear in mind the factors described in the following sections:

- “Avoid turning on modified data tags (MDTs) unnecessarily”
- “Use FRSET to reduce inbound traffic” on page 76
- “Do not send blank fields to the screen” on page 76
- “Use the MAPONLY option when possible” on page 77
- “Send only changed fields to an existing screen” on page 77
- “Design data entry operations to reduce line traffic” on page 78
- “Compress data sent to the screen” on page 78
- “Use nulls instead of blanks” on page 78
- “Use methods that avoid the need for nulls or blanks” on page 78

Avoid turning on modified data tags (MDTs) unnecessarily

The MDT is the bit in the attribute byte that determines whether a field should be transmitted on a READ MODIFIED command (the command used by CICS for all but copy operations).

The MDT for a field is normally turned on by the 3270 hardware when the user enters data into a field. However, you can also turn the tag on when you send a map to the screen, either by specifying FSET in the map or by sending an override attribute byte that has the tag on. Never set the tag on in this way for a field that is constant in the map, or for a field that has no label (and is not sent to the program that receives the map).

Also, you do not normally need to specify FSET for an ordinary input field. This is because, as already mentioned, the MDT is turned on automatically in any field in which the user enters data. This is then included in the next RECEIVE command. These tags remain on, no matter how many times the screen is sent, until explicitly turned **off** by the program (by the FRSET, ERASEAUP, or ERASE option, or by an override attribute with the tag off).

You can store information, between inputs, that the user did not enter on the screen. This is an intended reason for turning the MDT on by a program. However, this storage technique is appropriate only to small amounts of data, and is more suitable for local than for remote terminals, because of the transmission overhead involved.

For example, this technique is particularly useful for storing default values for input fields. In some applications, the user must complete a screen in which some fields already contain default values. A user who does not want to change a default just skips that field. The program processing the input has to be informed what these defaults are. If they are always the same, they can be supplied as constants in the program. If they are variable, however, and depend on earlier inputs, you can simply save them on the screen by turning the MDT on with FSET in the map that writes the screen. The program reading the screen then receives the default value from a user who does not change the field and the new value from a user who does.

Note: The saved values are not returned to the screen if the CLEAR, PA1, PA2, or PA3 key is pressed.

Use FRSET to reduce inbound traffic

If you have a screen with many input fields, which you may have to read several times, you can reduce the length of the input data stream by specifying FRSET when you write back to the screen in preparation for the next read. FRSET turns off the MDTs, so that fields entered before that write are not present unless the user reenters them the next time. If you are dealing with a relatively full screen and a process where there may be a number of error cycles (or repeat transmissions for some other reason), this can be a substantial saving. However, because only **changed** fields are sent on subsequent reads, the program must save input from each cycle and merge the new data with the old. This is not necessary if you are not using FRSET, because the MDTs remain on, and all fields are sent regardless of when they were entered.

Do not send blank fields to the screen

Sending fields to the screen that consist entirely of blanks or that are filled out on the right by trailing blanks usually wastes line capacity. The only case in which BMS requires you to do this is when you need to erase a field on the screen that currently contains data, or to replace it with data shorter than that currently on the screen, without changing the rest of the screen.

This is because, when BMS builds the data stream representing your map, it includes blanks (X'40') but omits nulls (X'00'). This makes the output data stream shorter. BMS omits any field whose first data character is null, regardless of subsequent characters in the field.

Although the contents of a BMS screen sent with nulls differ slightly from one sent with blanks, the screen displayed is the same. Also, when this screen is read in,

the input data stream is shorter because the control unit does not transmit the nulls. The only difference to the program processing the input is that it must be able to interpret a field of nulls correctly, by replacing it with blanks if subsequent program logic requires this. However, the savings in transmission usually justify the extra code that may be needed.

BMS requires you to initialize to nulls any area to be used to build a map. This is done by moving nulls (X'00') to the mapnameO field in the symbolic map structure. See "Initializing the output map" on page 244 for more information. BMS uses nulls in attribute positions and in the first position of data to indicate that no change is to be made to the value in the map. If you are reusing a map area in a program or in a TIOA, take special care to clear it in this way.

Use the MAPONLY option when possible

The MAPONLY option sends only the **constant** data in a map, and does not merge any variable data from the program. The resulting data stream is not always shorter, but the operation has a shorter path length in BMS. When you send a skeleton screen to be used for data entry, you can often use MAPONLY.

Send only changed fields to an existing screen

Sending only changed fields is important when, for example, a message is added to the screen, or one or two fields on an input screen are highlighted to show errors. In these situations, use the DATAONLY option to send a map that consists of nulls except for the changed fields. For fields in which only the attribute byte has changed, you need send only that byte, and send the remaining fields as nulls. BMS uses this input to build a data stream consisting of only the fields in question, and all other fields on the screen remain unchanged.

It may be tempting to ignore this advice and send an unnecessarily long data stream. For example, when a program that is checking an input screen for errors finds one, there are two options.

- It can simply add the error information to the input map (highlighted attributes, error messages, and so on) and resend it.
- It can build an entirely new screen, consisting of just the error and message fields.

The former is slightly easier to code (you do not need to have two map areas or move any fields), but it may result in very much longer transmissions because the output data stream contains the correct input fields as well as the error and message fields. In fact, it may even be longer than the original input stream because, if there were empty or short fields in the input, BMS may have replaced the missing characters with blanks or zeros.

With the 3270 hardware, if the input stream for a terminal exceeds 256 bytes, the terminal control unit automatically breaks it up into separate transmissions of 256 bytes maximum. This means that a long input stream may require several physical I/O operations. Although this is transparent to the application program, it does cause additional line and processor overhead. The **output** stream is generally sent in a single transmission.

Design data entry operations to reduce line traffic

Often, users are required to complete the same screen several times. Only the data changes on each cycle; the titles, field labels, instructions, and so on remain unchanged. In this situation, when an entry is accepted and processed, you can respond with a SEND CONTROL ERASEAUP command (or a map that contains only a short confirmation message and specifies the ERASEAUP option). This causes all the **unprotected** fields on the screen (that is, all the input data from the last entry) to be erased and to have their MDTs reset. The labels and other text, which are in protected fields, are unchanged, the screen is ready for the next data-entry cycle, and only the necessary data has been sent.

Compress data sent to the screen

When you send unformatted data to the screen, or create a formatted screen outside BMS, you can compress the data further by inserting set buffer address (SBA) and repeat-to-address (RA) orders into the data stream. SBA allows you to position data on the screen, and RA causes the character following it to be generated from the current point in the buffer until a specified ending address. SBA is useful whenever there are substantial unused areas on the screen that are followed by data. RA is useful when there are long sequences of the same character, such as blanks or dashes, on the screen. However, note that the speed with which RA processes is not uniform across all models of 3270 control units. Check how it applies to your configuration before use.

CICS provides an exit that is driven just before output is sent to a terminal (XTCOUT). You may want to add SBA and RA substitutions to this exit to compress the data stream using a general subroutine. This has the dual benefit of removing compression logic from your application program and of applying to all output data streams, whether they are produced by BMS or not.

Use nulls instead of blanks

Note that, outside BMS, nulls have no special significance in an **output** data stream. If you need a blank area on a screen, you can send either blanks or nulls to it; they take up the same space in the output stream. However, if the blank field is likely to be changed by the user and subsequently read, use nulls, because they are not transmitted back.

Use methods that avoid the need for nulls or blanks

For any **large** area of a screen that needs to be blank, consider methods other than transmitting blanks or nulls; for example, when using BMS, putting SBA and RA orders directly into the data stream, or using the ERASE and ERASEAUP options.

Page-building and routing operations

BMS page-building facilities provide a powerful and flexible tool for building and displaying long messages, sending messages to multiple destinations, and formatting a single message for several devices with different physical characteristics. However, as for any high-function tool, it requires a substantial overhead, as mentioned in “Browsing data sets” on page 74. You may need the page-building option (ACCUM) when:

- “Sending multipage output” on page 79
- “Sending messages to destinations other than the input terminal” on page 79

- “Sending pages built from multiple maps” on page 80
- “Using the BMS page-copy facility” on page 80

Sending multipage output

Transactions that produce very large output messages, consisting of many screen-size pages, tend to tax system resources. First, all the pages have to be created, which involves processor activity, execution of the CSPG transaction, and data set I/O activity. The pages must then be saved in temporary storage. If the terminal user looks at every page in a message, a large number of transactions are run to process the paging requests, each of which needs line and processor overhead. Obviously some overhead is caused by the size and complexity of the transaction, and it may be unavoidable. Indeed, if several users are scrolling rapidly through paged output at the same time, the transactions needed can monopolize a system.

If users really need to see all the pages, and need to scroll backward and forward frequently, it may be more efficient to produce all the pages at the same time and present them using “traditional” CICS paging services. However, if users need only a few of the pages, or can easily specify how far back or forward in the message they would like to scroll, there are two choices:

- Construct a pseudoconversational transaction to produce just one screen of output. The first time this transaction is run, it produces the first page of the many-page output. The output screen contains space for users to indicate the page they want next. The transaction always sets the next transaction identifier to point to itself, so that it can display the requested page when it is next run.

You probably want to give users some of the options that CICS provides (such as one page forward, one page back, and skip to a selected page) and some relevant to the application, such as a data set key at which to begin the next page of output.

- Page-build a multipage output message with the ACCUM option, but limit the number of pages in the message (for example, to five). Users page through the subset pages with the usual CICS page commands. On the last screen of the output, you add an indication that there is more output and a place for them to indicate whether they want to see the next segment. As in the first example, the next transaction identifier is set to the original transaction so that, if CICS does not receive a paging command, it invokes that transaction.

Sending messages to destinations other than the input terminal

If you need to send a message to a terminal other than the input terminal associated with a task, BMS routing may be the most efficient way of doing so. This is especially so if the message must be sent to multiple destinations or if it involves multiple pages. Routing is the recommended method if the message recipients need CICS paging commands to access it.

However, if neither of the above conditions apply, you have a choice of two other methods of delivering output to a terminal not associated with the transaction.

1. You can use an EXEC CICS START command, with the TERMID option, to specify the terminal to which you want to write and the FROM option to specify the data you want to send. Your own transaction is the started transaction. It issues an EXEC CICS RETRIEVE command for the message and then sends it

to its own terminal. See the *CICS Application Programming Reference* manual for programming information about the START command.

2. Similarly, you can put messages destined for a particular terminal on a transient data queue that has the same name as the terminal. The definition for the transient data queue must specify:

- DESTFAC=TERMINAL
- The terminal identifier
- A trigger level (usually 1)
- A transaction name

Your own transaction reads the transient data queue and sends the message to its terminal. It repeats this sequence until the queue is empty, and then terminates. The trigger level you specified means that it is invoked every time the specified number of messages are placed on the queue. The *CICS Sample Applications Guide* gives an example of a program that performs this function.

Note: Because of the overhead associated with routing messages (by whatever means), use facilities such as ROUTE=ALL with caution.

Sending pages built from multiple maps

Although you can easily build a screen gradually using different maps, you can sometimes avoid considerable overhead by not using page-building operations, especially where there is only one screen of output and no other need for paging. An example of this is an application whose output consists of a header map, followed by a variable number of detail segments, sent with a second map, and finally a trailer map following the detail. Suppose the average output screen for such an application contains eight (2-line) detail segments, plus header and trailer, and all this fits on a single screen. Writing this screen with page-building requires 11 BMS calls (header, details, trailer, and page-out) whereas, if the program builds the output screen internally, it needs only one call.

Using the BMS page-copy facility

Because the individual pages that make up an accumulated BMS message are saved in temporary storage, BMS enables the terminal user to copy individual pages to other terminals. However, if the ability to copy is the only reason for using page-building, consider using either the 3274 control unit copy facilities or the CICS **copy key** facility instead.

The 3274 copy facilities require no participation from CICS and no transmission, and are by far the most efficient method. The CICS copy key facility does have an overhead (see "Requests for printed output"), although of a different type from the BMS copy facility. It also has destination restrictions that do not apply to BMS copying.

Requests for printed output

A CICS print request asks CICS to copy what is on the requesting screen to the first available printer on the same control unit. The overhead involved depends on whether a printer is available, and whether the requesting terminal is remote or local to CICS.

If no printer is available, and the request is from a remote or a local device:

- CICS reads the buffer to the display terminal. This involves transmitting **every** position on the screen, including nulls.

For requests from a local device, the buffer is read at channel speeds, so that the large input message size does not increase response time too much, and does not monopolize the line.

- An error task is generated so that the terminal error program can dispose of the message. If a printer is available and the request is from a local device, this step is not needed.
- The 3270 print task (CSPP) is attached to write the entire buffer to the printer when it is available.

If a printer is available, and the request is from a remote device, CICS sends a very short data stream to the control unit asking for a copy of the requesting device buffer to be sent to the output device buffer.

Additional terminal control considerations

The following sections describe additional points to consider when using the CICS terminal control services:

- “Use only one physical EXEC CICS SEND command per screen”
- “On BTAM, avoid the WAIT option on an EXEC CICS SEND command”
- “Use the EXEC CICS CONVERSE command” on page 82
- “Limit the use of message integrity options” on page 82
- “Limit the use of the DEFRESP option on EXEC CICS SEND commands” on page 82
- “Avoid using unnecessary transactions” on page 82
- “Send unformatted data without maps” on page 82

Use only one physical EXEC CICS SEND command per screen

It is usually more efficient to create a screen with a single call to BMS, than to build the screen with a series of EXEC CICS SEND MAP ACCUM commands. It is important to send the screen in a single physical output to the terminal. It is **very** inefficient to build a screen in parts and send each part with a separate command, because of the additional processor overhead of using several commands and the additional line and access method overhead.

On BTAM, avoid the WAIT option on an EXEC CICS SEND command

Note: CICS Transaction Server for VSE/ESA Release 1 does not support BTAM. You can run BTAM transactions on CICS Transaction Server for VSE/ESA Release 1 if you initiate them from a system that supports BTAM transactions, and use transaction routing to CICS Transaction Server for VSE/ESA Release 1.

If your program is pseudoconversational, it has only one EXEC CICS SEND command, by definition. (See “Choosing between pseudoconversational and conversational design” on page 53.) Unless you require notification to this program of an error on the EXEC CICS SEND command, omit the WAIT option. This allows CICS task control to reclaim the control blocks and user storage for your program long before it would otherwise be able to do so. Indeed, use of the WAIT option reduces substantially the savings effected by pseudoconversational programming.

Use the EXEC CICS CONVERSE command

Use the EXEC CICS CONVERSE command rather than the EXEC CICS SEND and EXEC CICS RECEIVE commands (or an EXEC CICS SEND, WAIT, RECEIVE command sequence if your program is conversational). They are functionally equivalent, but the EXEC CICS CONVERSE command crosses the CICS services interface only once, which saves processor time.

Limit the use of message integrity options

Like specifying the WAIT option on the final EXEC CICS SEND command of a transaction, the MSGINTEG option of CEDA requires CICS to keep the transaction running until the last message has been delivered successfully.

The PROTECT option of the PROFILE definition implies message integrity and causes the system to log all input and output messages, which adds to I/O and processor overhead.

Limit the use of the DEFRESP option on EXEC CICS SEND commands

Avoid using the DEFRESP option on EXEC CICS SEND commands, unless the transaction must verify successful delivery of the output message. It delays termination of the transaction in the same way as MSGINTEG.

Avoid using unnecessary transactions

Avoid situations that may cause users to enter an incorrect transaction or to use the CLEAR key unnecessarily, thus adding to terminal input, task control processing, terminal output, and overhead. Good screen design and standardized PF and PA key assignments should minimize this.

Send unformatted data without maps

If your output to a terminal is entirely or even mostly unformatted, you can send it using terminal control commands rather than BMS commands (that is, using an EXEC CICS SEND command without the MAP or TEXT options).

Chapter 4. Sharing data across transactions

CICS facilities for sharing data across transactions include:

- “Common work area (CWA)”
- “TCTTE user area (TCTUA)” on page 87
- “COMMAREA in EXEC CICS RETURN commands” on page 88
- “Display screen” on page 89

Other facilities include:

- Temporary storage
- Intrapartition transient data
- Storage acquired by a GETMAIN command with SHARED option
- Your own data sets

These last four items provide more flexibility and function than the items in the first list, and therefore involve somewhat more overhead. You can also use any of these methods within transactions.

With the exception of COMMAREA and the display screen, data stored in these facilities is available to any transaction in the system. Subject to resource security and storage protection restrictions, any transaction may write to them and any transaction may read them.

The use of some of these facilities may cause intertransaction affinities. See Chapter 5, “Affinity among transactions” on page 91 for more information about transaction affinities.

Common work area (CWA)

The common work area (CWA) is a single control block that is allocated at system startup time and exists for the duration of that CICS session. The size is fixed, as specified in the system initialization parameter, WRKAREA. The CWA has the following characteristics:

- There is almost no overhead in storing or retrieving data from the CWA. Command-level programs must issue one EXEC CICS ADDRESS command to get the address of the area but, after that, they can access it directly.
- Data in the CWA is not recovered if a transaction or the system fails.
- It is not subject to resource security.
- CICS does not regulate use of the CWA. All programs in all applications that use the CWA must follow the same rules for shared use. These are usually set down by the system programmers, in cooperation with application developers, and require all programs to use the same “copy” module to describe the layout of the area.

You must not exceed the length of the CWA, because this causes a storage violation. Furthermore, you must ensure that the data used in one transaction does not overlay data used in another. One way to protect CWA data is to use the storage protection facility, which protects the CWA from being written to by user-key applications. See “Protecting the CWA” on page 84 for more information.

- The CWA is especially suitable for small amounts of data, such as status information, that are read or updated frequently by multiple programs in an application.
- The CWA is not suitable for large-volume or short-lived data because it is always allocated.

Protecting the CWA

The CWAKEY system initialization parameter allows you to specify whether the CWA is to be allocated from CICS-key or user-key storage. See the *CICS System Definition Guide* for details about the CWAKEY parameter.

If you want to restrict write access to the CWA, you can specify CWAKEY=CICS. This means that CICS allocates the CWA from CICS-key storage, restricting application programs defined with EXECCKEY(USER) on the RDO PROGRAM resource definition to read-only access to the CWA. The only programs allowed to write to a CWA allocated from CICS-key storage are those you define with EXECCKEY(CICS).

Because any program that executes in CICS key can also write to CICS storage, ensure that such programs are thoroughly tested to make sure that they do not overwrite CICS storage.

If you want to give preference to protecting CICS rather than the CWA, specify CWAKEY=USER for the CWA, and EXECCKEY(USER) for all programs that write to the CWA. This ensures that, if a program exceeds the length of the CWA, it does not overwrite CICS storage. For more information about storage protection, see Chapter 26, “Storage control and storage protection” on page 391.

Figure 8 illustrates a particular use of the CWA where the CWA itself is protected from user-key application programs by CWAKEY=CICS. In this illustration, the CWA is not used directly to store application data and constants. The CWA contains pairs of application identifiers and associated addresses, with the address fields containing the addresses of data areas that hold the application-related data. For protection, the CWA is defined with CWAKEY=CICS, therefore the program that, in this illustration, is a program defined in the program list table post initialization (PLTPI) list, and which loads the CWA with addresses and application identifiers. All PLT programs execute in CICS-key regardless of the value of the EXECKEY option on the RDO PROGRAM resource definition. Any application programs requiring access to the CWA should be defined with EXECKEY(USER), thereby ensuring the CWA is protected from overwriting by application programs. In Figure 8, one of the data areas is obtained from CICS-key storage, whereas the other is obtained from user-key storage.




Figure 8. Example of use of CWA in CICS-key storage.

In the sample code shown in Figure 9 on page 86, the program list table post-initialization (PLTPI) program is setting up the application data areas, with pointers to the data stored in the CWA.

```

ID DIVISION.
PROGRAM-ID. PLTPROG.
ENVIRONMENT DIVISION.
DATA DIVISION.
WORKING-STORAGE SECTION.
77 APPLID                PIC X(8)      VALUE SPACES.
77 SYSID                 PIC X(4)      VALUE SPACES.
01 COMM-DATA.
   03 AREA-PTR           USAGE IS POINTER.
   03 AREA-LENGTH        PIC S9(8)     COMP.
LINKAGE SECTION.
01 COMMON-WORK-AREA.
   03 APPL-1-ID          PIC X(4).
   03 APPL-1-PTR         USAGE IS POINTER.
   03 APPL-2-ID          PIC X(4).
   03 APPL-2-PTR         USAGE IS POINTER.
PROCEDURE DIVISION.
MAIN-PROCESSING SECTION.
* Obtain APPLID and SYSID values
  EXEC CICS ASSIGN APPLID(APPLID)
                SYSID(SYSID)

  END-EXEC.
* Set up addressability to the CWA
  EXEC CICS ADDRESS
                CWA(ADDRESS OF COMMON-WORK-AREA)

  END-EXEC.
* Get 12KB of CICS-key storage for the first application ('APP1')
  MOVE 12288    TO    AREA-LENGTH.
  EXEC CICS GETMAIN SET(AREA-PTR)
                FLENGTH(AREA-LENGTH)
                SHARED

  END-EXEC.
* Initialize CWA fields and link to load program
* for storage area 1.
  MOVE 'APP1'   TO    APPL-1-ID.
  SET APPL-1-PTR TO    AREA-PTR.
  EXEC CICS LINK PROGRAM('LOADTAB1')
                COMMAREA(COMM-DATA)

  END-EXEC.

```

Figure 9 (Part 1 of 2). Sample code for loading the CWA

```

* Get 2KB of user-key storage for the second application ('APP2')
  MOVE 2048      TO  AREA-LENGTH.
  EXEC CICS GETMAIN SET(AREA-PTR)
                FLENGTH(AREA-LENGTH)
                SHARED
                USERDATAKEY
  END-EXEC.
* Initialize CWA fields and link to load program
* for storage area 2.
  MOVE 'APP2'    TO  APPL-2-ID.
  SET APPL-2-PTR TO  AREA-PTR.
  EXEC CICS LINK PROGRAM('LOADTAB2')
                COMMAREA(COMM-DATA)
  END-EXEC.
  EXEC CICS RETURN
  END-EXEC.
  MAIN-PROCESSING-EXIT.
  GOBACK.

```

Figure 9 (Part 2 of 2). Sample code for loading the CWA. This example illustrates how to create global data for use by application programs, with addresses of the data stored in the CWA—for example, by a PLTPI program. The first data area is obtained from CICS-key storage, which is the default on a GETMAIN command issued by a PLTPI program, the second from user-key storage by specifying the USERDATAKEY option. The CWA itself is in CICS-key storage.

TCTTE user area (TCTUA)

The TCT user area (TCTUA) is an optional extension to the terminal control table entry (TCTTE). Each entry in the TCT specifies whether this extension is present and, if so, how long it is (by means of the USERAREALEN attribute of the RDO TYPETERM resource definition used for the terminal). See the *CICS Resource Definition Guide* for more information about the TYPETERM resource definition.

The system initialization parameters TCTUALLOC and TCTUAKEY specify the location and storage key for all TCTUAs.

- TCTUALLOC=BELOW or ANY specifies whether you want 24-or 31-bit addressability to the TCTUAs, and whether TCTCUAs must be stored below the 16MB line or may be either above or below the line.
- TCTUAKEY=USER or CICS specifies whether you want the TCTUAs allocated from user-key or CICS-key storage.

TCTUAs have the following characteristics in common with the CWA:

- Minimal processor overhead (only one EXEC CICS ADDRESS command needed)
- No recovery
- No resource security
- No regulation of use by CICS
- Fixed length
- Unsuitability for large-volume or short-lived data

Unlike the CWA, however, the TCTUA for a particular terminal is usually shared only among transactions using that terminal. It is therefore useful for storing small amounts of data of fairly standard length between a series of transactions in a pseudoconversational sequence. Another difference is that it is not necessarily permanently allocated, because the TCTUA exists only while the TCTTE is set up. For nonautoinstall terminals the TCTUA is allocated from system startup; for autoinstall terminals the TCTUA is allocated when the TCTTE is generated.

Using the TCTUA in this way does not require special discipline among using transactions, because data is always read by the transaction following the one that wrote it. However, if you use TCTUAs to store longer-term data (for example, terminal or operator information needed by an entire application), they require the same care as the CWA to ensure that data used in one transaction does not overlay data used in another. Do not exceed the length of the allocated TCTUA, because this produces a storage violation.

COMMAREA in EXEC CICS RETURN commands

The COMMAREA option of the EXEC CICS RETURN command is designed specifically for passing data between successive transactions in a pseudoconversational sequence. It is implemented as a special form of user storage, although the EXEC interface, rather than the application program, issues the GETMAIN and FREEMAIN requests.

The COMMAREA is allocated from the CICS shared subpool in main storage, and is addressed by the TCTTE, between tasks of a pseudoconversational application. The COMMAREA is freed unless it is passed to the next task.

The first program in the next task has automatic addressability to the passed COMMAREA, as if the program had been invoked by either an EXEC CICS LINK command or an EXEC CICS XCTL command (see “COMMAREA in EXEC CICS LINK and EXEC CICS XCTL commands” on page 66). You can also use the COMMAREA option of the EXEC CICS ADDRESS command to obtain the address of the COMMAREA.

For a COMMAREA passed between successive transactions in a pseudoconversational sequence in a distributed environment, VTAM® imposes a limit of 32KB on the size of the total data length. This limit applies to the entire transmitted package, which includes control data added by VTAM. The amount of control data increases if the transmission uses intermediate links.

To summarize:

- Processor overhead is low (equivalent to using COMMAREA with an XCTL command and approximately equal to using main temporary storage).
- It is not recoverable.
- There is no resource security.
- It is not suitable for very large amounts of data (because main storage is used, and it is held until the terminal user responds).
- It is available only to the first program in a transaction, unless that program explicitly passes the data or its address to succeeding programs.

Display screen

You can also store data between pseudoconversational transactions from a 3270 display terminal on the display screen itself. For example, if users make errors in data that they are asked to enter on a screen, the transaction processing the input usually points out the errors on the screen (with highlights or messages), sets the next transaction identifier to point to itself (so that it processes the corrected input), and returns to CICS.

The transaction has two ways of using the **valid** data. It can save it (for example, in COMMAREA), and pass it on for the next time it is run. In this case, the transaction must merge the changed data on the screen with the data from previous entries. Alternatively, it can save the data on the screen by not turning off the modified data tags of the keyed fields.

Saving the data on the screen is very easy to code, but has two limitations:

- Do not use it with screens that contain large amounts of data if the likelihood of errors is high. This is because of the additional network traffic needed to resend the unchanged data. It does not apply to locally-attached terminals.
- If the user presses the CLEAR key, the screen data is lost, and the transaction must be able to recover from this. You can avoid this by defining the CLEAR key to mean CANCEL or QUIT, if this is appropriate for the application concerned.

Data other than keyed data may also be stored on the screen. This data can be protected from changes (except those caused by CLEAR) and can be nondisplay, if necessary.

Chapter 5. Affinity among transactions

CICS transactions use many different techniques to pass data from one to another. Some of these techniques require that the transactions exchanging data must execute in the same CICS region, and therefore impose restrictions on the dynamic routing of transactions. If transactions exchange data in ways that impose such restrictions, there is said to be an affinity among them. This chapter describes:

- “What is transaction affinity?”
- “Programming techniques used by CICS application programs” on page 92
- “Safe programming techniques” on page 94
- “Unsafe programming techniques” on page 97
- “Suspect programming techniques” on page 103
- “Detecting intertransaction affinities” on page 112
- “Duration and scope of intertransaction affinities” on page 113
- “Recommendations” on page 117

What is transaction affinity?

There are two types of affinity that affect dynamic transaction routing:

- Intertransaction affinity
- Transaction-system affinity

Intertransaction affinity

Transaction affinity among two or more CICS transactions is caused by the transactions using techniques to pass information between one another, or to synchronize activity between one another, in a way that requires the transactions to execute in the same CICS region. This type of affinity is intertransaction affinity, where a set of transactions share a common resource and/or coordinate their processing. Intertransaction affinity, which imposes restrictions on the dynamic routing of transactions, can occur in the following circumstances:

- One transaction terminates, leaving “state data” in a place that a second transaction can access only by running in the same CICS region as the first transaction.
- One transaction creates data that a second transaction accesses while the first transaction is still running. For this to work safely, the first transaction usually waits on some event, which the second transaction posts when it has read the data created by the first transaction. This technique requires that both transactions are routed to the same CICS region.
- Two transactions synchronize, using either an event control block (ECB) or an enqueue (ENQ) mechanism. Because CICS has no function shipping support for these techniques, this type of affinity means the two transactions must be routed to the same CICS region.

Transaction-system affinity

There is another type of transaction affinity that is not an affinity among transactions themselves. It is an affinity between a transaction and a particular CICS region, where the transaction interrogates or changes the properties of that CICS region—transaction-system affinity.

Transactions with affinity to a particular system, rather than another transaction, are not eligible for dynamic transaction routing. In general, they are transactions that use EXEC CICS INQUIRE and SET commands, or have some dependency on global user exit programs, which also have an affinity with a particular CICS region.

Using INQUIRE and SET commands and global user exits

There is no remote (that is, function shipping) support for INQUIRE and SET commands, nor is there a SYSID option on them, hence transactions using these commands must be routed to the CICS regions that own the resources they refer to. In general, such transactions cannot be dynamically routed to any AOR, and therefore transactions that use INQUIRE and SET should be statically routed.

Global user exits running in different CICS regions cannot exchange data. It is unlikely that user transactions pass data or parameters by means of user exits, but if such transactions do exist, they must run in the same AOR as the global user exits.

Programming techniques used by CICS application programs

From the point of view of intertransaction affinity in a dynamic transaction routing environment, the programming techniques used by your application programs can be considered in three broad categories. The first category comprises those techniques that are generally safe and do not cause intertransaction affinities; the second comprises those techniques that are inherently unsafe; and the third comprises those that are suspect in that they may, or may not, create affinities depending on exactly how they are implemented.

Safe techniques

The programming techniques in the generally safe category are:

- The use of the communication area (COMMAREA), supported by the CICS API on a number of CICS commands. However, it is the COMMAREA option on the CICS RETURN command only that is of interest in a dynamic transaction routing environment with regard to transaction affinity, because it is the COMMAREA on a RETURN command that is passed to the next transaction in a pseudoconversational transaction.
- The use of a TCT user area (TCTUA) that is optionally available for each terminal defined to CICS.

For more information about the COMMAREA and the TCTUA, see “Safe programming techniques” on page 94.

Unsafe techniques

The programming techniques in the unsafe category are:

- The use of long-life shared storage:
 - The common work area (CWA)
 - EXEC CICS GETMAIN SHARED storage
 - Storage obtained via an EXEC CICS LOAD PROGRAM HOLD
- The use of task-lifetime local storage shared by synchronized tasks

It is possible for one task to pass the address of some task-lifetime storage to another task.

It may be safe for the receiving task to use the passed address, provided the owning task does not terminate. It is possible, but ill-advised, to use a CICS task-synchronization technique to allow the receiving task to prevent the sending task from terminating (or freeing the storage in some other way) before the receiver has finished with the address. However, such designs are not robust because there is a danger of the sending task being purged by some outside agency.

(See “Sharing task-lifetime storage” on page 100 for more details.)

- Synchronization or serialization of tasks using EXEC CICS commands:
 - WAIT EVENT / WAIT EXTERNAL / WAITCICS
 - ENQ / DEQ

For more information about unsafe programming techniques, see “Unsafe programming techniques” on page 97.

Suspect techniques

Some programming techniques may, or may not, create affinity depending on exactly how they are implemented. A good example is the use of temporary storage. Application programs using techniques in this category must be checked to determine whether they will work without restrictions in a dynamic transaction routing environment. The programming techniques in the suspect category are:

- The use of temporary storage queues with restrictive naming conventions
- Transient data queues and the use of trigger levels
- Synchronization or serialization of tasks using EXEC CICS commands:
 - RETRIEVE WAIT / START
 - START / CANCEL REQID
 - DELAY / CANCEL REQID
 - POST / CANCEL REQID
- INQUIRE and SET commands and global user exits.

For more information about suspect programming techniques, see “Suspect programming techniques” on page 103.

Safe programming techniques

Some techniques for passing data between transactions are generally safe in that they do not create intertransaction affinity. These involve either the use of a communication area (COMMAREA), or a terminal control table user area (TCTUA).

However, to remain free from affinity, COMMAREAs and TCTUAs must **not** contain addresses. Generally the storage referenced by such addresses would have to be long-life storage, the use of which is an unsafe programming technique in a dynamic transaction routing environment.

The use of the COMMAREA and TCTUA for passing data between transactions is discussed further in the following sections.

The COMMAREA

The use of the COMMAREA option on the RETURN command is the principal example of a safe programming technique that you can use to pass data between successive transactions in a CICS pseudoconversational transaction. CICS treats the COMMAREA as a special form of user storage, even though it is CICS that issues the GETMAIN and FREEMAIN requests for the storage, and not the application program.

CICS ensures that the contents of the COMMAREA specified on a RETURN command are always made available to the first program in the next transaction. This is true even when the sending and receiving transactions execute in different AORs. In a pseudoconversation, regardless of which AOR a dynamic routing program chooses to route the next transaction to, CICS ensures the COMMAREA specified on the previous RETURN command is made available in the target AOR. This is illustrated in Figure 10 on page 95.

Some general characteristics of a COMMAREA are:

- Processor overhead is low.
- It is not recoverable.
- The length of a COMMAREA on a RETURN command can vary from transaction to transaction, up to a theoretical upper limit of 32 763 bytes. (However to be safe, do not exceed 24KB, as recommended in the *CICS Application Programming Reference* manual, because of a number of factors that can reduce the limit from the theoretical maximum.)
- CICS holds a COMMAREA in CICS main storage until the terminal user responds with the next transaction. This may be an important consideration if you are using large COMMAREAs, because the number of COMMAREAs held by CICS relates to terminal usage, and not to the maximum number of tasks in a region at any one time.
- A COMMAREA is available only to the first program in the next transaction, unless that program explicitly passes the data to another program or a succeeding transaction.




Figure 10. The use of a COMMAREA by a pseudoconversation in a dynamic transaction routing environment

The COMMAREA used in a pseudoconversational transaction, as shown in Figure 10, can be passed from transaction to transaction across a CICSplex and, provided the COMMAREA contains only data and not addresses of storage areas, no intertransaction affinity is created.

The TCTUA

The TCTUA is an optional extension to the terminal control table entry (TCTTE), each entry specifying whether the extension is present, and its length. You specify that you want a TCTUA associated with a terminal by defining its length on the USERAREALEN parameter of a RDO TYPETERM resource definition. This means that the TCTUAs are of fixed length for all the terminals created using the same TYPETERM definition.

A terminal control table user area (TCTUA) is safe to use in a dynamic transaction routing environment as a means of passing data between successive transactions in a pseudoconversational transaction. Like the COMMAREA, the TCTUA is always accessible to transactions initiated at a user terminal, even when the transactions in a pseudoconversation are routed to different AORs. This is illustrated in Figure 11 on page 96. Some other general characteristics of TCTUAs are:

- Minimal processor overhead (only one CICS command is needed to obtain the address).
- It is not recoverable.
- The length is fixed for the group of terminals associated with a given TYPETERM definition. It is suitable only for small amounts of data, the maximum size allowed being 255 bytes.
- If the terminal is autoinstalled, the TCTUA lasts as long as the TCTTE, the retention of which is determined by the AILDELAY system initialization parameter. The TCTTE, and therefore any associated TCTUA, is deleted when the AILDELAY interval expires after a session between CICS and a terminal is ended.

If the terminal is defined to CICS by an explicit terminal definition, the TCTTE and its associated TCTUA are created when the terminal is installed and remain until the next cold start of CICS.

Note that the TCTUA is available to a dynamic transaction routing program in the TOR as well as application programs in the AOR. It can be used store information relating to the dynamic routing of a transaction. For example, you can use the TCTUA to store the name of the selected AOR to which a transaction is routed.




Figure 11. The use of a TCTUA by a pseudoconversation in a dynamic transaction routing environment

Using the TCTUA in an unsafe way

The EXEC CICS ADDRESS TCTUA(*ptr-ref*) provides direct addressability to the TCTUA, and this is how each task requiring access to a TCTUA should obtain the TCTUA address. If tasks attempt to pass the address of their TCTUAs in some other way, such as in a temporary storage queue, or to use the TCTUA itself to pass addresses of other storage areas, the TCTUA ceases to provide a safe programming technique for use in a dynamic transaction routing environment.

It is also possible for a task to obtain the TCTUA of a principal facility other than its own, by issuing an EXEC CICS INQUIRE TERMINAL command that names the terminal associated with another task (the EXEC CICS INQUIRE TERMINAL command returns the TCTUA address of the specified terminal). Using the TCTUA address of a terminal other than a task's own principal facility is another example of an unsafe use of the TCTUA facility. Depending on the circumstances, particularly in a dynamic transaction routing environment, the TCTUA of a terminal that is not the inquiring task's principal facility could be deleted after the address has been obtained. For example, in an AOR, an EXEC CICS INQUIRE TERMINAL command could return the TCTUA address associated with a surrogate terminal that is running a dynamically routed transaction. If the next transaction from the terminal is routed to a different AOR, the TCTUA address ceases to be valid.

Unsafe programming techniques

Some CICS application programming techniques, notably those that pass or obtain addresses to shared storage, create an affinity between transactions.

The programming techniques that are generally unsafe are described in the following sections.

Using the common work area

The CWA in a CICS region is created (optionally) during CICS initialization, exists until CICS terminates, and is not recovered on a CICS restart (warm or emergency). The EXEC CICS ADDRESS CWA(*ptr-ref*) command provides direct addressability to the CWA.

A good example of how the use of long-life shared storage such as the CWA can create affinity is when one task stores data in the CWA, and a later task reads the data from it. Clearly, the task retrieving the data must run in the same AOR as the task that stored the data, or it will reference a completely different storage area in a different address space. This restricts the workload balancing capability of the dynamic routing program, as shown in Figure 12.




Figure 12. Illustration of intertransaction affinity created by use of the CWA. The dynamic routing program needs to be aware of this CWA affinity, and ensure it routes TRN2 to the same AOR as TRN1.

However, if the CWA contains read-only data, and this data is replicated in more than one AOR, it is possible to use the CWA and continue to have the full benefits of dynamic transaction routing. For example, you can run a program during the post-initialization phase of CICS startup (a PLTPI program) that loads the CWA with read-only data in each of a number of selected AORs. In this way, all transactions routed to AORs loaded with the same CWA data have equal access to the data, regardless of which of the AORs the transactions are routed to. With CICS subsystem storage protection, you can ensure the read-only integrity of the CWA data by requesting the CWA from CICS-key storage, and define all the programs that read the CWA to execute in user key.

Using GETMAIN SHARED storage

Shared storage is allocated by an EXEC CICS GETMAIN SHARED command, and remains allocated until explicitly freed by the same, or by a different, task. Shared storage can be used to exchange data between any CICS tasks that run during the lifetime of the shared storage. Transactions designed in this way must execute in the same CICS region to work correctly. The dynamic transaction routing program should ensure that transactions using shared storage are routed to the same AOR.

Figure 13 illustrates the use of shared storage.




Figure 13. Illustration of intertransaction affinity created by use of shared storage. The dynamic transaction routing program needs to be aware of this affinity, and ensure it routes TRN2 to the same AOR as TRN1.

If the two transactions shown in Figure 13 are parts of a pseudoconversational transaction, the use of shared storage should be replaced by a COMMAREA (provided that the amount of storage fits within the COMMAREA size limits.)

Using the LOAD PROGRAM HOLD command

A program (or table) that CICS loads in response to an EXEC CICS LOAD PROGRAM HOLD command remains in directly addressable storage until explicitly released by the same, or by a different, task. Any CICS tasks that run while the loaded program (table) is held in storage can use the loaded program's storage to exchange data, provided that:

- The program is not loaded into read-only storage, or

- The program is not defined to CICS with RELOAD(YES) on it's RDO PROGRAM resource definition.

Although you could use a temporary storage queue to make the address of the loaded program's storage available to other tasks, the more usual method would be for other tasks to issue a LOAD PROGRAM command also, with the SET(ptr_ref) option so that CICS can return the address of the held program.

The nature of the affinity caused by the use of the EXEC CICS LOAD PROGRAM HOLD command is virtually identical to that caused by the use of GETMAIN SHARED storage (see Figure 13 on page 98 and Figure 14), and the same rule applies: to preserve the application design, the dynamic routing program must ensure that all transactions that use the address of the loaded program (or table) are routed to the same AOR.




Figure 14. Illustration of intertransaction affinity created by use of shared storage. The dynamic transaction routing program needs to be aware of this affinity, and ensure it routes TRN2 to the same AOR as TRN1.

Note: This rule applies also to programs defined with the RESIDENT option on the resource definition for the loaded program (whether or not the HOLD option is specified on the LOAD command). However, regardless of affinity considerations, it is unsafe to use the RESIDENT option to enable transactions to share data, because programs defined with RESIDENT are subject to EXEC CICS SET PROGRAM(program_name) NEWCOPY commands, and can therefore be changed.

The rule also applies to a nonresident, nonheld, loaded program where the communicating tasks are synchronized.

Sharing task-lifetime storage

The use of any task-lifetime storage belonging to one task can be shared with another task, provided the owning task can pass the address to the other task in the same CICS address space. This technique creates an affinity among the communicating tasks, and requires any task retrieving and using the passed address to execute in the same AOR as the task owning the task-lifetime storage.

For example, it is possible to use a temporary storage queue to pass the address of a PL/I automatic variable, or the address of a COBOL working-storage structure (see Figure 15 for an example).




Figure 15. Illustration of intertransaction affinity created by use of task-lifetime storage. TRN2 must execute in the same AOR as TRN1. Also, TRN1 must not terminate until TRN2 has finished using its task-lifetime storage.

For two tasks to share task-lifetime storage belonging to one of them requires the tasks to be synchronized in some way. See Table 9 on page 101 for commands that provide ways of suspending and resuming a task that passes the address of its local storage.

<i>Table 9. Methods for suspending and resuming (synchronizing) tasks</i>	
Suspending operation	Resuming operation
WAIT EVENT, WAIT EXTERNAL, WAITCICS	VSE POST or hand-post
RETRIEVE WAIT	START
DELAY	CANCEL
POST	CANCEL
START	CANCEL
ENQ	DEQ

Some of these techniques themselves require that the transactions using them must execute in the same AOR, and these are discussed later in this chapter. However, even in those cases where tasks running in different AORs can be synchronized, it is not safe to pass the address of task-lifetime storage from one to the other. Even without dynamic transaction routing, designs that are based on the synchronization techniques shown in Table 9 are fundamentally unsafe because it is possible that the storage-owning task could be purged.

Notes:


1. Using synchronization techniques, such as EXEC CICS RETRIEVE WAIT / START, to allow sharing of task-lifetime storage is unsafe in CICS/VSE® Version 2 because the task issuing, for example, the EXEC CICS RETRIEVE WAIT could be purged by a CEMT SET TASK(...) PURGE command. In CICS Transaction Server for VSE/ESA Version 1 the SPURGE parameter on the RDO TRANSACTION resource definition could be used to protect the first task, but even so the design is not recommended.
2. No intertransaction affinity is caused in those cases where the task sharing another task's task-lifetime storage is started by an EXEC CICS START command, except when the EXEC CICS START command is function-shipped to a remote system. This is true even if the started task is terminal-related (probably to a printer), because started tasks are always routed back to the AOR in which the START command was issued.

Using the WAIT EVENT command

The WAIT EVENT command is used to synchronize a task with the completion of an event performed by some other CICS or VSE task.

The completion of the event is signaled (posted) by the setting of a bit pattern into the event control block (ECB). Both the waiting task and the posting task must have direct addressability to the ECB, hence both tasks must execute in the same AOR. The use of a temporary storage queue is one way that the waiting task can pass the address of the ECB to another task.

This synchronization technique is illustrated in Figure 16 on page 102.



WAIT EXTERNAL command

Figure 16. Illustration of intertransaction affinity created by use of. TRN2 must execute in the same AOR as TRN1.

If TRN2 shown in Figure 16 executed in a different AOR from TRN1, the value of *ptr-ref* would be invalid, the post operation would have unpredictable results, and the waiting task would never be resumed. For this reason, a dynamic transaction routing program must ensure that a posting task executes in the same AOR as the waiting task to preserve the application design.

The same considerations apply to the use of EXEC CICS WAIT EXTERNAL and EXEC CICS WAITCICS commands for synchronizing tasks.

Using ENQ and DEQ commands

The ENQ and DEQ commands are used to serialize access to a shared resource. However, these commands work only for CICS tasks running in the same region, and cannot be used to serialize access to a resource shared by tasks in different regions.

The use of ENQ and DEQ for serialization is illustrated in Figure 17 on page 103.




Figure 17. Illustration of intertransaction affinity created by use of ENQ / DEQ commands. TRN2 must execute in the same AOR as TRN1.

If TRN2 shown in Figure 17 executed in a different AOR from TRN1, TRN2 would not be suspended while TRN1 accessed the shared resource. For this reason, a dynamic transaction routing program must ensure that all tasks that enqueue on a given resource name must execute in the same AOR to preserve the application design. TRN2 would, of course, be serialized with other CICS tasks that issue ENQ commands on the same resource name in its AOR.

Suspect programming techniques

Some CICS application programming techniques may create an affinity among transactions depending on how they are implemented.

The programming techniques that may be suspect are described in the sections that follow.

Using temporary storage

CICS application programs commonly use temporary storage (TS) queues to hold temporary application data, and to act as scratch pads.

Sometimes a TS queue is used to pass data between application programs that execute under one instance of a transaction (for example, between programs that pass control via an EXEC CICS LINK or EXEC CICS XCTL command in a multi-program transaction). Such use of a TS queue requires the queue to exist only for the lifetime of the transaction instance, and therefore it does not need to be shared between different AORs, because a transaction instance executes in one, and only one, AOR.

Note: This latter statement is not strictly true in the case of a multi-program transaction, where one of the programs is linked by a distributed program link command and the linked-to program resides in a remote system. In this case, the program linked by a DPL command runs under another CICS task in the remote region. The recommended method for passing data to a DPL program is via a COMMAREA, but if a TS queue is used for passing data in a DPL application, the queue must be shared between the two regions.

However, a DPL command does not create an affinity of the type described here, because it is not affected by dynamic transaction routing.

Sometimes a TS queue holds information that is specific to the AOR, or holds read-only data. In this case the TS queue can be replicated in each AOR, and no sharing of data between AORs is necessary.

However, in many cases a TS queue is used to pass data between transactions, in which case the queue must be globally accessible to enable the transactions using the queue to run in any dynamically selected AOR. It is possible to make a TS queue globally accessible by function shipping TS requests to a queue-owning region (QOR), provided the TS queue can be defined as remote.

In a pseudoconversational transaction, a better way is to change the program to use a COMMAREA to pass data between the phases of the conversation.

Naming conventions for remote queues

To define a queue as remote you must include an entry for the queue in a temporary storage table (TST). TS queue names are frequently generated dynamically, but they can also be unique fixed names. The TST naming convention allows for dynamic names by accepting generic names formed by a constant prefix, to which a CICS application program can add a variable suffix. (Generic names are formed from the leading characters of the 8-character queue names and can be up to seven characters long. Names in a TST entry using all eight characters specify unique TS queues.)

The usual convention is a 4-character prefix (for example, the transaction identifier) followed by a 4-character terminal identifier as the suffix. This generates queue names that are unique for a given terminal. Such generic queue names can be defined easily as remote queues that are owned, for example, by a QOR, thus avoiding transaction affinity problems. However, if the naming convention for dynamically named queues does not conform to this rule, the queue cannot be defined as remote, and all transactions that access the queue must be routed to the AOR where the queue was created. Furthermore, a TS queue name cannot be changed from a local to a remote queue name using the global user exits for TS requests.

See Figure 18 on page 105 for an illustration of the use of a remote queue-owning region.




Figure 18. Using remote queues to avoid intertransaction affinity relating to temporary storage. This example shows a combined file-owning and queue-owning region. Separate regions can be used, but these require special care during recovery operations because of 'in-doubt' windows that can occur when recovering data managed independently by file control and temporary storage control.

Exception conditions for globally accessible queues

When you eliminate intertransaction affinity relating to TS queues by the use of a global QOR, you must also take care to review exception condition handling. This is because some exception conditions can occur that previously were not possible when the transactions and the queue were local in the same region. This situation arises because the AOR and QOR can fail independently, causing circumstances where:

- The queue already exists, because only the AOR failed while the QOR continued.
- The queue is not found, because only the QOR failed while the AOR continued.

Using transient data

Another form of data queue that CICS application programs commonly use is the transient data queue (TD). The dynamic transaction routing considerations for TD queues have much in common with those for temporary storage. If transactions use a TD queue that needs to be shared, and if you want the transactions to be dynamically routed to an AOR, you must ensure that the TD queues are globally accessible.

All TD queues must be predefined in a destination control table (DCT) and these definitions can be changed easily to support a remote TD queue-owning region (QOR).

However, there is a restriction for TD queues that use the trigger function. The transaction to be invoked when the trigger level is reached must be defined as a local transaction in the region where the queue resides (in the QOR). Thus the trigger transaction must execute in the QOR. However, any terminal associated

with the queue need not be defined as a local terminal in the QOR. This does not create an intertransaction affinity.

See Figure 19 for an illustration of the use of a remote transient data queue-owning region.




Figure 19. Using remote queues to avoid intertransaction affinity relating to transient data. In the DCT defined for the AORs (suffix=XX), all the TD queues are remote, defined as owned by the QOR (SYSID=CICQ). In the DCT defined for the QOR (suffix=YY), all the TD queues are local, some with trigger levels.

Exception conditions for globally accessible queues

When you eliminate intertransaction affinity relating to TD queues by the use of a global QOR, there should not be any new exception conditions (other than SYSIDERR if there is a system definition error or failure).

Using the RETRIEVE WAIT and START commands

The use of some synchronization techniques permit the sharing of task-lifetime storage between two synchronized tasks. For example, the EXEC CICS RETRIEVE WAIT and EXEC CICS START commands could be used for this purpose, as illustrated in Figure 20 on page 107.

In this example, TRN1 is designed to retrieve data from an asynchronous task, TRN2, and therefore must wait until TRN2 makes the data available. Note that for this mechanism to work, TRN1 must be a terminal-related transaction.

The steps are as follows:

1. TRN1 writes data to a TS queue, containing its TRANSID and TERMID.

2. To cause itself to suspend, TRN1 issues an EXEC CICS RETRIEVE WAIT command, which causes CICS to suspend the task until the RETRIEVE can be satisfied, which is not until TRN2 issues a START command with data passed via the FROM parameter.
3. However, TRN2 can only issue a START command to resume TRN1 if it knows the TRANSID and TERMID of the suspended task (TRN1 in our example). Thus it reads the TS queue to obtain the information written by TRN1. Using a temporary storage queue is one way that this information can be passed by the suspending task.
4. Using the information from the TS queue, TRN2 issues the START command for TRN1, causing CICS to resume TRN1 by satisfying the outstanding RETRIEVE WAIT.




Figure 20. Illustration of task synchronization using EXEC CICS RETRIEVE WAIT and EXEC CICS START commands

In the example of task synchronization using EXEC CICS RETRIEVE WAIT and EXEC CICS START commands shown in Figure 20, the EXEC CICS START command that satisfies the RETRIEVE WAIT must:

- Be issued in same AOR as the transaction (TRN1 in our example) issuing the EXEC CICS RETRIEVE WAIT command, or
- Specify the SYSID of the AOR where the EXEC CICS RETRIEVE WAIT command was executed, or
- Specify a TRANSID (TRN1 in our example) that is defined as a remote transaction residing on the AOR that executed the EXEC CICS RETRIEVE WAIT command.

An application design based on the remote TRANSID technique works only for two AORs. An application design using the SYSID option on the START command works only for multiple AORs if all AORs have connections to all other AORs (which may not be desirable). In either case, the application programs need to be modified: there is no acceptable way to use this programming technique in a dynamic transaction routing environment except by imposing restrictions on the routing program. In general, this means that the dynamic transaction routing program has to ensure that TRN2 has to execute in the same region as TRN1 to preserve the application design.

Using the **START** and **CANCEL REQID** commands

When this CICS application programming technique is used, one transaction issues an EXEC CICS START command to start another transaction after a specified interval. Another transaction (not the one requested on the EXEC CICS START command) determines that it is no longer necessary to run the requested transaction (which is identified by the REQID parameter) and cancels the START request. Note that the cancel is effective only if the specified interval has not yet expired.

A temporary storage queue is one way that the REQID can be passed from task to task.

Figure 21 illustrates this programming technique.




Figure 21. Illustration of the use of the EXEC CICS START and EXEC CICS CANCEL REQID commands

When this application programming technique is used, the EXEC CICS CANCEL command that cancels the START request must:

- Specify the REQID option, but the EXEC CICS START command need not. This is because, provided the NOCHECK option is not specified on the EXEC CICS START command, CICS generates a REQID for the request and stores it in the EXEC interface block (EIB) in field EIBREQID.
- Be issued in same AOR that the EXEC CICS START command was executed in, or
- Specify the SYSID of the AOR where the EXEC CICS START command was executed, or
- Specify a TRANSID (TRNX in our example) that is defined as a remote transaction residing on the AOR where the EXEC CICS START command was executed.

Note: An EXEC CICS START command is not necessarily executed in the same region as the application program issuing the command. It can be function shipped and executed in a different CICS region. The above rules apply to the region where the START command is finally executed.

An application design based on the remote TRANSID technique works only for two AORs. An application design using the SYSID option on the cancel command works only for multiple AORs if all AORs have connections to all other AORs. In either case, the application programs need to be modified: there is no acceptable way to use this programming technique in a dynamic transaction routing environment except by imposing restrictions on the routing program.

In general, this means that the dynamic transaction routing program has to ensure that TRN2 executes in the same region as TRN1 to preserve the application design, and also that TRNX is defined as a local transaction in the same region.

Using the EXEC CICS DELAY and EXEC CICS CANCEL REQID commands

When this CICS application programming technique is used, one task can resume another task that has been suspended by an EXEC CICS DELAY command.

A DELAY request can be canceled only by a different task from the one issuing the EXEC CICS DELAY command, and the EXEC CICS CANCEL command must specify the REQID associated with DELAY command. Both the EXEC CICS DELAY command and the EXEC CICS CANCEL command must specify the REQID option to use this technique.

The steps involved in this technique using a temporary storage queue to pass the REQID are as follows:

1. A task (TRN1) writes a predefined DELAY REQID to a TS queue. For example:

```
EXEC CICS WRITEQ TS
      QUEUE('DELAYQUE')
      FROM(reqid_value)
```

2. The task waits on another task by issuing an EXEC CICS DELAY command, using the *reqid_value* as the REQID. For example:

```
EXEC CICS DELAY
      INTERVAL(1000)
      REQID(reqid_value)
```

3. Another task, TRN2, reads the REQID of the DELAY request from TS queue called 'DELAYQUE'.
4. TRN2 completes its processing, and resumes TRN1 by canceling the DELAY request.

The process using a TS queue is illustrated in Figure 22.

Another way to pass the REQID when employing this technique would be for TRN1 to start TRN2 using the EXEC CICS START command with the FROM and TERMID options. TRN2 could then obtain the REQID with the EXEC CICS RETRIEVE command, using the INTO option.




Figure 22. Illustration of the use of the EXEC CICS DELAY and EXEC CICS CANCEL REQID commands

When this application programming technique is used, the EXEC CICS CANCEL command that cancels the DELAY request must:

- Be issued in same AOR as the EXEC CICS DELAY command was executed in, or
- Specify the SYSID of the AOR where the EXEC CICS DELAY command was executed, or
- Specify a TRANSID (TRN1 in our example) that is defined as a remote transaction residing on the AOR where the EXEC CICS DELAY command was executed.

An application design based on the remote TRANSID technique works only for two AORs. An application design using the SYSID option on the cancel command

works only for multiple AORs if all AORs have connections to all other AORs. In either case, the application programs need to be modified: there is no acceptable way to use this programming technique in a dynamic transaction routing environment except by imposing restrictions on the routing program.

If the EXEC CICS CANCEL command is issued by a transaction that is initiated from a terminal, it is possible that the transaction could be dynamically routed to the wrong AOR.

Using the EXEC CICS POST and EXEC CICS CANCEL REQID commands

The EXEC CICS POST command is used to request notification that a specified time has expired. Another transaction (TRN2) can force notification, as if the specified time has expired, by issuing a CANCEL of the POST request.

The time limit is signaled (posted) by CICS by setting a bit pattern in the event control block (ECB). To determine whether notification has been received, the requesting transaction (TRN1) has either to test the ECB periodically, or to issue a WAIT command on the ECB.

A TS storage queue is one way that can be used to pass the REQID of the POST request from task to task.

Figure 23 illustrates this technique.




Figure 23. Illustration of the use of the EXEC CICS POST command

If this technique is used, the dynamic transaction routing program has to ensure that TRN2 executes in the same CICS region as TRN1 to preserve the application design.

The EXEC CICS CANCEL command that notifies the task that issued the POST must:

- Be issued in same AOR that the EXEC CICS POST command was executed in, or
- Specify the SYSID of the AOR where the EXEC CICS POST command was executed, or
- Specify a TRANSID (TRN1 in our example) that is defined as a remote transaction residing on the AOR where the EXEC CICS POST command was executed.

An application design based on the remote TRANSID technique works only for two AORs. An application design using the SYSID option on the cancel command works only for multiple AORs if all AORs have connections to all other AORs. In either case, the application programs need to be modified: there is no acceptable way to use this programming technique in a dynamic transaction routing environment except by imposing restrictions on the routing program.

In general, this means that the dynamic transaction routing program has to ensure that TRN2 executes in the same region as TRN1 to preserve the application design.

Clearly, there is no problem if the CANCEL is issued by the same task that issued the POST. If a different task cancels the EXEC CICS POST command, it must specify REQID to identify the particular instance of that command. Hence the EXEC CICS CANCEL command with REQID is indicative of an intertransaction affinity problem. However, REQID need not be specified on the EXEC CICS POST command because CICS will automatically generate a REQID and pass it to the application in EIBREQID.

Detecting intertransaction affinities

You can use one of the following methods to detect affinities, although you are strongly recommended to use the utility.

- Review application design, paying particular attention to the techniques used for intertransaction communication.
- Search the source of application programs, looking for instances of the EXEC CICS commands that can give rise to intertransaction affinity.
- Run a trace analysis program that can analyze CICS auxiliary trace. For example, if you run the CICS trace utility program, DFHTU410, with the ABBREV option to format CICS auxiliary trace output, you can analyze the resulting abbreviated trace data to find instances of suspect commands.

Intertransaction affinities caused by application generators

Application generators may give rise to particularly difficult problems of intertransaction affinity:

1. The affinity may be hidden from the application programmer.
2. The application generator may have a different concept of a transaction to CICS: it is affinity among CICS transactions that is of concern, because these are the entities that are dynamically routed.
3. Some application generators use a single transaction code for all transactions within an application, making it difficult for the router to select those instances of transactions that have affinities.

Duration and scope of intertransaction affinities

When planning your dynamic transaction routing strategy, and planning how to manage intertransaction affinities, it is important to understand the concepts of affinity relations and affinity lifetimes. The relations and lifetimes of intertransaction affinities must be taken into account when designing a dynamic transaction routing program, because they define the scope and duration of intertransaction affinities. Clearly, the ideal situation for a dynamic transaction routing program is for there to be no intertransaction affinities at all, which means there are no restrictions in the choice of available AORs for dynamic transaction routing. However, even when intertransaction affinities do exist, there are limits to the scope of these affinities, the scope of the affinity being determined by affinity relation and affinity lifetime.

Understanding the relations and lifetimes of transaction affinities is important in deciding how to manage them in a dynamic transaction routing environment.

Affinity transaction groups

In order to manage affinities within a dynamic transaction routing environment, you must first categorize transactions by their affinity. One way to do this is to place transactions in groups, where a group is a set of transactions that have intertransaction affinity. Each affinity transaction group (or affinity group, for short) thus represents a group of transactions that have an affinity with one another. Defining affinity groups is one way that a dynamic transaction routing program can determine which AOR a transaction should be routed to.

Clearly, the more intertransaction affinity you have in a given CICS workload, the less effective a dynamic transaction routing program can be in balancing the workload across a CICSplex. To minimize the impact of intertransaction affinity, affinities within an affinity group can be characterized by their relation and lifetime. These relation and lifetime attributes determine the scope and duration of an affinity.

Thus, ideally, an affinity transaction group consists of an affinity group identifier, a set of transactions that constitute the affinity group, with the affinity relation and affinity lifetime associated with the group.

Relations and lifetimes

When you create an affinity group, assign to the group the appropriate affinity relation and affinity lifetime attributes. The relation determines how the dynamic transaction routing program is to select a target AOR for a transaction instance associated with the affinity, and the lifetime determines when the affinity is ended.

There are three possible affinity relations that you can assign to your affinity groups:

- “The global relation”
- “The LUsername (terminal) relation” on page 115
- “The userid relation” on page 116

These are described in the following sections, together with the permitted lifetimes for each relation.

The global relation

A group of transactions whose affinity relation is defined as global is one where **all** instances of all transactions in the group that are initiated from any terminal must execute in the same AOR for the lifetime of the affinity. The affinity lifetime for global relations can be as follows:

System The affinity lasts for as long as the target AOR exists, and ends whenever the AOR terminates (at a normal, immediate, or abnormal termination).

Permanent The affinity extends across all CICS restarts. This is the most restrictive of all the intertransaction affinities.

An example of a global intertransaction affinity with a lifetime of permanent is where the transaction uses (reads and/or writes) a local, recoverable, temporary storage queue, and where the TS queue name is **not** derived from the terminal. (You can only specify that a TS queue is recoverable in the CICS region in which the queue is local.)

Generally, transactions in this affinity category are not suitable candidates for dynamic transaction routing. Consider making them statically routed transactions.

An example of a global relation is illustrated in Figure 24 on page 115.




Figure 24. Managing intertransaction affinity with global relation and permanent lifetime. In this example, the first instance of transid GGGG, from any terminal, starts a permanent-lifetime affinity. The first instance of GGGG can be routed to any suitable AOR (AOR1 thru AOR6), but all other instances, from any terminal, must be routed to whichever AOR is selected for GGGG.

The LUname (terminal) relation

A group of transactions whose affinity relation is defined as LUname is one where **all** instances of all transactions in the group that are initiated from the **same terminal** must execute in the same AOR for the lifetime of the affinity. The affinity lifetime for LUname relations can be as follows:

Pseudoconversation

The affinity lasts for the whole pseudoconversation, and ends when the pseudoconversation ends at the terminal.

Logon

The affinity lasts for as long as the terminal remains logged-on to CICS, and ends when the terminal logs off.

System

The affinity lasts for as long as the target AOR exists, and ends whenever the AOR terminates (at a normal, immediate, or abnormal termination).

Permanent The affinity extends across all CICS restarts.

A typical example of transactions that have an LUname relation are those that:

- Use a local TS queue to pass data between the transactions in a pseudoconversation, and
- The TS queue name is derived, in part, from the terminal name (see “Naming conventions for remote queues” on page 104 for information about TS queue names).

These types of transaction can be placed in an affinity group with a relation of terminal and lifetime of pseudoconversation. When the dynamic transaction routing

program detects the first transaction in the pseudoconversation initiated by a specific terminal (LUNAME), it is free to route the transaction to any AOR that is a valid candidate for that transaction. However, any subsequent transaction within the affinity group that is initiated at **the same terminal** must be routed to the same AOR as the transaction that started the pseudoconversation. When the affinity ends (at the end of the pseudoconversation on a given terminal), the dynamic transaction routing program is again free to route the first transaction to any candidate AOR.

This form of affinity is manageable and does not impose too severe a constraint on dynamic transaction routing, and may occur commonly in many CICSplexes. It can be managed easily by a dynamic transaction routing program, and should not inhibit the use of dynamic transaction routing.

This example is illustrated in Figure 25.




Figure 25. Managing intertransaction affinity with LUNAME relation and pseudoconversation lifetime. In this example, each instance of transid AAAA from a terminal starts a pseudoconversational-lifetime affinity. AAAA can be routed to any suitable AOR (AOR1 thru AOR6), but other transactions in the group from the same terminal (IGKS201 in this example) must be routed to whichever AOR is selected for AAAA.

The userid relation

A group of transactions whose affinity relation is defined as userid is one where **all** instances of the transactions that are initiated from a terminal and executed on behalf of the **same userid**, must execute in the same AOR for the lifetime of the affinity. The affinity lifetime for userid relations can be as follows:

Signon The affinity lasts for as long as the user is signed on, and ends when the user signs off. Note this lifetime is only possible in those situations where only one user per userid is permitted. Signon lifetime cannot be detected if multiple users are permitted to be signed on with the same userid at the same time (at different terminals).

System The affinity lasts for as long as the target AOR exists, and ends whenever the AOR terminates (at a normal, immediate, or abnormal termination).

Permanent The affinity extends across all CICS restarts.

A typical example of transactions that have a userid relation is where the userid is used dynamically to identify a resource, such as a TS queue. The least restrictive of the affinities in this category is one that lasts only for as long as the user remains signed on.

An example of an affinity group with the userid relation and a signon lifetime is shown in Figure 26.




Figure 26. Managing intertransaction affinity with userid relation and sign-on lifetime. In this example, any instance of a transaction from a terminal starts a sign-on lifetime affinity. It can be routed to any suitable AOR (AOR1 thru AOR6), but other transactions in the group for the same user (ANOTHER in this example) must be routed to whichever AOR is selected for the first instance of a transaction in the group.

Recommendations

The best way to deal with intertransaction affinity is to avoid creating intertransaction affinity in the first place.

Where it is not possible to avoid affinities:

- Make the intertransaction affinity easily recognizable, by using appropriate naming conventions, and
- Keep the lifetime of the affinities as short as possible

Even if you could avoid intertransaction affinities by changing your application programs, this is not absolutely necessary provided you include logic in your

dynamic transaction routing program to cope with the affinity. Finally, you can statically route the affected transactions.

Chapter 6. Intercommunication considerations

This chapter provides only a summary of what you need to consider when writing applications that communicate with other CICS systems. For further information, see the *CICS Intercommunication Guide*.

You can run application programs in a CICS intercommunication environment using one or more of the following functions, which are described in the rest of this chapter.

“Transaction routing” on page 120

Enables a terminal in one CICS system to run a transaction in another CICS system.

“Function shipping” on page 120

Enables your application program to access resources in another CICS system.

“Distributed program link (DPL)” on page 121

Enables an application program running in one CICS region to link to another application program running in a remote CICS system.

“Asynchronous processing” on page 132

Enables a CICS transaction to start another transaction in a remote system and optionally pass data to it.

“Distributed transaction processing (DTP)” on page 132

Enables a CICS transaction to communicate with a transaction running in another system. There are two interfaces available for DTP; command-level EXEC CICS and the SAA interface for DTP known as Common Programming Interface Communications (CPI Communications).

“Common Programming Interface (CPI) Communications” on page 132

Provides DTP on APPC connections and defines an API that can be used on multiple system platforms.

“External CICS interface (EXCI)” on page 133

Enables a non-CICS program running in VSE to allocate and open sessions to a CICS system, and to issue DPL requests on these sessions.

The intercommunication aspects of the CICS Front End Programming Interface (FEPI) are not discussed in this book. See the *CICS Front End Programming Interface User's Guide* for details about FEPI.

Design considerations

If your application program uses more than one of these facilities, you obviously need to bear in mind the design considerations for each one. Also, if your program uses more than one intersystem session for distributed transaction processing, it must control each session according to the rules for that type of session.

Programming language

Generally speaking, you can use COBOL, C, PL/I, or assembler language to write application programs that use CICS intercommunication facilities. There is, however, an exception. You can only use C or assembler language for DTP application programs that hold APPC unmapped conversations using the EXEC CICS GDS commands.

Transaction routing

Transactions that can be invoked from a terminal owned by another CICS system, or that can acquire a terminal owned by another CICS system during transaction initiation, must be able to run in a transaction routing environment.

Generally, you can design and code such a transaction just like one used in a local environment. However, there are a few restrictions related to basic mapping support (BMS), pseudoconversational transactions, and the terminal on which your transaction is to run. All programs, tables, and maps that are used by a transaction **must** reside on the system that owns the transaction. (You can duplicate them in as many systems as you need.)

Some CICS transactions are related to one another, for example, through common access to the CWA or through shared storage acquired using a GETMAIN command. When this is true, the system programmer must ensure that these transactions are routed to the same CICS system. Avoid (where possible) any techniques that might create intertransaction affinities that could adversely affect your ability to perform dynamic transaction routing. See Chapter 5, "Affinity among transactions" on page 91 for more information about transaction affinity.

When a request to process a transaction is transmitted from one CICS system to another, transaction identifiers can be translated from local names to remote names. However, a transaction identifier specified in a RETURN command is not translated when it is transmitted from the transaction-owning system to the terminal-owning system.

Function shipping

You code a program to access resources in a remote system in much the same way as if they were on the local system. You can use:

DL/I calls (EXEC DLI commands)

To access data associated with a remote CICS system.

File control commands

To access files on remote systems. Note that requests that contain the TOKEN keyword may not be function-shipped.

Interval control commands

To initiate transactions on other CICS systems, or other non-CICS systems that implement SNA LUTYPE 6 protocols, such as IMS™.

Temporary storage commands

To access data from temporary storage queues on remote systems.

Transient data commands

To access transient data queues on remote systems.

Three additional exception conditions can occur with remote resources. They occur if the remote system is not available (SYSIDERR), if a request is invalid (ISCINVREQ), or if the mirror transaction abends (ATNI for ISC connections and AZI6 for MRO).

Distributed program link (DPL)

The distributed program link function enables a CICS program (the client program) to call another CICS program (the server program) in a remote CICS region. There are several reasons why you might want to design your application to use distributed program link. Some of these are:

- To separate the end-user interface (for example, BMS screen handling) from the application business logic, such as accessing and processing data, to enable parts of the applications to be ported from host to workstation more readily
- To obtain performance benefits from running programs closer to the resources they access, and thus reduce the need for repeated function shipping requests
- To reuse existing business logic as you move it into client and server operations
- To offer a simple alternative, in many cases, to writing distributed transaction processing (DTP) applications

You can specify that the program to which an application is linking is remote, by specifying the remote system name:

- In an XPCREQ global user exit
- On an EXEC CICS LINK command
- On the installed PROGRAM resource definition²
- Using the program autoinstall user-replaceable module (URM)

The basic flow in distributed program link is described in the *CICS Intercommunication Guide*. The following terms, illustrated in Figure 27 on page 122, are used in the discussion of distributed program link:

Client region The CICS region running an application program that issues a link to a program in another CICS region.

Server region The CICS region to which a client region ships a link request.

Client program

The application program that issues a remote link request.

Server program

The application program that is specified on the link request and that is executed in the server region.

DPL subset

The subset of EXEC CICS commands to which a server program is restricted when it is executed remotely. Table 11 on page 131 lists the API commands that are prohibited in programs invoked by DPL.

² By **installed PROGRAM definition**, we mean a PROGRAM definition that has been installed statically, or dynamically using AUTOINSTALL, or by an EXEC CICS CREATE command.

In Figure 27 on page 122, a program (known as a **client program**) issues a program-control EXEC CICS LINK command for a program called PROG2 (known as the **server program**). The EXEC CICS LINK command includes a SYSID option that indicates that PROG2 is owned by a remote CICS system called CICY. If this option is not coded, CICS then searches the installed program definitions and can discover that PROG2 is owned by a remote system. CICS changes the EXEC CICS LINK request into a suitable transmission format, and then ships it to the remote system (CICY) for execution.

In CICY, the mirror transaction (AC20) is attached. The mirror program recreates the original request, issues it on CICY and, when the server program has run to completion, returns any communication-area data to CICX. There are two instances of transaction AC20; one in each system.




Figure 27. Illustration of distributed program link

Using the distributed program link function

The distributed program link function provides a number of options. You can specify:

- The name of the remote system (the server region).
- The name of the server program, if it is known by a different name in the server region.
- That you want to run the linked program locally, but restrict it to the distributed program link subset of the application programming interface (API) for testing purposes. (Server programs cannot use the entire CICS API when executed remotely; the restrictions are listed in Table 11 on page 131.)
- That the server program takes a syncpoint independently from the client.
- The name of the transaction you want the program to run under in the server region.

- The data length of the COMMAREA being passed.

A server program can itself issue a distributed program link and act as a client program with respect to the program it links to.

The options shown in Table 10 are used on the EXEC CICS LINK command and the program resource definition in support of the distributed program link facility.

<i>Table 10. Options on EXEC CICS LINK command and program resource definitions to support DPL</i>		
Where specified	Keyword	Description
EXEC CICS LINK command options	DATALENGTH	Specifies the length of the contiguous area of storage (from the start of the COMMAREA) that the application is sending to a server program.
	SYSID	Specifies the name of the connection to the server region to which you want the client region to ship the program link request. Note: SYSID specified on the EXEC CICS LINK command overrides the REMOTESYSTEM name specified on the program resource definition.
	SYNCONRETURN	Specifies that you want the server region to take a syncpoint on successful completion of the server program. Note: This option is unique to the EXEC CICS LINK command and cannot be specified on the program resource definition.
	TRANSID	Specifies the name of the transaction that the server region is to attach for execution of the server program. Note: TRANSID specified on the EXEC CICS LINK command overrides any TRANSID specified on the program resource definition.
Program resource definition options	REMOTESYSTEM	Specifies the name of the connection to the server region (SYSID) to which you want the client region to ship the program link request.
	REMOTENAME	Specifies the name by which the program is known in the server region (if different from the local name).
	EXECUTIONSET	Specifies whether the program is restricted to the distributed program link subset of the CICS API. Note: This option is unique to the program definition and cannot be specified on the EXEC CICS LINK command.
	TRANSID	Specifies the name of the transaction that the server region is to attach for execution of the server program.

Note: Programming information, including the full syntax of the EXEC CICS LINK command, is in the *CICS Application Programming Reference* manual but note that, for a distributed program link, you cannot specify the INPUTMSG or INPUTMSGLEN options.

Examples of distributed program link

A COBOL example of a distributed program link command is shown in Figure 28. The numbers down the right-hand side of the example refer to the numbered sections, following the figure, which give information about each option.

```
EXEC CICS LINK PROGRAM(' DPLPROG' )           1
          COMMAREA(DPLPROG-DATA-AREA)         ] 2
          LENGTH(24000)
          DATALENGTH(100)
          SYSID(' CI CR' )                   3
          TRANSID(' AC20' )                  4
          SYNCONRETURN                         5
END-EXEC.
```

Figure 28. COBOL example of a distributed program link

1. The program name of the server program

A program may have different names in the client and server regions. The program name you specify on the EXEC CICS LINK command depends on whether or not you specify the SYSID option, and also on whether the program definition in the client region specifies the REMOTENAME option.

If the program link command specifies a remote system on the SYSID option, CICS ships the link request to the server region without reference to the program resource definition in the client region. In this case, the program name specified on the EXEC CICS LINK command must be the name by which the program is known in the server region.

If you do not specify the SYSID option on an EXEC CICS LINK command, however, the local name of the program is used. CICS looks up the program resource definition in the local (client) region. If the program is remote, CICS ships the request to the region specified in REMOTESYSTEM using the name by which the program is known in the server region. This is the REMOTENAME if one was specified; otherwise the local name is used.

If the region name (SYSID) on the EXEC CICS LINK command is the same name as the client region, CICS initially processes the request as though it is a local link. CICS checks the local program resource definition and, if it specifies a remote system name, the client region ships the request to the region specified in the REMOTESYSTEM parameter. In this case, the remote system name on the program resource definition overrides the SYSID name on the program link command. If REMOTESYSTEM is not specified in the program resource definition, or it names the local region name, CICS invokes the program in the client region (a local link).

2. The communication data area (COMMAREA)

The length of the communication area (COMMAREA) to be made available to the remote program is defined by the LENGTH option, as in the case of an ordinary EXEC CICS LINK command.

To improve performance, you can specify the DATALENGTH option on the EXEC CICS LINK command. This allows you to specify the amount of COMMAREA data you want the client region to pass to the server program. Typically, you use this option when a large COMMAREA is required to hold data that the server program is to return to the client program, but only a small

amount of data needs to be sent to the server program by the client program, as in the example.

See the *CICS Application Programming Reference* manual for the full syntax of the EXEC CICS LINK command and details of the DATALENGTH option.

3. The remote system ID (SYSID)

The SYSID option on the EXEC CICS LINK command, or the REMOTESYSTEM option on the program resource definition, enables you to specify a 4-character name for the server region to which you want the application region to ship a program link request. This is the name of the connection definition installed in the client region defining the connection with the server region. (CICS uses the connection name in a table look-up to obtain the netname (VTAM APPLID) of the server region.) The name of the server region you specify on the SYSID or REMOTESYSTEM option can be the name of the client region in which case the program is run locally.

If the server region is unable to load or run the requested program (DPLPROG in our example), CICS returns the PGMIDERR condition to the client program in response to the link request. Note that EIBRESP2 values are not returned over the link for a distributed program link request where the error is detected in the server region. For errors detected in the client region, EIBRESP2 values are returned.

You can also specify, or modify, the name of a server region in an XPCREQ global user exit program. See the *CICS Customization Guide* for programming information about the XPCREQ global user exit point.

4. The remote transaction (TRANSID) to be attached

The TRANSID option is available on both the EXEC CICS LINK command and the program resource definition. This enables you to tell the server region the transaction identifier to use when it attaches the mirror task under which the server program runs. If you specify the TRANSID option either on the EXEC CICS LINK command or in the program resource definition, you must define the transaction in the server region, and associate it with the supplied mirror program, DFHMIRS. This option allows you to specify your own attributes on the transaction definition for the purpose of performance and fine tuning. For example, you could vary the task priority and transaction class attributes.

You are recommended to specify the transaction identifier of the client program as the transaction identifier for the server program. This enables any statistics and monitoring data you collect to be correlated correctly under the same transaction.

The transaction identifier (TRANSID) used on a distributed link program request is passed to the server program as follows:

- If you specify your own transaction identifier for the distributed link program request, this is passed to the server program in the EIBTRNID field of the EIB.
- If SYSID is omitted (or specifies the local system), and you do not specify a TRANSID or an EXEC CICS LINK request, any value specified on the client program resource definition is used.
- Otherwise, it defaults to the client's transaction code, which is the same value that is in the client's EIBTRNID.

5. The synchronization option for the server program

When you specify the SYNCONRETURN option, it means that the resources on the server are committed in a separate logical unit of work immediately before returning control to the client; that is, an implicit EXEC CICS SYNCPOINT is issued for the server just before the server returns control to the client.

Figure 29 on page 126 provides an example of using distributed program link with the SYNCONRETURN option. The SYNCONRETURN option is intended for use when the client program is not updating any recoverable resources, for example, when performing screen handling. However, if the client does have recoverable resources, they are not committed at this point. They are committed when the client itself reaches a syncpoint or in the implicit syncpoint at client task end. You must ensure that the client and server programs are designed correctly for this purpose, and that you are not risking data integrity.

For example, if your client program has shipped data to the server that results in the server updating a database owned by the server region, you only specify an independent syncpoint if it is safe to do so, and when there is no dependency on what happens in the client program. This option has no effect if the server program runs locally in the client region unless EXECUTIONSET(DPLSUBSET) is specified. In this case, the syncpoint rules governing a local link apply; that is, either the client or the server program can take a syncpoint and both the client program's and the server program's resources are committed at that syncpoint.

Without the SYNCONRETURN option, the client commits the logical unit of work for both the client and the server resources, with either explicit commands or the implicit syncpoint at task end. Thus, in this case, the server resources are committed at the same time as the client resources are committed.

Figure 30 on page 127 shows an example of using distributed program link without the SYNCONRETURN option.




Figure 29. Using distributed program link with the SYNCONRETURN option

Note: This includes three logical units of work: one for the client and two for the server. The client resources are committed separately from the server.




Figure 30. Using distributed program link without the SYNCONRETURN option

Note: The implicit or explicit syncpoint causes all client and server resources to be committed. There is only one logical unit of work because the client is responsible for determining when both the client and server resources are committed.

You need to consider the case when the client has a `HANDLE ABEND` command. When the client is handling abends in the server, the client gets control when the server abends. This is also true when the `SYNCONRETURN` option has been specified on the `EXEC CICS LINK` command. In this case, it is recommended that the client issues an abend after doing the minimum of cleanup. This causes both the client logical unit of work and the server logical unit of work to be backed out.

The server program can find out how it was started and therefore whether it is allowed to issue independent syncpoint requests, by issuing the `EXEC CICS ASSIGN STARTCODE` command. The server program can issue a syncpoint only if the `STARTCODE` option on `EXEC CICS ASSIGN` indicates that it was linked to with the `SYNCONRETURN` option. If the server program does issue a syncpoint, this does not commit changes made by the client program. For changes to be committed across the distributed unit of work, the client program must issue the syncpoint. The client program can also backout changes across the distributed unit of work, provided that the server program has not already committed its changes.

Programming considerations for distributed program link

When writing application programs that use distributed program link, consider the factors in the next few sections.

Issuing multiple distributed program links from the same client task

A client task cannot request distributed program links to a single CICS server region using more than one transaction code in a single client unit of work unless the SYNCONRETURN option is specified. It can issue multiple distributed program links to one CICS server system with the same or the default transaction code.

Sharing resources between the client program and server program

The server program does not have access to the lifetime storage of tasks on the client, for example, the TWA. Nor does it necessarily have access to the resources that the client program is using, for example, files, unless the file requests are being function shipped.

Mixing DPL and function shipping to the same CICS system

Great care should be taken when mixing function shipping and DPL to the same CICS system, from the same client task. These are some considerations:

- A client task cannot function ship requests and then use distributed program link with the SYNCONRETURN option in the same client logical unit of work. The distributed program link fails with an INVREQ response. In this case EIBRESP2 is set to 14.
- A client task cannot function ship requests and then use distributed program link with the TRANSID option in the same client logical unit of work. The distributed program link fails with an INVREQ response. In this case, EIBRESP2 is set to 15.
- Any function-shipped requests that follow a DPL request with the SYNCONRETURN option runs in a separate logical unit of work from the server logical unit of work.
- Any function-shipped requests running that follow a DPL request with the TRANSID option to the same server region runs under the transaction code specified on the TRANSID option, instead of under the default mirror transaction code. The function-shipped requests are committed as part of the overall client logical unit of work when the client commits.
- Any function-shipped requests running before or after a DPL request without the SYNCONRETURN or TRANSID options are committed as part of the overall client logical unit of work when the client commits.

See the *CICS Intercommunication Guide* for more information about function shipping.

Mixing DPL and DTP to the same CICS system

Care should be taken when using both DPL and DTP in the same application, particularly using DTP in the server program. For example, if you have not used the SYNCONRETURN option, you must avoid taking a syncpoint in the DTP partner which requires the DPL server program to syncpoint.




Figure 31. Example of mixing DPL and DTP

Restricting a program to the distributed program link subset

When a program executes as the result of a distributed program link, it is restricted to a subset of the full CICS API called the distributed program link subset. The commands that are prohibited in a server program are summarized in Table 11 on page 131.

You can specify, in the program resource definition only, that you want to restrict a program invoked by a local EXEC CICS LINK command to this subset with the EXECUTIONSET(DPLSUBSET) option. The use of any prohibited commands can then be detected before an application program is used in a distributed environment. The EXECUTIONSET(DPLSUBSET) option should be used for very early testing purposes only, and should never be used in production.

When the server program is running locally the following considerations apply:

- If EXECUTIONSET(DPLSUBSET) is specified on the server program then the SYNCONRETURN option causes an implicit syncpoint to be taken in the local server program, prior to returning control to the client program. In this case, because the server program is running locally, both the client and server resources are committed. However, it should be noted that SYNCONRETURN is intended for use when the client has no recoverable resources.
- If EXECUTIONSET(FULLAPI) is specified on the server program, the SYNCONRETURN option is ignored.
- The TRANSID and DATALENGTH options are ignored when processing the local link, but the format of the arguments is checked, for example, the TRANSID argument cannot be all blank.

Determining how a program was invoked

The 2-byte values returned on the STARTCODE option of the EXEC CICS ASSIGN command are extended in support of the distributed program link function enabling the server program to find out that it is restricted to the distributed program link subset. See the *CICS Application Programming Reference* manual for programming information about EXEC CICS commands.

Accessing user-related information with the EXEC CICS ASSIGN command

The values returned with the USERID and OPID keywords of the EXEC CICS ASSIGN command in the server program depend on the way the ATTACHSEC option is defined for the connection being used between the client CICS region and the server CICS region. For example, the system could be defined so that the server program could access the same USERID and OPID values as the client program or could access different values determined by the ATTACHSEC option.

If ATTACHSEC(LOCAL) is specified, the userid to which the OPID and USERID parameters correspond is one of the following, in the order shown:

1. The userid specified on the USERID parameter (for preset security) of the SESSIONS resource definition, if present
2. The userid specified on the SECURITYNAME parameter of the RDO CONNECTION resource definition, if present and no preset security userid is defined on the sessions
3. The userid specified on the DFLTUSER system initialization parameter of the server region, if neither the sessions nor connection definitions specify a userid

If any value other than LOCAL is specified for ATTACHSEC, the signed-on userid is the one received in the function management header (FMH5) from the client region.

See the *CICS Security Guide* for more information about link security and the ATTACHSEC parameter.

Another security-related consideration concerns the use of the CMDSEC and RESSEC options of the ASSIGN command. These are attributes of the transaction definition for the mirror transaction in the server region. They can be different from the definitions in the client region, even if the same TRANSID is used.

Exception conditions for the EXEC CICS LINK command

There are error conditions introduced in support of DPL.

Exception conditions returned to the client program: Condition codes returned to a client program describe such events as “remote system not known” or “failure to commit” in the server program. There are different reasons, identified by EIBRESP2 values, for raising the INVREQ and LENGERR conditions on an EXEC CICS LINK command. The ROLLEDBACK, SYSIDERR, and TERMERR conditions may also be raised. See the *CICS Application Programming Reference* manual for programming information about these commands.

The PGMIDERR condition is raised on the EXEC CICS HANDLE ABEND PROGRAM, EXEC CICS LOAD, EXEC CICS RELEASE, and EXEC CICS XCTL commands if the local program definition specifies that the program is remote. This exception is qualified by an EIBRESP2 value of 9.

Exception conditions returned to the server program: The INVREQ condition covers the use of prohibited API commands. INVREQ is returned, qualified by an EIBRESP2 value of 200, to a server program if it issues one of the prohibited commands summarized in Table 11 on page 131. If the server program does not handle the INVREQ condition, the default action is to abend the mirror transaction under which the server program is running with abend code ADPL.

For programming information about the DPL-related exception conditions, see the *CICS Application Programming Reference* manual.

<i>Table 11. API commands prohibited in programs invoked by DPL</i>	
Command	Options
ADDRESS	ACEE TCTUA
ALLOCATE	All
ASSIGN	ALTSCRNHT ALTSCRNWD APLKYBD APLTEXT BTRANS COLOR DEFSCRNHT DEFSCRNWD DELIMITER DESTCOUNT DESTID DESTIDLENG DS3270 DSSCS EWASUPP EXTDS FACILITY FCI GCHARS GCODES GMMI HILIGHT INPARTN KATAKANA LDCMNEM LDCNUM MAPCOLUMN MAPHEIGHT MAPLINE MAPWIDTH MSRCONTROL NATLANGINUSE NEXTTRANSID NUMTAB OPCLASS OPSECURITY OUTLINE PAGENUM PARTNPAGE PARTNS PARTNSET PS QNAME SCRNHT SCRNWD SIGDATA SOSI STATIONID TELLERID TERMCODE TERMPRIORITY TEXTKYBD TEXTPRINT UNATTEND USERNAME USERPRIORITY VALIDATION
CONNECT PROCESS	all
CONVERSE	all
EXTRACT ATTRIBUTES	all
EXTRACT PROCESS	all
FREE	all
HANDLE AID	all
ISSUE	ABEND ABORT ADD CONFIRMATION DISCONNECT END ENDOUTPUT EODS ERASE ERASEAUP ERROR LOAD NOTE PASS PREPARE PRINT QUERY RESET SEND SIGNAL RECEIVE REPLACE WAIT
LINK	INPUTMSG INPUTMSGLEN
PURGE MESSAGE	all
RECEIVE	all
RETURN	INPUTMSG INPUTMSGLEN
ROUTE	all
SEND	CONTROL MAP PARTNSET TEXT TEXT(MAPPED) TEXT(NOEDIT) PAGE
SIGNOFF	all
SIGNON	all
SYNCPOINT	Note: Can be issued in server region if SYNCONRETURN specified on EXEC CICS LINK command
WAIT CONVID	All
WAIT TERMINAL	all
XCTL	INPUTMSG INPUTMSGLEN

The following commands are also restricted but can be used in the server region if SYNCONRETURN is specified on EXEC CICS LINK:

- CALL DLI CHKP
- CALL DLI TERM
- CPIRR COMMIT

- CPIRR BACK
- EXEC DLI CHKP
- EXEC DLI TERM

Where only certain options are prohibited on the command, they are shown. All the APPC commands listed are prohibited only when they refer to the principal facility. One of these, the EXEC CICS CONNECT PROCESS command, causes an error even if it refers to the principal facility in a non-DPL environment. It is included here because, if an EXEC CICS CONNECT PROCESS command refers to its principal facility in a server program, the exception condition raised indicates a DPL error.

Asynchronous processing

The response from a remotely initiated transaction is not necessarily returned to the task that initiated the transaction, which is why the processing is referred to as **asynchronous**. Asynchronous processing is useful when you do not need or want to tie up local resources while having a remote request processed. For example, with online inquiry on remote databases, terminal operators can continue entering inquiries without having to wait for an answer to the first one.

You can start a transaction on a remote system using an EXEC CICS START command just like a local transaction. You can use the EXEC CICS RETRIEVE command to retrieve data that has been stored for a task as a result of a remotely issued EXEC CICS START, EXEC CICS CANCEL, EXEC CICS SEND, or EXEC CICS RECEIVE command, as if it were a local transaction.

Distributed transaction processing (DTP)

The main advantage of DTP is that it allows the two transactions to have exclusive control of a session and to “converse.” DTP is particularly useful when you need remote resources to be processed remotely or if you need to transfer data between systems. It also allows you to design very flexible and efficient applications. DTP can be used with either EXEC CICS or CPI Communications. You can use C and assembler language in DTP application programs that hold LU type 6.2 unmapped conversations using the EXEC CICS API as well as applications that use the CICS intercommunication facilities.

DTP can be used with a variety of partners, including both CICS and non-CICS platforms, as long as they support APPC. For further information about DTP, see the *CICS Family: Interproduct Communication* manual.

Common Programming Interface (CPI) Communications

CPI Communications provides an alternative API to existing CICS APPC support. CPI Communications provides DTP on APPC connections and can be used in COBOL, C, PL/I, and assembler language.

CPI Communications defines an API that can be used in APPC networks that include multiple system platforms, where the consistency of a common API is seen to be of benefit.

The CPI Communications interface can converse with applications on any system that provides an APPC API. This includes applications on CICS platforms. You may use EXEC CICS APPC API commands on one end of a conversation and CPI Communications commands on the other.

CPI Communications requires specific information (side information) to begin a conversation with partner program. CICS implementation of side information is achieved using the partner resource that your system programmer is responsible for maintaining.

The application's calls to the CPI Communications interface is resolved by link-editing it with the CICS CPI Communications stub (DFHCPLC). You can find information about how to do this in the *CICS System Definition Guide*.

The CPI Communications API is defined as a general call interface. The interface is described in *Common Programming Interface Communications Reference* manual.

External CICS interface (EXCI)

The external CICS interface is an application programming interface that enables a non-CICS program (a client program) running in VSE to call a program (a server program) running in a CICS TS for VSE/ESA region and to pass and receive data by means of a communications area. The CICS program is invoked as if linked-to by another CICS program.

This programming interface allows a user to allocate and open sessions (pipes) to a CICS system and to pass distributed program link (DPL) requests over them. CICS interregion communication (IRC) supports these requests and each pipe maps onto one MRO session.

For programming information about EXCI, see the *CICS External CICS Interface* manual.

A client program that uses the external CICS interface can operate multiple sessions for different users (either under the same or separate TCBs) all coexisting in the same VSE partition without knowledge of, or interference from, each other.

The external CICS interface provides two forms of programming interface:

- The EXCI CALL interface consists of six commands that allow you to:
 - Allocate and open sessions to a CICS system from non-CICS programs running under VSE/ESA.
 - Issue DPL requests on these sessions from the non-CICS programs.
 - Close and deallocate the sessions on completion of the DPL requests.
- The EXEC CICS interface provides:
 - A single composite command (LINK PROGRAM) that performs all six commands of the EXCI CALL interface in one invocation

The command takes the same form as the distributed program link command of the CICS command-level application programming interface.

Chapter 7. Recovery considerations

This chapter is about two techniques that are available to help recover or reconstruct events or data changes during CICS execution:

- “Journaling”
- “Syncpointing” on page 137

Journaling

CICS provides facilities for creating and managing **journals** during CICS processing. A journal is a set of special-purpose sequential data sets. Journals may contain any and all data the user needs to facilitate subsequent reconstruction of events or data changes. For example, a journal might act as an audit trail, a change-file of database updates and additions, or a record of transactions passing through the system (often referred to as a **log**). Each journal can be written from any task.

Journal control commands are provided to allow the application programmer to:

- Create a journal record (EXEC CICS WRITE JOURNALNUM command)
- Synchronize with (wait for completion of) journal output (EXEC CICS WAIT JOURNALNUM command)

Exception conditions that occur during execution of a journal control command are handled as described in Chapter 9, “Dealing with exception conditions” on page 141. (The earlier JFILEID option is supported for compatibility purposes only.)

Journal records

Data may be written to any journal specified in the journal control table (JCT), which defines the journals available during a particular CICS execution. The JCT may define one or more journals on direct access storage or magnetic tape. Each journal is identified by a number known as the journal identifier. This number may range from 2 through 99; the value 1 is reserved for a journal known as the system log.

When a journal record is built, the data is moved to the journal buffer area. All buffer space and other work areas needed for journal operations are acquired and managed by CICS. The user task supplies only the data to be written to the journal.

Journal records are built into blocks compatible with standard variable-blocked format. CICS uses the sequential access method of the host operating system to write the blocks to auxiliary storage.

Each journal record begins with a standard length field (LLbb), a user-specified identifier, and a system-supplied prefix. This data is followed in the journal record by any user-supplied prefix data (optional), and finally by the user-specified data. Journal control is designed so that the application programmer requesting output services does not have to be concerned with the detailed layout and precise

contents of journal records. The programmer has to know only which journal to use, what user data to specify, and which user-identifier to supply.

Journal output synchronization

When a synchronous journal record is created by issuing the WRITE JOURNALNUM command with the WAIT option, the requesting task can wait until the output has been completed. By specifying that this should happen, the application programmer ensures that the journal record is written on the external storage device associated with the journal before processing continues; the task is said to be **synchronized** with the output operation.

The application programmer can also request asynchronous journal output. This causes a journal record to be created in the journal buffer area and, optionally, initiates the data output operation from the buffer to the external device, but allows the requesting task to retain control and thus to continue with other processing. The task may check and wait for output completion (that is, synchronize) later by issuing the EXEC CICS WAIT JOURNALNUM command.

The basic process of building journal records in the buffer space of a given journal continues until one of the following events occurs:

- A request specifying the STARTIO option is made (from any task) from the EXEC CICS WAIT JOURNAL command, and is made for output of a journal record.
- A request is rejected because of insufficient journal buffer space.
- One second elapses after the last occasion on which any task issued an implied or explicit wait on records in this journal buffer.

When any one of these occurs, all journal records present in the buffer, including any deferred output resulting from asynchronous requests, are written to auxiliary storage as one block.

The advantages that may be gained by deferring journal output are:

- Transactions may get better response times by waiting less.
- The load of physical I/O requests on the host system may be reduced.
- Journal data sets may contain fewer but larger blocks for better utilization of auxiliary storage devices.

However, these advantages are achievable only at the cost of more buffer space and greater programming complexity. It is necessary to plan and program to control synchronizing with journal output. Additional decisions that depend on the data content of the journal record and how it is to be used must be made in the application program. In any case, the full benefit of deferring journal output is obtained only when the load on the journal is high.

The STARTIO option is used with EXEC CICS WRITE JOURNALNUM commands to specify that the journal output operation is to be initiated immediately. For asynchronous output requests, control returns directly to the requesting program. The STARTIO option should not be used unnecessarily because, if every journal request used STARTIO, no improvement over synchronous output requests, in terms of reducing the number of physical I/O operations and increasing the average block size, would be possible.

If the journal buffer space available at the time of the request is not sufficient to contain the journal record, the NOJBUFSP condition occurs. If no EXEC CICS HANDLE CONDITION command is active for this condition, the requesting task loses control, the contents of the current buffer are written out, and the journal record is built in the resulting freed buffer space before control returns to the requesting task.

If the requesting task is not willing to lose control (for example, if some housekeeping must be performed before other tasks get control), a HANDLE CONDITION command should be issued. If the NOJBUFSP condition occurs, no journal record is built for the request, and control is returned directly to the requesting program at the location provided in the HANDLE CONDITION command. The requesting program can perform any housekeeping needed before reissuing the journal output request.

Journal commands can cause immediate or deferred output to the journal. To distinguish between user journal record types and system journal record types, you must specify a journal identifier (JOURNALNUM) and a journal type identifier (JTYPEID). If the user journaling is to the system log, the journal type identifier (according to the setting of the high-order bit) also serves to cause recovery control to copy the records to the restart data set during its backward scan of the log as follows:

- For in-flight tasks only (high-order bit off)
- For all records encountered until the scan is terminated (high-order bit on)

See the *CICS Customization Guide* for programming information and details of the layout of journal records. See the section on emergency restart in the *CICS Recovery and Restart Guide* for background information and a description of the recovery process.

Syncpointing

To facilitate recovery in the event of abnormal termination of a CICS task or of failure of the CICS system, the system programmer can, during CICS table generation, define specific resources (for example, files) as recoverable. If a task is terminated abnormally, these resources are restored to the condition they were in at the start of the task, and can then be rerun. The process of restoring the resources associated with a task is termed **backout**.

If an individual task fails, backout is performed by the dynamic transaction backout program. If the CICS system fails, backout is performed as part of the emergency restart process. See the *CICS Recovery and Restart Guide*, which describes these facilities, which in general have no effect on the coding of application programs.

However, for long-running programs, it may be undesirable to have a large number of changes, accumulated over a period of time, exposed to the possibility of backout in the event of task or system failure. This possibility can be avoided by using the EXEC CICS SYNCPOINT command to split the program into logically separate sections known as units of work (UOWs); the end of an UOW is referred to as a synchronization point (**syncpoint**). For more information about syncpoints, see the *CICS Recovery and Restart Guide*.

If failure occurs after a syncpoint but before the task has been completed, only changes made after the syncpoint are backed out.

Alternatively, you can use the SAA Resource Recovery interface instead of the EXEC CICS SYNCPOINT command. This provides an alternative API to existing CICS resource recovery services. You may wish to use the SAA Resource Recovery interface in networks that include multiple SAA platforms, where the consistency of a common API is seen to be of benefit. In a CICS system, the SAA Resource Recovery interface provides the same function as the EXEC CICS API.³

The SAA Resource Recovery interface is implemented as a call interface, having two call types:

SRRCMIT Commit—Equivalent to SYNCPOINT command.

SRRBACK Backout—Equivalent to SYNCPOINT ROLLBACK command.

For further information about the SAA Resource Recovery interface, see the *SAA Common Programming Interface for Resource Recovery Reference* manual.

UOWs should be entirely logically independent, not merely with regard to protected resources, but also with regard to execution flow. Typically, a UOW comprises a complete conversational operation bounded by EXEC CICS SEND and EXEC CICS RECEIVE commands. A browse is another example of a UOW; an EXEC CICS ENDBR command must therefore precede the syncpoint.

In addition to a DL/I termination call being considered to be a syncpoint, the execution of an EXEC CICS SYNCPOINT command causes CICS to issue a DL/I termination call. If a DL/I PSB is required in a subsequent UOW, it must be rescheduled using a program control block (PCB) call or an EXEC CICS SCHEDULE command.

With distributed program link (DPL), it is possible to specify that a syncpoint is taken in the server program, to commit the server resources before returning control to the client. This is achieved by using the SYNCONRETURN option on the EXEC CICS LINK command. For programming information about the SYNCONRETURN option, see “Using the distributed program link function” on page 122 and the *CICS Application Programming Reference* manual.

A BMS logical message, started but not completed when an EXEC CICS SYNCPOINT command is processed, is forced to completion by an implied EXEC CICS SEND PAGE command. However, do not rely on this, because a logical message whose first page is incomplete is lost. Also code an explicit EXEC CICS SEND PAGE command before the EXEC CICS SYNCPOINT command or before termination of the transaction.

Consult your system programmer if syncpoints are to be issued in a transaction that is eligible for transaction restart.

³ Full SAA Resource Recovery provides some return codes that are not supported in its CICS implementation. (See the CICS appendix in the *SAA Common Programming Interface for Resource Recovery Reference* manual.)

Chapter 8. Minimizing errors

This chapter describes ways of making your applications error-free. Some of these suggestions apply not only to programming, but also to operations and systems.

What often happens is that, when two application systems that run perfectly by themselves are run together, performance goes down and you begin experiencing “lockouts” or waits. The scope of each system has not been defined well enough.

The key points in a well-designed application system are:

- At all levels, each function is defined clearly with inputs and outputs well-stated.
- Resources that the system uses are adequately-defined.
- Interactions with other systems are known.

Protecting CICS from application errors

There are various tools and techniques you can use to minimize errors in your application programs. In general:

- You can use the storage protection facility to prevent CICS code and control blocks from being overwritten accidentally by your application programs. You can choose whether you want to use this facility by means of CICS system initialization parameters. See Chapter 26, “Storage control and storage protection” on page 391 and the *CICS System Definition Guide* for more information about this facility.
- Consider using standards that avoid problems that may be caused by techniques such as the use of GETMAIN commands.

Testing applications

The following general rules apply to testing applications:

- Do not test on a production CICS system—use a test system, where you can isolate errors without affecting “live” databases.
- Have the testing done by someone other than the application developer, if possible.
- Document the data you use for testing.
- Test your applications several times. See Chapter 30, “Testing applications: the process” on page 415 for more information about testing applications.
- Use the CEDF transaction for initial testing. See “The CEBR commands” on page 444 for more information about using CEDF.
- Use stress or volume testing to catch problems that may not arise in a single-user environment. Teleprocessing Network Simulator (TPNS, licensed program number 5740-XT4) is a good tool for doing this.

TPNS is a telecommunications testing package that enables you to test and evaluate application programs before you install them. You can use TPNS for testing logic, user exit routines, message logging, data encryption, and device-dependencies, if these are used in application programs in your

organization. It is useful in investigating system performance and response times, stress testing, and evaluating TP network design. For further information, see the *TPNS General Information* manual.

- Test whether the application can handle **correct** data and **incorrect** data.
- Test against complete copies of the related databases.
- Consider using multiregion operation. (MRO is fully described in the *CICS Intercommunication Guide*.)
- Before you move an application to the production system, it is a good idea to run a **final** set of tests against a copy of the production database to catch any errors.

In particular, look for destroyed storage chains.

Assembler language programs (if not addressing data areas properly) can be harder to identify because they can alter something that affects (and abends) another transaction.

For more information about solving a problem, see the *CICS Problem Determination Guide*.

Chapter 9. Dealing with exception conditions

Every time you process an EXEC CICS command in one of your applications, CICS automatically raises a condition, or return code, to tell you what happened. You can choose to have this condition, which is usually NORMAL, passed back by the CICS EXEC interface program to your application. It is sometimes called a RESP value, because you may get hold of it by using the RESP option in your command. Alternatively, you may obtain this value by reading it from the EXEC interface block (EIB) field EIBRESP.

If something out of the ordinary happens, you get an **exception condition**, which simply means a condition other than NORMAL. By testing this condition, you can find out what has happened and, possibly, why.

Many exception conditions have an additional (RESP2) value associated with them, which gives further information. You may obtain this RESP2 value either by using the RESP2 option in your command in addition to the RESP option, or by reading it from the EIBRESP field.

Not all conditions denote an error situation, even if they are not NORMAL. For example, if you get an ENDFILE condition on an EXEC CICS READNEXT command during a file browse, it might be exactly what you expect. For information about all possible conditions and the commands on which they can occur, see the *CICS Application Programming Reference* manual.

The rest of this chapter covers the following topics:

- “Default CICS exception handling”
- “Handling exception conditions by in-line code” on page 142
- “Modifying the default CICS exception handling” on page 146

Default CICS exception handling

If your application is written in a language other than C, and you do not specify otherwise, CICS uses its built-in exception handling whenever an exception condition occurs. If your application is written in C, CICS itself takes no action when an exception condition occurs and it is left to the application to handle it. See “Handling exception conditions by in-line code” on page 142 for information on coping with exception conditions.

The most common action by CICS is to cause an abend of some type to happen. The particular behaviors for each condition and for each command are detailed in the *CICS Application Programming Reference* manual and the *CICS System Programming Reference* manual.

Sometimes you will be satisfied with the CICS default exception handling, in which case you need do nothing. More often you will prefer some other course of action.

These are the different ways of turning off the default CICS handling of exception conditions.

- Turn off the default CICS handling of exception conditions on a particular EXEC CICS command call by specifying the NOHANDLE option.

- Alternatively, turn off the default CICS handling of exception conditions by specifying the RESP option on the command. This, of itself, switches off the default CICS exception handling in the same way as NOHANDLE does. It also causes the variable named by the argument of RESP to be updated with the value of the condition returned by the command. This is described in more detail in “Handling exception conditions by in-line code.”
- Write your application program in C.

If the default CICS exception handling is turned off, ensure that your program copes with anything that may happen in the command call.

The traditional, but no longer recommended, way to specify some other course of action is available only if you are programming in a language other than C: it is to use combinations of the EXEC CICS HANDLE ABEND, EXEC CICS HANDLE CONDITION, and EXEC CICS IGNORE CONDITION commands to modify the default CICS exception handling. This is described in “Modifying the default CICS exception handling” on page 146.

Handling exception conditions by in-line code

This section describes the method of handling exception conditions which is recommended for new applications and is the only available choice if your programs are in C language. If your program is not written in C, it involves either using the NOHANDLE option or specifying the RESP option on EXEC CICS commands, which prevents CICS performing its default exception handling. Additionally, the RESP option makes the value of the exception condition directly available to your program, for it to take remedial action.

If your program is written in C, in-line code is the only means you have of handling exception conditions.

If you use the NOHANDLE or RESP options, ensure that your program can cope with whatever condition may arise in the course of executing the commands. The RESP value is available to enable your program to decide what to do, and more information that it may need to use is carried in the EXEC interface block (EIB). In particular, the RESP2 value is contained in one of the fields of the EIB. See the *CICS Application Programming Reference* manual for more information on the EIB. Alternatively, if your program specifies RESP2 in the command, the RESP2 value will be returned by CICS directly.

The DFHRESP built-in translator function makes it very easy to test the RESP value. It allows, you to examine RESP values symbolically. This is easier than examining binary values that are less meaningful to someone reading the code.

How to use the RESP and RESP2 options

The argument of RESP is a user-defined fullword binary data area (long integer). On return from the command, it contains a value corresponding to the condition that may have been raised. Normally its value is DFHRESP(NORMAL).

Use of RESP and DFHRESP in COBOL and PL/I

Here is an example of an EXEC CICS call in COBOL which uses the RESP option. A PL/I example would be similar, but would end in “;” instead of END-EXEC.

```
EXEC CICS WRITEQ TS FROM(abc)
        QUEUE(qname)
        NOSUSPEND
        RESP(xxx)
END-EXEC.
```

An example of using DFHRESP to check the RESP value is:

```
IF xxx=DFHRESP(NOSPACE) THEN ...
```

Use of RESP and DFHRESP in C

Here is an example of an EXEC CICS call in C which uses the RESP option, including the declaration of the RESP variable:

```
long response;
:
EXEC CICS WRITEQ TS FROM(abc)
        QUEUE(qname)
        NOSUSPEND
        RESP(response);
```

An example of using DFHRESP to check the RESP value is:

```
if (response == DFHRESP(NOSPACE))
{
:
}
```

Use of DFHRESP in assembler

An example of a test for the RESP value in assembler language is:

```
CLC   xxx,DFHRESP(NOSPACE)
BE    ...
```

An example of exception handling in C

The following example is a typical function that could be used to receive a BMS map and to cope with exception conditions:

```

int ReadAccountMap(char *mapname, void *map)
{
    long    response;
    int     ExitKey;

    EXEC CICS RECEIVE MAP(mapname)
           MAPSET("ACCOUNT")
           INTO(map)
           RESP(response);
    switch (response)
    {
    case DFHRESP(NORMAL):
        ExitKey = dfheiptr->eibaid;
        ModifyMap(map);
        break;

    case DFHRESP(MAPFAIL):
        ExitKey = dfheiptr->eibaid;
        break;

    default:
        ExitKey = DFHCLEAR;
        break;
    }

    return ExitKey;
}

```

Figure 32. An example of exception handling in C

The **ReadAccountMap** function has two arguments:

1. **mapname** is the variable that contains the name of the map that is to be received.
2. **map** is the address of the area in memory to which the map is to be written.

The RESP value will be returned in **response**. The declaration of **response** sets up the appropriate type of automatic variable.

The EXEC CICS statement asks for a map of the name given by **mapname**, of the mapset ACCOUNT, to be read into the area of memory to which the variable **map** points, with the value of the condition being held by the variable **response**.

The condition handling can be done by using if statements. However, to improve readability, it is often better, as here, to use a switch statement, instead of compound if ... else statements. The effect on program execution time is negligible.

Specific cases for two conditions:

1. A condition of NORMAL is what is normally expected. If a condition of NORMAL is detected in the example here, the function then finds out what key the user pressed to return to CICS and this value is passed to ExitKey. The program then makes some update to the map held in memory via the ModifyMap function, which need not concern us further.

2. A condition of MAPFAIL, signifying that the user has made no updates to the screen, is also fairly normal and is specifically dealt with here. In this case the program again updates ExitKey but does not call ModifyMap.

In this example, any other condition is held to be an error. The example sets ExitKey to DFHCLEAR—the same value that it would have set if the user had cleared the screen—which it then returns to the calling program. By checking the return code from ReadAccountMap, the calling program would know that the map had not been updated and that some remedial action is required.

An example of exception handling in COBOL

The following example is a typical function that could be used to receive a BMS map and to cope with exception conditions:

```
03  RESPONSE                PIC S9(8)  BINARY.
03  EXITKEY                  PIC X(4)   COMP-3.
:
EXEC CICS RECEIVE MAP(MAPNAME)
      MAPSET('ACCOUNT')
      INTO(MAP)
      RESP(RESPONSE)
      END-EXEC.

IF (RESPONSE NOT = DFHRESP(NORMAL)) AND
   (RESPONSE NOT = DFHRESP(MAPFAIL))
   MOVE DFHCLEAR TO EXITKEY
ELSE
   MOVE EIBAID TO EXITKEY
   IF RESPONSE = DFHRESP(NORMAL)
     GO TO MODIFYMAP.
   END-IF.
:
MODIFYMAP.
:
```

Figure 33. An example of exception handling in COBOL

MAPNAME is the variable that contains the name of the map that is to be received.

The RESP value will be returned in RESPONSE. RESPONSE is declared as a fullword binary variable in the data section.

The EXEC CICS statement asks for a map of the name given by **MAPNAME**, of the mapset ACCOUNT, to be read, with the value of the condition being held by the variable **RESPONSE**.

The condition handling is done by using IF ... statements. If the condition is neither NORMAL nor MAPFAIL the program behaves as if the user had cleared the screen.

If the condition is either NORMAL or MAPFAIL the program saves the value of the key that the user pressed to exit the screen in EXITKEY. In addition, if the condition is NORMAL, the program branches to MODIFYMAP to perform some additional function.

Modifying the default CICS exception handling

Section not applicable to C programs

CICS provides the following EXEC CICS commands that modify the default CICS exception handling and one that modifies the way CICS handles abends:

HANDLE CONDITION

Specify the label to which control is to be passed if a condition occurs.

IGNORE CONDITION

Specify that no action is to be taken if a condition occurs.

HANDLE ABEND

Activate, cancel, or reactivate an exit for abnormal termination processing.

An abend is the commonest way in which CICS handles exception conditions.

The current effect of IGNORE CONDITION, HANDLE ABEND and HANDLE CONDITION may be suspended by using PUSH HANDLE and reinstated by using POP HANDLE.

All the commands mentioned above are described in the *CICS Application Programming Reference* manual. You have two ways of passing control to a specified label:

1. Use a HANDLE CONDITION condition(label) command, where condition is the name of an exception condition
2. Use a HANDLE CONDITION ERROR(label) command

The EXEC CICS HANDLE CONDITION command sets up some CICS code to name conditions that interest you, and then uses this code to pass control to appropriate sections of your application if those conditions arise. So with an active HANDLE CONDITION command, control goes to whichever label you specified for that particular condition.

The **same** condition can arise, in some cases, on many different commands, and for a variety of reasons. For example, you can get an IOERR condition during file control operations, interval control operations, and others. One of your first tasks, therefore, is to sort out **which command** has raised a particular condition; only when you have discovered that, can you begin to investigate why it has happened. This, for many programmers, is reason enough to start using the RESP option in their new CICS applications. Although you need only one HANDLE CONDITION command to set your error-handling for several conditions, it can sometimes be awkward to pinpoint exactly which of several HANDLE CONDITION commands is currently active when a CICS command fails somewhere in your code.

If a condition that you have not named arises, CICS takes the default action, unless this is to abend the task, in which case it raises the ERROR condition. If you name the condition but leave out its label, any EXEC CICS HANDLE CONDITION command for that condition is deactivated, and CICS reverts to taking the default action for it, if and when it occurs.

The need to deal with **all** conditions is a common source of errors when using the EXEC CICS HANDLE CONDITION command. When using an unfamiliar command, read the *CICS Application Programming Reference* manual to find out which exception conditions are possible. Even if you then issue HANDLE commands for all of these, you may not finish all the error-handling code adequately. The outcome is sometimes an error-handling routine that, by issuing a RETURN command, allows incomplete or incorrect data changes to be committed.

The best approach is to use the EXEC CICS HANDLE CONDITION command, but to let the system default action take over if you cannot see an obvious way round a particular problem.

Bearing in mind the distinction between an error condition, a condition that merely causes a wait (see page 151 for examples of conditions that cause a wait), and the special case of EXEC CICS SEND MAP command overflow processing (see the *CICS Application Programming Reference* manual), an EXEC CICS HANDLE CONDITION command is active after an EXEC CICS HANDLE CONDITION condition(label), or EXEC CICS HANDLE CONDITION ERROR(label) command has been run in your application.

If no EXEC CICS HANDLE CONDITION command is active for a condition, but one is active for ERROR, control passes to the label for ERROR, if the condition is an error, not a wait.

If you use EXEC CICS HANDLE CONDITION commands, or are maintaining an application that uses them, do not include any commands in your error routine that can cause the same condition that gave you the original branch to the routine, because you will cause a loop.

Take special care not to cause a loop on the ERROR condition itself. You can avoid a loop by reverting temporarily to the system default action for the ERROR condition. Do this by coding an EXEC CICS HANDLE CONDITION ERROR command with no label specified. At the end of your error processing routine, you can reinstate your error action by including a HANDLE CONDITION ERROR command with the appropriate label. If you know the previous HANDLE CONDITION state, you can do this explicitly. In a general subroutine, which might be called from several different points in your code, the PUSH HANDLE and POP HANDLE command may be useful—see PUSH HANDLE and POP HANDLE on page 152.

Use of the EXEC CICS HANDLE CONDITION command

Use the EXEC CICS HANDLE CONDITION command to specify the label to which control is to be passed if a condition occurs. You must include the name of the condition and you must ensure that the EXEC CICS HANDLE CONDITION command is executed before the command that may give rise to the associated condition.

You cannot include more than 16 conditions in the same command. You must specify any additional conditions in further EXEC CICS HANDLE CONDITION commands. You can also use the ERROR condition within the same list to specify that all other conditions are to cause control to be passed to the same label.

The EXEC CICS HANDLE CONDITION command for a given condition applies only to the program in which it is specified. The EXEC CICS HANDLE CONDITION command:

- Remains active while the program is running, or until:
 - An EXEC CICS IGNORE CONDITION command for the same condition is met, in which case the EXEC CICS HANDLE CONDITION command is overridden
 - Another EXEC CICS HANDLE CONDITION command for the same condition is met, in which case the new command overrides the previous one
- Is temporarily deactivated by the NOHANDLE or RESP option on a command

When control passes to another program, via an EXEC CICS LINK or EXEC CICS XCTL command, the EXEC CICS HANDLE CONDITION commands that were active in the calling program are deactivated. When control returns to a program from a program at a lower logical level, the EXEC CICS HANDLE CONDITION commands that were active in the higher-level program before control was transferred from it are reactivated, and those in the lower-level program are deactivated. (See Chapter 25, “Program control” on page 377 for information about logical levels.)

The following example shows you how to handle conditions, such as DUPREC, LENGERR, and so on, that can occur when you use an EXEC CICS WRITE command to add a record to a data set. Suppose that you want DUPREC to be handled as a special case; that you want standard system action (that is, to terminate the task abnormally) to be taken for LENGERR; and that you want all other conditions to be handled by the error routine ERRHANDL. You would code:

```
EXEC CICS HANDLE CONDITION
      ERROR(ERRHANDL)
      DUPREC(DUPRTN) LENGERR
END-EXEC.
```

In a PL/I application program, a branch to a label in an inactive procedure or in an inactive begin block, caused by a condition, produces unpredictable results.

In an assembler language application program, if a HANDLE condition and the command that caused the condition are at the same logical level, the registers are restored to their values in the application program at the point where the command that caused the condition was issued. However, when the command that causes a condition occurs at a lower logical level, the registers are restored to the values saved in DFHEISTG when control is passed from the HANDLE CONDITION level.

Use of the EXEC CICS HANDLE CONDITION ERROR command

Figure 34 on page 149 shows the first of only two EXEC CICS HANDLE CONDITION commands used in program ACCT01:

```

PROCEDURE DIVISION.
*
*   INITIALIZE.
*   TRAP ANY UNEXPECTED ERRORS.
EXEC CICS HANDLE CONDITION
      ERROR(OTHER-ERRORS)
      END-EXEC.
*

```

Figure 34. Trapping the unexpected with the HANDLE CONDITION ERROR command

ACCT01 passes control to the paragraph at label OTHER-ERRORS if any condition arises for a command that does not specify NOHANDLE or RESP.

The HANDLE CONDITION ERROR command is the first command executed in the procedure division of this COBOL program. This is because a HANDLE CONDITION command must be processed before any CICS command is processed that can raise the condition being handled. Note, however, that your program does not see the effects when it processes the HANDLE CONDITION command; it only sees them later, if and when it issues a CICS command that actually raises one of the named conditions.

In this, and the other ACCT programs, you generally use the RESP option. All the commands specifying the RESP option have been written with a “catch-all” test (IF RESPONSE NOT = DFHRESP(NORMAL) GO TO OTHER-ERRORS) **after** any explicit tests for specific conditions. So any exceptions, other than those you might particularly “expect”, take control to the paragraph at OTHER-ERRORS in each program. Those relatively few commands that do not have RESP on them take control to exactly the same place if they result in any condition other than NORMAL because of this HANDLE CONDITION ERROR command.

Use of the EXEC CICS IGNORE CONDITION command

Just as you can arrange for control to pass to a particular label for a specific condition with a HANDLE CONDITION command, so you can have the program continue when a specific condition occurs. You do this by setting up an IGNORE CONDITION command to ignore one or more of the conditions that can potentially arise on a command. The IGNORE CONDITION command means that no action is to be taken if a condition occurs; control returns to the instruction following the command and return codes are set in the EIB. The following example ignores the MAPFAIL condition:

```

EXEC CICS IGNORE CONDITION MAPFAIL
      END-EXEC.

```

An IGNORE CONDITION command for a given condition applies only to the program you put it in, and it remains active while the program is running, or until a later HANDLE CONDITION command naming the same condition is met, in which case the IGNORE CONDITION command is overridden.

You can choose an IGNORE CONDITION command if you have a program reading records that are sometimes longer than the space you provided, but you do not consider this an error and do not want anything done about it. You might, therefore, code IGNORE CONDITION LENGERR before issuing READ commands.

You can also use an IGNORE CONDITION ERROR command to catch any condition considered as an error for which there is no currently active HANDLE CONDITION command that includes a label. When an error occurs, control is passed to the next statement and it is up to the program to check for return codes in the EIB. See page 151 for examples of conditions that are **not** considered as errors.

You cannot code more than 16 conditions in the same command. You must specify additional conditions in further IGNORE CONDITION commands.

Use of the EXEC CICS HANDLE ABEND command

Note to C programmers

Although HANDLE ABEND is supported in C when used with the PROGRAM option, it is not helpful in the context of this chapter because exception conditions in C programs do not cause abends.

The EXEC CICS HANDLE ABEND command activates or reactivates a program-level abend exit within your application program; you can also use this command to cancel a previously activated exit. For more information, see the *CICS Application Programming Reference* manual.

The EXEC CICS HANDLE ABEND command lets you supply your own code to be executed when an abend is processed. This means that your application can cope with the abnormal situation in an orderly manner and carry on executing. You provide the user exit programs and rely on CICS calling them when required.

The flow of control during abend processing is shown in Figure 35 on page 158.

EIBRESP and NOHANDLE options

You can temporarily deactivate the effect of any EXEC CICS HANDLE CONDITION command by using the EIBRESP or NOHANDLE option on a command. The way to use these options is described in “Handling exception conditions by in-line code” on page 142. If you do this, you lose the ability to use any system default action for that command. In other words, you have to do your own “catch-all” error processing.

You can also switch from ignoring a condition to handling it, or to using the system default action. For example, you could code:

```
*   MIXED ERROR PROCESSING
    EXEC CICS IGNORE CONDITION LENGERR
    END-EXEC.
:
    EXEC CICS HANDLE CONDITION DUPREC(DUPRTN)
    LENGERR
    ERROR(ERRHANDL)
    END-EXEC.
```

Because this code initially ignores condition LENGERR, nothing happens if the program raises a LENGERR condition; the application simply continues its processing. Of course, if the fact that LENGERR has arisen means that the application cannot sensibly continue, you have a problem.

Later in the code, you can explicitly set condition LENGERR to the system default action by naming it in an EXEC CICS HANDLE CONDITION command without a label. When this command has been executed, the program no longer ignores condition LENGERR, and if it subsequently occurs, it now causes the system default action. The point about mixing methods is that you can, and that each condition is treated separately.

How CICS keeps track of what to do

CICS has a table of the conditions referred to by HANDLE CONDITION and IGNORE CONDITION commands in your application. Each execution of one of these commands either updates an existing entry in this table, or causes CICS to make a new entry if this is the first time the condition has been quoted in such a command. Each entry tells CICS what to do by indicating one of the three exception-handling states your application can be in, namely:

1. **Let the program continue**, with control coming straight back from CICS to the next instruction following the command that has failed in your program. You can then find out what happened by testing, for example, the RESP value that CICS returns after executing a command. The result of this test enables you to decide what to do next. For details, see “Handling exception conditions by in-line code” on page 142.

This is the recommended method, which is the approach taken in the “File A” sample programs referred to in the *CICS Sample Applications Guide*. It is also the recommended approach for any new CICS applications. It lends itself to structured code and removes the need for implied GOTOs that CICS required in the past.

2. **Pass control to a specified label** if a named condition arises. You do this by using an EXEC CICS HANDLE CONDITION command or EXEC CICS HANDLE CONDITION ERROR command to name both the condition and the label of a routine in your code to deal with it. For details, see “Use of the EXEC CICS HANDLE CONDITION command” on page 147 and “Use of the EXEC CICS HANDLE CONDITION ERROR command” on page 148.
3. **Taking the CICS system default action**, where for most conditions, this is to terminate the task abnormally and means that you do nothing by way of testing or handling conditions.

For the conditions ENQBUSY, NOJBUFSP, NOSTG, QBUSY, SESSBUSY, and SYSBUSY, the normal default is to force the task to wait until the required resource (for example, storage) becomes available, and then resume processing the command. You can change this behavior to ignoring the condition by using the NOSUSPEND or NOQUEUE options. For the NOSPACE condition, the normal default is to wait if processing an EXEC CICS WRITEQ TS command, but to abend the task if processing a WRITEQ TD, WRITE, or REWRITE command. Coding the EXEC CICS WRITEQ TS command with the NOSUSPEND option makes it ignore any NOSPACE condition that arises. For further information on coding this command, see the *CICS Application Programming Reference* manual.

CICS keeps a table of these conditions for each link level. Essentially, therefore, each program level has its own HANDLE state table governing its own condition handling.

This behavior is modified by two more EXEC CICS commands:

EXEC CICS PUSH HANDLE and EXEC CICS POP HANDLE commands

The EXEC CICS PUSH HANDLE command suspends the current effect of EXEC CICS HANDLE CONDITION, IGNORE CONDITION, HANDLE ABEND and HANDLE AID commands.

The EXEC CICS POP HANDLE command reinstates the effect of EXEC CICS HANDLE CONDITION, IGNORE CONDITION, HANDLE ABEND and HANDLE AID commands to what they were before the previous PUSH HANDLE was called.

CICS also keeps a table of conditions for each EXEC CICS PUSH HANDLE command that has not been countermanded by a matching EXEC CICS POP HANDLE command. When each condition occurs, CICS performs the following sequence of tests:

1. If the command has the RESP or NOHANDLE option, control returns to the next instruction in your application program. Otherwise, CICS scans the condition table to see what to do.
2. If an entry for the condition exists, this determines the action.
3. If no entry exists and the default action for this condition is to suspend execution:
 - a. If the command has the NOSUSPEND or NOQUEUE option, control returns to the next instruction.
 - b. If the command does not have one of these options, the task is suspended.
4. If no entry exists and the default action for this condition is to abend, a second search is made, this time for the ERROR condition:
 - a. If found, this entry determines the action.
 - b. If ERROR cannot be found, the task is abended. You can choose to handle abends. For information about the HANDLE ABEND command, see the *CICS Application Programming Reference* manual.

Note: The OVERFLOW condition on a SEND MAP command is an exception to the above rules. See the *CICS Application Programming Reference* manual for more information.

The EXEC CICS ALLOCATE, ENQ, GETMAIN, WRITE JOURNALNUM, READQ TD, and WRITEQ TS commands can all raise conditions for which the default action is to suspend your application program until the specified resource becomes available. So, on these commands, you have the NOSUSPEND option to inhibit this waiting and return immediately to the next instruction in your application program.

Some conditions can occur during the execution of a number of unrelated commands. If you want the same action for all occurrences, code a single EXEC CICS HANDLE CONDITION command at the start of your program.

Note: Because using RESP implies NOHANDLE, be careful when using RESP with the RECEIVE command, because it overrides the HANDLE AID command as well as the HANDLE CONDITION command. This means that PF key responses are ignored, and is the reason for testing them earlier in the ACCT code. See "The HANDLE AID command" on page 261.

_____ End of Section not applicable to C programs _____

Chapter 10. Access to system information

You can write many application programs using the CICS command-level interface without any knowledge of, or reference to, the fields in the CICS control blocks and storage areas. However, you might need to get information that is valid outside the local environment of your application program. You use the EXEC CICS ADDRESS and EXEC CICS ASSIGN commands to access such information. For programming information about these commands, see the *CICS Application Programming Reference* manual.

When you use the EXEC CICS ADDRESS and EXEC CICS ASSIGN commands, various fields can be read but should not be set or used in any other way. This means that you must not use any of the CICS fields as arguments in CICS commands, because these fields may be altered by the EXEC interface modules.

System programming commands

The EXEC CICS INQUIRE, SET, and PERFORM commands allow application programs to access information about CICS resources. The application program can retrieve and modify information for CICS data sets, terminals, system entries, mode names, system attributes, programs, and transactions. These commands are primarily for the use of the system programmer. For programming information, see the *CICS System Programming Reference* manual.

EXEC interface block (EIB)

In addition to the usual CICS control blocks, each task in a command-level environment has a control block known as the EXEC interface block (EIB) associated with it. An application program can access all of the fields in the EIB by name. The EIB contains information that is useful during the execution of an application program, such as the transaction identifier, the time and date (initially when the task is started, and subsequently, if updated by the application program using ASKTIME), and the cursor position on a display device. The EIB also contains information that is helpful when a dump is used to debug a program. For programming information about EIB fields, see the *CICS Application Programming Reference* manual.

Chapter 11. Abnormal termination recovery

CICS provides a program-level abend exit facility so that you can write exits of your own that can receive control during abnormal termination of a task. The “cleanup” of a program that has started but not completed normally is an example of a function performed by such an abend exit.

Here are some causes of abnormal terminations:

- A user request by, for example:
`EXEC CICS ABEND ABCODE(...)`
- A CICS request as a result of an invalid user request. For example, an invalid `EXEC CICS FREEMAIN` request gives the transaction abend code `ASCF`.
- A program check, in which case the system recovery program (DFHSRP) is driven, and the task abends with code `ASRA`.
- An operating system abend, in which case DFHSRP is driven, and the task abends with code `ASRB`.
- A looping task, in which case DFHSRP is driven, and the task abends with code `AICA`.

See the *CICS Problem Determination Guide* for full details about identifying and fixing problems. See the *VSE/ESA Messages and Codes Volume 3* for information about the transaction abend codes for abnormal terminations that are initiated by CICS, their meanings, and your responses.

The rest of this chapter deals with the following topics:

- “The `EXEC CICS HANDLE ABEND` command”
- “Restrictions on retrying operations” on page 157
- “Exception handling” on page 158
- “Trace” on page 161
- “Monitoring” on page 162
- “Dumping” on page 163

The `EXEC CICS HANDLE ABEND` command

The `EXEC CICS HANDLE ABEND` command activates or reactivates a program-level abend exit within your application program; you can also use this command to cancel a previously activated exit.

When activating an exit, you must use the `PROGRAM` option to specify the name of a program to receive control, or (except for C and PL/I programs) the `LABEL` option to specify a routine label to which control branches when an abnormal termination condition occurs. Using an `ON ERROR` block in PL/I is the equivalent of using the `EXEC CICS HANDLE ABEND LABEL` command.

An `EXEC CICS HANDLE ABEND` command overrides any preceding such command in any application program at the same logical level. Each application program of a transaction can have its own abend exit, but only one abend exit at each logical level can be active. (Logical levels are explained in Chapter 25, “Program control” on page 377.)

When a task terminates abnormally, CICS searches for an active abend exit, starting at the logical level of the application program in which the abend occurred, and proceeding to successively higher levels. The first active abend exit found, if any, is given control. This procedure is shown in Figure 35 on page 158, which also shows how subsequent abend processing is determined by the user-written abend exit.

If CICS finds no abend exit, it passes control to the abnormal condition program to terminate the task abnormally. This program invokes the user replaceable program error program, DFHPEP. See the *CICS Customization Guide* for programming information about how to customize DFHPEP.

CICS deactivates the exit upon entry to the exit routine or program to prevent recursive abends in an abend exit. If you wish to retry the operation, you can branch to a point in the program that was in control at the time of the abend and issue an EXEC CICS HANDLE ABEND RESET command to reactivate the abend exit. You can also use this command to reactivate an abend exit (at the logical level of the issuing program) that was canceled previously by an EXEC CICS HANDLE ABEND CANCEL command. You can suspend the EXEC CICS HANDLE ABEND command by means of the EXEC CICS PUSH HANDLE and EXEC CICS POP HANDLE commands as described in PUSH HANDLE and POP HANDLE on page 152.

Note that when an abend is handled, the dynamic transaction backout program is not be invoked. If you need the dynamic transaction backout program, you take an implicit or explicit syncpoint or issue an EXEC CICS SYNCPOINT ROLLBACK or issue an EXEC CICS ABEND command.

Where the abend is the result of a failure in a transaction running in an IRC-connected system, for example AZI2, the syncpoint processing may abend ASP1 if it attempts to use the same IRC connection during its backout processing.

The EXEC CICS HANDLE ABEND command cannot intercept ASPx or APSJ abend codes.

If a program running under LE abends as a result of LE control blocks being corrupted, any CICS abend handler that is active for that rununit will not be driven.

Creating a program-level abend exit

You can either define abend exits by using RDO, EXEC CICS CREATE commands, or by using the program autoinstall exit. If you use the autoinstall method, the program definition is not available at the time of the HANDLE ABEND. This may mean that a program functions differently the first time it is invoked. If the program is not defined at the time the HANDLE ABEND is issued, and program autoinstall is active, the security check on the name of the program is the only one that takes place. Other checks occur at the time the abend program is invoked. If the autoinstall fails, the task abends APCT and control is passed to the next higher level.

Abend exit programs can be coded in any supported language, but abend exit routines must be coded in the same language as their program.

For abend exit routines, the addressing mode and execution key are set to the addressing mode and execution key in which the EXEC CICS HANDLE ABEND command has been issued.

Upon entry to an abend exit program, no addressability can be assumed other than that normally assumed for any application program coded in that language. There are no register values for C or PL/I languages because these languages do not support the EXEC CICS HANDLE ABEND LABEL command.

Upon entry to an abend exit routine, the register values are:

COBOL

Control returns to the EXEC CICS HANDLE ABEND command with the registers restored; a COBOL GOTO is then executed

Assembler

Reg 15 Abend label.

Reg 0-14 Contents at the time of the last CICS service request.

There are three means of terminating processing in an abend exit routine or program, as listed below. It is recommended that when abend routines and programs are called by CICS internal logic, they should terminate with an abend because further processing is likely to cause more problems.

- Using an EXEC CICS RETURN command to indicate that the task is to continue running with control passed to the program on the next higher logical level. If no such program exists, the task is terminated normally, and any recoverable resources are committed.
- Using an EXEC CICS ABEND command to indicate that the task is to be abnormally terminated with control passed either to an abend exit specified for a program on a higher logical level or, if there is not one, to the abnormal condition program for abnormal termination processing.

In the case of an abend caused by a timeout on an outstanding EXEC CICS RECEIVE command, it is important to let the CICS abend continue, so that CICS can cancel the RECEIVE.

The EXEC CICS ABEND command can also be used to request a dump.

- Branching to retry an operation. When you are using this method of retrying an operation, and you want to reenter the original abend exit routine or program if a second failure occurs, the abend exit routine or program should issue the HANDLE ABEND RESET command before branching. This is because CICS has disabled the exit routine or program to prevent it reentering the abend exit.

Restrictions on retrying operations

If an abend occurs during the invocation of a CICS service, be aware that issuing a further request for **the same service** may cause unpredictable results, because the reinitialization of pointers and work areas, and the freeing of storage areas in the exit routine, may not have been completed.

Do not try to recover from ATNI or ATND abends by attempting further I/O operations. Either of these abends results in a TERMERR condition, requiring the session to be terminated in all cases.

If intersystem communication is being used, an abend in the remote system may cause a branch to the specified program or label, but subsequent requests to use resources in the remote system fails.

If an abend occurs as a result of a BMS command, control blocks are not tidied up before control is returned to the BMS program, and results are unpredictable if the command is retried.




Figure 35. ABEND exit processing

Exception handling

Exception handling under CICS is similar to a C environment for `signal` and `raise` behavior. But the abend codes, the error messages, and the run-time option behavior are slightly different.

The C exception handler handles only program exceptions such as operation and protection; all other abends are handled by the CICS exception handler. Thus an `EXEC CICS ABEND ABCODE(...)` statement would not be processed by the C exception handler.

If the LE/VSE TRAP(OFF) run-time option has been specified, the C exception handler does not process any program exceptions that occur in the program; CICS handles the exception according to its exception handler. See the *LE/VSE Installation and Customization Guide* for more information about run-time options.

If the CICS command EXEC CICS HANDLE ABEND PROGRAM(*name*) was specified in the application, the specified program (*name*) receives control if a program exception occurs. This program always takes precedence over the C exception handler.

For example, if an EXEC CICS HANDLE ABEND PROGRAM(...) **and** a C user handler registered through a signal were specified in the application, the specified program (*name*) receives control whenever a program exception occurs. However, if the raise() function was specified, the corresponding registered C handler would receive control.

When a C application is invoked by an EXEC CICS LINK PROGRAM(...), this invoked program inherits any handlers registered with EXEC CICS HANDLE ABEND PROGRAM(...) in the invoking program. Any handlers registered in the invoked program override the inherited handlers. C signal handlers, however, are not inherited. If the invoked program abends, and no signal handler was specified in that program, the default action takes place.

The following code fragments show exception handling in a CICS-C environment. Figure 36 shows exception handling using only C functions. Figure 37 on page 160 shows exception handling using both C functions and CICS commands.

```
#pragma runopts(SPIE,STAE) 1
#include <signal.h>
void handle();
int main(void) {
    int a;
    .
    .
    .
    signal(SIGSEGV,handle); 2
    .
    .
    .
    <protection exception> 3
    a++; 4
    .
    .
    .
}
void handle() {
    .
    .
    .
    return;
}
```

Figure 36. Exception handling using only C functions

Notes:

1. The run-time options SPIE and STAE are the defaults indicating that the C exception handler should be invoked to handle an abend.
2. This `signal` function registers a user handler to be invoked in the event of a SIGSEGV abend.
3. A protection exception has occurred and control is returned to CICS. The C-CICS interface is invoked which, in turn, eventually invokes the application's handler, the `handle()` function.
4. After execution of the handler, control is returned to the statement after the abend.

```
#pragma runopts(SPIE,STAE)
#include <signal.h>
int main(void) {
    /*
     *
     */
    EXEC CICS HANDLE ABEND PROGRAM("PROGRAMB"); 1
    /*
     *
     */
    signal(SIGSEGV,handle); 2
    .
    .
    .
    raise(SIGSEGV); 3
    .
    .
    .
    /* <protection exception> */ 4
    .
    .
    .
}
void handle() {
    .
    .
    .
    return;
}
```

Figure 37. Exception handling using both C and CICS

Notes:

1. The EXEC CICS HANDLE ABEND PROGRAM registers PROGRAMB to CICS as the program to handle abends that occur in the task.
2. The `signal()` function registers the `handle()` function to be invoked in the case of a SIGSEGV abend.
3. The `raise()` function sends the signal SIGSEGV to the signal handler; in this example, the `handle()` function is invoked to handle the signal. After execution of the handler, control is returned to the statement after the call to `raise()`.
4. A protection exception has occurred and control is returned to CICS. Because an EXEC CICS HANDLE ABEND PROGRAM(...) takes precedence over the C signal handler, CICS invokes PROGRAMB and on completion does not return to the abending program. The return is determined by the rules for CICS abend handling.

Trace

CICS trace is a debugging aid for application programmers, system programmers, and IBM field engineers. It produces trace entries in response to trace commands. The trace entries can be sent to any trace destination that is currently active. The destinations are:

- Internal trace table
- Auxiliary trace data set

For information about trace destinations, see the *CICS Problem Determination Guide*.

You can:

- Specify user trace entry points (ENTER TRACENUM). (The earlier ENTER TRACEID command is supported for compatibility purposes. See the *CICS Migration Guide* for details.)
- Switch CICS internal trace on or off using the SET TRACEDEST, SET TRACEFLAG, and SET TRACETYPE commands. (TRACE ON and TRACE OFF are supported for compatibility only.)

Trace entry points

The points at which trace entries are produced during CICS operation are of four types: system trace entry points, user trace entry points, exception trace entry points, and user exception trace entry points. See the *CICS Problem Determination Guide* for more information about tracing.

System trace entry points

These are points within CICS at which trace control requests are made. The most important system trace entry points for application programmers are for the EXEC interface program. These produce entries in the trace table whenever a CICS command is processed.

Two trace entries are made: the first when the command is issued, and the second when CICS has performed the required function and is about to return control to your application program. Between them, these two trace entries allow you to trace the flow of control through an application, and to check which exception conditions,

if any, occurred during its execution. The EXEC CICS ABEND, EXEC CICS RETURN, EXEC CICS TRACEFLAG, and EXEC CICS XCTL commands produce single entries only. For programming information about these commands, see the *CICS Application Programming Reference* manual.

User trace entry points

These are additional points within your application program that you can include in the trace table to allow complete program debugging. For example, you could specify an entry for a program loop containing a counter value showing the number of times that the loop had been entered.

A trace entry is produced wherever the ENTER TRACENUM command is run. Each trace entry request, which can be given a unique identifier, causes data to be placed in the trace table.

Exception trace entry points

These are additional points where CICS has detected an exception condition. These are made from specific points in the CICS code, and data is taken from areas that might provide some information about the cause. Exception trace entry points do not have an associated “level” attribute; trace calls are only ever made from them when exception conditions occur.

User exception trace entry points

These are trace entries that are always written to the internal trace table (even if internal tracing is set off), but are written to other destinations only if they are active. You can identify them by the character string *EXCU in any formatted trace output produced by the CICS utility programs. See the *CICS Problem Determination Guide* for general information about user exception trace entry points; programming information is in the *CICS Application Programming Reference*.

Monitoring

CICS monitoring provides information about the performance of your application transactions.

Use the EXEC CICS MONITOR command for user event monitoring points.

In addition to the monitoring data collected from a system defined elsewhere, monitoring points (EMPs) within CICS, a user application program can contribute data to user fields within the CICS monitoring records. You can do this by using the EXEC CICS MONITOR POINT command to invoke user-defined EMPs. At each of these EMPs, you can add or change up to 16,384 bytes of your own data in each performance monitoring record. In those 16,384 bytes, you can have any combination of the following:

- In the range 0 through 256 counters
- In the range 0 through 256 clocks
- A single 8192-byte character string

For example, you could use these user EMPs to count the number of times a certain event occurs, or to time the interval between two events. For programming information about monitoring, see the *CICS Customization Guide*; for background information, see the *CICS Performance Guide*.

Dumping

CICS dump allows you to specify areas of main storage to be dumped, by means of the EXEC CICS DUMP TRANSACTION command, onto a sequential data set, which can be either on disk or tape. (The DUMP command is still supported for compatibility reasons.) See the *CICS Application Programming Reference* for details of the EXEC CICS DUMP TRANSACTION command.

The EXEC CICS PERFORM DUMP command allows you to request a VSE SDUMP. See the *CICS System Programming Reference* manual for programming information about EXEC CICS PERFORM DUMP.

You can format the contents of the dump data set and you can print them offline using the CICS dump utility program (DFHDU410 for transaction dumps) or the VSE/ESA Info/Analysis (INFOANA) program for system dumps. Instructions on using these programs are given in the *CICS Operations and Utilities Guide*.

Only one dump control command is processed at a time. If you issue additional dump control commands, while another task is taking a transaction dump, activity within the tasks associated with those commands is suspended until the dump is completed. Remaining dump commands are processed in the order in which they are made. Using the EXEC CICS DUMP TRANSACTION command causes some fields (for example, EIBFN and EIBRCODE) in the EIB and the TCA to be overwritten. See the *CICS Application Programming Reference* manual for programming information about EXEC CICS DUMP TRANSACTION.

Options on the EXEC CICS DUMP TRANSACTION command allow you to dump the following areas of main storage in various combinations:

- Task-related storage areas: selected main storage areas related to the requesting task. You would normally use a dump of these areas to test and debug your application program. (CICS automatically provides this service if the related task is terminated abnormally.)
- CICS control tables:
 - File control table (FCT)
 - Destination control table (DCT)
 - Program control table (PCT)
 - Processing program table (PPT)
 - System initialization table (SIT)
 - Terminal control table (TCT)

A dump of these tables is typically the first dump taken in a test in which the base of the test must be established; subsequent dumps are usually of the task-related storage type.

- It is sometimes appropriate during execution of a task to have a dump of both task-related storage areas and CICS control tables. Specifying one CICS control tables dump and a number of task-related storage dumps is generally more efficient than specifying a comparable number of complete dumps. However, do not use this facility excessively, because CICS control tables are primarily static areas.
- In addition, the EXEC CICS DUMP TRANSACTION command used with the three options, SEGMENTLIST, LENGTHLIST, and NUMSEGMENTS, allows you to dump a series of task-related storage areas simultaneously.

Program storage is not dumped for programs defined with the attribute RELOAD(YES).

You also get a list of the CICS nucleus modules and active PPT programs, indexed by address, at the end of the printed dump.

Part 3. Files and databases

Chapter 12. An overview of file control	167	The FROM option	180
VSAM data sets	167	The TOKEN option	181
Key-sequenced data set (KSDS)	167	Avoiding transaction deadlocks	181
Entry-sequenced data set (ESDS)	168	Rules for avoiding deadlocks	183
Relative record data set (RRDS)	168	KEYLENGTH option for remote data sets	183
Empty data sets	168		
VSAM alternate indexes	169	Chapter 13. File control—VSAM	
DAM data sets	169	considerations	185
CICS shared data tables	170	Record identification	185
Accessing data sets from CICS application		Key	185
programs	172	Relative byte address (RBA) and relative	
Reading records	172	record number (RRN)	185
Direct reading (using READ command)	172	CICS locking of VSAM records in	
Sequential reading (browsing)	173	recoverable files	186
Skip-sequential processing	176		
Updating records with the READ command	176	Chapter 14. File control—DAM	
Deleting records	177	considerations	189
Deleting single records	177	Record identification	189
Deleting groups of records (generic delete)	178	Block reference subfield	189
Read integrity	178	Physical key subfield	189
Adding records	178	Deblocking argument subfield	189
Adding to a KSDS	178	Updating records from DAM data sets	190
Adding to an ESDS	178	Browsing records from DAM data sets	190
Adding to an RRDS	178	Adding records to DAM data sets	191
Specifying record length	179	DAM exclusive control	192
Sequential adding of records (WRITE			
MASSINSERT command)	179	Chapter 15. Database control	193
Review of file control command options	179	DL/I databases	193
The RIDFLD option	179	Address restrictions	193
The INTO and SET options	180	DATABASE 2 (DB2) databases	193

Chapter 12. An overview of file control

CICS file control offers you access to data sets that are managed by:

- “VSAM data sets”
- “DAM data sets” on page 169

CICS file control lets you read, update, add, and browse data in VSAM and DAM data sets and delete data from VSAM data sets. You can also access CICS data tables using file control.

A CICS application program reads and writes its data in the form of individual records. Each read or write request is made by a CICS command.

To access a record, the application program must identify both the record and the data set that holds it. It must also specify the storage area into which the record is to be read or from which it is to be written.

The rest of this chapter contains the following topics:

- “CICS shared data tables” on page 170
- “Accessing data sets from CICS application programs” on page 172
- “Reading records” on page 172
- “Updating records with the READ command” on page 176
- “Deleting records” on page 177
- “Adding records” on page 178
- “Review of file control command options” on page 179
- “Avoiding transaction deadlocks” on page 181
- “KEYLENGTH option for remote data sets” on page 183

VSAM data sets

CICS supports access to the following types of data set:

- Key-sequenced data set (KSDS)
- Entry-sequenced data set (ESDS)
- Relative record data set (RRDS)

VSAM data sets are held on direct access storage devices (DASD) auxiliary storage. VSAM divides its data set storage into control areas (CA), which are further divided into control intervals (CI). Control intervals are the unit of data transmission between virtual and auxiliary storage. Each one is of fixed size and, in general, contains a number of records. A KSDS or ESDS can have records that extend over more than one control interval. These are called spanned records.

Key-sequenced data set (KSDS)

A **key-sequenced data set** has each of its records identified by a key. (The **key** of each record is simply a field in a predefined position within the record.) Each key must be unique in the data set.

When the data set is initially loaded with data, or when new records are added, the logical order of the records depends on the collating sequence of the key field. This also fixes the order in which you retrieve records when you browse through the data set.

To find the physical location of a record in a KSDS, VSAM creates and maintains an **index**. This relates the key of each record to the record's relative location in the data set. When you add or delete records, this index is updated accordingly.

Entry-sequenced data set (ESDS)

An **entry-sequenced data set** is one in which each record is identified by its relative byte address (RBA).

Records are held in an ESDS in the order in which they were first loaded into the data set. New records added to an ESDS always go after the last record in the data set. You may not delete records or alter their lengths. After a record has been stored in an ESDS, its RBA remains constant. When browsing, records are retrieved in the order in which they were added to the data set.

Relative record data set (RRDS)

A **relative record data set** has records that are identified by their relative record number (RRN). The first record in the data set is RRN 1, the second is RRN 2, and so on.

Records in an RRDS can be fixed or variable length records, and the way in which VSAM handles the data depends on whether the data set is a fixed or variable RRDS. A fixed RRDS has fixed-length slots predefined to VSAM, into which records are stored. The length of a record on a fixed RRDS is always equal to the size of the slot. VSAM locates records in a fixed RRDS by multiplying the slot size by the RRN (which you supply on the file control request), to calculate the byte offset from the start of the data set.

A variable RRDS, on the other hand, can accept records of any length up to the maximum for the data set. In a variable RRDS VSAM locates the records by means of an index.

A fixed RRDS generally offers better performance. A variable RRDS offers greater functionality.

Empty data sets

An empty data set is a data set that has not yet had any records written to it. VSAM imposes several restrictions on an empty data set such as not allowing you to have an alternate index in the upgrade set. CICS will allow you to use an empty data set in the same way as you use a normal data set. However, only attempt to write records to an empty data set.

When the first write request, or group of EXEC CICS WRITE MASSINSERT requests to the data set completes, CICS automatically closes the data set. This should be invisible to the application program. If you want to avoid the restrictions in the use of empty data sets, particularly the extension of upgrade sets, you may want to preload the data set using IDCAMS REPRO and BLDINDEX commands. See the *CICS System Definition Guide* for further information about using these commands. It is not in load mode when it is next opened.

VSAM alternate indexes

Sometimes you want to access the same set of records in different ways. For example, you may have records in a personnel data set that have as their key an employee number. No matter how many Smiths you have, each of them has a unique employee number. Think of this as the primary key. If you were producing a telephone directory from the data set, you would want to list people by name rather than by employee number.

You can identify records in a data set with a secondary (alternate) key instead of the primary key described above. So the primary key is the employee number, and the employee name is the **alternate key**. Alternate keys are just like the primary key in a KSDS—fields of fixed length and fixed position within the record. You can have any number of alternate keys per base file and, unlike the primary or base key, alternate keys need not be unique.

To continue the personnel example, the employee's department code might be defined as a further alternate key.

VSAM allows KSDS and ESDS (but not RRDS) data sets to have alternate keys. When the data set is created, one secondary or **alternate index** is built for each alternate key in the record and is related to the primary or base key. To access records using an alternate key, you must define a further VSAM object, an **alternate index path**. The path then behaves as if it were a KSDS in which records are accessed using the alternate key.

When you update a record by way of a path, the corresponding alternate index is updated to reflect the change. However, if you update the record directly by way of the base, or by a different path, the alternate index is updated only if it has been defined to VSAM (when created) to belong to the **upgrade set** of the base data set. For most applications, you probably want your alternate index (AIX) to be in the upgrade set.

A CICS application program disregards whether the file it is accessing is a path or the base. In a running CICS system, access to a single base data set can be made by way of the base and by any of the paths defined to it, if each access route is defined in the file control table (FCT).

DAM data sets

CICS supports access to keyed and nonkeyed DAM data sets. DAM support uses the physical nature of a record on a DASD device. DAM data sets consist of unblocked records with the following format:




Figure 38. Format of unblocked records in a DAM data set

Keyed DAM files have a physical key identifying the DAM record. The count area contains the physical key length, the physical data length, and the record's data location.

CICS can define a further structure on top of DAM data sets, introducing the concept of blocked-data sets:



Figure 39. Blocked-data set

The data portion of the physical record is viewed as a block containing logical records. CICS supports the retrieval of logical records from the data part of the physical record. CICS also supports unblocked records (where the structure reverts to the original DAM concept of one logical record per physical record).

To retrieve data from a physical record in a DAM file under CICS, a record identification field (RIDFLD) has to be defined to specify how the physical record should be retrieved. This may be done using the physical key, by relative address, or by absolute address.

If the data set is defined to CICS as being blocked, individual records within the block can be retrieved (deblocked) in two addressing modes: by key or by relative record.

Deblocking by key uses the key of the logical record (that is, the key contained in the logical record) to identify which record is required from the block. Deblocking by relative record uses the record number in the block, relative to zero, of the record to be retrieved.

You specify the key or relative record number used for deblocking in a subfield of the RIDFLD option used when accessing CICS DAM files. The addressing mode for CICS DAM files is set in the FCT using the RELTYPE keyword.

For more details about record identification and DAM record access, see Chapter 14, "File control—DAM considerations" on page 189.

CICS shared data tables

A CICS file is a representation of a data set on DASD. If you specify that the file is to use data table services, CICS copies the contents of the data set into a VSE/ESA data space when the file is opened and uses that copy whenever possible. Data tables offer a method of constructing, maintaining, and gaining rapid access to data records contained in tables held in virtual storage.

Because of the way that the data table services access the records, they can be used only with a VSAM key-sequenced data set (KSDS). Each data table is associated with a VSAM KSDS and is known as the **source data set**. The copy in virtual storage is called the **data table**. The process of copying the records is called **loading** the data table. The EXEC CICS file control commands can be used

to access data tables. Data tables also exploit the data in the memory facility that can dramatically improve performance of transactions whose primary access to VSAM data sets is read-only.

A table is defined using the RDO DEFINE FILE command or the DFHFCT macro. When a table is opened, CICS builds it by extracting data from the table's corresponding source data set and loading it into virtual storage.

When a file is read by a CICS command-level application, the record is normally retrieved from the data table. When the file is updated by a CICS application, the effect depends on the type of data table that you have defined for the file.

CICS supports two types of data table, as follows:

- **CICS-maintained tables**

All CICS file control commands can be used in applications that access a CICS-maintained data table. This means that you can immediately obtain the benefits of data tables without changing your existing applications. Full recovery support of the file is retained.

For update requests, CICS changes both the source data set and the data table. For most read requests, CICS needs to access only the data table. See the *CICS Shared Data Tables Guide* for details of which read requests cannot be satisfied by access to the data table. The CICS-maintained data table is best suited to applications that require high-speed, low-cost access to read-only data with some updating of the source data set, and with full recoverability after a system or transaction failure.

The full file control API appropriate to VSAM KSDS data sets is supported for CICS-maintained data tables. Requests that cannot be satisfied by reference to the data table result in calls to VSAM to access the source data set. Tables defined to be recoverable are supported with full integrity.

- **User-maintained tables**

This type of data table lets you optimize the benefits of using a data table by allowing you to eliminate activity on the source data set, for update requests as well as read requests. When you update the file, CICS changes only the data table.

There are some restrictions on which commands and options can be used. There are also some exceptional conditions that are unique to user-maintained data tables. These restrictions and conditions are described in detail in the *CICS Shared Data Tables Guide*. Because some file requests are not supported, a user-maintained data table is more suitable for new applications; existing application programs might need to be changed.

A user-maintained data table can be associated with a “dummy” KSDS, but can be loaded with non-VSAM data (from a DL/I database, for example). The records can then be accessed by applications as if they were KSDS records.

Unlike a CICS-maintained data table, if a request cannot be satisfied from a user-maintained data table, CICS does not access the source data set; it returns the NOTFND exceptional-condition response.

Records that were in the source data set when the data table was opened might be absent from the data table because they were not copied during

loading. This can be due to suppression by the user exit XDTRD or some abnormal event such as the data table becoming full.

A table defined as recoverable participates in dynamic transaction backout but is not recovered at restart or XRF takeover. For more information about data tables, see the *CICS Shared Data Tables Guide*.

You can use the user exits in data table services to put only the records that you need to access in the data table; there is no possibility of the source data set being accessed for those that you do not want.

Accessing data sets from CICS application programs

This section describes the facilities available to application programs for accessing data sets. Although VSAM data sets, are discussed, most of the facilities apply equally to DAM.

Reading records

A file can be defined in the file definition as containing either fixed-length or variable-length records. Fixed-length records should be defined only if:

- The definition of the VSAM data set (using access method services) specifies an average record size that is equal to the maximum record size

and

- All the records in the data set are of that length.

For direct reading and browsing, if the file contains fixed-length records, and if the application program provides an area into which the record is to be read, that area must be of the defined length. If the file contains variable-length records, the command must also specify the length of the area provided to hold them (which should normally be the maximum length of records in the file).

For fixed-length records and for records retrieved into storage provided by CICS (when you use the SET option), you need not specify the LENGTH argument. However, you may do so, to check whether the record being read is not too long for the available data area. If you specify LENGTH, CICS uses the LENGTH field to return the actual length of the record retrieved.

Direct reading (using READ command)

You read a record in the file with the EXEC CICS READ command. This must identify the record you want and say whether it is to be read into an area of storage provided by your application program (EXEC CICS READ INTO), or into CICS SET storage acquired by file control (EXEC CICS READ SET). If the latter, the address of the data in the CICS SET storage is returned to your program.

EXEC CICS SET storage normally remains valid until the next syncpoint, or the end of the task, whichever comes first.

Direct reading from a KSDS

When reading from a VSAM KSDS, you can identify the record you want uniquely by specifying its full key, or you can retrieve the first (lowest key) record whose key meets certain requirements. There are two options that qualify your key value; GENERIC and GTEQ.

The GENERIC option indicates that you require a match on only a part of the key; when you specify the GENERIC option, you also must specify the KEYLENGTH option, to say how many positions of the key, starting from the left, must match. For the READ command, CICS uses only the first KEYLENGTH option characters.

The GTEQ option indicates that you want the first record whose key is “greater than or equal to” the key you have specified. You can use GTEQ with either a full or a generic key.

The opposite of the GTEQ option is the EQUAL option (the default), which means that you want only a record whose key matches exactly that portion (full or generic) of the key that you have specified.

Direct reading from an ESDS

When reading from a VSAM ESDS, the individual record you want is identified by an RBA. Because the RBA of a record in an ESDS cannot change, your application program can keep track of the values of the RBAs corresponding to the records it wants to access. An RBA must always point to the beginning of a record. There is no equivalent to the GENERIC or GTEQ options that you can use to position approximately in a KSDS.

Direct reading from an RRDS

When reading from a VSAM RRDS, the record to be retrieved is identified by its relative record number. The application program must know the RRN values of the records it wants. There is no equivalent to the GENERIC or GTEQ options that you can use to position approximately in a KSDS.

Direct reading by way of a path

If a KSDS or an ESDS has an alternate index and an alternate index path (and an appropriate entry in the FCT), you can retrieve a record in the file by using the alternate key that you set up in the alternate index. The GENERIC option and the GTEQ (greater than or equal to) option still work in the same way as for a read from a KSDS using the primary key.

If the alternate key in an EXEC CICS READ command is not unique, the first record in the file with that key is read and you get the DUPKEY condition. To retrieve other records with the same alternate key, you have to start a browse operation at this point.

Sequential reading (browsing)

You start a browse with the EXEC CICS STARTBR command, identifying a particular record in the same way as for a direct read. However, the EXEC CICS STARTBR command only identifies the starting position for the browse; it does not retrieve a record.

The EXEC CICS READNEXT command reads records sequentially from this starting point. On completion of each EXEC CICS READNEXT command, CICS

returns the full key of the record it retrieved in the field specified in the RIDFLD option. (Be sure to provide a field as long as the full key, even if you use a STARTBR command with a shorter generic key.)

As in the case of a direct read, the record may be read into an area supplied by the application program (when you use the INTO option), or into storage provided by CICS (when you use the SET option). In the latter case, the CICS storage remains valid until the browse ends, syncpoint, end of task, whichever occurs first.

For VSAM, but not for DAM, you can also browse backward in the file, by using EXEC CICS READPREV commands instead of EXEC CICS READNEXT commands, and you can switch from one direction to the other at any time. The READPREV command is like the READNEXT command, except that the records are read sequentially **backward** from the current position. As you switch from one direction to the other, you retrieve the same record twice, unless you reposition.

After the browse has started, you can change the current browse position either by using a RESETBR command, or a READNEXT command, or a READPREV command. The RESETBR command can also be used for other purposes, however.

For VSAM, but not for DAM, you can reposition simply by varying the value in RIDFLD when you issue the next EXEC CICS READNEXT or READPREV command. When you change RIDFLD, the record identifier must be in the same form as on the previous STARTBR or RESETBR command (key, RBA, or RRN). In addition, you can change the length of a generic key by specifying a KEYLENGTH in your READNEXT or READPREV command, which is different from the current generic key length and not equal to the full length. If you change the length of a generic key in this way, you reposition to the generic key specified by RIDFLD option.

RESETBR command must be used to change the browse position in a DAM file. If you wish to change the form of the key from key to RBA or vice versa, you must use a RESETBR command. You must also use a RESETBR command to switch between generic and full keys or between “equal to” and “greater than or equal to” searches. You can also only use X'FF' characters to point to the last record in the file if you are using a RESETBR or STARTBR command.

Under certain conditions, CICS uses VSAM skip-sequential processing when you change the key in this way, as explained in “Skip-sequential processing” on page 176.

Browsing through a KSDS

You can use a generic key on the STARTBR command when browsing through a KSDS. However, the browse can only continue forward through the file. If you process a READPREV command during such a browse, you get the INVREQ condition.

You can use the options “key equal to” and “key greater than or equal to” on the STARTBR command and they have the same meaning as on the READ command. However, the STARTBR command assumes that you want to position at the key specified or the first key greater if that key does not exist. That is, option GTEQ is the default for the STARTBR command, whereas EQUAL is the default for the READ command.

You can start a forward browse through a VSAM KSDS at the start of the file by specifying a key of hexadecimal zeros, or by specifying options GENERIC, GTEQ, and KEYLENGTH(0) on the STARTBR, RESETBR, READNEXT, or READPREV command. (In the latter case, you need the RIDFLD keyword although its value is not used and, after the command completes, CICS is using a generic key length of one.)

You can start from the end of the data set by specifying a complete key of X'FF' characters on the STARTBR or RESETBR command. This points to the last record in the file ready for a backward browse.

A STARTBR, RESETBR, or READNEXT command having the option KEYLENGTH(0) is always treated as if KEYLENGTH(1) and a partial key of one byte of binary zeros have been specified.

Browsing through an ESDS

Positioning for a browse in an ESDS is identical to that for reading. If you want to begin reading at the beginning of the data set, use an RBA of low values (X'00'), and to begin at the end, use high values (X'FF').

Browsing through an RRDS

You can use the GTEQ option on an EXEC CICS STARTBR command when browsing through an RRDS. It is the default, even though on a direct READ this option has no effect. A direct read command with the GTEQ option that specifies an RRN that does not exist returns the NOTFND condition, because only the EQUAL option is taken. However, a STARTBR GTEQ command using the same RRN completes successfully, and sets a pointer to the relevant position in the data set for the start of the browse. The first record in the file is identified using an RRN of 1, and the last record by high values (X'FF').

Browsing using a path

Browsing can also use an alternate index path to a VSAM KSDS or an ESDS. The browse is just like that for a VSAM KSDS, but using the alternate key. The records are retrieved in alternate key order.

If the alternate key is not unique, the DUPKEY condition is raised for each retrieval operation except the last occurrence of the duplicate key. For example, if there are three records with the same alternate key, DUPKEY is raised on the first two, but not the third. The order in which records with duplicate alternate keys are returned is the order in which they were written to the data set. This is true whether you are using a READNEXT or a READPREV command. For this reason, you cannot retrieve records with the same alternate key in reverse order.

Ending the browse

Trying to browse past the last record in a file raises the ENDFILE condition. Stop a browse with the ENDBR command. You must issue the ENDBR command before performing an update operation on the same file (a READ UPDATE, DELETE with RIDFLD, or WRITE command). If you do not, you get unpredictable results, possibly including deadlock within your own transaction.

Simultaneous browse operations

CICS allows a transaction to perform more than one browse on the same file at the same time. You distinguish between browse operations by including the REQID option on each browse command.

Skip-sequential processing

When possible, CICS uses VSAM “skip-sequential” processing to speed browsing. Skip-sequential processing applies only to forward browsing of a file when RIDFLD is specified in key form. CICS uses it when you increase the key value in RIDFLD on your READNEXT command and make no other key-related specification, such as KEYLENGTH. In this situation, VSAM locates the next desired record by reading forward through the index, rather than repositioning from scratch. This method is faster if the records you are retrieving are relatively close to each other but not necessarily adjacent; it can have the opposite effect if the records are very far apart in a large file. If you know that the key you are repositioning to is much higher in the file, and that you may incur a long index scan, you may wish to consider using a RESETBR command that forces a reposition from scratch.

Updating records with the READ command

To update a record, you must first retrieve it using an EXEC CICS READ command with the UPDATE option. The record is identified in exactly the same way as for a direct read. In a VSAM KSDS or ESDS, the record may (as with a direct read) be accessed either by way of a FCT entry that refers to the base, or to a path defined over it.

After modification by the application program, the record is written back to the data set using the EXEC CICS REWRITE command, which does not identify the record being rewritten.

In a task, each REWRITE command should be associated with a previous READ UPDATE by a common keyword (TOKEN). Attempts to perform multiple concurrent update requests within a task, on the same data set without the use of TOKENs, are prevented by CICS. If you want to release the string held by a READ UPDATE without rewriting or deleting the record, use the UNLOCK command. The UNLOCK command releases any CICS storage acquired for the READ command. If TOKEN is specified with the UNLOCK command, CICS attempts to match this with an outstanding READ UPDATE whose TOKEN has the same value.

If more than one file definition refers to the same base data set, and this has been denied to the VSAM with SHAREOPTIONS(1) or SHAREOPTIONS(2), only a single file definition can be open for output at any time unless data set name sharing is used. For information about data set name sharing, see the *CICS Performance Guide*. Multiple concurrent updates to a data set through different file definitions can be made using data set name sharing, which preserves data integrity. If the data set is defined with SHAREOPTIONS(3) or SHAREOPTIONS(4), CICS cannot ensure that data integrity is preserved.

The EXEC CICS REWRITE command does not identify the record being rewritten because, in any one transaction, CICS allows only a single update to a given data set to be in progress at any time. To avoid deadlock, no other operations against the same file should occur between the EXEC CICS READ UPDATE and EXEC CICS REWRITE. If you want to release the string held by an EXEC CICS READ

UPDATE without rewriting or deleting the record, use the EXEC CICS UNLOCK command. EXEC CICS UNLOCK releases any CICS storage acquired for the EXEC CICS READ and releases VSAM resources held by the EXEC CICS READ.

For both update and nonupdate commands, you must identify the record to be retrieved by the record identification field specified in the RIDFLD option. Immediately upon completion of an EXEC CICS READ UPDATE command, the RIDFLD data area is available for reuse by the application program.

A record retrieved as part of a browse operation may not be updated during the browse. The application program must end the browse, read the desired record with an EXEC CICS READ UPDATE, and perform the update. Failure to do this may cause a deadlock.

The record to be updated may (as in the case of a direct read) be read into an area of storage supplied by the application program, or into CICS-provided SET storage. For an EXEC CICS READ UPDATE command, CICS SET storage remains valid until the next EXEC CICS REWRITE, EXEC CICS UNLOCK, EXEC CICS DELETE (without RIDFLD), or EXEC CICS SYNCPOINT, whichever is encountered first.

For a VSAM KSDS, the base key in the record must not be altered when the record is modified. Similarly, if the update is being made by way of a VSAM path, the alternate key used to identify the record must not be altered either, although other alternate keys may be altered. If the file definition allows variable-length records, the length of the record may be changed.

The length of records in an ESDS, an RRDS, and a fixed-length KSDS must not be changed on update.

For a file defined to CICS as containing fixed-length records, the length of record being rewritten **must equal the original length**. For variable-length records, you must specify the length option with both the EXEC CICS READ UPDATE and the EXEC CICS REWRITE commands. The length must not be greater than the maximum defined to VSAM.

Deleting records

Records can never be deleted from an ESDS.

Deleting single records

You delete a single record in a VSAM KSDS or RRDS in one of two ways:

- Retrieve it for update with a READ UPDATE command, and then issue a DELETE command without specifying the RIDFLD option.
- Issue a DELETE command specifying the RIDFLD option.

If a full key is provided with the DELETE command, a single record with that key is deleted. So, if the data set is being accessed by way of an alternate index path that allows nonunique alternate keys, only the first record with that key is deleted. After the deletion, you know whether further records exist with the same alternate key, because you get the DUPKEY condition if they do.

Deleting groups of records (generic delete)

You can use a generic key with the DELETE command. Then, instead of deleting a single record, all the records in the file whose keys match the generic key are deleted with the single command. However, this cannot be used if the KEYLENGTH value is equal to the length of the whole key (even if duplicate keys are allowed). The number of records deleted is returned to the application program if the NUMREC option is included with the command. If access is by way of an alternate index path, the records deleted are all those whose alternate keys match the generic key.

Read integrity

The NOSUSPEND option discussed in “The NOSUSPEND option” on page 63, also applies to the CICS browse commands when you are using them to update a file.

Adding records

Add new records to a file with the WRITE command. They must always be written from an area provided by the application program.

Adding to a KSDS

When adding a record to a VSAM KSDS, the base key of the record identifies the position in the data set where the record is to be inserted. Although the key is part of the record, CICS also requires the application program to specify the key separately using the RIDFLD option on the WRITE command.

A record added to a KSDS by way of an alternate index path is also inserted into the data set in the position determined by the base key. However, the command must also include the alternate index key as the record identifier.

Adding to an ESDS

A record added to an ESDS is always added to the end of the file. You cannot insert a record in an ESDS between existing records. After the operation is completed, the relative byte address in the file where the record was placed is returned to the application program.

When adding a record to an ESDS by way of an alternate index path, the record is also placed at the end of the data set. The command must include the alternate index key in the same way as for a KSDS path.

Adding to an RRDS

To add a record to an RRDS, include the relative record number as a record identifier on the WRITE command. The record is then stored in the file in the position corresponding to the RRN.

Specifying record length

When writing to a fixed-length file, the record length must match the value specified at the time the file was created. In this case, you need not include the length with the command, although you may do so to check whether the length agrees with that originally defined to VSAM. If the file is defined as containing variable-length records, the command must always include the length of the record.

Sequential adding of records (WRITE MASSINSERT command)

MASSINSERT on a WRITE command offers potential improved performance where there is more than one record to add to a KSDS, ESDS, or path. The performance improvement is obtained only when the keys in successive WRITE MASSINSERT requests are in ascending order.

A MASSINSERT is completed by the UNLOCK command. This ensures that all the records are written to the file and the position is released. Always issue an UNLOCK command before performing an update operation on the same data set (read update, delete with RIDFLD, or write). If you do not, you may get unpredictable results, possibly including a deadlock.

Without an UNLOCK command, the MASSINSERT is completed when a syncpoint is issued, or at task termination.

Note: A READ command does not necessarily retrieve a record that has been added by an incomplete MASSINSERT operation.

See “VSAM data sets” on page 72 for more information about MASSINSERT.

Review of file control command options

Some of the file control command options you may specify are:

- RIDFLD
- INTO or SET
- FROM
- TOKEN

Use of the LENGTH option varies, depending on how you use the other options.

The RIDFLD option

Whatever you do to a record (read, add, delete, or start a browse), you identify the record by the RIDFLD option, except when you have read the record for update first. However, there is no RIDFLD for ENDBR, REWRITE, and UNLOCK commands. Further, during a browse using READNEXT or READPREV commands, you must include the RIDFLD option to give CICS a way to return the identifier of each record retrieved.

The RIDFLD option identifies a field containing the record identification appropriate to the access method and the type of file being accessed.

The RIDFLD option by itself is not always enough to identify a specific record in the file. So, when retrieving records from a KSDS, or from a KSDS or ESDS by way of an alternate index path, or when setting a starting position for a browse in this type

of data set, you can have one or both of the further options GTEQ and GENERIC with your command.

With READNEXT or READPREV commands, the application program would not usually set the RIDFLD field. After each command, CICS updates this field with the actual identifier of the record retrieved. (You can alter the RIDFLD value to set a new position from which to continue the browse.)

The INTO and SET options

With the READ, READNEXT, or READPREV command, the record is retrieved and put in main storage according to your INTO and SET options.

The INTO option specifies the main storage area into which the record is to be put.

For fixed-length records, you need not include the LENGTH option. If you do, the length specified must exactly match the defined length; otherwise, you get the LENGERR condition.

For variable-length records, always specify (in the LENGTH option) the longest record your application program accepts (which must correspond to the value defined as the maximum record size when the data set was created); otherwise, you get the LENGERR condition. LENGERR occurs if the record exceeds the length specified, and the record is then truncated to that length. After the record retrieval, if you include the LENGTH option, the data area specified in it is set to the actual record length (before any truncation occurs).

The SET option specifies a pointer to the buffer in main storage acquired by CICS to hold the record. When using the SET option, you need not include the LENGTH option. If you do include it, the data area specified is set to the actual record length after the record has been retrieved.

The FROM option

When you add records (using the EXEC CICS WRITE command), or update records (using the REWRITE command), specify the record to be written with the FROM option.

The FROM option specifies the main storage area that contains the record to be written. In general, this area is part of the storage owned by your application program. With the REWRITE command, the FROM area is usually (but not necessarily) the same as the corresponding INTO area on the READ UPDATE command. The length of the record can be changed when rewriting to a KSDS with variable-length records.

Always include the LENGTH option when writing to a file with variable-length records. If the value specified exceeds the maximum allowed in the cluster definition, LENGERR is raised when the command is executed. LENGERR is also raised if the LENGTH option is omitted when accessing a file with variable-length records.

When writing to a file with fixed-length records, CICS uses the length specified in the cluster definition as the length of the record to be written, so you need not have the LENGTH option. If you do, its value is checked against the defined value and you get a LENGERR condition if the values do not match.

The TOKEN option

The TOKEN option is a unique value within a task that is supplied by CICS during a READ UPDATE and is returned with an associated REWRITE, DELETE, or UNLOCK command. However, for each data set that is being updated by a task, there may be a single update request outstanding at any one time, which does not include the TOKEN option with it.

You can perform multiple concurrent updates on the same data set using the same task by including the TOKEN option with READ UPDATE, REWRITE, DELETE, or the UNLOCK command. Note that a set of concurrent updates fails if more than one record is being updated in the same CI, irrespective of the TOKEN associated with the request.

You cannot function ship a command containing a TOKEN option.

Avoiding transaction deadlocks

Design your applications so as to avoid transaction deadlocks. A deadlock occurs if each of two transactions (for example, A and B) needs exclusive use of some resource (for example, a particular record in a data set) that the other already holds. Transaction A waits for the resource to become available. However, if transaction B is not in a position to release it because it, in turn, is waiting on some resource held by A, both are deadlocked and the only way of breaking the deadlock is to cancel one of the transactions, thus releasing its resources.

A transaction may have to wait for a resource for several reasons while executing file control commands:

- For both VSAM and DAM data sets, any record that is being modified is held in exclusive control by the access method for the duration of the request. (With VSAM files, not just the record but the complete control interval containing the record is held in exclusive control.)
- If a transaction has modified a record in a recoverable file, CICS locks that record to the transaction even after the request that performed the change has completed. The transaction can continue to access and modify the same record; other transactions must wait until the transaction releases the lock, either by terminating or by issuing a syncpoint request. For more information, see “Syncpointing” on page 137.

Whether a deadlock actually occurs depends on the relative timing of the acquisition and release of the resources by different concurrent transactions. Application programs may continue to be used for some time before meeting circumstances that cause a deadlock; it is important to recognize and allow for the possibility of deadlock early in the application program design stages.

Here are examples of different types of deadlock found in recoverable data sets:

- Two transactions running concurrently are modifying records within a single recoverable file, through the same FCT entry, as follows:

Transaction 1		Transaction 2	
READ UPDATE	record 1	DELETE	record 2
UNLOCK	record 1		
WRITE	record 2	READ UPDATE	record 1
		REWRITE	record 1

Transaction 1 has acquired the record lock for record 1 (even though it has completed the READ UPDATE command with an UNLOCK command). Transaction 2 has similarly acquired the record lock for record 2. The transactions are then deadlocked because each wants to acquire the CICS lock held by the other. The CICS lock is not released until syncpoint.

- Two transactions running concurrently are modifying two recoverable files as follows:

Transaction 1		Transaction 2	
READ UPDATE	file 1, record 1	READ UPDATE	file 2, record 2
REWRITE	file 1, record 1	REWRITE	file 2, record 2
READ UPDATE	file 2, record 2	READ UPDATE	file 1, record 1
REWRITE	file 2, record 2	REWRITE	file 1, record 1

Here, the record locks have been acquired on different files as well as different records; however, the deadlock is similar to the first example.

- Two transactions running concurrently are modifying a single recoverable KSDS, through the same FCT entry, as follows:

Transaction 1		Transaction 2	
READ UPDATE	record 1	DELETE	record 3
WRITE	record 3	READ UPDATE	record 2

Suppose records one and two are held in the same control interval (CI). The first READ UPDATE has acquired VSAM exclusive control of the CI holding record one. The DELETE operation has completed and acquired the CICS record lock on record three. The WRITE operation must wait for the lock on record three to be released before it can complete the operation. Finally, the last READ UPDATE must wait for the VSAM exclusive control lock held by transaction one to be released.

- A transaction is browsing through a VSAM file that uses shared resources (LSRPOOLID not equal to NONE in the file resource definition). Before the ENDBR command, the transaction issues a further request to update the current record or another record that happens to be in the same control interval. Because VSAM already holds shared control of the control interval on

behalf of the first request, the second request wants exclusive control of the control interval and therefore enters deadlock.

The transaction about to enter a deadlock is abended with the abend code AFCE if it is deadlocked by another transaction, or with abend code AFCEG if it has deadlocked itself.

Rules for avoiding deadlocks

You can avoid deadlocks by following these rules:

- All applications that update (modify) multiple resources should do so in the same order. For instance, if a transaction is updating more than one record in a data set, it can do so in ascending key order. A transaction that is accessing more than one file should always do so in the same predefined sequence of files.
- An application that issues a READ UPDATE command should follow it with a REWRITE, DELETE without RIDFLD, or UNLOCK command to release the position before doing anything else to the file.
- A sequence of WRITE MASSINSERT commands must terminate with the UNLOCK command to release the position. No other operation on the file should be performed before the UNLOCK command has been issued.
- An application must end all browses on a file by means of ENDBR commands (thereby releasing the VSAM lock) before issuing a READ UPDATE, WRITE, or DELETE with RIDFLD command, to the file.

KEYLENGTH option for remote data sets

In general, file control commands need the RIDFLD and KEYLENGTH options. The KEYLENGTH option can be specified explicitly in the command, or determined implicitly from the file definition.

For remote files for which the SYSID option has been specified, the KEYLENGTH option must be specified if the RIDFLD option is passing a key to file control. If the remote file is being browsed, the KEYLENGTH option is not required for the READNEXT or READPREV command.

For a remote DAM file, where the DEBKEY or DEBREC options have been specified, KEYLENGTH (when specified explicitly) should be the total length of the key (that is, all specified subfields).

Chapter 13. File control—VSAM considerations

Here we shall see how you identify records in a VSAM data set, and look in more detail at how CICS locks VSAM records in recoverable data sets.

Record identification

You have three ways to identify records in data sets:

- Key
- Relative byte address (RBA)
- Relative record number (RRN)

Key

Generally, if you use a key, you can specify either a complete key or a generic (partial) key. The exceptions to this rule are when you write a record to a KSDS or when you write a record via an alternate index path. In either of these cases you **must** specify the complete key in the RIDFLD option of the command.

When you use a generic key, you must specify its length in the KEYLENGTH option and you must specify the GENERIC option on the command. A generic key cannot have a key length equal to the full key length. You must define it to be shorter than the complete key.

You can also specify the GTEQ option on certain commands, for both complete and generic keys. The command then positions at, or applies to, the record with the next higher key if a matching key cannot be found. When accessing a data set by way of an alternate index path, the record identified is the one with the next higher alternate key when a matching record cannot be found.

Even when using generic keys, always use a storage area for the record identification field that is equal in length to the length of the complete key. During a browse operation, after retrieving a record, CICS copies into the record identification area the actual identifier of the record retrieved. CICS returns a complete key to your application, even when you specified a generic key on the command. For example, a generic browse through a KSDS returns the complete key to your application on each READNEXT and READPREV command.

Relative byte address (RBA) and relative record number (RRN)

You can use the RBA and RRN options on most commands that access data sets. In effect, they define the format of the record identification field (RIDFLD). Unless you specify either the RBA or the RRN, the RIDFLD option should hold a key to be used for accessing a KSDS (or a KSDS or ESDS by way of an alternate index).

RBA

RBA specifies that the record identification field contains the relative byte address of the record to be accessed. A relative byte address is used to access an ESDS, and it may also be used to access a KSDS. All file control commands that refer to an ESDS base, and specify the RIDFLD option, must also specify the RBA option.

Note: If a KSDS is accessed in this way, the RBA of the record may change during the transaction as a result of another transaction adding records to, or deleting records from, the same data set.

RRN

RRN specifies that the record identification field contains the relative record number of the record to be accessed. The first record in the data set is number one. All file control commands that refer to an RRDS, and specify the RIDFLD option, must also specify the RRN option.

CICS locking of VSAM records in recoverable files

In “Avoiding transaction deadlocks” on page 181, we described the prevention of transaction deadlocks in terms of the record locks acquired whenever records in a recoverable file are modified. The locks are held on behalf of the transaction doing the change until it issues a syncpoint request or terminates (at which time a syncpoint is automatically performed). For VSAM recoverable file processing, note the following:

- Whenever a VSAM record is obtained for modification or deletion, CICS file control locks the record with an ENQUEUE request using the primary record identifier as the enqueue argument.

If a record is modified by way of a path, the enqueue uses the base key or the base RBA as an argument. So CICS permits only one transaction at a time to perform its request, the other transactions having to wait until the first has reached a syncpoint.

- For the READ UPDATE and REWRITE-related commands the record lock is acquired as soon as the READ UPDATE command has been issued.

For a DELETE command that has not been preceded by a READ UPDATE command, or for a WRITE command, the record lock is acquired at the time the VSAM command is executed.

For a WRITE MASSINSERT command (which consists of a series of WRITE commands), a separate record lock is acquired at the time each individual WRITE command is performed. Similarly, for a DELETE GENERIC command, each record deleted acquires a separate lock on behalf of the transaction issuing the request.

The record locks referred to above are known as update locks, because they are acquired whenever a record is updated (modified). A further type of lock (a delete lock) may also be acquired by file control whenever an EXEC CICS DELETE, an EXEC CICS WRITE, or an EXEC CICS WRITE MASSINSERT command is being performed for a recoverable KSDS or a recoverable path over a KSDS. A delete operation therefore may acquire two separate locks on the record being deleted.

The separate delete lock is needed because of the way file control does its write operations. Before executing a WRITE MASSINSERT command to a KSDS or RRDS, file control finds and locks the empty range into which the new record or records are to go. The empty range is locked by identifying the next existing record in the data set and acquiring its delete lock.

The empty range is locked to stop other requests simultaneously adding records into it. Moreover, the record defining the end of the empty range cannot be

removed during the add operation. If another transaction issues a request to add records into the empty range or to delete the record at the end of the range, the delete lock makes the transaction wait until the WRITE or WRITE MASSINSERT command is complete. The record held with a delete lock may, however, be **updated** by another transaction during the write operation if it is in another CI.

Unlike an update lock, a delete lock is held only for the duration of a delete or write operation, or a sequence of WRITE MASSINSERT commands terminated by an UNLOCK command. An WRITE MASSINSERT command that adds records to the file into more than one empty range releases the previous delete lock as it moves into a new empty range.

The CICS enqueueing mechanism is only for updates and deletes and does not prevent a read request being satisfied before the next syncpoint. The integrity of a READ command in these circumstances is not guaranteed.

Chapter 14. File control—DAM considerations

This chapter explains how to identify records in a DAM data set, and looks at DAM exclusive control.

Record identification

You identify records in DAM data sets by a **block reference**, a **physical key** (keyed data set), or a **deblocking argument** (blocked-data set). The record identification (in the RIDFLD option) has a subfield for each item. These subfields, when used, must be in the above order.

Block reference subfield

This is one of the following:

- Relative track and record (hexadecimal format): 2-byte TT, 1-byte R (RELTYPE=HEX).

The 2-byte TT begins at relative track zero. The 1-byte R begins at relative record one.

- Relative track and record (zoned decimal format): 6-byte TTTTTT, 2-byte RR (RELTYPE=DEC).
- Actual (absolute) address: 8-byte MBCCCHHR (RELTYPE operand omitted).

The system programmer must specify the type of block reference you are using in the RELTYPE operand of the DFHFCT TYPE=FILE macro that defines the data set.

Physical key subfield

You need this only if the data set has been defined to contain recorded keys. If used, it must immediately follow the block reference. Its length must match the length specified in the BLKKEYL operand of the DFHFCT TYPE=FILE macro that defines the data set.

Deblocking argument subfield

You need this only if you want to retrieve specific records from a block. If used, it must immediately follow the physical key (if present) or the block reference. If you omit it, you retrieve an entire block.

The deblocking argument can be a key or a relative record number. If it is a key, specify the DEBKEY option on a READ or STARTBR command and make sure its length matches that specified in the KEYLEN operand of the DFHFCT TYPE=FILE macro. If it is a relative record number, specify the DEBREC option on a READ or STARTBR command. It is a 1-byte binary number (first record=zero).

The examples in Figure 40 assume a physical key of four bytes and a deblocking key of three bytes.

Byte Number																
0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	
T	T	R														Search by relative track and record and key; deblock by key
M	B	B	C	C	H	H	R									Search by actual address; deblock by relative record
T	T	T	T	T	T	R	R									Search by zoned decimal relative track and record and key; deblock by key
T	T	R														Search by relative track and record; deblock by key

Figure 40. Examples of DAM record identification

Updating records from DAM data sets

You cannot change the record length of a variable blocked or unblocked DAM record on a REWRITE command that specifies deblocking. You cannot change the record length of an undefined format DAM record on a REWRITE command either.

Browsing records from DAM data sets

The record identification field must contain a block reference (for example, TTR or MBBCCHHR) that conforms to the addressing method defined for the data set. Processing begins with the specified block and continues with each subsequent block until you end the browse.

If the data set contains blocked records, processing begins at the first record of the first block and continues with each subsequent record, regardless of the contents of the record identification field.

In other words, CICS uses only the information held in the TTR or MBBCCHHR subfield of the RIDFLD to identify the record. It ignores all other information, such as physical key and relative record, or logical key.

After the READNEXT command, CICS updates the RIDFLD with the complete identification of the record retrieved. For example, assume a browse is to be started with the first record of a blocked, keyed data set, and deblocking by logical key is to be performed. Before issuing the STARTBR command, put the TTR (assuming that is the addressing method) of the first block in the record identification field. After the first READNEXT command, the record identification field might contain X'0000010504', where X'000001' represents the TTR value, X'05' represents the block key, (of length 1), and X'04' represents the logical record key.

Now assume that a blocked, nonkeyed data set is being browsed using relative record deblocking and the second record from the second physical block on the third relative track is read by a READNEXT command. Upon return to the application program, the record identification field contains X'00020201', where

X'0002' represents the track, X'02' represents the block, and X'01' represents the number of the record in the block relative to zero.

Note: Specify the options DEBREC and DEBKEY on the STARTBR command when browsing blocked-data sets. This enables CICS to return the correct contents in the RIDFLD. Specifying DEBREC on the STARTBR command causes the relative record number to be returned. Specifying DEBKEY on the STARTBR command causes the logical record key to be returned.

Do not omit DEBREC or DEBKEY when browsing a blocked file. If you do, the logical record is retrieved from the block, the length parameter is set equal to the logical record length, **but** the RIDFLD is not updated with the full identification of the record. **Do not use this method.**

Compare this with what happens if you omit the DEBREC or DEBKEY option when reading from a blocked DAM data set. In this case, you retrieve the **whole** block, and the length parameter is set equal to the length of the block.

Adding records to DAM data sets

When adding records to a DAM data set, bear in mind the following:

- When adding undefined or variable-length records (keyed or nonkeyed), you must specify the track on which each new record is to be added. If space is available on the track, the record is written following the last previously written record, and the record number is put in the “R” portion of the record identification field of the record. The track specification may be in any format except relative block. If you use zoned-decimal relative format, the record number is returned as a 2-byte zoned decimal number in the seventh and eighth positions of the record identification field.

The extended search option allows the record to be added to another track if no space is available on the specified track. The location at which the record is added is returned to the application program in the record identification field being used.

When adding records of undefined length, use the LENGTH option to specify the length of the record. When an undefined record is retrieved, the application program must find out its length.

- When adding keyed fixed-length records, you must first format the data set with dummy records or “slots” into which the records may be added. You signify a dummy record by a key of X'FF's. The first byte of data contains the record number.
- When adding nonkeyed fixed-length records, give the block reference in the record identification field. The new records are written in the location specified, destroying the previous contents of that location.
- When adding keyed fixed-length records, track information only is used to search for a dummy key and record which, when found, is replaced by the new key and record. The location of the new record is returned to the application program in the block reference subfield of the record identification field.

For example, for a record with the following identification field:

```
0 3 0 ALPHA  
T T R KEY
```

the search starts at relative track three. When control is returned to the application program, the record identification field is:

```
0 4 6 ALPHA
```

showing that the record is now record six on relative track four.

- When adding variable-length blocked records you must include a 4-byte record description field (RDF) in each record. The first two bytes specify the length of the record (including the 4-byte RDF); the other two bytes consist of zeros.

DAM exclusive control

When a blocked record is read for update, CICS maintains exclusive control of the containing block. An attempt to read a second record from the block before the first is updated (by a REWRITE command), or before exclusive control is released (by an UNLOCK command), causes a deadlock.

Chapter 15. Database control

DL/I databases give you more data independence than file control does. DATABASE 2™ databases also provide data independence, offering a logical view of the database as a series of tables that you can interrelate more or less as you wish.

DL/I databases

You get a logical view of the database in terms of a hierarchy of segments. DL/I lets you manipulate these segments without needing to know how they are organized. DL/I databases are processed by the IBM licensed program , DL/I VSE.

CICS has two programming interfaces to DL/I. We recommend that you use the higher-level EXEC DLI interface. It is simple, works with EDF, and can fully exploit a 31-bit environment. The lower-level DL/I programming interface is the DL/I CALL interface.

This book does not discuss EXEC DLI commands. See the *CICS Application Programming Reference* for a brief description of the EXEC CICS DL/I commands. You can find a full description in the *DL/I Application Programming: High Level Programming Interface* manual.

Address restrictions

For EXEC DLI programs, both the program and the data can reside above the 16MB line.

For CALL-level programs that use a local DL/I call interface, the parameter list and all storage areas that are referred to in the parameter list of the DL/I call must reside below the 16MB line.

DATABASE 2 (DB2) databases

DATABASE 2 databases also provide data independence, offering a logical view of the database as a series of tables that you can interrelate more or less as you wish. DB2 lets you manipulate these tables without needing to know how they are organized. DB2 databases are processed by DB2 for VSE/ESA 6.1.

CICS has one interface to DB2—the EXEC SQL interface, which offers powerful statements for manipulating sets of tables—thus relieving the application program of record-by-record (segment-by-segment, in the case of DL/I) processing.

CICS applications that process DB2 tables can also access DL/I databases. Any CICS resources (files, transient data, and so on), DL/I, and DB2 can be accessed from within one transaction.

For information about SQL commands, which are not discussed in this book, see the *SQL/DS SQL Reference* manual.

Part 4. Data communication

This part of the manual covers CICS facilities for communicating with the terminals through which the users access CICS applications.

CICS supports communication with other applications as well as with terminals, including other CICS regions, IMS/DC, and any program that understands one of the standard protocols that CICS uses for the purpose. Communications between applications are different in character from those with a terminal, however, and require a different application programming interface. CICS provides several sets of commands for this purpose. They are not covered in this manual; you will find them instead in the *CICS Distributed Transaction Programming Guide*.

Chapter 16. Introduction to data communication	199	Full BMS	227
Basic CICS terms	199	A BMS output example	228
How tasks get started	200	Creating the map	231
Automatic task initiation (ATI)	200	Defining the map: DFHMDDI	234
Unsolicited input	201	Defining the map set: DFHMDDI	234
Which transaction?	201	Rules for writing BMS macros	235
CICS APIs for terminals	203	Assembling the map	237
Topics elsewhere in this book	204	Physical and symbolic map sets	238
Where to find more information	205	The SDF II alternative	239
		Grouping maps into map sets	239
Chapter 17. The 3270 family of terminals	207	Complex fields	239
Introduction	207	Composite fields: the GRPNAME option	239
Background	208	Repeated fields: the OCCURS option	240
Screen fields	208	Sending mapped output: basics	241
Personal computers	209	The SEND MAP command	242
The 3270 buffer	210	Acquiring and defining storage for the maps	242
Writing to a 3270 terminal	211	Initializing the output map	244
3270 write commands	211	Moving the variable data to the map	244
Write control character	212	Setting the display characteristics	244
3270 display data: defining 3270 fields	212	Changing the attributes	246
Display characteristics	212	Attribute value definitions: DFHBMSCA	246
3270 field attributes	213	Control options on the EXEC CICS SEND MAP command	247
Extended attributes	215	Other BMS SEND options: WAIT and LAST	248
Orders in the data stream	215	Options for merging the symbolic and physical maps	248
Outbound data stream sample	218	MAPONLY option	248
Input from a 3270 terminal	221	DATAONLY option	248
Attention keys	221	The EXEC CICS SEND CONTROL command	249
Reading from a 3270 terminal	222	Summary: what appears on the screen	249
Inbound field format	223	What you start with	250
Input example	223	What is sent	250
Unformatted mode	224	Where the values come from	250
		Outside the map	252
Chapter 18. Basic mapping support	225	Using GDDM® and BMS	252
Introduction to BMS	225		
BMS support levels	226		
Minimum BMS	226		
Standard BMS	227		

Positioning the cursor	252	Formatting text output	288
Sending invalid data and other errors	253	The SEND TEXT command	289
Receiving data from a display	253	Text logical messages	289
An input-output example	254	Page format for text messages	290
The symbolic input map	256	How BMS breaks text into lines	291
Programming simple mapped input	257	Header and trailer format for text	
The EXEC CICS RECEIVE MAP command	257	messages	292
Getting storage for mapped input: INTO and SET	258	SEND TEXT extensions: SEND TEXT	
Reading from a formatted screen: what comes in	259	MAPPED and SEND TEXT NOEDIT	293
Modified data	259	Message routing: the ROUTE command	294
Upper case translation	260	How routing works	294
Other information from EXEC CICS		Specifying destinations for a routed	
RECEIVE MAP	260	message	295
The attention identifier: what caused		Route list format	297
transmission	260	Delivery conditions	299
The HANDLE AID command	261	Undeliverable messages	300
Processing the mapped input	262	Temporary storage and routing	300
Handling input errors	263	Programming considerations with routing	301
Mapped output after mapped input	264	Using SET	302
MAPFAIL and other exceptional conditions	265	Partition support	304
Formatting other input	266	Uses for partitioned screens	305
Support for non-3270 terminals	267	How to define partitions	306
Output considerations for non-3270		3290 character size	307
devices	267	Programming considerations	307
Differences on input	268	Establishing the partitioning	308
Special options for non-3270 terminals	268	Partition options for BMS SEND	
Device-dependent maps: map suffixes	269	commands	309
The MAPPINGDEV facility.	272	Determining the active partition	309
SEND MAP with the MAPPINGDEV		Partition options for BMS RECEIVE	
option.	272	commands	309
RECEIVE MAP with the MAPPINGDEV		ASSIGN options for partitions	310
option.	273	Partitions and logical messages	310
Sample assembler MAPPINGDEV		Partitions and routing	310
application.	274	New attention identifiers and exception	
Block data	276	conditions	311
Sending mapped output: additional facilities	276	Terminal sharing	311
Output disposition options: TERMINAL,		Restrictions on partitioned screens	312
SET, and PAGING	277	Logical device components	312
BMS logical messages	277	Defining logical device components	312
Terminal operator paging: the CSPG		Sending data to a logical device	
transaction	280	component	313
Changing your mind: the PURGE		LDCs and logical messages	313
MESSAGE command	281	LDCs and routing	313
Logical message recovery	282	BMS support for other special hardware	314
Page formation: the ACCUM option	282	10/63 magnetic slot reader	314
Floating maps: how BMS places maps		Field selection features	315
using ACCUM	282	Cursor- and pen-detectable fields	316
Page breaks: BMS overflow processing	283	Outboard formatting	318
Map placement rules	284	Chapter 19. Terminal control	319
ASSIGN options for cumulative processing	286	Access method support	319
Input from a composite screen	286	Basic Telecommunications Access Method	
Performance considerations	287	(BTAM)	319
		Terminal control commands	320

Data transmission commands	320	Non-3270 CICS printers	346
Send/receive mode	321	Determining the characteristics of a CICS printer	347
Speaking out of turn	322	CICS printers: getting the data to the printer	349
Interrupting	323	Printing with an EXEC CICS START command	349
Terminal waits	323	Printing with transient data	350
What you get on an EXEC CICS RECEIVE	323	Printing with BMS routing	351
Control commands	324	Non-CICS printers	352
Finding the right commands	325	Formatting for non-CICS printers	352
Finding out about your terminal	330	Non-CICS printers: Delivering the data	352
EIB feedback on terminal control operations	332	CICS API considerations	352
VTAM considerations	333	Notifying the print application	354
Chaining input data	333	Printing display screens	354
Chaining output data	334	CICS print key	355
Handling logical records	334	EXEC CICS ISSUE PRINT and EXEC CICS ISSUE COPY	355
Response protocol	335	Hardware print key	356
Using function management headers	335	BMS screen copy	356
Preventing interruptions (bracket protocol)	336	Chapter 21. CICS interface to POWER	357
Sequential terminal support	336	CICS interface to VSE/POWER	357
Coding considerations for sequential terminals	337	Input to CICS from VSE/POWER	357
Batch data interchange	338	Output from CICS using VSE/POWER	358
Destination selection and identification	340	Typical use of EXEC CICS SPOOL commands	359
Definite response	340	Programming note for spool commands	362
Waiting for function completion	340	Chapter 22. Report Controller	363
Chapter 20. CICS support for printing	341	Reports	364
Formatting for CICS printers	341		
3270 printers	342		
Options for 3270 printers	343		

Chapter 16. Introduction to data communication

This chapter introduces the subject of data communications, preparing the ground for the remaining chapters in this part of the book. It contains:

- “Basic CICS terms”
- “How tasks get started” on page 200
- “Which transaction?” on page 201
- “CICS APIs for terminals” on page 203
- “Topics elsewhere in this book” on page 204
- “Where to find more information” on page 205

Basic CICS terms

Here are several terms that are used throughout this part of the book:

Terminal A hardware device from which data can enter or leave CICS over a communication channel. It is usually some combination of keyboard, display screen and print mechanism. It may also be a control unit or a processor emulating a terminal, but it is not another application. The material in this part, together with descriptions of the commands cited here, tells you how to program for the wide variety of terminals that CICS supports. The commands are described in the *CICS Application Programming Reference* manual.

Logical unit (LU)

The VTAM term for the end point of a data transmission. That is, VTAM transmits data from one LU to another. (VTAM is an access method for terminals; see “Access method support” on page 319.) An LU may be a terminal, but it may also be a host application (CICS or IMS/DC, among others). VTAM distinguishes between different categories of end points by defining seven types of logical unit. Types 1, 2, and 3 are generally called “terminals”. Type 0 may be a terminal, or it may be a control unit or processor using a very simple communications protocol. Type 4 is a programmable control unit. Type 6.1 and Type 6.2 (APPC) are reserved for program-to-program communications (the type not covered in this chapter). In this section, the term **logical unit** is used to mean specifically a terminal connected under VTAM, and **terminal** when the access method is not important.

Transaction

When CICS is described as an online transaction processing system, **transaction** is used in the ordinary sense of an interaction between two participants. Making an airline reservation, mailing a payment, and sending a bill are all examples.

However, **transaction** also has a specific meaning in CICS, which applies to the material in this section and throughout the manual. It means the processing executed for one specific **type** of request. A request type may represent a whole process, like making an airline reservation, or a subcomponent of that process, like selecting a seat. Your application design defines what constitutes a request type, and you describe each one to CICS with an RDO TRANSACTION resource definition. This definition tells CICS several things about the work to be done; the key one is what program to invoke first. (You only have to tell

CICS where to start in the TRANSACTION definition; execution flow is controlled by the programs themselves thereafter.)

Transactions are identified by a **transaction identifier** (or, often, **transaction code**), a one- to four-character code by which both users and programs indicate the type of processing to be performed.

Task You will also see the word **task** used extensively here. This word also has a specific meaning in CICS: it is **one instance** of the execution of a particular transaction type. That is, one execution of a transaction, with a particular set of data, usually on behalf of a particular user at a particular terminal.

Principal facility

If you have looked at the programming information for the commands for terminals in the *CICS Application Programming Reference* manual, you may have noticed that there is no way to indicate **which** terminal you are talking about. That is because CICS allows a task to communicate directly with only one terminal, namely its **principal facility**. CICS assigns the principal facility when it initiates the task, and the task “owns” the facility for its duration. No other task can use that terminal until the owning task ends. If a task needs to communicate with a terminal other than its principal facility, it must do so indirectly, by creating another task that has the terminal as its principal facility. This requirement arises most commonly in connection with printing, and how you can create such a task is explained in “CICS printers: getting the data to the printer” on page 349.⁴

Alternate facility

Although a task may communicate directly with only one terminal, it can also establish communications with one or more remote systems. It does this by asking CICS to assign a conversation with that system to it as an **alternate facility**. The task “owns” its alternate facilities in the same way that it owns its principal facility. Ownership lasts from the point of assignment until task end or until the task releases the facility.

How tasks get started

Work gets started in CICS—that is, tasks get initiated—in one of two ways:

- By automatic task initiation (ATI)
- From unsolicited input

Automatic task initiation (ATI)

Automatic task initiation occurs when:

- An existing task asks CICS to create another one. The START command, the IMMEDIATE option on a RETURN command (discussed in “EXEC CICS RETURN IMMEDIATE” on page 322), and the SEND PAGE command (in “Ending a logical message: the SEND PAGE command” on page 279) all do this.

⁴ BMS routing, described in “Message routing: the ROUTE command” on page 294, is another form of indirect access to other terminals via queues.

- CICS creates a task to process a transient data queue (see “Automatic transaction initiation (ATI)” on page 405).
- CICS creates a task to deliver a message sent via a BMS ROUTE request (see “Message routing: the ROUTE command” on page 294). The CSPG tasks you see after using the CICS-supplied transaction CMSG are an example of this. CMSG uses a ROUTE command that creates a CSPG transaction for each target terminal in your destination list.

Unsolicited input

The primary mechanism for initiating tasks, however, is unsolicited input. When a user transmits input from a terminal that is not the principal facility of an existing task, CICS creates a task to process it. The terminal that sent the input becomes the principal facility of the new task.

Unsolicited inputs from other systems are handled in the same way: CICS creates a task to process the input, and assigns the conversation over which the input arrived as the principal facility. (Thus a conversation with another system may be either a principal or alternate facility. In the case where a task in one CICS region initiates a conversation with another CICS region, the conversation is an alternate facility of the initiating task, but the principal facility of the partner task created by the receiving system. By contrast, a terminal is always the principal facility.)

Not all tasks have a principal facility. Tasks that result from unsolicited input always do, by definition, but a task that comes about from automatic task initiation may or may not need one. When it does, CICS waits to initiate the task until the requested facility is available for assignment to the task.

Which transaction?

Having received an unsolicited input, how does CICS decide what to do with it? That is, what transaction should the task created to process it execute? The short answer is that the previous task with the same principal facility usually tells CICS what transaction to execute next just before it ends, via the TRANSID option on its final RETURN. This is almost always the case in a pseudoconversational transaction sequence, and usually in menu-driven applications as well. Failing that, and in any case to get a sequence started, CICS interprets the first few characters of the input as a transaction code. However, it is more complicated than that; the exact process goes as follows. The step numbers indicate the order in which the tests are made and refer to Figure 41 on page 202, a diagram of this logic.

0. On the very first input from a terminal, CICS sometimes schedules a preliminary task before creating one to process the input. This task executes the CICS-supplied “query” transaction, CQRY, which causes the terminal to transmit an encoded description of some of its hardware characteristics—extended attributes, character sets, and so on.

CQRY allows the system programmer to simplify maintenance of the terminal network by omitting these particulars from the terminal definitions. It occurs only if the terminal definition so specifies, and has no effect on the subsequent determination of what transaction to use to process the input, which goes as follows.




Figure 41. Determining which transaction to execute

1. If the terminal is a 3270 and the input is the “print request key”, the CICS-supplied transaction that prints the contents of the screen, CSPP, is initiated. See “CICS print key” on page 355 for more information about this

feature. For this purpose, a “3270 logical unit” or any other device that accepts the 3270 data stream counts as a 3270.

2. If full BMS support is present, the terminal is of a type supported by BMS terminal paging, and the input is a paging command, the CICS-supplied transaction CSPG is initiated to process the request. BMS support levels are explained in “BMS support levels” on page 226, and the same section contains a list of the terminals that BMS supports. The PGRET, SKRxxxx, PGCHAIN, PGCOPY and PGPURGE system initialization parameters define the paging commands. As paging requires full BMS, this step is skipped if the CICS system contains less than that level.
3. If the terminal definition indicates that a specific transaction should be used to process all unsolicited inputs from that terminal, the indicated transaction is executed. (If present, this information appears in the TRANSACTION attribute of the RDO TERMINAL resource definition.)
4. If the previous task at the terminal specified the TRANSID option of the RETURN command that ended it, the transaction named is executed.
5. If an attach function management header is present in the input, the attach names in the header are converted to a four-character CICS transaction identifier, and that transaction is executed.
6. If the terminal is a 3270, and the attention identifier is defined as a transaction, that transaction is executed. “Attention keys” on page 221 explains attention identifiers. You define one as a transaction identifier with the TASKREQ attribute of the corresponding RDO TRANSACTION resource definition.
7. If all of the preceding tests fail, the initial characters of the input are used to identify the transaction to be executed. The characters used are the first ones (up to four) after any control information in the data stream and before the first field separator character. Field separators are defined in the FLDSEP system initialization parameter (the default is a blank).

If there are no such characters in the input, as occurs when you use the CLEAR key, for example, or if there is no transaction definition that matches the input, CICS cannot determine what transaction to execute and sends an “invalid transaction identification” message to the terminal.

Note: This logic for deciding which transaction to execute applies only to tasks initiated to process unsolicited inputs. For automatic transaction initiation, the transaction is always known. You specify it in the TRANSID option when you create a task with a START or RETURN IMMEDIATE. Similarly, you specify what transaction should be used to process a transient data queue in the queue definition. Tasks created to route messages always execute the CICS-supplied transaction CSPG.

CICS APIs for terminals

The CICS application programming interface contains two sets of commands for communicating with terminals:

- Terminal control commands
- Basic Mapping Support (BMS)

Terminal control is the more basic of the two. It gives you flexibility and function, at the cost of more programming. In particular, if you code at the terminal control level, you need to build the device data stream in your application.

BMS lets you communicate with a terminal at a much higher language level. It formats your data, and you do not need to know the details of the data stream. It is thus easier to code initially and easier to maintain, especially if your application has to support new types of terminal. However, BMS pathlengths are longer (BMS itself uses terminal control), and BMS does not support all the terminal types that terminal control does. BMS is the subject of Chapter 18, “Basic mapping support” on page 225, and Chapter 19, “Terminal control” on page 319 covers terminal control.

Finally, you can use CPI-C “sockets” calls to communicate with terminals or other systems. This interface is covered in Chapter 6, “Intercommunication considerations” on page 119.

Topics elsewhere in this book

Even though BMS and terminal control are discussed separately, some of the discussion in the BMS chapter applies to terminal control as well, and vice-versa. These topics have been covered where they naturally arise. Be aware of them if you are using only one interface. They include:

- DFHBMSCA, a useful set of attribute byte definitions in a CICS-supplied copybook (see “Attribute value definitions: DFHBMSCA” on page 246).
- What happens when you send an invalid data stream, in “Sending invalid data and other errors” on page 253.
- Translation of mixed case input to uppercase, in “Upper case translation” on page 260.
- “The HANDLE AID command” on page 261. HANDLE AID lets you specify program flow based on the key used to transmit the input.
- Saving input data between tasks in a pseudoconversational sequence, discussed in “Saving the good input” on page 263.
- “Performance considerations” on page 287.
- ASSIGN command options specific to terminals. You can use these to find out the characteristics of the principal facility for your task. See “Finding out about your terminal” on page 330.
- EIB (execute interface block) fields that contain information specific to terminal operations, discussed in “EIB feedback on terminal control operations” on page 332.
- Send-receive conventions, in “Send/receive mode” on page 321 and “Preventing interruptions (bracket protocol)” on page 336.

If you are not familiar with 3270s and you plan to code for one or to use BMS on any terminal, read Chapter 17, “The 3270 family of terminals” on page 207. You do not have to read the whole chapter; we have noted shortcuts for BMS users. There is also material on special features of 3270s in the BMS chapter, in “BMS support for other special hardware” on page 314, and 3270 printers are covered in Chapter 20, “CICS support for printing” on page 341. Also read “Personal

computers” on page 209 if you are using a personal computer as a terminal or as the client in a client-server configuration.

Where to find more information

The commands cited in the chapters that follow are described fully in the *CICS Application Programming Reference* manual, and that manual should be used in conjunction with this one. In addition, the CICS device guides provide specific information for many families of devices.

Chapter 17. The 3270 family of terminals

This chapter describes the IBM 3270 family of terminals. It contains:

- “Introduction”
- “Background” on page 208
- “The 3270 buffer” on page 210
- “3270 display data: defining 3270 fields” on page 212
- “Input from a 3270 terminal” on page 221
- “Unformatted mode” on page 224

Introduction

The 3270 is a family of display and printer terminals, with supporting control units, that share common characteristics and use the same encoded data format to communicate between terminal and host processor. This data format is known as the **3270 data stream**.

This chapter helps you to understand 3270 facilities and operation, so that you can use these terminals to best advantage in creating the end-user interface for your application. Some appreciation of the 3270 is also crucial to understanding BMS, because so many facilities of BMS exploit features of the 3270.

The 3270 is a complex device with many features and capabilities. Only basic operations are covered here and the emphasis is on the way CICS supports the 3270. For a comprehensive discussion of 3270 facilities, programming and data stream format, see the *IBM 3270 Information Display System Data Stream Programmer's Reference* manual. BMS support for a few special features is discussed in the BMS chapter (see page 314).

The first section in this chapter gives some background information that may help you understand the 3270 better. If you have experienced this history first-hand, you can move directly to “The 3270 buffer” on page 210.

The sections that follow first describe how to write to a 3270 terminal, and then how to read from one. In the process, it gives sufficient information on the 3270 data stream to allow you to understand the capabilities of the 3270. Programmers using terminal control commands still need to consult the *IBM 3270 Information Display System Data Stream Programmer's Reference* manual for details. The *IBM CICS/DOS/VS 3270 Data Stream Device Guide* also contains much important information. It is primarily intended for programmers using terminal control, but contains information that may be helpful for BMS programmers as well.

Although the discussion in this chapter is focused on display terminals, most of the material applies equally to 3270 printers. A 3270 printer accepts the same data stream as a 3270 display and simply delivers the screen image in hardcopy form. Most of the differences relate to input, which is (mostly) lacking on printers.

However, additional formatting facilities are available for use with printers, and there are special considerations in getting your printed output to the desired printer. All of these issues are covered in the chapter devoted to printing, Chapter 20, “CICS support for printing” on page 341.

Background

The development of the 3270 coincided with, and in part caused, the explosive growth of online transaction processing that began in the late 1960s. Consequently, the 3270 was a major influence in the design of transaction processing systems such as CICS.

The earliest terminal devices for online processing were adaptations of the teletype, the original and most basic computer terminal. Output was typed, and structure in the input typed by the operator was determined entirely by program convention, without any assists from the hardware. Cathode-ray tube terminals brought a revolutionary improvement in output speed, allowing a complexity of application not previously possible, but formatting on early CRTs was not much more sophisticated than on their hard-copy predecessors.

Screen fields

The 3270 transformed the user interface by introducing the concept of **fields** on a display screen. Each field on the screen has a starting position and individual attributes, such as display intensity, color, and whether or not you can key data into it. Fields introduce structure into the communication between program and terminal operator in the same way that fields in a file record provide structure for interaction between programs and data.

```
Billing information on customer:
Reference Number      KRK123456
Full Name             Phileas Arthur Fogg
Amount Owed          $40.07
```

Figure 42. Part of a formatted screen, showing fields. Each block of text on the screen is a separate field. The fields on the left were filled in by program; those on the right were completed by an operator.

Organizing a screen display into fields has many advantages:

- The screen is easier to read, because fields can have different display characteristics.
- Data entry is enhanced by providing clear visual and keyboard cues about the order and format of the information required. The screen can be as explicit as a standard “fill-in-the-blanks” paper form. (Keyboard facilities reinforce the structure imposed by the fields. The keyboard locks if the operator tries to key into the wrong place. There are keys that tab from one field to the next, another that erases just the current field, and so on.)
- The length of the outbound data stream is reduced, because you send only nonblank (that is, nonspacer) data.
- The inbound data stream is also reduced, because the host normally reads only the changed fields.

Personal computers

The advent of personal computers (**PCs**) and intelligent workstations brought a second revolution in terminal display function. These terminals differ from 3270s in two important respects:

- They are generally “all points addressable”. That is, you can address any point on the display raster, just as you can on a television screen. A typical display might contain a grid of 640 by 480 points in the space normally used to display a single character on an earlier display. Moreover, a whole palette of colors and intensities is available at each point.

In contrast, a 3270 screen is divided into an array of character positions, typically 24 down and 80 across. Each position consists of an array of raster points, but you cannot address them individually. You can only select a character, from a set of about 190 choices, for each position. Some terminals allow you to select from several character sets and to load new sets, allowing a rudimentary form of graphics, but essentially you are working with a terminal that displays text, numbers and symbols. You get some control of how the characters are displayed, but the choices are very limited in comparison with a PC display.

- The second difference is what makes the first possible. Personal computers and intelligent workstations contain a processor, memory, and programming (that is, “intelligence”) that make it possible to communicate with this very much more complex hardware through a relatively simple programming interface and minimum long-distance transmission of data.

These characteristics make possible a much higher-function end-user interface than that of the 3270. You can draw pictures, select from a variety of fonts, scale images in size, and so on. If you are writing a new application, and all of your users will access it from such terminals, you may want to take advantage of this function to create the most efficient end-user interface possible for your application.

CICS cannot provide this type of function directly, but it does provide a number of ways for a task to communicate with a workstation, so that you can use a software package tailored for your particular workstation in combination with CICS. One popular approach is to use one of these packages, executing on the PC, to build your screens and handle the interactions with your user—that is, to implement the “front end” of your application. This code can then communicate with the part of your application that does the actual processing—the “back end” or “business logic” part—executing under CICS on the host. Communication between the two parts of the application can be done in several ways, depending on what your workstation supports:

- You can use one of the SNA application-to-application protocols, such as APPC.
- You can use the CPI-C “sockets” interface (see Chapter 6, “Intercommunication considerations” on page 119).
- You can use CICS on the workstation and use CICS facilities to communicate, or even distribute the business logic between the host and the workstation. CICS runs on many of these platforms, including OS/2®, AIX®, OS/400®, and others.

When you do this, you can execute specific commands on the host (file operations, for example), or whole programs, or whole tasks. Executing

commands remotely is called **function shipping**, executing a program remotely is called a **distributed program link**, and executing the whole task remotely is called **transaction routing**. See the *CICS Intercommunication Guide* for a full discussion of the possibilities, and the *CICS Distributed Transaction Programming Guide* for implementation details.

- You can use the terminal in emulation mode, a technique explained in “PCs as 3270s.”

If some of your users have 3270s or other nonprogrammable terminals, on the other hand, or if you are modifying an existing 3270 application, you need to use one of the CICS APIs for terminals. See “CICS APIs for terminals” on page 203 for information on which to base your choice.

PCs as 3270s

Although there is a different programming interface for a PC display, you can use PCs as “3270” terminals. Almost all PCs have programs available that emulate a 3270. These programs convert output in 3270 data stream format into the set of PC instructions that produces the same display on the screen, and similarly convert keyboard input into the form that would have come from a 3270 with the same screen contents.

Under an emulator, the PC display has essentially the same level of function as a real 3270. This limits your access to the more powerful PC hardware, although an emulator program often gives you a means to switch easily from its control to other programs that use the display in full function mode. Moreover, the hardware on a particular PC does not always permit exact duplication of 3270 function (the keyboard may be different, for example). Consequently, your PC may not always behave precisely as described in this chapter or in the *IBM 3270 Information Display System Data Stream Programmer's Reference* manual, although the differences are usually minor.

The 3270 buffer

Communication with a 3270 device occurs through its **character buffer**, which is a hardware storage mechanism similar to the memory in a processor. Output to the 3270 is sent to the buffer. The buffer, in turn, drives the display of a display terminal and the print mechanism of a printer terminal.

Conversely, keyboard input reaches the host through the buffer, as explained in “Input from a 3270 terminal” on page 221.

Each position on the screen corresponds to one in the buffer, and the contents of that buffer position determine what is displayed on the screen. When the screen is formatted in fields, the first position of each field is used to store certain display characteristics of the field and is not available to display data (it appears blank). In the original models of the 3270, this byte was sufficient to store all of the display characteristics. In later models, which have more types of display characteristics, the additional information is kept in an area of buffer storage not associated with a fixed position on the screen. There is more about display characteristics on page 212.

Writing to a 3270 terminal

To create a 3270 display, you send a stream of data that consists of:

- A write command (one byte)
- A write control character or **WCC** (one byte)
- Display data (variable number of bytes)

The WCC and display data are not always present; the write command determines whether a WCC follows and whether data may or must be present.

When you use BMS, CICS builds the entire data stream for you. The WCC is assembled from options in the SEND command, and the write command is selected from other SEND options and information in the RDO PROFILE of the transaction being executed. Display data is built from map or text data that you provide, which BMS translates into 3270 format for you.

When you use terminal control commands, CICS still supplies the write command, built from the same information. However, you provide the WCC and you must express the display data in 3270 format.

3270 write commands

Even though CICS supplies the write command, you need to know the possibilities, so that you can select the options that produce the one you want. There are five 3270 commands that send data or instructions to a terminal:

- Write
- Erase/write
- Erase/write alternate
- Erase all unprotected fields
- Write structured fields

The 3270 **write** command sends the data that follows it to the 3270 buffer, from which the screen (or printer) is driven. **Erase/write** and **erase/write alternate** also do this, but they erase the buffer first (that is, they set it entirely to null values). They also determine the buffer size (the number of rows and columns on the screen), if the terminal has a feature called **alternate screen size**.

Terminals with this feature have two sizes, **default size** and **alternate size**. The erase/write command causes the default size to be used in subsequent operations (until the next erase/write or erase/write alternate command), and erase/write alternate selects the alternate size, as the names suggest.

CICS uses the plain write command to send data unless you include the ERASE option on your SEND command. If you specify ERASE DEFAULT on your SEND, CICS uses erase/write instead (setting the screen to default size), and ERASE ALTERNATE causes CICS to use erase/write alternate (setting alternate size). If you specify ERASE without DEFAULT or ALTERNATE, CICS looks at the PROFILE definition associated with the transaction you are executing to decide whether to use erase/write or erase/write alternate.

The **erase unprotected to address** command causes a scan of the buffer for unprotected fields (these are defined more precisely in “3270 field attributes” on page 213). Any such fields that are found are set to nulls. This selective erasing is useful in data entry operations, as explained in “The EXEC CICS SEND

CONTROL command” on page 249. No WCC or data follows this command; you send only the command.

Write structured fields causes the data that follows to be interpreted as 3270 structured fields. **Structured fields** are required for some of the advanced function features of the 3270. They are not covered here, but you can write them with terminal control SEND commands containing the STRFIELD option. See the *IBM CICS/DOS/VS 3270 Data Stream Device Guide* if you wish to do this.

Write control character

The byte that follows a 3270 write, erase/write or erase/write alternate command is the **write control character** or **WCC**. The WCC tells the 3270 whether or not to:

- Sound the audible alarm
- Unlock the keyboard
- Turn off the modified data tags
- Begin printing (if terminal is a printer)
- Reset structured fields
- Reset inbound reply mode

In BMS, CICS creates the WCC from the ALARM, FREEKB, FRSET and PRINT options on your SEND MAP command. If you use terminal control commands, you can specify your WCC explicitly, using the CTLCHAR option. If you do not, CICS will generate one that unlocks the keyboard and turns off the modified data tags (these are explained shortly, in “Modification” on page 214).

3270 display data: defining 3270 fields

Display data consists of a combination of characters to be displayed and instructions to the device on how and where to display them. Under ordinary circumstances, this data consists of a series of field definitions, although it is possible to write the screen without defining fields, as explained in “Unformatted mode” on page 224.

After a write command that erases, you need to define every field on the screen. Thereafter, you can use a plain write command and send only the fields you want to change.

To define a field, you need to tell the 3270:

- How to display it
- What its contents are
- Where it goes on the screen (that is, its starting position in the buffer)

Display characteristics

Each field on the screen has a set of display characteristics, called **attributes**. Attributes tell the 3270 **how** to display a field, and you need to understand what the possibilities are whether you are using BMS or terminal control commands. Attributes fall into two categories:

Field attributes

These include:

- Protection (whether the operator can modify the field or not)
- Modification (whether the operator **did** modify the field)

- Display intensity

All 3270s support field attributes; “3270 field attributes” explains your choices for them.

Field attributes are stored in the first character position of a field. This byte takes up a position on the screen and not only stores the field attributes, but marks the beginning of the field. The field continues up to the next attributes byte (that is, to the beginning of the next field). If the next field does not start on the same line, the current one wraps from the end of the current line to the beginning of the next line until another field is encountered. A field that has not ended by the last line returns to the first.

Extended field attributes

(Usually shortened to **extended attributes**). These are not present on all models of the 3270. Consequently, you need to be aware of which ones are available when you design your end-user interface. Extended attributes include special forms of highlighting and outlining, the ability to use multiple symbol sets and provision for double-byte character sets. Table 12 on page 215 lists the seven extended attributes and the values they can take.

3270 field attributes

As noted above, the field attributes byte holds the protection, modification and display intensity attributes of a field. Your choices for each of these attributes are described here using the terms that BMS uses in defining formats. If you use terminal control commands, you need to set the corresponding bits in the attributes byte to reflect the value you choose. (See the *IBM 3270 Information Display System Data Stream Programmer's Reference* manual for the bit assignments. See also “Attribute value definitions: DFHBMSCA” on page 246 for help from CICS in this area.)

Protection

There are four choices for the protection attribute, using up two bit positions in the attributes byte. They are:

Unprotected

The operator can enter any data character into an unprotected field.

Numeric-only

The effect of this designation depends on the keyboard type of the terminal. On a data entry keyboard, a numeric shift occurs, so that the operator can key numbers without shifting. On keyboards equipped with the “numeric lock” special feature, the keyboard locks if the operator uses any key except one of the digits 0 through 9, a period (decimal point), a dash (minus sign) or the DUP key. This prevents the operator from keying alphabetic data into the field, although the receiving program must still inspect the entry to ensure that it is a number of the form it expects. Without the numeric lock feature, numeric-only allows any data into the field.

Protected

The operator cannot key into a protected field. Attempting to do so locks the keyboard.

Autoskip The operator cannot key into an autoskip field either, but the cursor behaves differently. (The cursor indicates where the operator's next keystroke will go; for more information about this, see "Input from a 3270 terminal" on page 221.) Whenever the cursor is being advanced to a new field (either because the previous field filled or because a field advance key was used), the cursor skips over any autoskip fields in its path and goes to the first field that is either unprotected or numeric-only.

Modification

The second item of information in the field attributes byte occupies only a single bit, called the **modified data tag** or **MDT**. The MDT indicates whether the field has been modified or not. The hardware turns on this bit automatically whenever the operator makes any change to the field contents. The MDT bit is very important because, for the read command that CICS normally uses, it determines whether the field is included in the inbound data or not. If the bit is on (that is, the field was changed), the 3270 sends the field; if not, the field is not sent.

You can also turn the MDT on by program, when you send a field to the screen. Using this feature ensures that a field is returned on a read, even if the operator cannot or does not change it. The FRSET option on BMS SEND commands allows you to turn off the tags for all the fields on the screen by program; you cannot turn off individual tags by program. If you are using terminal control commands, you turn on a bit in the WCC to turn off an individual tag.

Intensity

The third characteristic stored in the attributes byte is the display intensity of the field. There are three mutually exclusive choices:

- Normal intensity** The field is displayed at normal brightness for the device.
- Bright** The field is displayed at higher than normal intensity, so that it appears highlighted.
- Nondisplay** The field is not displayed at all. The field may contain data in the buffer, and the operator can key into it (provided it is not protected or autoskip), but the data is not visible on the screen.

Two bits are used for display intensity, which allows one more value to be expressed than the three listed above. For terminals that have either of the associated special hardware features, these same two bits are used to determine whether a field is light-pen detectable or cursor selectable. Because there are only two bits, not all combinations of intensity and selectability are possible. The compromise is that bright fields are always detectable, nondisplay fields are never detectable, and normal intensity fields may be either. "Cursor- and pen-detectable fields" on page 316 contains more information about these features.

Base color

Some terminals support **base color** without, or in addition to, the **extended colors** included in the extended attributes. There is a mode switch on the front of such a terminal, allowing the operator to select base or default color. Default color shows characters in green unless field attributes specify bright intensity, in which case they are white. In base color mode, the protection and intensity bits are used in combination to select among four colors: normally white, red, blue and green; the protection bits retain their protection functions as well as determining color.

Extended attributes

In addition to the field attributes just described, some 3270 terminals have extended attributes as well. Table 12 lists the types of extended attributes in the first column and the possible values for each type in the second column.

Attribute type	Values
Extended color	Blue, red, pink, green, turquoise, yellow, neutral
Extended highlighting	Blinking, reverse video, underscoring
Field outlining	Lines over, under, left and right, in any combination
Background transparency	Background transparent, background opaque
Field validation	Field must be entered; field must be filled; field triggers input
Programmed symbol sets	Number identifying the symbol set Note: The control unit associated with a terminal contains a default symbol set and can store up to five additional ones. To use one of these others, you need to load the symbol set into the controller prior to use. You can use a terminal control SEND command to do this.
SO/SI creation	Shift characters indicating double-byte characters may be present; shift characters are not present

The *IBM 3270 Information Display System Data Stream Programmer's Reference* manual contains details about extended attributes and explains how default values are determined. You can use EXEC CICS ASSIGN and INQUIRE commands to determine which extended attributes your particular terminal has. These commands are described in "Finding out about your terminal" on page 330.

Some models of the 3270 also allow you to assign extended attribute values to individual characters within a field that are different from the value for the field as a whole. Generally, you need to use terminal control commands to do this, because BMS does not make explicit provision for character attributes. However, you can insert the control sequences for character attributes in text output under BMS, as explained in "How BMS breaks text into lines" on page 291. "The set attribute order" on page 218 describes the format of such a sequence.

Orders in the data stream

The next several sections tell you how to format outbound data to express the attributes, position, and contents of a field. You need to know this information if you are writing to a 3270 using terminal control commands. If you are using BMS, all this is done for you, and you can move on to "Input from a 3270 terminal" on page 221.

When you define a field in the 3270 data stream, you begin with a **start field (SF)** or a **start field extended (SFE)** order. **Orders** are instructions to the 3270. They tell it how to load its buffer. They are one byte long and usually are followed by data in a format specific to the order.

The start field order

The SF order is supported on all models and lets you specify the field attributes and the display contents of a field, but not extended attributes. To define a field with SF, you insert a sequence in the data stream as in Figure 43.




Figure 43. Field definition using SF order

If you need to specify extended attributes, and your terminal supports them, you use the start field extended order instead. SFE requires a different format, because of the more complex attribute information. Extended attributes are expressed as byte pairs. The first byte is a code indicating which type of attribute is being defined, and the second byte is the value for that attribute. The field attributes (the first byte in an SF sequence) are treated collectively as an additional attribute type and also expressed as a byte pair. Immediately after the SFE order, you give a one-byte count of the attribute pairs, then the attribute pairs, and finally the display data. The whole sequence is shown in Figure 44.




Figure 44. Field definition using SFE order

The modify field order

When a field is on the screen, you can change it with a command almost identical in format to SFE, called **modify field (MF)**. The only differences from SFE are:

- The field must already exist.
- The command code is X'2C' instead of X'29'.

- You send only the attributes you want to change from their current values, and you send display data only if you want to change it.
- A null value sets an attribute back to its default for your particular terminal (you accomplish the same thing in an SFE order by omitting the attribute).

For example, to change the “menu” field of earlier examples back to the default color for the terminal and underscore it, you would need the sequence in Figure 45.




Figure 45. Changing field attributes within an MF order

The set buffer address order

The SF and SFE orders place the field they define at the current position in the buffer, and MF modifies the field at this position. Unless the field follows the last character sent (that is, begins in the current buffer position), you need to precede these orders with a **set buffer address (SBA)** order to indicate where you want to place or change a field. To do this, you send an SBA order followed by a two-byte address, as in Figure 46.



Figure 46. SBA sequence

The address in the figure is a “12-bit” address for position 112 (X'70'), which is row 2, column 33 on an 80-column screen. There are two other addressing schemes used: “14-bit” and “16-bit”. Buffer positions are numbered sequentially in all of them, but in 12- and 14-bit addressing, not all the bits in the address are used, so that they do not appear sequential. (The X'70' (B'1110000') in the figure appears as B'110000' in the low-order six bits of the rightmost byte of the address and B'000001' in the low-order six bits of the left byte.) The *IBM 3270 Information Display System Data Stream Programmer's Reference* manual explains how to form addresses.

After an SF, SFE, or MF order, the current buffer address points to the first position in the buffer you did not fill—right after your data, if any, or after the field attributes byte if none.

The set attribute order

To set the attributes of a single character position, you use a **set attribute (SA)** order for each attribute you want to specify. For example, to make a character blink, you need the sequence in Figure 47.




Figure 47. SA sequence to make a character blink

The attributes you specify with SA orders are assigned to the current buffer position, in the same way that field definitions are placed at the current buffer position, so you generally need to precede your SAs with SBA sequences.

Outbound data stream sample

This section shows you an annotated example of the data stream required to paint a particular 3270 screen, to reinforce the explanation of how the data stream is built.

Figure 48 shows an example screen that is part of an application that keeps track of cars used by the employees at a work site, and is used to record a new car. The only inputs are the employee identification number, the license plate (tag) number, and, if the car is from out-of-state, the licensing state.

Car Record

Employee No: _____ Tag No: _____ State: ____

Figure 48. Example of a data-entry screen

Note: This is an unrealistically simple screen, designed to keep the explanation manageably short. It does not conform to generally accepted standards of screen design, so do not use it as a model.

There are eight fields on this screen:

1. Screen title, “Car Record”, on line 1, column 26
2. Label field, “Employee No:” (line 3, column 1), indicating what the operator is to enter into the next field
3. An input field for the employee number (line 3, column 14), six positions long
4. Label field, “Tag. No:”, at line 3, column 21
5. An input field (tag number) at line 3, column 31, eight positions long

6. Label field, "State:", at line 3, column 40
7. An input field (state), at line 3, column 49, two positions long
8. A field to mark the end of the previous (state) input field, at line 3, column 52

Table 13 shows the outbound data stream:

<i>Table 13 (Page 1 of 2). 3270 output data stream</i>		
Bytes	Contents	Notes
1	X'F5'	The 3270 command that starts the data stream, in this case erase/write.
2	X'C2'	WCC; this value unlocks the keyboard, but does not sound the alarm or reset the MDTs.
3	X'11'	SBA order to position first field at ...
4-5	X'40D6'	Address of line 1, column 23 on 24 by 80 screen, using 12-bit addressing.
6	X'1D'	SF order to begin first field definition.
7	X'F8'	Field attributes byte; this combination indicates a field that is autoskip and bright, with the MDT initially off.
8-17	'Car record'	Display contents of the field.
18-20	X'11C260'	SBA sequence to reset the current buffer position to line 3, column 1 for second field.
21	X'1D'	SF order for second field.
22	X'F0'	Field attributes byte: autoskip, normal intensity, MDT off.
23-34	'Employee No:'	Display contents of field.
35	X'29'	SFE order to start fourth field. SFE is required, instead of SF, because you need to specify extended attributes. This field starts immediately after the previous one left off, so you do not have to precede it with an SBA sequence.
36	X'02'	Count of attribute types that are specified (two here: field outlining and field attributes).
37	X'41'	Code indicating attribute type of extended highlighting.
38	X'F4'	Extended highlighting value indicating underscoring.
39	X'C0'	Code indicating attribute type of field attributes.
40	X'50'	Field attributes value for numeric-only, normal intensity, MDT off. Any initial data for this field would appear next, but there is none.
41	X'13'	Insert cursor (IC) order, which tells the 3270 to place the cursor at the current buffer position. We want it at the start of the first field that the operator has to fill in, which is the current buffer position.
42-44	X'11C2F4'	SBA sequence to position to line 3, column 21, to leave the six positions required for an employee number. The beginning of the "Tag No" label field marks the end of the employee number input field, so that the user is aware immediately if he tries to key too long a number.

<i>Table 13 (Page 2 of 2). 3270 output data stream</i>		
Bytes	Contents	Notes
45	X'1D'	SF order to start field.
46	X'F0'	Field attributes byte: autoskip, normal intensity, MDT off.
47-55	' Tag No:'	Display data. We attach two leading blanks to the label for more space between the fields. (We could have used a separate field, but this is easier for only a few characters.)
56	X'29'	SFE (the next field is another input field, where we want field outlining, so we use SFE again).
57	X'02'	Count of attribute types.
58-59	X'41F4'	Code for extended highlighting with value of underscoring.
60-61	X'C040'	Code for field attributes and attributes of unprotected, normal intensity, MDT off.
62-64	X'11C3C7'	SBA sequence to reposition to line 3, column 40, leaving eight positions for the tag.
65	X'1D'	SF to start field.
66	X'F0'	Field attributes byte: autoskip, normal intensity, MDT off.
67-74	' State:'	Field data (two leading blanks again for spacing).
75-80	X'290241F4C040'	SFE order and attribute specifications for state input field (attributes are identical to those for tag input field).
81-82	X'0000'	The (initial) contents of the state field. We could have omitted this value as we did for other input fields, but we would need an SBA sequence to move the current buffer position to the end of the field, and this is shorter.
83	X'1D'	SF. The last field indicates the end of the previous one, so that the user does not attempt to key more than two characters for the state code. It has no initial data, just an attributes byte. This kind of field is sometimes called a "stopper" field.
84	X'F0'	Field attributes byte: autoskip, normal intensity, MDT off.

Note: If you use terminal control commands and build your own data stream, the data you provide in the FROM parameter of your EXEC CICS SEND command starts at byte 3 in the table above; CICS supplies the write command and the WCC from options on your EXEC CICS SEND command.

Input from a 3270 terminal

As explained earlier, keyboard input reaches the host through the buffer. There are many different keyboard arrangements available for 3270 terminals, but in any arrangement, a key falls into one of three categories:

- Data key
- Keyboard control key
- Attention key

The data keys include all the familiar letters, numbers, punctuation marks and special characters. Depressing a data key simply changes the content of the buffer (and therefore the screen) at the point indicated by the **cursor**. The cursor is a visible pointer to the position on the screen (that is, in the buffer) where the next data keystroke will be stored. As the operator keys data, the cursor advances to the next position on the screen, skipping over fields defined with the autoskip attribute on screens that have been formatted.

Keyboard control keys move the cursor to a new position, erase fields or individual buffer positions, cause characters to be inserted, or otherwise change where or how the keyboard modifies the buffer.

Attention keys

None of the keys in either of these two groups causes any interaction with the host; they are handled entirely by the device and its control unit. An attention key, on the other hand, signals that the buffer is ready for transmission to the host. If the host has issued a read to the terminal, the usual situation in CICS, transmission occurs at this time.

There are five types of attention key:

- ENTER
- PF (program function) key
- CLEAR
- PA (program attention) key
- CNCL (cancel key, present only on some keyboard models)

In addition to pressing an attention key, there are other operator actions that cause transmission:

- Using an identification card reader
- Using a magnetic slot reader or hand scanner
- Selecting an attention field with a light pen or the cursor select key
- Moving the cursor out of a trigger field

Trigger field capability is provided with extended attributes on some terminal models, but all the other actions listed above require special hardware, and in most cases the screen (buffer) must be set up appropriately beforehand. We talk about these features in “BMS support for other special hardware” on page 314. For this chapter, we concentrate on standard features.

The AID

The 3270 identifies the key that causes transmission by an encoded value in the first byte of the inbound data stream. This value is called the **attention identifier** or **AID**.

Ordinarily, the key that the terminal operator chooses to transmit data is dictated by the application designer. The designer assigns specific meanings to the various attention keys, and the user must know these meanings in order to use the application. (Often, there are only a few such keys in use: ENTER for normal inputs, one PF key to exit from control of the application, another to cancel a partially completed transaction sequence, for example. Where there are a number of choices, you may want to list the key definitions on the screen, so that the user does not have to memorize them.)

There is an important distinction between two groups of attention keys, which the application designer must keep in mind. The ENTER and PF keys transmit data from the buffer when the host issues a “read modified” command, the command normally used by CICS. CLEAR, CNCL, and the PA keys do not, although you do get the AID (that is, the identity of the key that was used). These are called the **short read** keys. They are useful for conveying simple requests, such as “cancel”, but not for those that require accompanying data. In practice, many designers use PF keys even for the nondata requests, and discard any accompanying data.

Note: The CLEAR key has the additional effect of setting the entire buffer to nulls, so that there is literally no data to send. CLEAR also sets the screen size to the default value, if the terminal has the alternate screen size feature, and it puts the screen into unformatted mode, as explained in “Unformatted mode” on page 224.

Reading from a 3270 terminal

There are two basic read commands for the 3270:

- **Read buffer**
- **Read modified**

For either command, the inbound data stream starts with a three-byte **read header** consisting of:

- Attention identifier (AID), one byte
- Cursor address, two bytes

As noted in the previous section, the AID indicates which action or attention key causes transmission. The cursor address indicates where the cursor was at the time of transmission. CICS stores this information in the EIB, at EIB AID and EIB POSN, on the completion of any RECEIVE command.

The read buffer command brings in the entire buffer following the read header, and the receiving program is responsible for extracting the information it wants based on position. It is intended primarily for diagnostic and other special purposes, and CICS uses it in executing a RECEIVE command only if the BUFFER option is specified. CICS never uses read buffer to read unsolicited terminal input, so the BUFFER option cannot be used on the first RECEIVE of a transaction initiated in this way.

With read modified, the command that CICS normally uses, much less data is transmitted. For the short read keys (CLEAR, CNCL and PAs), only the read

header comes in. For other attention keys (ENTER and PFs), the fields on the screen that were changed (those with the MDT on, to be precise) follow the read header. We describe the format in the next section. When transmission occurs because of a trigger field, light pen detect, or cursor select, the amount and format of the information is slightly different; these special formats are described in “BMS support for other special hardware” on page 314. Input from a **program attention** key on an SCS printer is also an exception; see “SCS input” on page 347 for a description of that datastream.

Inbound field format

The next several sections describe the format in which the 3270 transmits data, which you need to understand if you are using terminal control commands. If you are using BMS, you can skip to “Unformatted mode” on page 224, because BMS translates the input for you.

Each modified field comes in as follows:

- SBA order
- Two-byte address of the first **data** position of field
- SF order
- Field contents

Only the nonnull characters in the field are transmitted; nulls are skipped, wherever they appear. Thus if an entry does not fill the field, and the field was initially nulls, only the characters keyed are transmitted, reducing the length of the inbound data. Nulls (X'00') are not the same as blanks (X'40'), even though they are indistinguishable on the screen. Blanks get transmitted, and hence you normally initialize fields to nulls rather than to blanks, to minimize transmission.

A 3270 read command can specify that the terminal should return the attribute values along with the field contents, but CICS does not use this option. Consequently, the buffer address is the location of the first byte of field data, not the preceding attributes byte (as it is in the corresponding outbound data stream).

Note: Special features of the 3270 for input, such as the cursor select key, trigger fields, magnetic slot readers, and so on, produce different input formats. See “Field selection features” on page 315 for details.

Input example

To illustrate an inbound data stream, we assume that an operator using the screen shown in Figure 48 on page 218 did the following:

- Put “123456” in the employee identifier field
- Put “ABC987” in the tag number
- Pressed ENTER, without filling in the state field

Here is the resulting inbound data stream:

Bytes	Contents	Notes
1	X'7D'	AID, in this case the ENTER key.
2-3	X'C3C5'	Cursor address: line 3, column 38, where the operator left it after the last data keystroke.
4	X'11'	SBA, indicating that a buffer address will follow.
5-6	X'C26E'	Address of line 3, column 15, which is the starting position of the field to follow.
7-12	'123456'	Input, the employee number entered by the operator.
13-15	X'11C3D1'	SBA sequence indicating a buffer address of line 3, column 32.
16	X'1D'	SF, indicating another input field follows.
17-22	'ABC987'	Input field: plate number. Notice that only six characters came in from a field that was eight long, because operator left the remaining positions null.

Note that the third input field (the state code) does not appear in the input data stream. This is because its MDT did not get turned on; it was set off initially, and the operator did not turn it on by keying into the field. Note also that no SF is required at Byte 7 because CICS normally issues a Read Modified All.

Unformatted mode

Even though the high function of the 3270 revolves around its field structure, it is possible to use the 3270 without fields, in what is called **unformatted mode**. In this mode, there are no fields defined, and the entire screen (buffer) behaves as a single string of data, inbound and outbound, much like earlier, simpler terminals.

When you write in unformatted mode, you define no fields in your data, although you can include SBA orders to direct the data to a particular positions on the screen. Data that precedes any SBA order is written starting at the current position of the cursor. (If you use an erase or write command, the cursor is automatically set to zero, at the upper left corner of the screen.)

When you read an unformatted screen, the first three bytes are the read header (the AID and the cursor address), just as when you read a formatted screen. The remaining bytes are the contents of the entire buffer, starting at position zero. There are no SBA or SF orders present, because there are no fields. If the read command was read modified, the nulls are suppressed, and therefore it is not always possible to determine exactly where on the screen the input data was located.

You cannot use a BMS RECEIVE MAP command to read an unformatted screen. BMS raises the MAPFAIL condition on detecting unformatted input, as explained in "MAPFAIL and other exceptional conditions" on page 265. You can read unformatted data only with a terminal control RECEIVE command in CICS.

Note: The CLEAR key puts the screen into unformatted mode, because it sets the buffer to nulls, thereby erasing all the attribute bytes that demarcate fields.

Chapter 18. Basic mapping support

This chapter describes basic mapping support (BMS), and contains the following topics:

- “Introduction to BMS”
- “BMS support levels” on page 226
- “A BMS output example” on page 228
- “Creating the map” on page 231
- “Attribute value definitions: DFHBMSCA” on page 246
- “Assembling the map” on page 237
- “Complex fields” on page 239
- “Sending mapped output: basics” on page 241
- “Initializing the output map” on page 244
- “Moving the variable data to the map” on page 244
- “Setting the display characteristics” on page 244
- “Control options on the EXEC CICS SEND MAP command” on page 247
- “Options for merging the symbolic and physical maps” on page 248
- “Summary: what appears on the screen” on page 249
- “Using GDDM® and BMS” on page 252
- “Positioning the cursor” on page 252
- “Sending invalid data and other errors” on page 253
- “Receiving data from a display” on page 253
- “An input-output example” on page 254
- “The symbolic input map” on page 256
- “Programming simple mapped input” on page 257
- “The EXEC CICS RECEIVE MAP command” on page 257
- “An input-output example” on page 254
- “Getting storage for mapped input: INTO and SET” on page 258
- “Reading from a formatted screen: what comes in” on page 259
- “Other information from EXEC CICS RECEIVE MAP” on page 260
- “Support for non-3270 terminals” on page 267
- “The MAPPINGDEV facility.” on page 272
- “Sending mapped output: additional facilities” on page 276
- “Using GDDM® and BMS” on page 252
- “Page formation: the ACCUM option” on page 282
- “Floating maps: how BMS places maps using ACCUM” on page 282
- “Formatting text output” on page 288
- “Message routing: the ROUTE command” on page 294
- “Using SET” on page 302
- “Partition support” on page 304
- “Logical device components” on page 312
- “BMS support for other special hardware” on page 314

Introduction to BMS

Basic mapping support (BMS) is an application programming interface between CICS programs and terminal devices. As noted on page 203, BMS is one of two sets of commands for this purpose. The other one, terminal control, is described in Chapter 19, “Terminal control” on page 319.

For many applications, BMS has several advantages. First, BMS removes device dependencies from the application program. It interprets **device-independent** output commands and generates **device-dependent** data streams for specific terminals. It also transforms incoming device-dependent data into device-independent format. These features eliminate the need to learn complex device data streams. They also allow you to use the same program for a variety of devices, because BMS determines the device information from the terminal definition, not from the application program.

Second, BMS separates the design and preparation of formats from application logic, reducing the impact of one on the other. Both of these features make it easier to write new programs and to maintain existing code.

This chapter describes the services BMS provides, and how to use them. We start with the simplest situation and build from there in this order:

1. Example of sending a simple formatted screen
2. Defining screen and report formats
3. Sending formatted data (the basics: one format, one page)
4. Reading formatted data, including an example
5. Multi-format screens and multi-page messages
6. Routing messages to terminals other than your own
7. A variant on the process for text data
8. Support for special hardware features

We do not cover every detail of BMS here. See the *CICS Application Programming Reference* manual for details on the syntax and operation of BMS commands. If you are printing with BMS, read Chapter 20, “CICS support for printing” on page 341 too. And, because BMS is designed around the 3270 family of terminals, read Chapter 17, “The 3270 family of terminals” on page 207 if you are unfamiliar with this hardware and the associated terminology. Finally, some situations require terminal control commands. These are described in Chapter 19, “Terminal control” on page 319.

BMS support levels

There are three levels of BMS support: minimum, standard, and full. Most installations use full BMS. If yours does, you can use all the features we describe in this chapter and not concern yourself with levels. If your installation uses minimum or standard BMS, note the features that require levels beyond your's. They are summarized here, and they are noted again whenever a facility that is not in minimum BMS is covered. The *CICS Application Programming Reference* manual also classifies BMS commands and options by level.

Minimum BMS

Minimum BMS supports all the basic functions for 3270 terminals, including everything described in our example and in the discussion of simple mapped output and mapped input.

Note: Minimum BMS has a substantially shorter path length than standard or full BMS. It is included in the larger versions and invoked as a kind of “fast path” on commands do not require function beyond what it provides. Specifically, it is used for SEND MAP and SEND CONTROL commands without the ACCUM, PAGING, SET, OUTPARTN, ACTPARTN, LDC, MSR, or REQID options, and for RECEIVE

MAP commands, when your principal facility is a 3270 display or printer whose definition does not include outboard formatting. You can tell whether a particular BMS request used the fast path by looking at the CICS trace table. When fast path is used, the trace table contains duplicate entries for the BMS entry and exit code.

Standard BMS

Standard BMS adds:

- Support for terminals other than 3270s
- Text output commands
- Support for special hardware features: partitions, logical devices codes, magnetic slot readers, outboard formatting, and so on
- Additional options on the SEND command: NLEOM and FMHPARM

Standard BMS supports these terminals:

- Sequential terminals (composed of card readers, line printers, tape or disk)
- TWX Model 33/35
- 1050
- 2740-1 (no buffer receive), 2740-2, 2741
- 2770
- 2780
- 2980, models 1, 2 and 4
- 3270
- 3600 (3601) LU
- 3650 (3653 and 3270 host conversational LUs)
- 3650 interpreter LU
- 3767/3770 interactive LU
- 3770 batch LU
- 3780
- LU type 4

Full BMS

Full BMS is required for:

- Sending BMS output other than directly to your own terminal (the SET and PAGING options, and BMS routing)
- Messages built cumulatively, with multiple BMS SEND commands (the ACCUM and PAGING options)

Some CICS platforms do not support all the features of BMS. Table 15 on page 228 shows the approximate level of support in each, for general guidance. However, there are differences among platforms even at the same level, usually imposed by differences in execution environment. These are described in detail, by function, in the *CICS Family: API Structure* manual. If your application may eventually move to another platform, or there is a chance that the end-user interface part of it may get distributed to one, consult that manual.

Platform	BMS support
CICS OS/2™	Minimum plus SEND TEXT of standard
CICS/400®	Minimum plus SEND TEXT of standard
CICS/6000®	Minimum plus SEND TEXT of standard
CICS/VSE®	Full
CICS/ESA®	Full

A BMS output example

To create a formatted screen, BMS takes a list of data items from a program and displays them on the screen (or printed page) according to a predefined format. It merges variable data supplied by the program with constant data in the format (titles, labels for variable fields, default values for these fields). It builds the data stream for the terminal to which you are writing, to show this merged data in the designated screen positions, with the proper attributes (color, highlighting, and so on). You do not have to know anything about the data stream, and you do not need to know much about the format to write the required CICS commands.

Note: For simplicity, this chapter is mainly concerned with display screens, but most of it applies equally to printers. Chapter 20, “CICS support for printing” on page 341 discusses differences between displays and printers and covers additional considerations that apply to printing. Furthermore, the examples and discussion assume a standard 3270 terminal because BMS is designed to support the features of the 3270. Other terminals are discussed in “Support for non-3270 terminals” on page 267.

You define the formats, called **maps**, separately from the programs that use them. This allows you to reposition fields, change their attributes, and change the constant text without modifying your programs. If you add or remove variable data, of course, you need to change the programs that use the affected fields.

The basics of how this works are explained by an atypically simple example. In real life, requirements are always more complex, but this will give you the essentials without too much confusing detail. There are more realistic and complete BMS examples among the CICS sample applications. These programs are included in source form in the VSE/ESA sublibrary, PRD1.BASE, and are described in the *CICS Sample Applications Guide*.

This example assumes that you need to write the code for a transaction used in a department store that checks a customer’s balance before a charge sale is completed. The transaction is called a “quick check”, because all it does is check that the customer’s account is open and that the current purchase is permissible, given the state of the account. The program for the **output** part of this transaction gets an account number as input, and produces the screen shown in Figure 49 on page 229 in response:

```
QCK                                Quick Customer Account Check
Account:    0000005
Name:      Thompson             Chris
Max charge: $500.00
```

Figure 49. Normal “quick check” output screen

The program uses the input account number to retrieve the customer’s record from the account file. From the information in this record, it fills in the account number and customer name in the map, and computes the maximum charge allowed from the credit limit, outstanding balance, and purchases posted after the last billing period. If the amount comes out negative, you are supposed to show a value of zero and add an explanatory message. You also need to alert the clerk if the charge card is listed as lost, stolen or canceled with a message as shown in Figure 50.

```
QCK                                Quick Customer Account Check
Account:    0000005
Name:      Thompson             Chris
Max charge:  $0.00
STOLEN CARD - SECURITY NOTIFIED
```

Figure 50. “Quick check” output screen with warning message

This message is to be highlighted, to draw the clerk’s attention to it.

The first thing you must do is define the screen. We explain how to do so for this particular map in “Creating the map” on page 231. For the moment, however, let us assume that one of the outputs of this process is a data structure like the one in Figure 51. (We show the COBOL-coded version of the structure, because we are using COBOL to code our examples. However, BMS will produce the structure in any language that CICS supports.) The map creation process stores this source code in a library from which you copy it into your program.

```

01 QCKMAPO.
   02 FILLER PIC X(12).
   02 FILLER PICTURE X(2).
   02 ACCTNOA PICTURE X.
   02 ACCTNOO PIC X(7).
   02 FILLER PICTURE X(2).
   02 SURNAMEA PICTURE X.
   02 SURNAMEO PIC X(15).
   02 FILLER PICTURE X(2).
   02 FNAMEA PICTURE X.
   02 FNAMEO PIC X(10).
   02 FILLER PICTURE X(2).
   02 CHGA PICTURE X.
   02 CHGO PIC $,$$0.00
   02 FILLER PICTURE X(2).
   02 MSGA PICTURE X.
   02 MSGO PIC X(30).

```

Figure 51. Symbolic map for “quick check”

The data names in this structure come from the map definition. You assign names to the fields that the program may have to change in any way. For our example, this category includes the fields where you will display the account number, last name, first name, maximum charge, and explanatory message. It does not include any of the field labels or screen titles that never change, such as “Quick Customer Account Check” and “Account”.

Each field that you name on the screen generates several fields in the data structure, which are distinguished by a one-character suffix added to the name you assigned in the map. Two appear here: the “A” suffix for the field attributes byte and the “O” suffix for the output data. If we were creating a map to use special device features like color and highlighting, or were using the map for input as well as output, there would be many more. We tell you about these other fields in “Setting the display characteristics” on page 244 and “Receiving data from a display” on page 253.

The key fields for this particular exercise are the ones suffixed with “O”. These are where you put the data that you want displayed on the screen. You use the “A” subfields if you want to change how the data is displayed. In our example, we use MSGA to highlight the message if our customer is using a dubious card.

Here is an outline of the code that is needed for the example. You have to copy in the data structure (Figure 51) produced by creating the map, and the COPY QCKSET statement in the third line does this. (Ordinarily, you would use a copy statement for the account record format too. We show it partly expanded here so that you can see its contents.)

```

WORKING-STORAGE SECTION.
C COPY IN SYMBOLIC MAP STRUCTURE.
01 COPY QCKSET.
01 ACCTFILE-RECORD.
    02 ACCTFILE-ACCTNO PIC S9(7).
    02 ACCTFILE-SURNAME PIC X(15).
    02 ACCTFILE-FNAME PIC X(10).
    02 ACCTFILE-CREDIT-LIM PIC S9(7) COMP-3.
    02 ACCTFILE-UNPAID-BAL PIC S9(7) COMP-3.
    02 ACCTFILE-CUR-CHGS PIC S9(7) COMP-3.
    02 ACCTFILE-WARNCODE PIC X.
:
PROCEDURE DIVISION.
:
EXEC CICS READ FILE (ACCT) INTO (ACCTFILE-RECORD) RIDFLD (CKNO)
... END-EXEC.
MOVE ACCTFILE-ACCTNO TO ACCTNOO.
MOVE ACCTFILE-SURNAME TO SURNAMEO.
MOVE ACCTFILE-FNAME TO FNAMEO.
COMPUTE CHGO = ACCTFILE-CREDIT-LIM - ACCTFILE-UNPAID-BAL
- ACCTFILE-CUR-CHGS.
IF CHGO < ZERO, MOVE ZERO TO CHGO
MOVE 'OVER CHARGE LIMIT' TO MSGO.
IF ACCTFILE-WARNCODE = 'S', MOVE DFHMBRY TO MSGA
MOVE 'STOLEN CARD - SECURITY NOTIFIED' TO MSGO
EXEC CICS LINK PROGRAM('NTFYCOPS') END-EXEC.
EXEC CICS SEND MAP ('QCKMAP') MAPSET ('QCKSET') END-EXEC.
EXEC CICS RETURN END-EXEC.

```

Creating the map

BMS provides three assembler language macro instructions (**macros**) for defining maps. This method of map definition is still widely used, and we are about to explain how to do it. However, there are also other products for creating maps that exploit the facilities of the display terminal to make the map creation process easier. They produce the same outputs as the BMS macros, generally with less programmer effort.

One of these is the Screen Definition Facility II for VSE. SDF II allows you to build your screens directly from a display station, testing the appearance and usability as you go. You can find out more about SDF II in the *Screen Definition Facility II for VSE: General Information* manual.

The three assembler macros used to define BMS maps are:

- **DFHMDF**—defines an individual field on a screen or page.
- **DFHMDI**—defines a single map as a collection of fields.
- **DFHMSD**—groups single maps into a map set.

The explanation of this process begins by telling you how to define individual fields. Then we explain how to go from the fields to a complete map, and from a map to a map set (the assembly unit). BMS is designed principally for 3270-type terminals, although it supports nearly all types. See Chapter 17, “The 3270 family of terminals” on page 207 for information on 3270 terms.

You should design the layout of your screen before you attempt to code any macros. After you have done that, you define each field on the screen (page) with a DFHMDF macro. In it, you indicate:

- The position of the field on the screen
- The length of the field
- The default contents (unless you always intend to provide them in the program)
- The **field** display attributes, governing whether and what the operator can key into the field, whether the cursor will stop there, the intensity of the characters, and the initial state of the modified data tag
- For some terminals, **extended** display attributes, such as color, underlining, highlighting
- The name by which you refer to the field in your program, if you ever modify its contents or attributes

Fields that are referenced by the application must be allocated field names. The length of the field name and the characters that may be used to form field names must conform to the following rules. (Note that these rules apply to currently-supported compilers and assemblers.)

The characters used must be valid for names of assembler ordinary symbols. This character set consists of the alphabetic characters A - Z (upper or lower case), \$, #, @, numeric digits 0 - 9, and the underscore (_) character for High Level Assembler.

There is one exception to this rule. The hyphen (-) character may be used in field names provided that:

- The mapset is only used by application programs written in COBOL.
- The mapset is generated using the High Level Assembler.

The first character of the field name must be alphabetic, but the other characters can be any from the character set described above.

In addition, the characters used in field names must conform to the character set supported by the programming language of the application using the map. For example, if the application language is COBOL, you cannot use the @ character. You should refer to the appropriate Language Reference manual for information about these character sets.

For High Level Assembler, the DFHMDF macro allows the length of field names to be from one through 30 characters. DFHMDF derives additional variable names by appending one of several additional characters to the defined name to generate a symbolic description map. These derived names may therefore be up to 31 characters in length. The assembler, PL/1, and C languages all support variable names of at least 31 characters. However the COBOL language only allows up to 30 characters, which means that field names used in maps must not exceed 29 characters for COBOL applications. For example, the following field definition is valid for all languages except COBOL:

```
ThisIsAnExtremelyLongFieldName DFHMDF LENGTH=10,POS=(2,1)
```

and the following field definition is only valid for COBOL:

Must-Not-Exceed-29-Characters DFHMDF LENGTH=10,POS=(2,1)

Not all the options for field definition are described here; you can read about them in the *CICS Application Programming Reference* manual.

Figure 52 shows the field definitions for the the map we considered in Figure 50 on page 229.

```
DFHMDF POS=(1,1),LENGTH=3,ATTRB=(ASKIP,BRT),INITIAL='QCK'
DFHMDF POS=(1,26),LENGTH=28,ATTRB=(ASKIP,NORM),                X
INITIAL='Quick Customer Account Check'
DFHMDF POS=(3,1),LENGTH=8,ATTRB=(ASKIP,NORM),                  X
INITIAL='Account:'
ACCTNO DFHMDF POS=(3,13),LENGTH=7,ATTRB=(ASKIP,NORM)
DFHMDF POS=(4,1),LENGTH=5,ATTRB=(ASKIP,NORM),INITIAL='Name:'
SURNAME DFHMDF POS=(4,13),LENGTH=15,ATTRB=(ASKIP,NORM)
FNAME DFHMDF POS=(4,30),LENGTH=10,ATTRB=(ASKIP,NORM)
DFHMDF POS=(5,1),LENGTH=11,ATTRB=(ASKIP,NORM),                X
INITIAL='Max charge:'
CHG DFHMDF POS=(5,13),ATTRB=(ASKIP,NORM),PICOUT='$,$$.00'
MSG DFHMDF LENGTH=20,POS=(7,1),ATTRB=(ASKIP,NORM)
```

Figure 52. BMS map definitions

1. The **POS** (position) parameter indicates the row and column position of the field, relative to the upper left corner of the map, position (1,1). It must be present. Remember that every field begins with a field attributes byte; POS defines the location of this byte; the contents of the field follow immediately to the right.
2. The **LENGTH** option tells how many characters long the field is. The length does **not** include the attributes byte, so each field occupies one more column than its LENGTH value. In the case of the first field in our map, for example, the attributes byte is in row 1, column 1, and the display data is in columns 2-4. Fields can be up to 256 characters long and can wrap from one line to another. (Take care with fields that wrap if your map is smaller than your screen. See “Outside the map” on page 252 for further information.)
3. The **ATTRB** (attributes) option sets the **field attributes** of the field, which we discussed in “3270 field attributes” on page 213. It is not required; BMS uses a default value of (ASKIP, NORM)—autoskip protection, normal intensity, modified data tag off—if you omit it. There are other options for each of the extended attributes, none of which was used in this map; these are described in “Setting the display characteristics” on page 244.
4. The **INITIAL** value for the field is not required either. You use it for label and title fields that have a constant value, such as ‘QCK’, and to assign a default value to a field, so that the program does not always have to supply a value.
5. The **PICOUT** option on the definition of the field CHG tells BMS what sort of PICTURE clause to generate for the field. It lets you use the edit facilities of COBOL or PL/I directly, as you move data into the map. If you omit PICOUT, and also the numeric (NUM) attribute, BMS assumes character data. Figure 51 on page 230 shows the effects of the PICOUT option for CHG and, in the other fields, its absence. You can omit the LENGTH option if you use PICOUT, because BMS infers the length from the picture.

6. The **GRPNAME** and **OCCURS** options do not appear in our simple example, because they are for more complex problems. **GRPNAME** allows you to subdivide a map field within the program for processing, and **OCCURS** lets you define adjacent, like map fields so that you can treat them as an array in the program. These options are explained in “Complex fields” on page 239 after some further information about maps.

Defining the map: DFHMDI

After all the fields on your map are defined, you tell BMS that they form a single map by preceding them with a DFHMDI macro. This macro tells BMS:

- The name of the map
- The size, in rows and columns
- Where it appears on the screen (you can put several maps on one screen)
- Whether it uses 3270 extended display attributes and, if so, which ones
- The defaults for these extended attributes for fields where you have not assigned specific values on the DFHMDF macro
- Device controls associated with sending the map (such as whether to sound the alarm, unlock the keyboard)
- The type of device the map supports, if you intend to create multiple versions of the map for different types of devices (see “Device-dependent maps: map suffixes” on page 269)

The map name and size are the critical information on a DFHMDI macro but, for documentation purposes, specify your other options explicitly rather than letting them default. The DFHMDI macro for our example might be:

```
QCKMAP  DFHMDI  SIZE=(24,80),LINE=1,COLUMN=1,CTRL=ALARM
```

We have named the map QCKMAP. This is the identifier we use in SEND MAP commands. It is 24 lines long, 80 columns wide, and starts in the first column of the first line of the display. We have also indicated that we want to sound the alarm when the map is displayed.

Defining the map set: DFHMSD

You need one more macro to create a map: DFHMSD, which defines a map set. Maps are assembled in groups called map sets. Typically you group all the maps used by a single transaction or several related transactions. (We discuss reasons for grouping maps further in “Grouping maps into map sets” on page 239.) A map set need not contain more than one map, incidentally, and in our simple example, the map set consists of just the “quick check” map.

One DFHMSD macro is placed in front of all the map definitions in the map set. It gives:

- The name of the map set
- Whether you are using the maps for output, input, or both
- Defaults for map characteristics that you did not specify on the DFHMDI macros for the individual maps
- Defaults for extended attributes that you did not specify in either the field or map definitions

- Whether you are creating physical or symbolic maps in the current assembly (see “Physical and symbolic map sets” on page 238)
- The programming language of programs that use the maps
- Information about the storage that is used to build the maps

Here’s the DFHMSD macro we need at the beginning of our example:

```
QCKSET DFHMSD TYPE=MAP,STORAGE=AUTO,MODE=OUT,LANG=COBOL,TIOAPFX=YES
```

This map set definition tells BMS that the maps in it are used only for output, and that the programs using them are written in COBOL. It assigns the name QCKSET to the map set. TIOAPFX=YES causes inclusion of a 12-byte “prefix” field at the beginning of each symbolic map (you can see the effect in the second line in Figure 51 on page 230). You always need this filler in command language programs, so specify it explicitly, because the default is sometimes omission. We explain MAP and STORAGE in a moment, in “Sending mapped output: basics” on page 241.

You need another DFHMSD macro at the end of your map definitions, to tell the assembler that it has reached the end of last map in the map set:

```
DFHMSD TYPE=FINAL
```

Rules for writing BMS macros

Because a BMS macro is an assembler language statement, you have to follow assembler syntax rules. We do not try to explain those in full here; you can find them in the *High Level Assembler for MVS, VM and ESA: Language Reference* manual. Instead, we give you a set of rules that work, although they are more restrictive than the actual rules.

1. Start names in column 1. Map and map set names may be up to seven characters long. The maximum length for field names (the DFHMDF macro) depends on the programming language. BMS creates labels by adding one-character suffixes to your field names. These labels must not be longer than the target language allows, because they get copied into the program. Hence the limit for a map field name is 29 characters for COBOL, 30 for PL/I and High Level Assembler. For C, it is 30 if the map is copied into the program as an internal data object, and 6 if it is an external data object (see “Acquiring and defining storage for the maps” on page 242 for more information about copying the map).
2. Start the macro identifier in column 10, or leave one blank between it and the name if the name exceeds eight positions. For field definitions, the identifier is always DFHMDF; for map definitions, DFHMDI; and for the map set macros that begin and end the map set, DFHMSD.
3. The rest of the field description consists of keywords (like POS for the position parameter) followed by values. The *CICS Application Programming Reference* manual lists the possible keywords and tells you how to express the values. Sometimes a keyword does not have a value, but if it does, an equals sign (=) always separates the keyword from its value.
4. Leave one blank after your macro identifier and then start your keywords. They can appear in any order.
5. Separate keywords by one comma (no blanks), but do not put a comma after the last one.

6. Keywords can extend through column 71. If you need more space, stop after the comma that follows the last keyword that fits entirely on the line and resume in column 16 of the next line.
7. Initial values (the INITIAL, XINIT, and GINIT keywords) are exceptions to the rule, because they may not fit even if you start on a new line. Except when double-byte characters are involved, you can split them at any point after the first character of the initial value itself. When you split in this way, use all of the columns through 71 and continue in column 16 of the next line. Double-byte character set (DBCS) data is more complicated to express than ordinary single-byte (SBCS) data. See Step 12 if you have DBCS initial values.
8. Surround initial values by single quote marks. If you need a single quote **within** your text, use two successive single quotes (the assembler will remove the extra one). Ampersands also have special significance to the assembler, and you use the same technique: use two ampersands where you want one, and the assembler will remove the extra.
9. If you use more than one line for a macro, put a character (any one except a blank) in column 72 of all lines **except the last**.
10. If you want comments in your map, use comment lines between macros, not among the lines that make up a single macro. Comment lines have an asterisk in column 1 and a blank in column 72. Your comments can appear anywhere among columns 2-71.
11. Use upper case only, except for values for the INITIAL parameter and in comments.
12. **For initial values containing DBCS.** If you have initial data that is **entirely** DBCS, use the GINIT keyword for your data and specify the keyword PS=8 as well. If your data contains both DBCS and SBCS characters, that is, if it is **mixed**, use INITIAL and specify SOSI=YES. (We need to explain a third alternative, XINIT, because you may find it in code you are maintaining. Use GINIT and INITIAL if possible, however, because XINIT is more difficult to use and your data is not validated as completely. XINIT can be used for either pure or mixed DBCS. XINIT with PS=8 follows the rules for GINIT, and XINIT with SOSI=YES follows those for INITIAL (mostly, at least). The main difference is that you express your data in hexadecimal with XINIT, but you use ordinary characters for GINIT and INITIAL.)

This is how you write DBCS initial values:

- You enclose your data with single quotes, as you do with the ordinary INITIAL parameter.
- You use two ordinary characters for each DBCS character in your constant (two pairs of hexadecimal digits with XINIT) and one for each SBCS character (one pair with XINIT).
- You bracket each DBCS character string with a shift-out (SO) character immediately preceding and a shift-in (SI) character immediately after. SO is hexadecimal X'0E', which appears as '<' on most keyboards, and SI is X'0F' ('>'). (XINIT with PS=8 is an exception; the SO/SI brackets are implied and you do not key them.) For example, all of these define the same initial value, which is entirely DBCS. (Ignore the LENGTH values for the moment; we explain those in a moment.)

```
GINIT='<D1D2D3D4D5>',PS=8,LENGTH=10
INITIAL='<D1D2D3D4D5>',SOSI=YES,LENGTH=12
XINIT='C4F1C4F2C4F3C4F4C4F5',PS=8,LENGTH=10
XINIT='0EC4F1C4F2C4F3C4F4C4F50F',SOSI=YES,LENGTH=12
```

- SBCS and DBCS sequences can follow each other in any combination with INITIAL (and XINIT with SOSI=YES). If we add 'ABC' in front of the DBCS string in the previous example, and 'def' following the string, we have:

```
INITIAL='ABC<D1D2D3D4D5>def',SOSI=YES,LENGTH=18
XINIT='C1C2C30EC4F1C4F2C4F3C4F4C4F50F848586',SOSI=YES,LENGTH=18
```

- To calculate the length of your initial value, count two for each DBCS character and one for each SBCS character, whether you express them in ordinary characters or hexadecimal pairs. With GINIT (and XINIT with PS=8), you do not count the SO and SI characters, but with INITIAL (and XINIT with SOSI=YES), you add one for each SO and for each SI. (Note the different LENGTH values for the same constants in the examples above.) In all cases, your LENGTH value must not exceed 256.
- For GINIT and INITIAL, if your constant will not fit on one line, you use "extended" continuation rules, which are a little different from the ones described earlier. With extended continuation, you can stop after any full character (SBCS character, DBCS pair, or the SI ending a DBCS string) within your initial value. If you are in the middle of a DBCS string, add an SI (the SOs and SIs on one line must balance). Then fill out the line through column 72 with a continuation character. Any character will do, so long as it is different from the last meaningful character on the line. If you have stopped within a DBCS string, put an SO character in column 16 of the next line and resume in 17; otherwise just resume in 16, thus:

```
GXMPL1 DFHMDF POS=(02,21),LENGTH=20,PS=8,GINIT='<D1D2D3D4D5D6>*****
<D7D8D9D0>'
IXMPL1 DFHMDF POS=(02,21),LENGTH=23,PS=8,INITIAL='abc<D1D2D3D4>ABC**
DEFGHIJ'
```

You cannot use extended continuation with XINIT; use the rules described in Step 7.

- If your LENGTH specification exceeds the length of the initial value you provide, the value will be filled out on the right with DBCS blanks to your LENGTH value if you have used GINIT (or XINIT with PS=8). If you have used INITIAL, the fill character will be an SBCS blank if the last part of the constant was SBCS, a DBCS blank if the last part was DBCS. If you use XINIT with SOSI=YES, the fill character is always an SBCS blank.

Assembling the map

Before you start coding, you need to assemble and link edit your map set. You usually have to assemble twice, to create the map set in two different forms. The TYPE option in the DFHMSD macro tells the assembler the form to produce in any particular assembly.

Physical and symbolic map sets

A TYPE=MAP assembly, followed by a link-edit, produces a load module called the **physical map set**. The physical map set contains format information in encoded form. CICS uses it at execution time for constant fields and to determine how to merge in the variable data from the program. The physical map set normally is stored in the same library as your application programs, and it requires an RDO MAPSET resource definition within CICS, just as a program requires an RDO PROGRAM resource definition.

The output of a TYPE=DSECT assembly is a series of data structures, collectively called the **symbolic map set**, coded in the source language specified in the LANG option. There is a structure for each map used for input, called the **symbolic input map**, and one for each map used for output, called the **symbolic output map**.

Symbolic map sets are used at compile (assembly) time. You copy them into your program, and they allow you to refer to the fields in the maps by name and to pass the variable data in the form dictated by the physical map set. We have already shown you an example of a symbolic output map in COBOL (Figure 51 on page 230) and used it in the example code. Symbolic map sets are usually stored in the library your installation defines for source code that gets copied into programs. Member names are usually the same as the map set names, but they need not be.

You need the TYPE=DSECT assembly before you compile or assemble your program. You can defer the TYPE=MAP assembly and link-edit until you are ready to test, because the physical map set is not used until execution time. However, because you must do both eventually, many installations provide a catalogued procedure to do this automatically; the procedure copies the source file for the map set and processes it once using TYPE=MAP and again using TYPE=DSECT. You also can use the SYSPARM option in your assembler procedure to override the TYPE value in a particular assembly. See the *High Level Assembler for MVS, and VM and VSE: Language Reference* manual for a description of SYSPARM in connection with map assemblies, and “Preparing BMS maps” on page 16 for more information about assembling maps.

Notes:

1. The fact that symbolic map sets are coded in a specific language does not prevent you from using the same map in programs coded in different languages. You simply assemble with TYPE=DSECT for each LANG value you need, taking care to store the outputs in different libraries or under different names. The LANG value does not affect the TYPE=MAP assembly, which need be done only once.
2. If you modify an existing map in a way that affects the symbolic map, you **must recompile (reassemble)** any programs using it, so that the compilation uses the symbolic structure that corresponds to the new physical structure. Changes to unnamed map fields do not affect the symbolic map, but addition, deletion, rearrangement, and length changes of named fields do. (Rearrangement refers to the DFHMDF macros; the order of the fields on the screen does not affect the symbolic map, although it is more efficient to have the DFHMDF macros in same order as the fields on the screen.) So make changes to the DSATTS option in the map definition—this option states the extended attributes you may want to change by program. It is always safest to recompile, of course.

The SDF II alternative

None of these assembly or link-edit steps are required if you use the IBM licensed program Screen Definition Facility II Release 6 for VSE (SDF II) . SDF II creates both the symbolic map set and the physical map set in the final step of the interactive map creation process. See the SDF II VSE manuals "Introducing SDF II for VSE Release 6" and "SDF II General Introduction" for more information.

Grouping maps into map sets

Because they are assembled together, all of the physical maps in a map set constitute a single load module. BMS gains access to all of them with a single load request, issued on the first use of the map set in the task. No further loads are required unless you request a map in a different set, in which case BMS releases the old map set and loads the new one. If you go back to the first map set subsequently, it gets loaded again. Loading and deleting does not necessarily involve I/O, but consider the pathlength when grouping your maps into map sets. Generally, if maps are used together, they should be in the same map set; those not used together should be in different map sets.

The limit to the number of maps in a set is 9 998, but also keep the size of any given load module reasonable. So you might keep infrequently used maps separate from those normally used in a given process.

Similarly, all of the symbolic maps for a map set are in a single symbolic structure. This affects the amount of storage you need while using the maps, as explained in "BASE and STORAGE options" on page 243. Depending on the programming language, it also may affect high-level names, and this may be a reason for separating or combining maps as well.

Complex fields

The symbolic maps we have shown so far consisted of a fixed set of fields for each named map field (the A and O subfields, and so on, in Figure 51 on page 230). Such fields are the most common, but BMS provides two options for field definition which produce slightly different structures, to account for two common programming situations.

Composite fields: the GRPNAME option

Sometimes, you have to refer to subfields within a single field on the display. For example, you may have a date field that appears on the screen like this:

```
03-17-92
```

It is one field on the screen (with one attributes byte, just before the digit "0"), but you need to be able to manipulate the month, day, and year components separately in your program.

You can do this with a "group field", using the GRPNAME option of the DFHMDF macro. To create one, you code a DFHMDF macro for each of the component subfields; each definition names the same group field in the GRPNAME option. To define the date above as a group field starting at the beginning of line 10, for example, we would write:

```

MO      DFHMDF POS=(10,1),LENGTH=2,ATTRB=BRT,GRPNAME=DATE
SEP1    DFHMDF POS=(10,3),LENGTH=1,GRPNAME=DATE,INITIAL='- '
DAY     DFHMDF POS=(10,4),LENGTH=2,GRPNAME=DATE
SEP2    DFHMDF POS=(10,6),LENGTH=1,GRPNAME=DATE,INITIAL='- '
YR      DFHMDF POS=(10,7),LENGTH=2,GRPNAME=DATE

```

These definitions produce the following in the symbolic output map:

```

02 DATE.
   03 FILLER PICTURE X(2).
   03 MOA    PICTURE X.
   03 MOO   PIC X(2).
   03 SEP1  PIC X(1).
   03 DAO   PIC X(2).
   03 SEP2  PIC X(1).
   03 YRO   PIC X(2).

```

Several rules must be observed when using a group field:

- There is only one attributes byte; it precedes the whole group field and applies to the whole field. You specify it just once, on the DFHMDF macro for the first subfield, MO here.
- Because there is only one attributes byte, the cursor behaves as if the group field were a single field. In our example, the cursor will not move from the last position of month to the first of day, or day to year, skipping over the hyphens. This is because the group really is a single field as far as the hardware goes; it is subdivided only for program access to the component subfields.
- Although subfields after the first do not have an attributes byte, you define the POS option as if they did, as shown in the example. That is, POS points to one character before the subfield begins, and can overlap the last character of the previous subfield, as occurs in our example.
- Although all the component subfields are adjacent in this example, they do not have to be. There can be gaps between the subfields, provided you do not define any other field in the gap. The group field spans all the columns from its first subfield to its last, and you must put the component DFHMDF macros in the order the subfields appear on the screen. The group ends with the first DFHMDF macro that does not specify its name.
- You must assign a field name to every subfield, even if you do not intend to refer to it (as we did in the SEP1 and SEP2 subfields in the example).
- You cannot use the OCCURS option (explained in the next section) for a group field or any of its components.

Repeated fields: the OCCURS option

Sometimes a screen contains a series of identical fields that you want to treat as an array in your program. Suppose, for example, that you need to create a display of 40 numbers, to be used when a clerk assigns an unused telephone number to a new customer. (The idea is to give the customer some choice.) You also want to highlight numbers that have been in service recently, to warn the customer of the possibility of calls to the previous owner.

You can define the part of your screen that shows the telephone numbers with a single field definition:

```

TELNO    DFHMDF POS=(7,1),LENGTH=9,ATTRB=NORM,OCCURS=40

```

This statement generates 40 contiguous but separate display fields, starting at position (7,1) and proceeding across the rows for as many rows as required (five, in our case). We have chosen a length that (with the addition of the attributes byte) divides the screen width evenly, so that our numbers appear in vertical columns and are not split across row boundaries. The attributes you specify, and the initial value as well, apply to each field.

The description of these fields in the symbolic map looks like this in COBOL:

```
02 TELNOG    OCCURS 40.  
   03 FILLER PICTURE X(2).  
   03 TELNOA  PICTURE X.  
   03 TELNOO  PIC X(9).
```

This structure lets you fill the map from an array in your program (or any other source) as follows:

```
        PERFORM MOVENO FOR I FROM 1 THROUGH 40.  
        ...  
MOVENO.  
    MOVE AVAIL-NO (I) TO TELNOO (I).  
    IF DAYS-SINCE-USE (I) < 90, MOVE DFHMBRY to TELNOA (I).
```

(DFHMBRY is a CICS-supplied constant for setting the field intensity to bright; we explain more in “Attribute value definitions: DFHBMSCA” on page 246.)

Labels for OCCURS fields vary slightly for the different languages that CICS supports, but the function is the same.

Each element of an array created by the OCCURS option is a single map field. If you need to repeat a series of fields (an array of structures, in other words), you cannot use OCCURS. To use such an array in a program, you must define all of the fields individually, without OCCURS, to produce the necessary physical map. Then you can modify the resulting symbolic map, replacing the individual field definitions with an array whose elements are the structure you need to repeat. You must ensure that the revised symbolic map has exactly the same field structure as the original, of course. An alternative is to use SDF II, which allows you to define such an array directly.

Sending mapped output: basics

When you have assembled your symbolic map set, you are ready to code. We have explained by example how you get data from an application program to a map. We discuss that process in greater detail now, describing all the steps that must be performed, and telling you more about the options you have.

You need to do the following steps to produce mapped output:

1. Acquire storage in which to build the map.
2. Copy the symbolic map set so that it defines the structure of this storage.
3. Initialize it.
4. Move the output data into the map structure.
5. Set the field attributes.
6. Write the map to the screen with an EXEC CICS SEND MAP command, adding any device control information required.

We tell you about the final step—the EXEC CICS SEND MAP command itself—first, because you need to know what it does in order to understand what you need to do beforehand. Then we cover the tasks you need to complete before you issue it.

The SEND MAP command

The SEND MAP command tells BMS:

- Which map to use (MAP option), and where to find that map (the MAPSET option)
- Where to find the variable data for the map (FROM option) and how to merge it with the values from the map (MAPONLY and DATAONLY)
- Which device controls to include in the data stream, and other control options
- Where to put the cursor, if you want to override the position in the map definition (the CURSOR option)
- Whether the message is complete or will be built cumulatively (the ACCUM option)
- What to do with the formatted output (TERMINAL, SET, and PAGING options)

The MAP and MAPSET options are self-explanatory, and we cover most of the rest as we describe the programming steps that precede a simple EXEC CICS SEND MAP command. The last two topics require a knowledge of BMS logical message facilities, which we take up in “Output disposition options: TERMINAL, SET, and PAGING” on page 277.

Until we get to that point, we assume the defaults: that each EXEC CICS SEND MAP creates one message, and we are sending that message to our own terminal. For full details of the EXEC CICS SEND MAP command, see the *CICS Application Programming Reference* manual.

Acquiring and defining storage for the maps

The first step in creating mapped output is to provide storage in which to arrange the variable map data that your program passes to BMS. If you place the map structure in working storage, CICS does the allocation for you. (CICS allocates a private copy of working storage for each execution of a program, so that data from one task does not get confused with that from another, as explained in “Program storage” on page 67.) To use working storage, copy the symbolic map set there with the language statement provided for the purpose:

```
COPY          in COBOL and assembler
%INCLUDE     in PL/I
#include      in C
```

Working storage is the WORKING-STORAGE SECTION in COBOL, automatic storage in PL/I and C, and DFHEISTG in a CICS assembler program. For example:

```
WORKING-STORAGE SECTION.
...
01  COPY QCKSET.
...
```

Alternatively, you can obtain and release map set storage as you need it, using EXEC CICS GETMAIN commands. (GETMAIN is discussed in Chapter 26,

“Storage control and storage protection” on page 391.) In this case you copy the map into storage addressed by a pointer variable (the LINKAGE SECTION in COBOL, based storage in PL/I and C, a DSECT in assembler). On return from the GETMAIN, you use the address returned in the SET option to associate the storage with the data structure, according to the facilities of the programming language.

We used working storage in the example back on page 228, but we could have used a GETMAIN. If we had, the code we just showed you would change to:

```
LINKAGE SECTION.  
...  
01 COPY QCKSET.  
...  
PROCEDURE DIVISION.  
...  
MOVE LENGTH OF QCKMAPO TO LL.  
EXEC CICS GETMAIN SET(ADDRESS OF QCKMAPO)  
        LENGTH(LL) END-EXEC.  
...
```

The length you need on your GETMAIN command is the length of the variable whose name is the map name suffixed by the letter “O”. In COBOL, PL/I, and C, you can use language facilities to determine this length, as in the example above. In assembler, it is defined in an EQUate statement whose label is the map name suffixed by “L”.

BASE and STORAGE options

Two options on the DFHMSD map set definition macro affect how storage for maps is defined: BASE and STORAGE=AUTO (the STORAGE option always has the value AUTO). You can use either one or neither, so there are **three** possibilities. If you specify neither for a map set containing several maps, the symbolic structures for the maps are defined so that they overlay one another. If you specify STORAGE=AUTO, they do not; each occupies separate space. Thus STORAGE=AUTO requires more storage.

However, when you use maps that overlay one another in a single program, you must use them serially or compensate for the reuse of storage by programming. Unless storage is a major issue, STORAGE=AUTO simplifies programming and reduces the risk of error.

In PL/I and C, STORAGE=AUTO has the additional effect of defining the map as automatic storage (storage that CICS will allocate); the absence of STORAGE=AUTO causes these compilers to assume based storage, for which you generally incur the overhead of an additional GETMAIN. BMS assigns the name BSMAPBR to the associated pointer variable, unless you specify another name with the BASE option.

The third possibility, BASE, lets you use the same storage for all the maps in multiple map sets. Its effect varies slightly with the programming language, but essentially, all the maps in map sets with the same BASE value overlay one another. In COBOL, BASE=xxxx causes the 01 levels (that is, each individual map) to contain a REDEFINES xxxx clause. In PL/I and C, it designates each map as storage based on the pointer variable xxxx. BASE cannot be used when the programming language is assembler.

Initializing the output map

Before you start building your output, make sure that the map storage is initialized to nulls, so that data left there by a previous process is not used inadvertently. If you have read data input data using this same map, or one that overlays it, you need to ensure that you have processed or saved this data first, of course. The relationship between input and output maps is discussed in “The symbolic input map” on page 256, and using the same map you used for input is discussed in “Mapped output after mapped input” on page 264.

You initialize by moving nulls (X'00') into the structure. The symbolic map structures are defined so that you can refer to the whole map area by the name of the map suffixed by the letter **O**. You can see this in Figure 51 on page 230, and, in fact, the statement:

```
MOVE LOW-VALUES TO QCKMAPO.
```

would clear the area in which we built the map in the “quick check” example. If you are using the map for both input and output, it may be easier to clear the map one field at a time, as you edit the input (see “Handling input errors” on page 263).

When you obtain map storage with an EXEC CICS GETMAIN command, you can initialize using the INITIMG option.

Moving the variable data to the map

Having obtained storage for your map, established the relationship of the map structure to the storage, and initialized, you are finally ready to create your output. There are two parts to it: the data itself and its display attributes. We tell you about the data first and get to the attributes right after.

In the usual case, an output display consists of some constant or default data (provided by the physical map) and some variable data (provided by the program). For each field that you want to supply by program, you move the data into the field in the symbolic map whose name is the name assigned in the map suffixed by the letter **O**. See the code on page 231 for an example.

If you do not supply a value for a field (that is, you leave it null, as initialized), BMS ordinarily uses the initial value assigned in the map, if any. Constants (that is, fields without names) also get the initial values specified in the map. However, the DATAONLY and MAPONLY options on the SEND MAP command modify the way in which program and map data are merged; we explain these options in “Options for merging the symbolic and physical maps” on page 248 and summarize the exact rules in “Summary: what appears on the screen” on page 249.

Setting the display characteristics

Display attributes are the second component of the output data. (See “3270 field attributes” on page 213 for information about attributes.) In the “quick check” example on page 231, we show how 3270 field attributes for a map field are defined with the ATTRB option, and how BMS generates the “A” subfield to let you override the map value by program if you name the field.

BMS always provides the **A** subfield, because all 3270 devices support field attributes. Many 3270s also have some of the extended attributes shown in Table 16 on page 245. BMS supports each of these attributes individually in much the same way that it does field attributes collectively. You can assign attribute values in your DFHMDf field definitions, and, if you name the field, BMS generates a subfield in the symbolic map, so that you can override the map-assigned value in your program. There is a separate subfield for each type of extended attribute.

The only difference is that while BMS assumes field attributes, you must specify the other types in your map definition. The DSATTS option tells BMS which types of attributes you may modify by program. Each attribute listed causes BMS to generate the corresponding subfield for all named fields in the map. The MAPATTS option lists attribute types to which values are assigned in the DFHMDf field definitions. You must also include any attribute types that appear in DSATTS, even if they do not appear in any DFHMDf macro.

Table 16. BMS attribute types. The columns show the types of attributes, the name of the associated MAPATTS and DSATTS value, and the suffix of the associated subfields in the symbolic map.

Attribute type	MAPATTS, DSATTS value	Subfield suffix
Field attributes	None (default)	A
Color	COLOR	C
Highlighting	HIGHLIGHT	H
Outlining	OUTLINE	U
Background transparency	TRANSP	T
Validation	VALIDN	V
Double-byte character capability	SOSI	M
Programmed symbols	PS	P

Note: If you use programmed symbols, you need to ensure that a suitable symbol set has been sent to the device first, unless you choose one that is permanently loaded in the device. You can use a terminal control SEND command to do this (see “Data transmission commands” on page 320). The *IBM 3270 Information Display System Data Stream Programmer’s Reference* manual describes what to send.

The types of attributes that apply depend on the features of your principal facility at the time of execution. If you specify a value for an attribute that the terminal does not possess, BMS ignores it. If you are supporting different terminal types, however, you may need to use different techniques to get the same visual clarity. You can find out what kind of terminal you are using with the EXEC CICS ASSIGN and INQUIRE commands, explained in “Finding out about your terminal” on page 330. There are also provisions in BMS for keeping your program independent of the terminal type; see “Device-dependent maps: map suffixes” on page 269.

Changing the attributes

Here is an example of how this works. Suppose that the terminals in our “quick check” application have color and highlighting capabilities. We might decide to show the maximum charge allowed in a different color from the rest of the screen, because this field is of most interest to the clerk. We might also make the warning message red, because when it appears at all, it is important for the clerk to notice it. And when we really want to get the clerk’s attention, because the card is stolen, we could change the attributes in the program to make the message blink. To add these features, we need to change our map definition as follows:

```
QCKMAP  DFHMDI SIZE=(24,80),..., X
        MAPATTS=(COLOR,HILIGHT),COLOR=GREEN,HILIGHT=OFF, X
        DSATTS=HILIGHT
```

The MAPATTS option tells BMS that we specify color and highlighting in the map (or in the program, because any attribute listed in DSATTS must be included in MAPATTS as well). The COLOR and HILIGHT values indicate that fields with no color assigned should be green and that highlighting should be off if not specified. The only field definitions that we need to change are the ones that **are not** green or **are** highlighted:

```
CHG      DFHMDF POS=(5,13),LENGTH=8,ATTRB=(ASKIP,NORM),PICOUT='$,$$0.00', X
        COLOR=WHITE
MSG      DFHMDF LENGTH=20,POS=(7,1),ATTRB=(ASKIP,NORM),COLOR=RED
```

The DSATTS option tells BMS that we want to alter the highlighting of some fields at execution time, and therefore it should produce “H”-suffix subfields in the symbolic map to let us do that. Each named field gets the extra subfield; the message field, for example, expands from the current three lines in Figure 51 on page 230 to:

```
02 FILLER PICTURE X(2).
02 MSGH   PICTURE X.
02 MSGA   PICTURE X.
02 MSGO   PIC X(30).
```

The program statement we need to produce the blinking is:

```
MOVE DFHBLINK to MSGH.
```

In general, BMS takes attribute values from the program if you supply them and from the map if you do not (that is, if you leave the program value null, as initialized). However, the MAPONLY and DATAONLY options on the SEND MAP command affect attribute values as well as field data, as explained in “Where the values come from” on page 250.

Attribute value definitions: DFHBMSCA

The one-byte values required to set attribute values are bit combinations defined by 3270 hardware. They are hard to remember and, in some languages, clumsy to express. To solve this problem, CICS provides source code that you can copy into your program. The code, named DFHBMSCA, defines all the commonly used values for all attributes and assigns meaningful names to each combination. DFHBLINK in the line of code above is an example. To define DFHBLINK, we simply copy DFHBMSCA into our working storage, thus:

```
WORKING-STORAGE SECTION.
...
01 COPY DFHBMSCA.
```

There is a separate version of DFHBMSCA for each programming language, but the value names are the same in all versions. The *CICS Application Programming Reference* manual lists and defines all the value names. If you need an attribute combination not included in DFHBMSCA, you can determine the value by referring to the *IBM 3270 Information Display System Data Stream Programmer's Reference* manual. If you make frequent use of the value, you may want to modify DFHBMSCA to include it.

Note: In assembler language only, the values are defined with EQUates, so you use MVI rather than MVC instructions.

Control options on the EXEC CICS SEND MAP command

There are many control options for the BMS SEND commands. Some apply only to particular devices or special features of BMS, and we defer describing these until we get to the associated device support or feature. The following device control options, however, apply generally:

- **ERASE, ERASEAUP, and FRSET** all modify the contents of the device buffer, if the terminal has one, before writing your output into it. ERASE sets the entire buffer to nulls (X'00'). If the terminal has the alternate screen size feature, ERASE also sets the buffer size. Therefore, the first EXEC CICS SEND MAP command in a task normally specifies the ERASE option, both to clear the buffer and to select the buffer size. (See "3270 write commands" on page 211 for more information about alternate screen size.)

ERASEAUP (erase all unprotected fields) sets the contents of all fields in the buffer that are unprotected (that is, fields that the operator can change) to nulls. This is useful for data entry, as we explain in "DATAONLY option" on page 248.

FRSET (field reset) turns off the modified data tag of all fields in the buffer. ("Saving the good input" on page 263 and "Modification" on page 214 explain more about this option).

- **FREEKB** (free keyboard) unlocks the keyboard when the output is sent to the terminal. You usually want to do this on a display terminal.
- **ALARM** sounds the audible alarm, if the terminal has one.
- **FORMFEED, PRINT, L40, L64, L80, and HONEOM** are specific to printing and are explained in "Options for 3270 printers" on page 343. **NLEOM** also is used mainly in printing, and is explained in the same section. NLEOM requires **standard BMS**.

Some of these options can also be specified in the map itself, in particular, the options that are expressed in the 3270 write control character and coded in the CTRL option of the DFHMDI or DFHMSD macros: PRINT, FREEKB, ALARM, FRSET, L40, L64, L80, HONEOM.

Note: CTRL options are always treated as a group, so if you include any of them on your EXEC CICS SEND MAP command, BMS ignores the values for **all** of them in your map definition and uses only those in the command.

As we noted earlier, you can also send device control options separate from your map data, using an EXEC CICS SEND CONTROL command. You can use any option on EXEC CICS SEND CONTROL that you can use on EXEC CICS SEND MAP, except those that relate expressly to data, such as NLEOM.

Other BMS SEND options: WAIT and LAST

When a task writes to a terminal with a BMS or terminal control SEND command CICS normally schedules the transmission and then makes the task ready for execution again. Actual transmission occurs somewhat later, depending on terminal type, access method and other activity in the system. If you want to ensure that transmission is complete before your program regains control, use the WAIT option.

WAIT can increase response time slightly, because it prevents overlap between processing and output transmission for a task. (Overlap occurs only until a subsequent EXEC CICS SEND, RECEIVE, or end of task, however, because CICS finishes one terminal operation completely before starting another.)

You can improve response time slightly for some terminals by using the LAST option. LAST indicates that the output you are sending is the last output for the task. This knowledge allows CICS to combine transmission of the data with the VTAM end-of-bracket flow that occurs at end of task.

Options for merging the symbolic and physical maps

So far, we have assumed that every display consists of some constant data (provided by the physical map) and some variable data (provided by the program and structured according to the symbolic map). Sometimes, however, one or more of these components is missing.

MAPONLY option

For example, a menu map may not need any data supplied by program. In such a case, you code the MAPONLY option in place of the FROM option on your EXEC CICS SEND MAP command. BMS then takes all the information from the physical map, sending the initial values for both the constant (unnamed) and named fields. You do not need to copy the symbolic map set into a program that always sends the map with MAPONLY, and, in fact, you can skip the TYPE=DSECT map set assembly if all programs use all the maps in the set in this way.

MAPONLY is also the way you get an input-only map to the screen.

DATAONLY option

The opposite situation is also possible: the program can supply all the data and not need any constant or default values from the map. This happens on the second and subsequent displays of a map in many situations: data entry applications, inquiry applications where the operator browses through a series of records displayed in identical format, and screens that are redisplayed after detection of an error in the input.

BMS takes advantage of this situation if you indicate it with the DATAONLY option. You still need to tell BMS which map and map set you are using for positioning information, but BMS sends only those fields that have nonnull attribute or data values in the symbolic map. Other fields and attribute values are left unchanged.

The EXEC CICS SEND CONTROL command

There are also occasions when you do not need to send data at all, but you do need to send device controls. For example, you might need to erase the screen or sound the alarm. You do this with an EXEC CICS SEND CONTROL command listing the options you need,

Here is an example of a program that uses all three of these techniques. It processes input in a data entry application. When first initiated, it displays the data entry map to format the screen with the input fields, the associated labels, screen headings and instructions. This first EXEC CICS SEND MAP command specifies MAPONLY, because the program sends no variable data. Thereafter, the program accepts one set of data input. If the input is correct, the program files it and requests another. It still does not need to send any variable data. What it needs to do is to erase the input from the screen and unlock the keyboard, to signal the operator to enter the next record.

```
EXEC CICS SEND CONTROL ERASEAUP FREEKB END-EXEC
```

does this. (See “Control options on the EXEC CICS SEND MAP command” on page 247 for a description of these and other device control options.)

If there are errors, the program does need to send variable data, to tell the operator how to fix the problem. This one changes the attributes of the fields in error to highlight them and sends a message in a field provided for the purpose. Here, our program uses the DATAONLY option, because the map is already on the screen. (We tell you more about handling input errors in “Handling input errors” on page 263.)

Use MAPONLY, DATAONLY, and EXEC CICS SEND CONTROL when they apply, especially when response time is critical, as it is in a data entry situation. MAPONLY saves pathlength, DATAONLY reduces the length of the outbound data stream, and EXEC CICS SEND CONTROL does both.

Summary: what appears on the screen

The interaction of physical map definition options, EXEC CICS SEND MAP options, program data and merge options is sufficiently complex that a summary of the rules for determining what appears on the screen after an EXEC CICS SEND MAP is in order.

The contents of the screen (buffer) are determined by:

- What was there before your EXEC CICS SEND MAP command
- The fields (field attributes, extended attributes, and display data) that get sent from your SEND MAP command
- Where the several values for these field elements come from

We discuss the possibilities in that order.

What you start with

The first thing that happens on an EXEC CICS SEND MAP command is that the entire screen (buffer) is cleared to nulls if the ERASE option is present, regardless of the size or origin of your map. On terminals that have the alternate screen size feature, the screen size is set as well, as explained in “3270 write commands” on page 211. The screen is in unformatted state, with no fields defined and no display data. If ERASEAUP is present, all of the unprotected fields on the screen are erased, but the field structure and attributes of all fields and the contents of protected fields are unchanged.

ERASE and ERASEAUP are honored **before** your EXEC CICS SEND MAP data is loaded into the buffer. If neither of these options appears on the EXEC CICS SEND MAP, the screen buffer starts out as it was left after the previous write operation, modified by whatever the operator did. In general, the positions of the fields (that is, of the attributes bytes) and their attributes are unchanged, but the data content of unprotected fields may be different. Furthermore, if the operator used the CLEAR key, the whole buffer is cleared to nulls and the screen is in unformatted state, just as if you had included the ERASE option.

What is sent

Secondly, BMS changes only those positions in the buffer within the confines of your map. Outside that area, the contents of the buffer are unchanged, although it is possible for areas outside your map to change in appearance, as explained in “Outside the map” on page 252.

Within the map area, what is sent depends on whether the DATAONLY option is present. In the usual case, where it is not present, BMS sends every component (field attributes, extended attributes, display data) of every field in your map. This creates a field at the position specified in the POS operand and overlays the number of bytes specified in the LENGTH field from POS. Buffer positions between the end of the display data and the next attributes byte (POS value) are left unchanged. (There may or may not be fields (attributes bytes) in these intervening spaces if you did not ERASE after the last write operation that used a different map.)

The values for these field elements come from the program, the map or defaults, as explained in the next section.

If DATAONLY is present, on the other hand, BMS sends only those fields, and only those components for them, that the program provides. Other screen data is unchanged.

Where the values come from

The values that determine screen contents may come from four sources:

- Program
- Map
- Hardware defaults
- Previous screen contents

BMS considers each component of each map field separately, and takes the value from the program, provided:

- The MAPONLY option has not been used.

- The field has a name in the map, so that the symbolic output map contains the corresponding set of subfields from which to get the data. The field attributes value comes from the program subfield whose name is the map field name suffixed by A. The display data comes from the subfield of the same name suffixed by O, and the extended attribute values come from the same-named subfields suffixed by the letter that identifies the attribute (see Table 16 on page 245). In the case of the extended attributes, the attribute must also appear among DSATTS in order for the symbolic map to contain the corresponding subfield.
- A value is present. The definition of “present” varies slightly with the field component:
 - For field attributes bytes, the value must not be null (X'00') or one of the values that can be left over from an input operation (X'80', X'02', or X'82').
 - For extended attribute bytes, the value must not be null.

Note: BMS sends only those extended attribute values that the terminal is defined as supporting. Values for other extended attributes are omitted from the final data stream.
 - For display data, the first character of the data must not be null.

If any of these conditions are not met, the next step depends on whether DATAONLY is present. With DATAONLY, BMS stops the process here and sends only the data it got from the program. BMS does this in such a way that components not changed by program are not changed on the screen. In particular, extended attributes values are not changed unless you specify a new value or ask for the hardware default. (A value of X'FF' requests the hardware default for all extended attributes except background transparency, for which you specify X'F0' to get the hardware default.)

Without DATAONLY, if one of the conditions above is not met, BMS takes the data from the map, as follows:

- For field attributes, it takes the value in the ATTRB option for the field. If none is present, BMS assumes an ATTRB value of (ASKIP,NORM).
- For extended attributes, BMS takes the value from:
 - The corresponding option in the DFHMDf field definition
 - If it is not specified there, the value is taken from the corresponding option in the DFHMDI map definition
 - If it is not there either, the value is taken from the corresponding option in the DFHMSD map set definition

(If no value is specified anywhere, BMS does not send one, and this causes the 3270 to use its hardware default value.)
- For display data, from the initial value in the map (the INITIAL, XINIT or GINIT option). If there is no initial value, the field is set to nulls.

Outside the map

We have assumed, so far, that your map is the same size as your screen or printer page. It need not be. Your application may use only a part of the screen area, or you may want to build your output incrementally, or both.

BMS logical messages allow you to build a screen from several maps, sending it with a single terminal write. You use the ACCUM option to do this, which we cover in “BMS logical messages” on page 277. Even without using ACCUM, you can build a screen from several maps if the terminal is a 3270-like device with a buffer. You do this with multiple EXEC CICS SEND MAP commands written to different areas of the screen (buffer), not erasing after the first command. Each EXEC CICS SEND MAP causes output and may produce a momentary “blink” at a display device. For this reason, and to eliminate the pathlength of extra I/O, you may prefer to use logical messages for such composite screens.

Outside the map just sent, the contents of the buffer are unchanged, except for the effects of ERASE and ERASEAUP cited earlier. In general, this means that the corresponding areas of the screen are unchanged. However, a screen position outside the map may get its attributes from a field within the map. This occurs if you do not define a field (using a different map) beyond the boundary of the map and before the position in question. If you change the attributes of the field inside your map governing this position outside, the appearance of the position may change, even though the contents do not.

Using GDDM® and BMS

One use of the buffer overlay technique we just described is the creation of screens containing a mixture of BMS and GDDM output. You generally write the BMS output first, followed by the GDDM. You can leave space in the BMS map for the GDDM output, or you can create a “graphic hole” in any display by writing a map with no fields in it to the position where you want the hole. Such a map is called a “null map,” and its size (height and width) correspond to the size of the hole.

If you use GDDM to combine graphics with BMS output, you need to include a GDDM PSRSRV call to prevent GDDM from corrupting programmed symbol sets that BMS may be using.

Positioning the cursor

Usually, you set the initial position for the cursor in the map definition by including “insert cursor” (IC) in the ATTRB values of the field where you want it. (Cursor position is not important for the output-only maps we have been discussing, but it becomes important as soon as you use a map for input too.)

The CURSOR option on the EXEC CICS SEND MAP command allows you to override this specification, if necessary, when the map is displayed. If you specify CURSOR(value), BMS places the cursor in that absolute position on the screen. Counting starts in the first row and column (the zero position), and proceeds across the rows. Thus, to place the cursor in the fourth column of the third row of an 80-column display, you code CURSOR(163).

Specifying CURSOR without a value signals BMS that you want “symbolic cursor positioning”. You do this by setting the length subfield of the field where you want

the cursor to minus one (-1). Length subfields are not defined on output-only maps, so you must define your map as INOUT to use symbolic cursor positioning. (We tell you about length subfields in “Reading from a formatted screen: what comes in” on page 259, and about INOUT maps in “Receiving data from a display.”) If you mark more than one field in this way, BMS uses the first one it finds.

Symbolic cursor positioning is particularly useful for input-output maps when the terminal operator enters incorrect data. If you validate the fields, setting the length of any in error to -1, BMS places the cursor under the first error when you redisplay. “Processing the mapped input” on page 262 shows this technique.

You can position the cursor with a SEND CONTROL command also, but only by specifying an absolute value for CURSOR; if you omit CURSOR on SEND CONTROL, the cursor is not moved.

Sending invalid data and other errors

The exceptional conditions that can occur on EXEC CICS SEND MAP and SEND CONTROL commands are listed with the descriptions of these commands in the *CICS Application Programming Reference* manual. Most of them apply only to the advanced BMS options: logical messages, partitions, and special devices.

However, it is also possible to send invalid data to a terminal. BMS does not check the validity of attribute and data values in the symbolic map, although it does not attempt to send an extended attribute, like color, to a terminal not defined to support that attribute.

The effects of invalid data depend on both the particular terminal and the nature of the incorrect data. Sometimes invalid data can be interpreted as a control sequence, so that the device accepts the data but produces the wrong output; sometimes the screen displays an error indicator; and sometimes an ATNI abend occurs. The point at which your task is notified of an ATNI depends on whether or not you specified the WAIT option (see “Other BMS SEND options: WAIT and LAST” on page 248).

Receiving data from a display

Formatted screens are as important for input as for output. Data entry applications are an obvious example, but most other applications also use formatted input, at least in part. On input, BMS does for you approximately the reverse of what it does on output: it removes device control characters from the data stream and moves the input fields into a data structure, so that you can address them by name.

Maps can be used exclusively for input, exclusively for output (the case we have already covered), or for both. Input-only maps are relatively rare, and we cover them as a special case of an input-output map, pointing out differences where they occur.

An input-output example

Before we explain the details of the input structure, let us re-examine the “quick check” example. Suppose that it is against our policy to let a customer charge up to the limit over and over again between the nightly runs when new charges are posted to the accounts. We want a new transaction that helps “quick check” processing by keeping a running total for the day.

In addition, we want to use the same screen for both input and output, so that there is only one screen entry per customer. In the new transaction, “quick update,” the clerk enters both the account number and the charge at the same time. The normal response is:

```
QUP                      Quick Account Update
Current charge okay; enter next
Account:  _____
Charge:  $  _____
```

Figure 53. Normal “quick update” response

When we reject a transaction, we leave the input information on the screen, so that the clerk can see what was entered along with the description of the problem:

```
QUP                      Quick Account Update
Charge exceeds maximum; do not approve
Account:    482554
Charge:    $ 1000.00
```

Figure 54. “Quick update” error response

(Here again, we are oversimplifying to keep our maps short for ease of explanation.)

The map definition we need for this exercise is:

```

QUPSET DFHMSD TYPE=MAP,STORAGE=AUTO,MODE=INOUT,LANG=COBOL,TERM=3270-2
QUPMAP DFHMDI SIZE=(24,80),LINE=1,COLUMN=1,CTRL=FRECKB
        DFHMDF POS=(1,1),LENGTH=3,ATTRB=(ASKIP,BRT),INITIAL='QUP'
        DFHMDF POS=(1,26),LENGTH=20,ATTRB=(ASKIP,NORM),                X
            INITIAL='Quick Account Update'
MSG     DFHMDF LENGTH=40,POS=(3,1),ATTRB=(ASKIP,NORM)
        DFHMDF POS=(5,1),LENGTH=8,ATTRB=(ASKIP,NORM),                X
            INITIAL='Account:'
ACCTNO  DFHMDF POS=(5,14),LENGTH=6,ATTRB=(UNPROT,NUM,IC)
        DFHMDF POS=(5,21),LENGTH=1,ATTRB=(ASKIP),INITIAL=' '
        DFHMDF POS=(6,1),LENGTH=11,ATTRB=(ASKIP,NORM),                X
            INITIAL='Charge: $ '
CHG     DFHMDF POS=(6,13),ATTRB=(UNPROT,NORM),PICIN='99999V99'
        DFHMDF POS=(6,21),LENGTH=1,ATTRB=(ASKIP),INITIAL=' '
        DFHMSD TYPE=FINAL

```

You can see that the map field definitions for this input-output map are very similar to those for the output-only “quick check” map, if we allow for changes to the content of the screen. The differences to note are:

- The MODE option in the DFHMSD map set definition is INOUT, indicating that the maps in this map set are used for both input and output. INOUT causes BMS to generate a symbolic structure for input as well as for output for every map in the map set. If this had been an input-only map, we would have said MODE=IN, and BMS would have generated only the input structures.
- We put names on the fields from which we want input (ACCTNO and CHG) as well as those to which we send output (MSG). As in an output-only map, we avoid naming constant fields to save space in the symbolic map.
- The input fields, ACCTNO and CHG, are unprotected (UNPROT), to allow the operator to key data into them.
- IC (insert cursor) is specified for ACCTNO. It positions the cursor at the start of the account number field when the map is first displayed, ready for the first item that the operator has to enter. (You can override this placement when you send the map; IC just provides the default position.)
- Just after the ACCTNO field, there is a constant field consisting of a single blank, and a similar one after the CHG field. These are called “stopper” fields. Normally, they are placed after each input field that is not followed immediately by some other field. They prevent the operator from keying data beyond the space you provided, into an unused area of the screen.

If you define the stopper field as “autoskip”, the cursor will jump to the next unprotected field after the operator has filled the preceding input field. This is convenient if most of the input fields are of fixed length, because the operator does not have to advance the cursor to get from field to field.

If you define the stopper field as “protected,” but not “autoskip,” the keyboard locks if the operator attempts to key beyond the end of the field. This choice may be preferable if most of the input fields are of variable length, where one usually has to use the cursor advance key anyway, because it alerts the operator to the overflow immediately. Whichever you choose, use the same choice throughout the application if possible, so that the operator sees a consistent interface.

- The CHG field has the option PICIN. PICIN produces an edit mask in the symbolic map, useful for COBOL and PL/I, and implies the field length. See

the *CICS Application Programming Reference* manual for details on using PICIN.

Figure 55 shows the symbolic map set that results from shows the symbolic map set that results from this INOUT map definition.




Figure 55. Symbolic map for “quick update”

The second part of this structure, starting at QUPMAPO, is the symbolic **output** map—the structure required to send data back to the screen. Apart from the fields we redefined, it looks almost the same as the one you would have expected if we had specified MODE=OUT instead of MODE=INOUT. See Figure 49 on page 229 for a comparison. The main difference is that the field attributes (A) subfield appears to be missing, but we will explain this in a moment.

The symbolic input map

The first part of the structure, under the label QUPMAPI, is new. This is the symbolic **input** map—the structure required for reading data from a screen formatted with map QUPMAP. For each named field in the map, it contains three subfields. As in the symbolic output map, each subfield has the same name as the map field, suffixed by a letter indicating its purpose. The suffixes and subfields related to input are:

- L** the length of the input in the map field
- F** the flag byte, which indicates whether the operator erased the field and whether the cursor was left there
- I** the input data itself

The input and output structures are defined so that they overlay one another field by field. That is, the input (I) subfield for a given map field always occupies the same storage as the corresponding output (O) subfield. Similarly, the input flag (F)

subfield overlays the output attributes (A) subfield. (For implementation reasons, the order of the subfield definitions varies somewhat among languages. In COBOL, the definition of the A subfield moves to the input structure in an INOUT map, but it still applies to output, just as it does in an output-only map. In assembler, the input and output subfield definitions are interleaved for each map field.)

BMS uses dummy fields to leave space in one part of the structure for subfields that do not occur in the other part. For example, there is always a two-byte filler in the output map to correspond to the length (L) subfield in the input map, even in output-only maps. If there are output subfields for extended attributes, such as color or highlighting, BMS generates dummy fields in the input map to match them. You can see examples of these fields (FILLERS in COBOL) in both Figure 49 on page 229 and Figure 55 on page 256.

The correspondence of fields in the input and output map structures is very convenient for processes in which you use a map for input and then write back in the same format, as you do in data entry transactions or when you get erroneous input and have to request a correction from the operator.

Programming simple mapped input

The programming required for mapped input is similar to that for mapped output, except, of course, that the data is going in the opposite direction. You define your maps and assemble them first, as for mapped output. In the program or programs reading from the terminal, you:

1. Acquire the storage to which the symbolic map set corresponds
2. Copy the symbolic map set to define the structure of this storage
3. Format the input data with an EXEC CICS RECEIVE MAP command
4. Process the input

We tell you more about these tasks and related topics in the paragraphs that follow, starting with the EXEC CICS RECEIVE MAP command. We also develop the code for the “quick update” transaction.

If the transaction also calls for mapped output, as “quick update” and most other transactions do, you simply continue with the steps outlined before, in “Sending mapped output: basics” on page 241. Some considerations and shortcuts for mapped input are described in “Mapped output after mapped input” on page 264.

The EXEC CICS RECEIVE MAP command

The EXEC CICS RECEIVE MAP command causes BMS to format terminal input data and make it accessible to your application program. It tells BMS:

- Which map to use in formatting the input data stream—that is, what format is on the screen and what data structure the program expects (the MAP option)
- Where to find this map (MAPSET option)
- Where to get the input (TERMINAL or FROM option)
- Whether to suppress translation to upper case (ASIS option)
- Where to put the formatted input data (the INTO and SET options)

The MAP and MAPSET options together tell BMS which map to use, and they work exactly as they do on an EXEC CICS SEND MAP command.

BMS gets the input data to format from the terminal associated with your task (its principal facility), unless you use the FROM option. FROM is an alternative to TERMINAL, the default, used in relatively unusual circumstances (see “Formatting other input” on page 266).

BMS also translates lower case input to upper case automatically in some cases; we explain how to control translation in “Upper case translation” on page 260.

You tell BMS where to put the formatted input with the INTO or SET option, which we cover in the next section. For the full syntax of the EXEC CICS RECEIVE MAP command, see the *CICS Application Programming Reference* manual.

Getting storage for mapped input: INTO and SET

When you issue a RECEIVE MAP command, BMS needs storage in which to build the input map structure. You can provide this space yourself, either in the working storage of your program or with an EXEC CICS GETMAIN command. These are the same choices you have for allocating storage in which to build an output map, and you use them the same way (see “Acquiring and defining storage for the maps” on page 242 for details and examples). For either, you code the INTO option on your RECEIVE command, naming the variable into which the formatted input is to be placed. For our “quick update”, for example, the required command is:

```
EXEC CICS RECEIVE MAP('QUPMAP') MAPSET('QUPSET')
      INTO(QUPMAPI) END-EXEC.
```

Usually, the receiving variable is the area defined by the symbolic input map, to which BMS assigns the map name suffixed by the letter “I”, as shown above. You can specify some other variable if you wish, however.

For input operations, you have a third choice for acquiring storage. If you code the SET option, BMS acquires the storage for you at the time of the RECEIVE command and returns the address in the pointer variable named in the SET option. So we could have coded the RECEIVE MAP command in “quick update” like this:

```
LINKAGE SECTION.
...
01 QUPMAP COPY QUPMAP.
...
PROCEDURE DIVISION.
...
EXEC CICS RECEIVE MAP('QUPMAP') MAPSET('QUPSET')
      SET(ADDRESS OF QUPMAPI) END-EXEC.
...
```

Storage obtained in this way remains until task end unless you issue an EXEC CICS FREEMAIN command to release it (see Chapter 26, “Storage control and storage protection” on page 391 for further information).

Reading from a formatted screen: what comes in

As we noted earlier, we explain receiving input from a terminal in terms of 3270 devices. Also read “Support for non-3270 terminals” on page 267 if you are writing for non-3270 terminals.

CICS normally reads a 3270 screen with a “read modified” command⁵. The data transmitted depends on what the operator did to cause transmission:

- The ENTER key or a PF key
- CLEAR, CNCL or a PA key (the “short read” keys)
- Field selection: cursor select, light pen detect or a trigger field

You can tell which event occurred, if you need to know; we explain how in “The attention identifier: what caused transmission” on page 260. You can also find more detail on 3270 input operations in “Input from a 3270 terminal” on page 221.

The short read keys transmit only the attention identifier (the identity of the key itself). No field data comes in, and there is nothing to map. For this reason, short read keys can cause the MAPFAIL condition, as explained on page 265. Field selection features transmit field data, but in most cases not the same data as the ENTER and PF keys, which we describe in the paragraphs that follow. See “BMS support for other special hardware” on page 314 for the exceptions if you plan to use these features.

Most applications are designed for transmission by the ENTER key or a PF key. When one of these is used to transmit, all of the fields on the screen that have been modified, and **only** those fields, are transmitted.

Modified data

As we explained in “Modification” on page 214, a 3270 screen field is considered modified only if the “modified data tag” (MDT), one of the bits in the field attributes byte, is on. The terminal hardware turns on this bit if the operator changes the field in any way—entering data, changing data already there, or erasing. You can also turn it on by program when you send the map, by including MDT among the ATTRB values for the field. You do this when you want the data in a particular field to be returned even if the operator does not change it.

You can tell whether there was input from a particular map field by looking at the corresponding length (L) subfield. If the length is zero, no data was read from that field. The associated input (I) subfield contains all nulls (X'00'), because BMS sets the entire input structure to nulls before it performs the input mapping operation. The length is zero either if the modified data tag is off (that is, the field was sent with the tag off and the operator did not change it) or if the operator erased the field. You can distinguish between these two situations, if you care, by inspecting the flag (F) subfield. It has the high-order bit on if the field contains nulls but the MDT is on (that is, the operator changed the field by erasing it). See “Finding the cursor” on page 261 for more information about the flag subfield.

⁵ CICS provides an option, BUFFER, for the terminal control RECEIVE command, with which you can capture the entire contents of a 3270 screen. See “Reading from a 3270 terminal” on page 222 if you need to do this.

If the length is nonzero, data was read from the field. Either the operator entered some, or changed what was there, or the field was sent with the MDT on. You may find the data itself in the corresponding input (I) subfield. The length subfield tells how many characters were sent. A 3270 terminal sends only nonnull characters, so BMS knows how much data was keyed into the field. Character fields are filled out with blanks on the right and numeric fields are filled on the left with zeros unless you specify otherwise in the JUSTIFY option of the field definition. BMS assumes that a field contains character data unless you indicate that it is numeric with ATTRB=NUM. (See the *CICS Application Programming Reference* manual for details of how these options work.)

Upper case translation

CICS converts lower case input characters to upper case automatically under some circumstances. The definition of the terminal and the transaction together determine whether translation occurs. See the UCTRAN option of the RDO PROFILE and the RDO TYPETERM resource definitions in the *CICS Resource Definition Guide* for how these specifications interact.

You can suppress this translation by using the ASIS option on your EXEC CICS RECEIVE MAP command, **except** on the first RECEIVE in a task initiated by terminal input. (The first RECEIVE may be either an EXEC CICS RECEIVE MAP (without FROM) or a terminal control RECEIVE.) CICS has already read and translated this input, and it is too late to suppress translation. (Its arrival caused the task to be invoked, as explained in “How tasks get started” on page 200.) Consequently, ASIS is ignored entirely in pseudoconversational transaction sequences, where at most one EXEC CICS RECEIVE MAP (without FROM) occurs per task, by definition. For the same reason, you cannot use ASIS with the FROM option (see “Formatting other input” on page 266).

Other information from EXEC CICS RECEIVE MAP

In addition to the data on the screen, the EXEC CICS RECEIVE MAP command tells you where the operator left the cursor and what key caused transmission. This information becomes available in the EIB on completion of the EXEC CICS RECEIVE MAP command. EIBAID identifies the transmit key (the “attention identifier” or AID), and EIBCURSR tells you where the cursor was left.

The attention identifier: what caused transmission

This information is part of the input in many applications, and you may also need it to interpret the input correctly.

For example, in the “quick update” transaction, we need some method for allowing the clerk to exit our transaction, and we have not yet provided for this. Suppose that we establish the convention that pressing PF12 causes you to leave control of the transaction. We would then code the following after our EXEC CICS RECEIVE MAP command:

```
IF EIBAID = DFHPF12
    EXEC CICS SEND CONTROL FREEKB ERASE END-EXEC
    EXEC CICS RETURN END-EXEC.
```

This would end the transaction without specifying which one should be executed next, so that the operator would regain control. The EXEC CICS SEND CONTROL

command that precedes the EXEC CICS RETURN unlocks the keyboard and clears the screen, so that the operator is ready to enter the next request.

The hexadecimal values that correspond to the various attention keys are defined in a copy book called DFHAID. To use these definitions, you simply copy DFHAID into your working storage, in the same way that you copy DFHBMSCA to use the predefined attributes byte combinations (see “Attribute value definitions: DFHBMSCA” on page 246). The contents of the DFHAID copy book are listed in the *CICS Application Programming Reference* manual.

The HANDLE AID command

You can also use an EXEC CICS HANDLE AID command to identify the attention key used (unless you are writing in C, which does not support EXEC CICS HANDLE AID commands). EXEC CICS HANDLE AID works like other HANDLE commands; you issue it before the first EXEC CICS RECEIVE command to which it applies, and it causes a program branch on completion of subsequent RECEIVES if a key named in the EXEC CICS HANDLE AID is used.

For example, an alternative to the “escape” code just shown would be:

```
EXEC CICS HANDLE AID PF12(ESCAPE) END-EXEC.  
...  
EXEC CICS RECEIVE MAP('QUPMAP') MAPSET('QUPSET') ...  
...  
    ESCAPE.  
EXEC CICS SEND CONTROL FREEKB ERASE END-EXEC  
EXEC CICS RETURN END-EXEC.
```

EXEC CICS HANDLE AID applies only to EXEC CICS RECEIVE commands in the same program. The specification for a key remains in effect until another EXEC CICS HANDLE AID in the same program supersedes it by naming a new label for the key or terminates it by naming the key with no label. A RESP, RESP2, or NOHANDLE option on an EXEC CICS RECEIVE command exempts that particular command from the effects of EXEC CICS HANDLE AID specifications, but they remain in effect otherwise.

If you have a HANDLE active for an AID received during an input operation, control goes to the label specified in the EXEC CICS HANDLE AID, regardless of any exceptional condition that occurs and whether or not an EXEC CICS HANDLE CONDITION is active for that exception. EXEC CICS HANDLE AID can thus mask an exceptional condition if you check for it with EXEC CICS HANDLE CONDITION. For this reason you may prefer to use an alternative test for the AID or exceptional conditions or both. You can check EIBAID for the AID and use the RESP option or check EIBRESP for exceptions. You need to be especially aware of MAPFAIL in this respect, as noted on page 265.

Finding the cursor

In some applications, you need to know where the operator left the cursor at the time of sending. There are two ways of finding out. If your map specifies CURSLOC=YES, BMS turns on the seventh (X'02') bit in the flag subfield (F) subfield of the map field where the cursor was left. This works only, of course, if the cursor is left in a map field to which you assigned a name.

Also, because the flag subfield is used to indicate both cursor presence and field erasure, you need to test the bits individually if you are looking for one in particular:

the X'80' bit for field erasure and the X'02' bit for the cursor. If you are using a language in which it is awkward to test bits, you can test for combinations. A value of X'80' or X'82' signals erasure; either X'02' or X'82' indicates the cursor. The DFHBMSCA definitions described in the *CICS Application Programming Reference* manual include all of these combinations.

You can also determine the position of the cursor from the EIBCPOSN field in the EIB. This is the absolute position on the screen, counting from the upper left (position zero) and going across the rows. Thus a value of 41 on a screen 40 characters wide would put the cursor in the second row, second column. Avoid this method if possible, because it makes your program sensitive to the placement of fields on the screen and to the terminal type.

Processing the mapped input

To illustrate how the input subfields are used, we return to “quick update”. After we have the input, we need to do some checks on it before continuing. First, we require that the charge be entered (that is, that the input length be greater than zero), and be positive and numeric.

```
IF CHGL = 0, MOVE -1 TO CHGL
    MOVE 1 TO ERR-NO
ELSE IF CHGI NOT > ZERO OR CHGI NOT NUMERIC,
    MOVE DFHUNIMD TO CHGA,
    MOVE -1 TO CHGL
    MOVE 2 TO ERR-NO.
```

The 'MOVE -1' statements here and following put the cursor in the first field in error when we redisplay the map, as explained in “Positioning the cursor” on page 252. The message number tells us what message to put in the message area; 1 is “enter a charge”, and so on through 6, for “charge is over limit”. We do these checks in roughly ascending order of importance, to ensure that the most basic error is the one that gets the message. At the end of the checking, we know that everything is okay if ERR-NO is zero.

An account number must be entered, as well as the charge. If we have one (whatever the condition of the charge), we can retrieve the customer’s account record to ensure that the account exists:

```
IF ACCTNOL = 0, MOVE -1 TO ACCTNOL
    MOVE 3 TO ERR-NO
ELSE EXEC CICS READ FILE (ACCT) INTO (ACCTFILE-RECORD)
    RIDFLD (ACCTNOI) UPDATE RESP(READRC) END-EXEC
    IF READRC = DFHRESP(NOTFOUND), MOVE 4 TO ERR-NO,
        MOVE DFHUNIMD TO ACCTNOA
        MOVE -1 TO ACCTNOL
    ELSE IF READRC NOT = DFHRESP(NORMAL) GO TO HARD-ERR-RTN.
```

If we get this far, we continue checking, until an error prevents us from going on. We need to ensure that the operator gave us a good account number (one that is not in trouble), and that the charge is not too much for the account:

```

IF ERR-NO NOT > 2
  IF ACCTFILE-WARNCODE = 'S', MOVE DFHBM BRY TO MSGA
  MOVE 5 TO ERR-NO
  MOVE -1 TO ACCTNOL
  EXEC CICS LINK PROGRAM('NTFYCOPS') END-EXEC
  ELSE IF CHGI > ACCTFILE-CREDIT-LIM - ACCTFILE-UNPAID-BAL
    - ACCTFILE-CUR-CHGS
    MOVE 6 TO ERR-NO
    MOVE -1 TO ACCTNOL.
IF ERR-NO NOT = 0 GO TO REJECT-INPUT.

```

Handling input errors

As illustrated in “quick update,” above, whenever you have operator input to process, there is almost always a possibility of incorrect data, and you must provide for this contingency in your code. Usually, what you need to do when the input is wrong is:

- Notify the operator of the errors. Try to diagnose all of them at once; it is annoying to the operator if you present them one at a time.
- Save the data already entered, so that the operator does not have to rekey anything except corrections.
- Arrange to recheck the input after the operator makes corrections.

Flagging errors

In the preceding code for the “quick update” transaction, we used the message field to describe the error (the first one, anyway). We highlighted all the fields in error, provided there was any data in them to highlight, and we set the length subfields to -1 so that BMS would place the cursor in the first bad field. We send this information using the same map, as follows:

```

REJECT-INPUT.
  MOVE LOW-VALUES TO ACCTNOO CHGO.
  EXEC CICS SEND MAP('QUPMAP') MAPSET('QUPSET') FROM(QUPMAPO)
    DATAONLY END-EXEC.

```

Notice that we specify the DATAONLY option. We can do this because the constant part of the map is still on the screen, and there is no point in rewriting it there. We cleared the output fields ACCTNOO and CHGO, to avoid sending back the input we had received, and we used a different attributes combination to make the ACCTNO field bright (DFHUNIMD instead of DFHBM BRY). DFHUNIMD highlights the field and leaves the modified data tag on, so that if the operator resends without changing the field, the account number is retransmitted.

Saving the good input

The next step is to ensure that whatever good data the operator entered gets saved. One easy technique is to store the data on the screen. You do not have to do anything additional to accomplish this; once the MDT in a field is turned on, as it is the first time the operator touches the field, it remains on, no matter how many times the screen is read. Tags are not turned off until you erase the screen, turn them off explicitly with the FRSET option on your EXEC CICS SEND, or set the attributes subfield to a value in which the tag is off.

The drawback to saving data on the screen is that all the data is lost if the operator uses the CLEAR key. If your task is conversational, you can avoid this hazard by moving the input to a safe area in the program before sending the error information

and asking for corrections. In a pseudoconversational sequence, where the component tasks do not span interactions with the terminal, the equivalent is for the task that detects the error to pass the old input forward to the task that processes the corrected input. You can forward data through a COMMAREA on the EXEC CICS RETURN command that ends a task, by writing to temporary storage, or in a number of other ways (see Chapter 4, “Sharing data across transactions” on page 83 for possibilities).

In addition to avoiding the CLEAR key problem, this method reduces inbound transmission time, because you transmit only changed fields on the error correction cycles. (You must specify FRSET when you send the error information to prevent the fields already sent and not corrected from coming in again.) You can also avoid repeating field audits because, after the first time, you need to audit only if the user has changed the field or a related one.

However, these gains are at the expense of extra programming and complexity, and therefore the savings in line time or audit pathlength must be considerable, and the probability of errors high, to justify this choice. You must add code to merge the new input with the old, and if you have turned off the MDTs, you need to check both the length and the flag subfield to determine whether the operator has modified a map field. Fields with new data have a nonzero length; those that had data and were subsequently erased have the high-order bit in the flag subfield on.

A good compromise is to save the data both ways. If the operator clears the screen, you use the saved data to refresh it; otherwise you simply use the data coming in from the screen. You do not need any merge logic, but you protect the operator from losing time over an unintended CLEAR.

For our “quick update” code, with its minimal audits and transmissions, we choose the “do nothing” approach and save the information on the screen.

Rechecking

The last requirement is to ensure that the input data is rechecked. If your task is conversational, this simply means repeating the audit section of your code after you have received (and merged, if necessary) the corrected input. In a pseudoconversational sequence, you usually repeat the transaction that failed. In the example, because we saved the data on the screen in such a way that corrected data is indistinguishable from new data, all we need to do is arrange to execute the same transaction against the corrected data, thus:

```
EXEC CICS RETURN TRANSID('QUPD') END-EXEC.
```

where 'QUPD' is the identifier of the “quick update” transaction.

Mapped output after mapped input

If your transaction makes it through its input audits and the attendant hazards, the processing specific to mapped input is complete. The next step, frequently, is to prepare and send the transaction output. In general, if the output is to be mapped, you follow the steps outlined in “Sending mapped output: basics” on page 241. However, the acquisition of storage for building the map may be affected by the input mapping you have already done. If the output and input maps are different, but in the same map set or in map sets defined to overlay one another, you have already done the storage acquisition during your input mapping process. If your output and input maps overlay one another, you need to ensure that you save any

map input you still need and clear the output structure to nulls before you start building the output map. If this is awkward, you may want to define the maps so that they do not overlay one another. (See “BASE and STORAGE options” on page 243 for your choices in this regard.)

Your transaction may also call for using the same map for output as input. This is routine in code that handles input errors, as we have already seen, and also in simple transactions like “quick update”. One-screen data-entry transactions are another common example.

When you are sending new data with a map already on the screen, you can reduce transmission with the DATAONLY option, and you may need only the EXEC CICS SEND CONTROL command. See “Options for merging the symbolic and physical maps” on page 248 for a discussion of these options.

For the “quick update” transaction, however, we need to fill in the message field with our “go” response (and update the file with the charge to finish our processing):

```
MOVE 'CURRENT CHARGE OKAY; ENTER NEXT' TO MSGO
ADD CHGI TO ACCTFILE-CUR-CHGS
EXEC CICS REWRITE FILE('ACCT') FROM (ACCTFILE-RECORD)....
```

We also need to erase the input fields, so that the screen is ready for the next input. We have to do this both on the screen (the ERASEAUP option erases all unprotected fields) and in the output structure (because the output subfield overlays the input subfield and the input data is still there).

```
MOVE LOW-VALUES TO ACCTNOO CHGO.
EXEC CICS SEND MAP('QUPMAP') MAPSET('QUPSET') FROM(QUPMAPO)
      DATAONLY ERASEAUP END-EXEC.
```

Finally, we can return control to CICS, specifying that the same transaction is to be executed for the next input.

```
EXEC CICS RETURN TRANSID('QUPD') END-EXEC.
```

MAPFAIL and other exceptional conditions

The exceptional conditions that can occur on an EXEC CICS RECEIVE command are all described in the *CICS Application Programming Reference* manual, and most are self-explanatory. One of them warrants discussion, however, because it can result from a simple operator error. This is MAPFAIL, which occurs when no usable data is transmitted from the terminal or when the data transmitted is unformatted (in the 3270 sense—see “Unformatted mode” on page 224). MAPFAIL occurs on an EXEC CICS RECEIVE MAP if the operator has used the CLEAR key or one of the PA keys. It also occurs if the operator uses ENTER or a PF key from a screen where:

- No fields defined in the map have the modified data tag set on (this means the operator did not key anything and you did not send any fields with the tags already set, so that no data is returned on the read), and
- The cursor was not left in a field defined in the map and named, or the map did not specify CURSLOC=YES.

Pressing ENTER prematurely or a “short read” key accidentally is an easy mistake for the operator to make. In the interest of user friendliness, you may want to refresh the screen after MAPFAIL instead of ending the transaction in error.

MAPFAIL also occurs if you issue an EXEC CICS RECEIVE MAP without first formatting with a SEND MAP or equivalent in the current or a previous task, and can occur if you use a map different from the one you sent. This might signal an error in logic, or it might simply mean that your transaction is in its startup phase. For instance, in our “quick update” example, we have not made any provision for getting started—that is, for getting an empty map onto the screen so that the operator can start using the transaction. We could use a separate transaction to do this, but we might as well take advantage of the code we need to refresh the screen after a MAPFAIL. What we need is:

```
IF RCV-RC = DFHRESP(MAPFAIL)
  MOVE 'PRESS PF12 TO QUIT THIS TRANSACTION' TO MSGO
  EXEC CICS SEND MAP('QUPMAP') MAPSET('QUPSET')
    FROM(QUPMAPO) END-EXEC.
```

We are reminding the operator how to escape, because attempts to do this may have caused the MAPFAIL in the first place. If we had not wanted to send this message, or if it was the default in the map, we could have used the MAPONLY option:

```
EXEC CICS SEND MAP('QUPMAP') MAPSET('QUPSET') MAPONLY END-EXEC.
```

When MAPFAIL occurs, the input map structure is not cleared to nulls, as it is otherwise, so it is important to test for this condition if your program logic depends on this clearing.

You can issue a HANDLE CONDITION command to intercept MAPFAIL, as you can other exception conditions. If you do, and you also have a HANDLE AID active for the AID you receive, however, control goes to the label specified for the AID and not that for MAPFAIL, as explained in “The HANDLE AID command” on page 261. In this situation you will be unaware of the MAPFAIL, even though you issued a HANDLE for it, unless you also test EIBRESP.

EOC condition

EOC is another condition that you encounter frequently using BMS. It occurs when the end-of-chain (EOC) indicator is set in the request/response unit returned from VTAM. EOC does not indicate an error, and the BMS default action is to ignore this condition.

Formatting other input

Although the data that you format with an EXEC CICS RECEIVE MAP command normally comes from a terminal, you can also format data that did not come from a terminal, or that came indirectly. For example, you might not know which map to use until you receive the input and inspect some part of it. This can happen when you use special hardware features like partitioning or logical device codes, and also in certain logic situations. You might also need to format data that was read from a formatted screen by an intermediate process (without mapping) and later passed to your transaction.

The FROM option of the EXEC CICS RECEIVE MAP command addresses these situations. FROM tells BMS that the data has already been read, and only the translation from the native input stream to the input map structure is required.

Because the input has already been read, you need to specify its length if you use FROM, because BMS cannot get this information from the access method, as it

does normally. If the data came originally from an EXEC CICS RECEIVE command in another task, the length on the EXEC CICS RECEIVE MAP FROM command should be the length produced by that original RECEIVE, including the 12-byte TIOA prefix.

For the same reason, you cannot suppress translation to upper case with the ASIS option when you use FROM. Moreover, BMS does not set EIBAID and EIBCURSR after an EXEC CICS RECEIVE FROM command.

And finally, BMS does not know from what device the input came, and it assumes that it was your current principal facility. (You cannot even use EXEC CICS RECEIVE FROM without a principal facility, even though no input/output occurs.) If the data came from a different type of device, you have to do the mapping in a transaction with a similar principal facility to get the proper conversion of the input data stream.

Note: You cannot map data read with a terminal control RECEIVE with the BUFFER option, because the input data is unformatted (in the 3270 sense). If you attempt to issue an EXEC CICS RECEIVE MAP FROM command for such input, MAPFAIL occurs.

Support for non-3270 terminals

Minimum BMS supports only 3270 displays and printers. This category includes the 3178, 3290, 8775 and 5520, LU type 2 and LU type 3 devices, and any other terminal that accepts the 3270 data stream. The *IBM 3270 Information Display System Data Stream Programmer's Reference* manual contains a full list.

Standard BMS expands 3270 support to SCS printers (3270 family printers **not** using the 3270 data stream) and all of the terminal types listed in Table 17 on page 270.

Because of functional differences among these terminal types, it is not possible to make BMS work in exactly the same way for each of them. The sections that follow outline the limitations in using BMS on devices that lack the hardware basis for certain features.

Output considerations for non-3270 devices

Because BMS separates the device-dependent content of the output data stream from the logical content, there are only a few differences between 3270 and non-3270 devices that you need to consider in creating BMS output.

The primary difference between 3270 and non-3270 devices is that the 3270 is **field-oriented**, and most others are not. Consequently, there are neither field attributes nor extended attributes associated with output fields sent to non-3270 terminals. BMS can position the output in the correct places, field by field, but the field structure is not reflected in the data stream. BMS can even emulate some field attributes on some terminals (it may underline a highlighted field, for example), but there is no modified data tag, no protection against keying into the field, and so on.

If you specify attributes on output that the terminal does not support, BMS simply ignores them. You do not need to worry about them, provided the output is understandable without the missing features.

Differences on input

The absence of field structure has more impact on input operations, because many of the features of BMS depend on the ability to read—by field—only those fields that were modified. If the hardware does not provide the input field-by-field with its position on the screen, you must provide equivalent information.

You can do this in either of two ways. The first is to define a field-separator sequence, one to four characters long, in the FLDSEP option of the map definition. You place this sequence between each field of your input and supply the input fields in the same order as they appear on the screen or page. You must supply every field on the screen, up to the last one that contains any data. If there is no input in a field, it can consist of only the terminating field-separator sequence. On hardcopy devices, input cannot overlay the output because of paper movement. On displays that emulate such terminals, the same technique is generally used. Input fields are entered in an area reserved for the purpose, in order, separated by the field-separator sequence.

The second method is to include control characters in your input. If you omit the FLDSEP option from your map, BMS uses control characters to calculate the position of the data on the “page” and maps it accordingly. The control characters that BMS recognizes are:

NL	new line	X'15'
IRS	interchange record separator	X'1E'
LF	line feed	X'25'
FF	form feed	X'0C'
HT	horizontal tab	X'05'
VT	vertical tab	X'0B'
CR	carriage return	X'0D'
RET	return on the TWX	X'26'
ETB	end text block	X'26'
ESC	escape, for 2780	X'27'

When you read data of this kind with a RECEIVE MAP command, there are some differences from true 3270 input:

- The flag byte (F subfield) is not set; it contains a null. You cannot determine whether the operator erased the field or whether the cursor was left in the field.
- You cannot preset a modified data tag on output to ensure that a field is returned on input.

Special options for non-3270 terminals

BMS provides some additional formatting options for non-3270 devices, to take advantage of device features that shorten the data stream. These include:

- Vertical and horizontal tabs. You can position your output with horizontal and vertical tab orders if the device supports them. The tab characters are defined by the HTAB and VTAB options in the map set definition. When you want to position to the next horizontal tab, you include the HTAB character in your data; you position to the next vertical tab by supplying the VTAB character in your data. BMS translates these characters to the tab sequence required by your particular principal facility.

Before you use tabs in BMS output, your task or some earlier task at the same terminal must have set the tabs in the required positions. This is usually done

with a terminal control SEND command, described in “Data transmission commands” on page 320.

- Outboard formatting. Some logical units can store format information and participate in the formatting process. This allows BMS to send much less data (essentially the symbolic map contents) and delegate the work of merging the physical and symbolic maps to the logical unit. See “Outboard formatting” on page 318 for details.
- NLEOM (new line, end of message). Standard BMS also gives you the option of requesting that BMS format your output with blanks and new-line (NL) characters rather than 3270 buffer control orders. This technique gives you more flexibility in page width settings on printers, as explained in “NLEOM option” on page 345.

Device-dependent maps: map suffixes

Because the position, default attributes, and default contents of map fields appear only in the physical map and not in the symbolic map, you can use a single program to build maps that contain the same variable information but different constant information in different arrangements on the screen. This is very convenient if the program you are writing must support multiple devices with different characteristics.

Suppose, for example, that some of the clerks using the “quick update” transaction use 3270 Model 2s (as we have assumed until now), and the rest use a special-purpose terminal that has only 3 rows and 40 columns. The format we designed for the big screen will not do for the small one, but the information will fit if we rearrange it:

```

QUP Quick Account Update:
Current charge okay; enter next
Acct: _____ Charge: $ _____
  
```

Figure 56. “Quick update” for the small screen

We need the following map definition:

```

QUPSET DFHMSD TYPE=MAP,STORAGE=AUTO,MODE=INOUT,LANG=COBOL,SUFFIX=9
QUPMAP DFHMDI SIZE=(3,40),LINE=1,COLUMN=1,CTRL=FREEKB
        DFHMDF POS=(1,1),LENGTH=24,ATTRB=(ASKIP,BRT), X
        INITIAL='QUP Quick Account Update'
MSG     DFHMDF LENGTH=39,POS=(2,1),ATTRB=(ASKIP,NORM)
        DFHMDF POS=(3,1),LENGTH=5,ATTRB=(ASKIP,NORM), X
        INITIAL='Acct:'
ACCTNO  DFHMDF POS=(3,11),LENGTH=6,ATTRB=(UNPROT,NUM,IC)
        DFHMDF POS=(3,18),LENGTH=1,ATTRB=(ASKIP),INITIAL=' '
        DFHMDF POS=(3,20),LENGTH=9,ATTRB=(ASKIP,NORM), X
        INITIAL='Charge: $'
CHG     DFHMDF POS=(3,30),LENGTH=7,ATTRB=(UNPROT,NORM),PICIN='99999v99'
        DFHMDF POS=(3,38),LENGTH=1,ATTRB=(ASKIP),INITIAL=' '
        DFHMSD TYPE=FINAL
  
```

The symbolic map set produced by assembling this version of the map is identical to the one shown in “An input-output example” on page 254, because the fields with names have the same names and same lengths, and they appear in the same

order in the map definition. (They do not need to appear in the same order on the screen, incidentally; you can rearrange them, provided you keep the definitions of named fields in the same order in all maps.) Thus we can use the same code to build the map. Furthermore, with some preparation, BMS picks the physical map that matches the terminal in the current execution when you issue EXEC CICS SEND MAP or EXEC CICS RECEIVE MAP.

Device-dependent support: DDS

The BMS feature that does this is called “device-dependent support” (DDS). DDS is an installation option that works as follows.

When you assemble your map sets, you specify the type of terminal the maps are for in the TERM option. This causes the assembler to store the physical map set under the MAPSET name suffixed by the character for that type of terminal⁶. When you issue EXEC CICS SEND MAP or EXEC CICS RECEIVE MAP with DDS active, BMS adds a one-character suffix to the name you supply in the MAPSET option. It chooses the suffix based on the definition of your terminal, and thus loads the physical map that corresponds to the terminal for any given execution.

BMS defines the suffixes used for the common terminal types. A 3270 Model 2 with a screen size of 24 rows and 80 columns is assigned the letter ‘M,’ for example. The type is determined from the TYPETERM definition if it is one of the standard types shown in Table 17.

<i>Table 17 (Page 1 of 2). Terminal codes for BMS</i>	
Code	Terminal or logical unit
A	CRLP (card reader input, line printer output) or TCAM terminal
B	Magnetic tape
C	Sequential disk
D	TWX Model 33/35
E	1050
F	2740-1, 2740-2 (without buffer receive)
G	2741
H	2740-2 (with buffer receive)
I	2770
J	2780
K	3780
L	3270-1 displays (40-character width)
M	3270-2 displays (80-character width), LU type 2s
N	3270-1 printers
O	3270-2 printers, LU type 3s
P	All interactive LUs, 3767/3770 Interpreter LU, 3790 full function LU, SCS printer LU
Q	2980 Models 1 and 2

⁶ You also can use JCL or the link-edit NAME statement to control the member name under which a map set is stored.

<i>Table 17 (Page 2 of 2). Terminal codes for BMS</i>	
Code	Terminal or logical unit
R	2980 Model 4
U	3600 (3601) LU
V	3650 Host Conversational (3653) LU
W	3650 Interpreter LU
X	3650 Host Conversational (3270) LU
Y	3770 Batch LU, 3770 and 3790 batch data interchange LUs, LU type 4s
blank	3270-2 (default if TERM omitted)

An installation can also define additional terminal types, such as the miniature screen described above. The system programmer does this by assigning an identifier to the terminal type and specifying it in the ALTSUFFIX option of the TYPETERM definition for the terminals. When you create a map for such a terminal, you specify this identifier in the SUFFIX option, instead of using the TERM option. Transactions using the map must also point to a PROFILE that specifies alternate screen size, so that ALTSUFFIX is used.

You might use the same technique to distinguish among standard terminals used for special purposes. For example, if an application were used by both English and French speakers, you could create two sets of physical maps, one with the constants in French and the other in English. You would assign each a suffix, and specify the English suffix as the ALTSUFFIX value in the definitions of the English terminals and the French suffix for French terminals. Transactions using the map would point to a PROFILE that specified alternate screen size. Then when you sent the map, BMS would pick the version with the suffix that matched the terminal (that is, in the appropriate language).

With DDS, the rules BMS uses for selecting a physical map are:

- BMS adds the ALTSUFFIX value in the terminal definition to your map set name, provided that definition specifies both ALTSUFFIX and ALTSCREEN, and provided that the screen size for the transaction is the alternate size (either because the transaction PROFILE calls for alternate size, or because the default and alternate sizes are the same).
- If these conditions are not met, or if BMS cannot find a map with that suffix, it attempts to find one with the suffix that corresponds to the terminal type in the terminal definition.
- If BMS cannot find that map either, it looks for one with no suffix. (A blank suffix indicates an all-purpose map, suitable for any terminal that might use it.)

Without DDS, BMS always looks first (and only) for an unsuffixed map.

Device-dependent support is an installation option for BMS, set by the system programmer in the system initialization table. Be sure that it is included in your system before taking advantage of it; know whether it is present, even if you are supporting only one device type.

With DDS in the system, there is an efficiency advantage in creating suffixed map sets, even if you are supporting only one device type, because you prevent BMS

from attempting to load a map set that does not exist before defaulting to the universal one (the blank suffix).

Without DDS, on the other hand, it is unnecessary to suffix your maps, because BMS looks for the universal suffix (a blank) and fails to locate suffixed maps.

Finding out about your terminal

Because of the overall design of BMS, and device-dependent support in particular, you generally do not need to know much about your terminal to format for it. However, if you need to know the characteristics of your principal facility, you can use the EXEC CICS ASSIGN and INQUIRE commands. You can tell, for example, whether your terminal supports a particular extended attribute, what national language is in use, screen size and so on. This type of information applies whether you are using BMS or terminal control to communicate with your terminal. You need it more often for terminal control, and so we describe the options that apply in that chapter, in “Finding out about your terminal” on page 330.

There are also ASSIGN options specific to BMS, but you need them most often when you use the ACCUM option, and so we describe them later, in “ASSIGN options for cumulative processing” on page 286.

The MAPPINGDEV facility.

Minimum BMS function assumes that the principal facility of your task is the **mapping device** that performs input and output mapping operations for the features and status that is defined in the TCTTE (Terminal Control Table entry).

The principal facility for transactions using BMS function should have a device type supported by BMS. However, the MAPPINGDEV facility is an extension of minimum BMS that allows you to perform mapping operations for a device that is **not** the principal facility. When the MAPPINGDEV request completes, the mapped data is returned to the application. BMS does not have any communication with the MAPPINGDEV device.

You can specify the MAPPINGDEV option on the RECEIVE MAP command and the SEND MAP command, but not on any other BMS command. See the *CICS Application Programming Reference* manual for further information about RECEIVE MAP and SEND MAP commands.

The TERMIID specified in the MAPPINGDEV option must represent a device in the 3270 family supported by BMS. If the device is partitioned, it is assumed to be in base state. Outboard formatting is ignored.

Data is mapped in exactly the same way as for minimum BMS, and there is no need to change mapset definitions or to re-generate mapsets.

SEND MAP with the MAPPINGDEV option.

Your SEND MAP commands that have the MAPPINGDEV option must also specify the SET option. (The SET option provides BMS with a pointer that sets the address of the storage area that contains the mapped output datastream.)

If you have storage protection active, the data is returned in storage in the key specified in the TASKDATAKEY option of the transaction definition. The storage is

located above or below the line depending on which TASKDATALOC option of the transaction definition you have specified.

The storage area is in task-related user storage but in the format of a TIOA (Terminal Input/Output Area). The application can reference the storage area using the DFHTIOA copybook. The TIOATDL field, at offset 8, contains the length of the datastream that starts at TIOADBA, at offset 12, in the storage area. The length value placed in TIOATDL does not include the length of the 4-byte page control area, which contains information such as the extended attributes that have been used in the datastream and can be referenced using the DFHPGADS copybook.

The storage area has usually a greater length than the datastream because the storage area is allocated before the exact length of the output datastream is determined. This storage area is in a form that can be used in a SEND TEXT MAPPED command and is described in “SEND TEXT extensions: SEND TEXT MAPPED and SEND TEXT NOEDIT” on page 293.

If you are familiar with using the SET option without the MAPPINGDEV option, (see “Output disposition options: TERMINAL, SET, and PAGING” on page 277 for details) you know that the datastream is returned to the application indirectly by a list of pages. However, when MAPPINGDEV is specified, a direct pointer to the storage area containing the datastream is returned to your application.

When the SEND MAP MAPPINGDEV command completes its processing, the storage area is under the control of the application and remains allocated until the end of the transaction unless your application FREEMAINS it. You are advised to FREEMAIN these storage areas, for long-running transactions but CICS frees these areas when the task terminates.

RECEIVE MAP with the MAPPINGDEV option.

You must specify the FROM option when using the MAPPINGDEV option on the RECEIVE MAP command. BMS needs the FROM option to supply a formatted 3270 input datastream that is consistent with the datastream returned via a Terminal Control RECEIVE command (that is, a normal input 3270 datastream). The only difference is that it does not start with an AID and input cursor address because this information is removed from the input datastream by terminal control, but there are options on the RECEIVE MAP command that allow you to specify an AID value and input cursor position when the MAPPINGDEV option is specified. If the datastream contains an AID and input cursor address, they are ignored by BMS.

When neither option is specified, BMS assumes that the input data operation was terminated with the ENTER key, and returns the appropriate AID value to the application from the EIBAID field. BMS also assumes that the input cursor was positioned at the home address and returns a value of zero to the application from the EIBCPOSN field.

The new AID option of the RECEIVE MAP command allows your application to specify an AID value which, if specified, overrides the default value of ENTER. Whether provided by the application, or defaulted by BMS, the AID value that you established causes control to be passed, when applicable, to the routine registered by a previous HANDLE AID request issued by the application.

The new CURSOR option of the RECEIVE MAP command allows your application to specify an input cursor position which, if specified, overrides the default value of zero. Whether provided by the application, or defaulted by BMS, the input cursor value is used in cursor location processing when you define the map with CURSLOC=YES.

As with the minimum BMS RECEIVE MAP command, the mapped data is returned to your application by the INTO or SET option. If neither option is specified, the CICS translator attempts to apply a default INTO option by appending the character 'I' to the map name.

When you use the SET option with the MAPPINGDEV option, it must provide a pointer variable that BMS sets with the address of the storage area containing the mapped input datastream. The data is returned in task-related user storage. If storage protection is active, the data is returned in storage in the key specified by the TASKDATAKEY option of the transaction definition. The storage is located above or below the line depending on the TASKDATALOC option of your transaction definition.

When the RECEIVE MAP MAPPINGDEV command completes its processing successfully, the storage area is returned by the SET option and is under the control of the application and remains allocated until the end of the transaction unless your application FREEMAINs it. You are advised to FREEMAIN these storage areas, for long-running transactions but CICS frees these areas when the task terminates.

Sample assembler MAPPINGDEV application.

Figure 57 on page 275 is a modification of the FILEA operator instruction sample program, and uses the same mapset named DFH\$AGA.

This application is only intended to demonstrate how to code the keywords associated with the MAPPINGDEV facility, and as a means of testing this function. It is not offered as a recommended design for applications that make use of the MAPPINGDEV facility.

```

DFH$AMNX CSECT
*
      DFHREGS
DFHEISTG DSECT
OUTAREA DS 0CL512
        DS CL8
OUTLEN  DS H
        DS H
OUTDATA DS CL500
INLEN   DS H
INAREA  DS CL256
PROOF   DS CL60
        COPY DFH$AGA
        COPY DFHBMSCA
DFH$AMNU CSECT
EXEC CICS HANDLE AID PF3(PF3_ROUTINE)
*
XC     DFH$AGAS(DFH$AGAL),DFH$AGAS
MVC    MSGO(L'APPLMSG),APPLMSG
EXEC   CICS SEND MAP('DFH$AGA') FROM(DFH$AGAO) ERASE      x
        MAPPINGDEV(EIBTRMID) SET(R6)
MVC    OUTAREA(256),0(R6)
MVC    OUTAREA+256(256),256(R6)
EXEC   CICS SEND TEXT MAPPED FROM(OUTDATA) LENGTH(OUTLEN)
*
EXEC   CICS RECEIVE INTO(INAREA) LENGTH(INLEN)             x
        MAXLENGTH(MAXLEN)
*
EXEC   CICS RECEIVE MAP('DFH$AGA') SET(R7) LENGTH(INLEN)  x
        MAPPINGDEV(EIBTRMID) FROM(INAREA)                 x
        CURSOR(820) AID(=C'3')
*
XC     PROOF,PROOF
MVC    PROOF(25),=C'You just keyed in number '
MVC    PROOF+25(6),KEYI-DFH$AGAI(R7)
FINISH DS 0H
EXEC   CICS SEND TEXT FROM(PROOF) LENGTH(60) ERASE FREEKB
TM     MSGF-DFH$AGAI(R7),X'02'
BNO    RETURN
XC     PROOF,PROOF
MVC    PROOF(33),=C'Input cursor located in MSG field'
EXEC   CICS SEND TEXT FROM(PROOF) LENGTH(60) ERASE FREEKB
*
*     THE RETURN COMMAND ENDS THE PROGRAM.
*
RETURN DS 0H
EXEC   CICS RETURN
*
PF3_ROUTINE DS 0H
        XC     PROOF,PROOF
        MVC    PROOF(30),=C'RECEIVE MAP specified AID(PF3)'
        B      FINISH
MAXLEN  DC H'256'
APPLMSG DC C'This is a MAPPINGDEV application'
END

```

Figure 57. ASM example of a MAPPINGDEV application

Block data

BMS provides an alternate format for the symbolic map, called block data format, that may be useful in specific circumstances. In block data format, the symbolic output map is an image of the screen or page going to the terminal. It has the customary field attributes (A) and output value (O) subfields for each named map field, but the subfields for each map field are separated by filler fields such that their displacement in the symbolic map structure corresponds to their position on the screen. There are no length subfields, and symbolic cursor positioning is unavailable as a consequence.

A symbolic description map is generated for block data format by specifying DATA=BLOCK on the DFHMSD or DFHMDI map definition macros.

For example, the symbolic map for the “quick check” screen in Figure 50 on page 229 would look like this in block data format (assuming a map 80 columns wide). Compare this with the normal “field data” format (in Figure 51 on page 230) from the same map definition.

```
01 QCKMAPO.
02 FILLER PIC X(12).           <---TIOAPFX still present
02 FILLER PICTURE X(192).     <---Spacer to
02 ACCTNOA PICTURE X.         <---Position (3,13)
02 ACCTNOO PIC X(7).
02 FILLER PICTURE X(72).     <---Spacer to
02 SURNAMEA PICTURE X.       <---Position (4,13)
02 SURNAMEO PIC X(15).
02 FNAMEA PICTURE X.         <---Position (4,30), no
02 FNAMEO PIC X(10).         preceding spacer required
02 FILLER PICTURE X(52).     <---Spacer to
02 CHGA PICTURE X.           <---Position (5,13)
02 CHGO PIC $,$$0.00
02 FILLER PICTURE X(139).    <---Spacer to
02 MSGA PICTURE X.           <---Position (7,1).
02 MSGO PIC X(30).
```

Figure 58. Symbolic map for “quick check” in block data format.

You can set only the field attributes in the program; BMS ignores the DSATTS option in the map and does not generate subfields for the extended attributes in block data format. You can use block data for input as well. The input map is identical in structure to the output map, except that the flag (F) replaces the field attributes (A) subfield, and the input (I) replaces the output (O) subfield, as in field format.

Block data format may be useful if the application program has built or has access to a printer page image that it needs to display on a screen. For most situations, however, the normal field data format provides greater function and flexibility.

Sending mapped output: additional facilities

In our examples so far, each SEND MAP produces one output message, of one screen or page, delivered immediately to your principal facility. In this section, we describe features of BMS that let you go beyond these basic messages, including:

- Additional disposition options for the message
- Messages of more than one page

- Composite pages, built with multiple BMS SEND commands
- Routing your message to terminals other than your own

We also cover the second type of BMS output, text, in the process, because these additional facilities apply to text as well as mapped output.

Output disposition options: **TERMINAL, SET, and PAGING**

The only disposition option we have described up to this point is TERMINAL, which sends the output to the principal facility of your task. TERMINAL is the default value that you get if you do not specify another disposition. There are, however, two other possibilities:

- BMS can return the formatted output stream to the task rather than sending it to the terminal. You use the SET disposition option to request this. You might do so to defer transmission or to modify the data stream to meet special requirements. “Using SET” on page 302 explains how and when to use SET.
- You can ask BMS to store and manage your output in CICS temporary storage for subsequent delivery to your terminal. This option, PAGING, implies that your message may contain more than one screen or page, and is particularly useful when you want to send a message to a display terminal that exceeds its screen capacity. BMS saves the entire message in temporary storage until you indicate that it is complete. Then it provides facilities for the operator to page through the output at the terminal. You can use PAGING for printers as well as displays, although you do not need the operator controls, and sometimes TERMINAL is just as satisfactory.

When you use PAGING, the output still goes to your principal facility, though indirectly, as just described. Full BMS also provides a feature, routing, that lets you send your message to another terminal, or several, in place of or in addition to your own. We tell you about routing in “Message routing: the ROUTE command” on page 294, after we cover the prerequisites.

Note: Both PAGING and SET and related options require **full** BMS. TERMINAL is the only disposition available in minimum and standard BMS.

BMS logical messages

The disposition options do not affect the correspondence between EXEC CICS SEND MAP commands and pages of output. You get one page for each EXEC CICS SEND MAP command, unless you also use a second feature of full BMS, the ACCUM option. ACCUM allows you to build pages piecemeal, using more than one map, and like PAGING, it allows your message to exceed a page. You do not have to worry about page breaks or about tailoring your output to a specific page or screen capacity. BMS handles these automatically, giving you control at page breaks if you wish. Details on cumulative page building are in “Page formation: the ACCUM option” on page 282.

As soon as you create an output message of more than one page, or a single page composed of several different maps, you are doing something BMS calls **cumulative** mapping. PAGING implies multiple pages, and ACCUM implies both multiple and composite pages, and so at the first appearance of either of these options, BMS goes into cumulative mapping mode and begins a **logical message**. The one-to-one correspondence between SEND commands and messages ends, and subsequent EXEC CICS SEND MAP commands simply add to the current logical message. Individual pages within the message are still disposed of as soon

as they are complete, but they all belong to the same logical message, which continues until you tell BMS to end it.

Rules for logical messages

When you start a logical message, you need to observe a number of rules:

- You can build only one logical message at a time. If you are routing this message, BMS may create more than one logical message internally, but in terms of content, there is only one. After you complete the message and dispose of it, you can build another in the same task, using different options if you wish.
- Options related to message management must be the same on all commands that build the message. These are:
 - The disposition option: PAGING, TERMINAL, or SET
 - The option governing page formation: ACCUM should be present on all commands or absent on all
 - The identifier for the message in CICS temporary storage: the REQID option value.

Switching options mid-message results in the INVREQ condition or, in the case of REQID, the IGREQID condition.

- The ERASE, ERASEAUP, NLEOM, and FORMFEED options are honored if they are used on **any** of the BMS commands that contribute to the page.
- The values of the CURSOR, ACTPARTN, and MSR options for the page are taken from the most recent EXEC CICS SEND MAP command, if they are specified there, and from the map if not.
- The 3270 write control character (WCC) from the most recent EXEC CICS SEND MAP command is used. The WCC is assembled from the ALARM, FREEKB, PRINT, FRSET, L40, L64, L80, and HONEOM options in the command whenever **any** of them is specified. Otherwise, it is built from the same options in the map; options from the command are never mixed with those in the map.
- The FMHPARMS from all commands used to build the message are included.
- You can use both EXEC CICS SEND MAP and EXEC CICS SEND CONTROL commands to build a logical message, as long as the options noted above are consistent. You can also build a logical message with a combination of EXEC CICS SEND TEXT and EXEC CICS SEND CONTROL commands. (EXEC CICS SEND TEXT is an alternative to EXEC CICS SEND MAP for formatting text output, covered in “The SEND TEXT command” on page 289.) However, you cannot mix EXEC CICS SEND MAP and EXEC CICS SEND TEXT in the same message unless you are using partitions or logical device codes, subjects covered in “Partition support” on page 304 and “Logical device components” on page 312 respectively.

There are also two special forms of EXEC CICS SEND TEXT which allow combined mapping and text output, but to which other restrictions apply. See “SEND TEXT extensions: SEND TEXT MAPPED and SEND TEXT NOEDIT” on page 293 for details.

- While you are building a logical message, you can still converse with your terminal. You cannot use BMS commands to write to the terminal unless you are also routing, but you can use BMS EXEC CICS RECEIVE MAP commands and terminal control EXEC CICS SEND and RECEIVE commands.

Ending a logical message: the SEND PAGE command

When you have completed a logical message, you notify BMS with an EXEC CICS SEND PAGE command. If you used the ACCUM option, EXEC CICS SEND PAGE causes BMS to complete the current page and dispose of it according to the disposition option you established, as it has already done for any previous pages. If your disposition is TERMINAL, this last page is written to your principal facility; if SET, it is returned to the program; and if PAGING, it is written to temporary storage. If your disposition was PAGING, BMS also arranges delivery of the entire message to your principal facility. Options on the EXEC CICS SEND PAGE command govern how this is done, as explained in “PAGING options: RETAIN and RELEASE.”

A SYNCPOINT command or the end of your task also ends a logical message, implicitly rather than explicitly. Where possible, BMS behaves as if you had issued SEND PAGE before your SYNCPOINT or RETURN, but you lose the last page of your output if you used the ACCUM option. Consequently, code EXEC CICS SEND PAGE explicitly.

Finally, you can also send a logical message by deleting it with the PURGE MESSAGE command, described on page 281.

PAGING options: RETAIN and RELEASE

When you complete a logical message with a disposition of PAGING, BMS arranges to deliver the entire logical message, which it has accumulated in temporary storage. The display or printing of pages can be done inline, immediately after the EXEC CICS SEND PAGE command, but it is more common to schedule a separate task for the purpose. In either case, CICS supplies the programs required. These programs allow a terminal operator to control the display of the message, paging back and forth, displaying particular pages, and so on. When a separate task is used, it executes pseudoconversationally under transaction code CSPG. When the display is inline, the work is done (by the same CICS-supplied programs) within the task that created the message, which becomes conversational as a result.

You indicate how and when the message is sent by specifying RETAIN, RELEASE, or neither on your EXEC CICS SEND PAGE command. The most common choice, and the default, is neither. It causes CICS to schedule the CICS-supplied transaction CSPG to display the message and then return control to the task. The CSPG transaction is queued with any others waiting to be executed at your terminal, which execute in priority sequence as the terminal becomes free. In the ordinary case, where no other tasks waiting, CSPG executes as soon as the creating task ends.

Note: The terminal must be defined as allowing automatic transaction initiation for CICS to start CSPG automatically (ATI(YES) in the associated TYPETERM definition). If it is not, the operator has to enter the transaction code CSPG or one of the paging commands to get the process started when neither RELEASE nor RETAIN is specified.

The RELEASE option works similarly, but your task does not regain control after EXEC CICS SEND PAGE RELEASE. Instead, BMS sends the first page of the message to the terminal immediately. It then ends your task, as if an EXEC CICS RETURN had been executed in the highest level program, and starts a CSPG transaction at your terminal so that the operator can display the rest of the pages.

The CSPG code executes pseudoconversationally, just as it does if you specify neither RELEASE nor RETAIN, and the original task remains pseudoconversational if it was previously.

There are two other distinctions between RELEASE and using neither option:

- RELEASE allows you to specify the transaction identifier for the next input from the terminal, after the operator is through displaying the message with CSPG.
- RELEASE also permits the terminal operator to chain the output from multiple transactions (see “Terminal operator paging: the CSPG transaction”).

EXEC CICS SEND PAGE RETAIN causes BMS to send the message immediately. When this process is complete, your task resumes control immediately after the EXEC CICS SEND PAGE command. When the terminal is a display, BMS provides the same operator facilities for paging through the message as the CSPG transaction does, but within the framework of your task. The code that BMS uses for this purpose issues EXEC CICS RECEIVE commands to get the operator's display requests, and this causes your task to become conversational.

Note: If an error occurs during the processing of an EXEC CICS SEND PAGE command, the logical message is not considered complete and no attempt is made to display it. BMS discards the message in its cleanup processing, unless you arrange to regain control after an error. If you do, you can either delete the logical message with an EXEC CICS PURGE command or retry the EXEC CICS SEND PAGE. Do not retry unless the error that caused the failure has been remedied.

The AUTOPAGE option

Your EXEC CICS SEND PAGE command also tells BMS how to deliver the pages to the terminal. For display terminals, you want CSPG to send one page at a time, at the request of the terminal operator. For printers, you want to send one page after another. You control this with the AUTOPAGE or NOAUTOPAGE option on your EXEC CICS SEND PAGE command. NOAUTOPAGE lets the terminal operator control the display of pages; AUTOPAGE sends the pages in ascending sequence, as quickly as the device can accept them. If you specify neither, BMS determine which is appropriate from the terminal definition.

Note: If your principal facility is a printer, you can sometimes use a disposition of TERMINAL rather than PAGING, because successive sends to a printer do not overlay one another as they do on a display. TERMINAL has less overhead, especially if you do not need ACCUM either, and thus avoid creating a logical message.

Terminal operator paging: the CSPG transaction

The CICS-supplied paging transaction, CSPG, allows a user at a terminal to display individual pages of a logical message by entering page retrieval requests. Your systems staff define the transaction identifiers for retrieval and other requests supported by CSPG in the system initialization table; sometimes program function keys are used to minimize operator effort.

Retrieval can be sequential (next page or previous page) or random (a particular page, first page, last page). In addition to page retrieval, CSPG supports the following requests:

Page copy

Copy the page currently on display to another terminal. BMS reformats the page if the target terminal has a different page size or different formatting characteristics, provided the terminal is of a type supported by BMS.

Message query

List the messages waiting to be displayed at the terminal with CSPG. The list contains the BMS-assigned message identifier and, for a routed message, the message title, if the sender provided one.

Purge message

Delete the logical message.

Page chaining

Suspend the current CSPG transaction after starting to display a message, execute one or more other transactions, and then resume the original CSPG display. An intervening transaction may itself produce BMS or terminal output. If this output is a BMS logical message created with the PAGING and RELEASE or RETAIN options, this message is “chained” to the original one, and the operator can switch between one and the other.

Switch to autopage

Switch from NOAUTOPAGE display mode to AUTOPAGE mode. For terminals that combine a keyboard and hard copy output, this allows an operator to purge or print a message based on inspection of specific pages.

The process of examining pages continues until the operator signals that the message can be purged. CSPG provides a specific request for this purpose, as noted above. If the EXEC CICS SEND PAGE command contained the option OPERPURGE, this request is the only way to delete the message and get control back from CSPG.

If OPERPURGE is not present, however, any input from the terminal that is not a CSPG request is interpreted as a request to delete the message and end CSPG. If the message was displayed with the RETAIN option, the non-CSPG input that terminates the display can be accessed with a BMS or terminal control RECEIVE when the task resumes execution. See the *CICS-Supplied Transactions* manual for detailed information about the CSPG transaction.

Changing your mind: the PURGE MESSAGE command

You also can delete an incomplete logical message if for some reason you decide not to send it. You use the EXEC CICS PURGE MESSAGE command in place of EXEC CICS SEND PAGE. EXEC CICS PURGE MESSAGE causes BMS to delete the current logical message and associated control blocks, including any pages already written to CICS temporary storage. You can create other logical messages subsequently in the same task, if you wish.

Logical message recovery

Logical messages created with a disposition of PAGING are kept in CICS temporary storage between creation and delivery. BMS constructs the temporary storage queue name for a message from the 2-character REQID on the EXEC CICS SEND commands, followed by a 6-position number to maintain uniqueness. If you do not specify REQID, BMS uses a value of two asterisks (**).

Temporary storage can be a recoverable resource, and therefore logical messages with a disposition of PAGING can be recovered after a CICS abend. In fact, because CICS bases the recoverability of temporary storage on generic queue names, you can make some of your messages recoverable and others not, by your choice of REQID for the message. The conditions under which logical messages are recoverable are covered in more detail in the *CICS Recovery and Restart Guide*.

Routed messages are eligible for recovery, as well as messages created for your principal facility. We explain routing in “Message routing: the ROUTE command” on page 294.

Page formation: the ACCUM option

The ACCUM option allows you to build your output cumulatively, from any number of EXEC CICS SEND MAP commands and less-than-page-size maps. Without it, each EXEC CICS SEND MAP command corresponds to one page (if the disposition is PAGING), or a whole message (if TERMINAL or SET). With ACCUM, however, BMS formats your output but does not dispose of it until either it has accumulated more than fits on a page or you end the logical message. You can intercept page breaks if you wish, or you can let BMS handle them automatically. “Floating maps: how BMS places maps using ACCUM” and “Page breaks: BMS overflow processing” on page 283 explain how BMS arranges maps on a page with ACCUM and what you can do at a page break.

Page size is determined by the PAGESIZE or ALTPAGE value in the terminal definition. PAGESIZE is used if the PROFILE under which your transaction is running specifies the default screen size, and ALTPAGE is used if it indicates alternate screen size. (Unlike screen size, page size is not affected by the DEFAULT and ALTERNATE options that you can include with the EXEC CICS ERASE command.)

Floating maps: how BMS places maps using ACCUM

In our example map on page 234, we described placing maps on a screen or page absolutely, by specifying the number of the line and column for the upper left corner. However, maps can **float**. That is, they can be positioned relative to maps already written to the same page and to any edge of the page. Floating maps save program logic when you need to support multiple screen sizes or build pages piecemeal out of headers, detail lines and trailers, where the number of detail lines depends on the data.

The BMS options that allow you to do this are:

- JUSTIFY
- HEADER and TRAILER
- Relative values (NEXT and SAME) for the LINE and COLUMN options

When you are building a composite screen with the ACCUM option, the position on the screen of any particular map is determined by:

- The space remaining on the screen at the time it is sent
- The JUSTIFY, LINE, and COLUMN option values in the map definition

The space remaining on the page, in turn, depends on:

- Maps already placed on the current page
- Whether you are participating in “overflow processing”, that is, the processing that occurs at page breaks. If you are, the sizes of the trailer maps in your map sets are also a factor.

The placement rules we are about to list apply even if you do not specify ACCUM, although JUSTIFY values of FIRST and LAST are ignored. However, without ACCUM, each EXEC CICS SEND MAP corresponds to a separate page, and thus the space remaining on the page is always the whole page.

Page breaks: BMS overflow processing

When you build a mapped logical message, you can ask BMS to notify you when a page break is about to occur; that is, when the map you just sent does not fit on the current page. This is very useful when you are forming composite pages with ACCUM. It allows you to put trailer maps at the bottom of the current page and header maps at the top of the next one, number your pages, and so on.

BMS gives your program control at page breaks if either:

- You have issued an EXEC CICS HANDLE CONDITION command naming a label for the OVERFLOW condition, or
- You specify the NOFLUSH option with either the RESP or the NOHANDLE option on your EXEC CICS SEND MAP commands.

Either of these actions has two effects:

- The calculation of the space remaining on the page changes. Unless the map you are sending is itself a trailer map, BMS assumes that you will eventually want one on the current page. It therefore reserves space for the largest trailer in the same map set. (The largest trailer map is the one containing the TRAILER option that has the most lines.) If you do not intercept page breaks (or if you send a trailer map), BMS uses the true end of the page to determine whether the current map fits.
- The flow of control changes if the map does not fit on the current page. On detecting this situation, BMS raises the OVERFLOW condition. Then it returns control to your task **without** processing the EXEC CICS SEND MAP command that caused the overflow. Control goes to the location named in the EXEC CICS HANDLE CONDITION command if you used one. With NOFLUSH, control goes to the statement after the EXEC CICS SEND MAP as usual, and you need to test the RESP value or EIBRESP in the EIB to determine whether overflow occurred.

When your program gets control after overflow, it should:

- Add any trailer maps that you want on the current page. BMS will have left room for the one with the most lines in the map set you just used. If this is not the right number of lines to reserve, or if you are using several map sets, you

can ensure the proper amount by including a dummy map in any map set that may apply. The dummy map must specify TRAILER and contain the number of lines you wish to reserve; you do not need to use it in any EXEC CICS SEND MAP commands.

- Write any header maps that you want at the top of the next page.
- Resend the map that caused the overflow in the first place. Keep track of the data and map name at the time the overflow occurs; BMS does not save this information for you. Note that if you have several EXEC CICS SEND MAP commands which might cause overflow, the program logic required to determine which one you need to reissue is more complex if you use EXEC CICS HANDLE CONDITION OVERFLOW than if you use NOFLUSH.

Once OVERFLOW is turned on, BMS suspends returning control to your program when the output does not fit on the current page, although it still uses overflow rules for calculating the remaining space. OVERFLOW remains on until BMS processes the first EXEC CICS SEND MAP naming a map which is **not** a header or a trailer. This allows you to send your trailers and headers without disabling your EXEC CICS HANDLE CONDITION for OVERFLOW or changing your response code tests, and reinstates your overflow logic as soon as you return to regular output. (Resending the map that originally caused overflow is usually the event that turns off the overflow condition.)

If you do not intercept overflows, BMS does not notify your program when a page break occurs. Instead, it disposes of the current page according to the disposition option you have established and starts a new page for the map that caused the overflow.

Map placement rules

The primary placement of maps on the screen is from top to bottom. You can place maps side-by-side where space permits, provided you maintain the overall flow from top to bottom. The precise rules for a given EXEC CICS SEND MAP ACCUM command are as follows:

1. The highest line on which the map might start is determined as follows:
 - a. If the map definition contains JUSTIFY=FIRST, BMS goes immediately to a new page (at Step 5), unless the only maps already on the page are headers placed there during overflow processing. In this case, BMS continues at Step 1c.
 - b. If the map specifies JUSTIFY=LAST, BMS starts the map on the lowest line that allows it to fit on the page. If the map is a trailer map or you are not intercepting overflows or you are already in overflow processing, BMS uses all the space on the page. Otherwise, BMS places the map as low on the page as it can while still retaining room for the largest trailer map. If the map fits vertically using this starting line, processing continues at Step 3 (the LINE option is ignored if JUSTIFY=LAST); if not, overflow occurs (Step 5).

Note: JUSTIFY=BOTTOM is the same as JUSTIFY=LAST for output operations with ACCUM. (There are differences without ACCUM and for input mapping; see "Input from a composite screen" on page 286 and the *CICS Application Programming Reference* manual).
 - c. If there is no vertical JUSTIFY value (or after any overflow processing caused by JUSTIFY=FIRST has been completed), the LINE operand is

checked. If an absolute value for LINE is given, that line is used, provided it is at or below the starting line of the map most recently placed on the page. If the value is above that point, BMS goes to a new page at Step 5.

If LINE=NEXT, the first completely unused line (below all maps currently on the page) is used. If LINE=SAME, the starting line of the map sent most recently is used.

2. BMS now checks that the map fits vertically on the screen, given its tentative starting line. Here again, BMS uses all of the space remaining if the map is a trailer map, if you are not intercepting overflows or if you are already in overflow processing. Otherwise, BMS requires that the map fit and still leave space for the largest trailer map. If the map does not fit vertically, BMS starts a new page (Step 5).

3. Next, BMS checks whether the map fits horizontally, assuming the proposed starting line. In horizontal positioning, the JUSTIFY option values of LEFT and RIGHT come into play. LEFT is the default, and means that the COLUMN value refers to the left-hand side of the map. A numeric value for COLUMN tells where the left edge of the map should start, counting from the left side of the page. COLUMN=NEXT starts the map in the first unused column from the left on the starting line. COLUMN=SAME means the left-hand column of the map most recently placed on the screen which also specified JUSTIFY=LEFT and which was not a header or trailer map.

JUSTIFY=RIGHT means that the COLUMN value refers to the right-hand edge of the map. A numeric value tells where the right edge of the map should start, **counting from the right**. COLUMN=NEXT means the first available column from the right, and COLUMN=SAME is the right-hand column of the map most recently placed which had JUSTIFY=RIGHT and was not a header or trailer.

If the map does not fit horizontally, BMS adjusts the starting line downward, one line at a time, until it reaches a line where the map does fit or overflow occurs. Processing resumes with the vertical check (Step 2) after each adjustment of the starting line.

4. If the map fits, BMS adds it to the current page and updates the available space, using the following rules:

- Lines above the first line of the map are completely unavailable.
- If the map specifies JUSTIFY=LEFT, the columns from the left edge of the page through the right-most column of the map are unavailable on the lines from the top of the map through the last line on the page that has anything on it (whether from this map or an earlier one).
- If the map specifies JUSTIFY=RIGHT, the columns between the right-hand edge of the page and the left-hand edge of the map are unavailable on the lines from the top of the map through the last line of the page that has anything on it.

Figure 59 on page 286 shows how the remaining space is reduced with each new map placed.

5. When the current map does not fit on a page, BMS determines whether it should return control to your program. If you have asked for control at overflow and you are not already in overflow processing, BMS returns control as described in "Page breaks: BMS overflow processing" on page 283. Otherwise, BMS disposes of the current page according to the disposition

option you have established, starts a new page, and resumes processing for the map that would not fit at Step 1.




Figure 59. Successive placement of maps on a page, showing space made unavailable by each

ASSIGN options for cumulative processing

To help you manage the complexities of building a composite screen, CICS provides EXEC CICS ASSIGN command options that relate specifically to cumulative processing:

- MAPCOLUMN
- MAPHEIGHT
- MAPLINE
- MAPWIDTH

All apply to the map most recently sent. MAPHEIGHT and MAPWIDTH are the size (number of rows and columns) and MAPLINE and MAPCOLUMN are the origin of the map (the position of the upper left corner).

Input from a composite screen

You can read mapped input from a screen built from multiple maps, but there are restrictions. First, you can specify only one map in your EXEC CICS RECEIVE MAP command, whereas the screen may have been written with several.

Second, BMS cannot know how to position a floating map for input and assumes that the map in your EXEC CICS RECEIVE MAP command was written to an empty screen. Thus LINE or COLUMN values of NEXT or SAME are interpreted differently on input than on output. JUSTIFY=LAST is ignored altogether; use JUSTIFY=BOTTOM if you want to place a map at the bottom of the screen and

read data back from it. See the *CICS Application Programming Reference* manual for the exact rules.

Performance considerations

There are three components to the overall efficiency of the part of your application that the end user sees: processor pathlength, communications line utilization, and user time. Pathlength and line time used to be paramount, and much design and programming effort has been invested in minimizing them.

As online systems have evolved, however, the emphasis has shifted steadily to making things as easy, pleasant and quick for the user as possible, at the expense of the other factors if necessary. Also, as processors have become cheaper, designers have been willing to expend cycles to reduce programming and maintenance effort as well as to make the user interface better.

We have already given you references on elements of good design for the user interface, in “Personal computers” on page 209, and usually these should be your overriding considerations. In this section, we point out some ways that you can reduce pathlength and line time as well. You need to judge for yourself whether and to what extent the savings justify extra programming effort or a less general design.

Minimizing pathlength

Ordinarily, the number of instructions executed in a single CICS command is large in comparison to the number of instructions in the application program that invoked it. Consequently, the pathlength for a given task ordinarily depends more on the number and type of CICS commands than on anything else, and commands are the most fertile area for tuning. Commands vary by type, of course, and pathlength for any given command may vary considerably with circumstances.

For BMS, some recommendations are:

- Build screens (pages) with a single command when practical. Avoid building a composite screen with the ACCUM feature when a modest amount of additional programming accomplishes the same function, and avoid building a composite screen by multiple physical writes, as described in “Outside the map” on page 252, except in unusual circumstances.
- Avoid producing more output at one time than the user is likely to inspect. Some transactions—inquiries, especially—produce many pages of output for certain input values. When this happens, the user usually narrows the search and repeats the inquiry, rather than page through the initial output. To avoid the pathlength of producing output that is never viewed, you can limit it to some reasonable number of pages, inform the user on the last page that there is more, and save the information required to restart the search from that point if the user requests it. The extra programming is minimal; see Chapter 4, “Sharing data across transactions” on page 83 for ways to save the restart data.
- Use commands that are on the BMS “fast path” if possible. (See “BMS support levels” on page 226 for the commands and terminal types that qualify.)
- Use terminal control commands for very simple inputs and outputs, where you do not need BMS formatting or other function. If pathlength is critical, you may want to use terminal control entirely. However, the advantages of BMS over

terminal control in terms of flexibility, initial programming effort and maintainability are significant, and usually outweigh the pathlength penalty.

Reducing message lengths

You can take advantage of 3270 hardware to reduce the length of both inbound and outbound messages. If the bandwidth in any link between the terminal and the processor is constrained, you will get better response overall with shorter messages. However, the time for any given transmission depends on the behavior of other users of those links at the time, and so you may not see improvement directly. Here are some of the possibilities for reducing the length of a 3270 datastream:

- Avoid turning on MDTs unnecessarily when you send a screen, because they cause the associated input fields to be transmitted on input. Ordinarily, you do not need to set the tag on, because the hardware does this when the user enters input into the field. The tag remains on, no matter how many times the screen is transmitted, until explicitly turned off by program (by FRSET, ERASEAUP, or ERASE, or by an override attribute byte with the tag off). The only time you need to set it on by program is when you want to store data on the screen in a field that the user does not modify, or when you highlight a field in error and you want the field returned whether or not the user changes it. In this case you need to turn on the MDT as well as the highlighting.
- Use FRSET to reset the MDTs when you do not want input on a screen retransmitted (that is, when you have saved it and the user does not need to change it on a subsequent transmission of the same screen). (See “Saving the good input” on page 263 for more.)
- Do not initialize input fields to blanks when you send the screen because, on input, blanks are transmitted and nulls are not. Hence the data stream is shortened by the unused positions in each modified field if you initialize with nulls. The appearance on the screen is the same, and the data returned to the program is also the same, if you map the input.
- For single-screen data entry operations, use ERASEAUP to clear data from the screen, rather than resending the screen.
- If you are updating a screen, send only the changed fields, especially if the changes are modest, as when you highlight fields in error or add a message to the screen. In BMS, you can use the DATAONLY option, both to shorten the data stream and reduce the pathlength (see “DATAONLY option” on page 248). To highlight a field, in fact, you send only the new attribute byte; the field data will remain undisturbed on the screen.
- If you are using terminal control commands, format with set buffer address (SBA) and repeat-to-address (RA) orders, rather than spacing with blanks and nulls. (BMS does this for you.)

Formatting text output

If the output you are sending to the terminal is simply text, and you do not need to format the screen for subsequent input, you do not need to create a map. BMS provides a different command expressly for this purpose: EXEC CICS SEND TEXT, which formats without maps.

When you use EXEC CICS SEND TEXT, BMS breaks your output into pages of the proper width and depth for the terminal to which it is directed. Lines are broken at

word boundaries, and you can add header and trailer text to each page if you wish. Page size is determined as it is for other BMS output (see “Ending a logical message: the SEND PAGE command” on page 279).

The SEND TEXT command

Except for the different type of formatting performed, the EXEC CICS SEND TEXT command is very similar to EXEC CICS SEND MAP. You specify the location of the text to be formatted in the FROM option and its length in the LENGTH option. Nearly all the options that apply to mapped output apply to text output as well, including:

Device controls

FORMFEED, ERASE, PRINT, FREEKB, ALARM, CURSOR.

Formatting options

NLEOM, L40, L64, L80, HONEOM.

Disposition options

TERMINAL, PAGING, SET.

Page formation option

ACCUM.

In general, these options have the same meaning on an EXEC CICS SEND TEXT command as they do on an EXEC CICS SEND MAP command, but always look up the *CICS Application Programming Reference* manual for the precise meanings. The EXEC CICS SEND TEXT command itself requires **standard** BMS; options like ACCUM, PAGING and SET that require **full** BMS in a mapped environment also require full BMS in a text environment.

Options for EXEC CICS SEND TEXT correspond to functions associated with the map in an EXEC CICS SEND MAP context. These are HEADER, TRAILER, JUSTIFY, JUSFIRST, and JUSLAST. We explain how they work in “Page format for text messages” on page 290.

Two EXEC CICS SEND MAP options that do not carry over to EXEC CICS SEND TEXT are ERASEAUP and NOFLUSH. ERASEAUP does not apply because text uses fields only minimally, and NOFLUSH does not apply because BMS does not raise the OVERFLOW condition on text output.

Text logical messages

The presence of either the ACCUM or PAGING option on an EXEC CICS SEND TEXT command signals BMS that you are building a logical message, just as it does in an EXEC CICS SEND MAP command. Text logical messages are subject to the same rules as mapped logical messages (see page 278). In particular, you can use both EXEC CICS SEND TEXT and EXEC CICS SEND CONTROL commands to build your message, but you cannot mix in EXEC CICS SEND MAP commands, except as noted there. You also end your message in the same way as a mapped message (see “BMS logical messages” on page 277).

Page format for text messages

Page formation with EXEC CICS SEND TEXT is somewhat different from page formation with EXEC CICS SEND MAP. First, a single EXEC CICS SEND TEXT command can produce more output than will fit on a screen or a printer page (EXEC CICS SEND MAP never does this). BMS sends the whole message, which means that you can deliver a multi-page message to a printer without using logical facilities. You cannot use the same technique for displays, however, because even though BMS delivers the whole message, the component screens overlay one another, generally too quickly for anyone to read.

If you specify ACCUM, BMS breaks the output into pages for you, and the second difference is that unless you specify a disposition of SET, your task does not get control at page breaks. Instead, when the current page has no more room, BMS simply starts a new one. It adds your header and trailer, if any, automatically, and does not raise the OVERFLOW condition. This is true whether you produced the pages with a single EXEC CICS SEND TEXT command or you built the message piecemeal, with several. The only situation in which your task is informed of a page break is when the disposition is SET. In this case, BMS raises the RETPAGE condition to tell you that one or more pages are complete, as explained in "Using SET" on page 302.

Here are the details of how BMS builds pages of text with ACCUM:

1. Every message starts on page 1, which is initially empty.
2. If you specify the HEADER option, BMS starts every page with your header text. BMS numbers your pages in the header or trailer if you wish. (Header format and page numbering are explained on page 292.)
3. If you specify one of the justification options (JUSTIFY, JUSFIRST, JUSLAST), BMS starts your text on the indicated line. JUSFIRST begins your text on the first line after the header, or the top line if there is no header. JUSTIFY=n starts your text on line n, and JUSLAST starts it on the lowest line that allows both it and your trailer (if any) to fit on the current page. If the contents of the current page prevent BMS from honoring the justification option there, BMS goes to a new page first, at step 6.

Justification applies only to the start of the data for each EXEC CICS SEND TEXT command; when the length of your data requires an additional page, BMS continues your text on it in the first available position there.

4. If you do not specify justification, BMS starts your text in the first position available. On the first EXEC CICS SEND TEXT of the message, this works out the same as JUSFIRST. Thereafter, your text follows one character after the text from the previous EXEC CICS SEND TEXT of the current logical message. (The intervening character is an attributes byte on 3270 terminals, a blank on others).
5. Having determined the starting position, BMS continues down the page, breaking your data into lines as explained in "How BMS breaks text into lines" on page 291, until it runs out of space or data. If you have specified a trailer, the available space is reduced by the requirement for the trailer. If the data is exhausted before the space, processing ends at this point. The message is completed when you indicate that you are finished with an EXEC CICS SEND PAGE or EXEC CICS PURGE MESSAGE command.

6. If your text does not fit on the current page, BMS completes it by adding your trailer text, if any, at the bottom and disposes of it according to the disposition option you have established (TERMINAL, PAGING, or SET), just as it does for a mapped logical message. The trailer is optional, like the header; you use the TRAILER option to specify it (see “Header and trailer format for text messages” on page 292).
7. BMS then goes to a new page and repeats from step 2 with the remaining data.

How BMS breaks text into lines

In breaking the text into lines, BMS uses the following rules:

1. Ordinarily, each line starts with what appears to be a blank. On a 3270 device, this is the attributes byte of a field that occupies the rest of the line on the screen or printed page. For other devices, it is simply a blank or a carriage control character.

An exception occurs if the task creating the output is running under a PROFILE that specifies PRINTERCOMP(YES) and the output device is a 3270 printer. In this case, no character is reserved at the beginning of each line. See “PRINTERCOMP option” on page 346.

2. BMS copies your text character for character, including all blanks, with two exceptions that occur at line end:
 - If a line ends in the middle of a word, BMS fills out the current line with blanks and places the word that would not fit in the first available position of the next line. For this purpose, a “word” is any string of consecutive nonblank characters.
 - If two words are separated by a single blank, and the first one fits on the current line without leaving room for the blank, the blank is removed and the next line starts at the beginning of the second word.
3. You can embed new-line (NL) characters and other print format orders as well as blanks to control the format, if the destination terminal is a printer. NLS and tabs are particularly useful with columnar data, and BMS does not filter or even interpret these characters. However, print format orders do not format displays; see “3270 printers” on page 342 for more information about using them.
4. You can also include set attribute (SA) order sequences in your output. (Each one sets the attributes of a single character in the data stream, as explained in “The set attribute order” on page 218.) BMS abends the task unless SA sequences are exactly three bytes long and represent a valid attribute type. However, if you use a valid SA sequence to a terminal that does not support the attribute, BMS removes the SA sequence and then sends the message. Attributes set with SA orders remain until overridden by subsequent orders or until another EXEC CICS SEND TEXT command, which resets them to their default values.

Do not include 3270 orders other than SA in your text. BMS treats them as display data and they do not format as intended; they may even cause a terminal error.

Header and trailer format for text messages

To place a header on the pages of a text message, you point to a block of data in the following format in the HEADER option:

refid=frame=none place=inline.



You use the same format for trailer text, but you point to it with the TRAILER option. Here:

LL

is the length of the header (trailer) data, not including the four bytes of LL, P, and C characters. LL should be expressed in halfword binary form.

P

is the page-number substitution character (see PNFLD below). Use a blank if you do not want page numbers.

C

is a reserved 1-byte field.

TEXT

is the header (trailer) text to be placed at the top (bottom) of each page of output. Use new-line characters (X'15') to indicate where line breaks should occur if you want multiple lines.

PNFLD

is the page number field within your header (trailer) text. If you want to number the pages of your output, choose a character that does not otherwise appear in your header (trailer) text. Place this character in the positions where the page number is to appear. You can use from one to five adjacent positions, depending on how large you expect your page numbers to get (32,767 is the maximum BMS allows). Place the same character in the P field above, to tell BMS where to make the substitution. Do not use X'0C', X'15', X'17', X'26', or X'FF' for P; these values are reserved for other purposes. If you do not want page numbering, simply place a blank (X'40') in P.

When you are building a logical message, repeat your HEADER and TRAILER options on each EXEC CICS SEND TEXT command, so that they are present when the page breaks occur, and specify the trailer again on the EXEC CICS SEND PAGE command that terminates the message.

Here is an example of a COBOL definition for a header that simply numbers the pages, leaving room for a number up to 99.

```
EXEC CICS SEND TEXT FROM (OUTPUT-AREA)
      HEADER(HEADER-TEXT) PAGING ACCUM END-EXEC.
```

where:

```
01 HEADER-TEXT
   02 HEADER-LL          PIC S9(4) COMP VALUE +11.
   02 HEADP             PIC X VALUE '@'.
   02 FILLER            PIC X VALUE LOW-VALUE.
   02 HEADING          PIC X(11) VALUE 'PAGE NO. @@'.
```

Screens built with EXEC CICS SEND TEXT are not designed for extensive input from the terminal operator. However, you can interpret the attention identifier and read simple inputs—such as those used in the CSPG transaction to control the page display—if the field structure on the screen is suitable and the operator knows or can see what is expected. (A new field starts at each line, as well as at the first character of the text sent with each EXEC CICS SEND TEXT command that made up the message. The fields defined are unprotected, alphameric and normal intensity, so that the operator can key into them.) Normally a terminal control RECEIVE is used in this situation; you can use EXEC CICS RECEIVE MAP only if you can build a map with a field structure matching that of the screen. However, you can interpret the attention identifier and read simple inputs—such as those used in the CSPG transaction to control the page display—if the field structure on the screen is suitable and the operator knows or can see what is expected. (A new field starts at each line, as well as at the first character of the text sent with each EXEC CICS SEND TEXT command that made up the message. The fields defined are unprotected, alphameric and normal intensity, so that the operator can key into them.) Normally a terminal control RECEIVE is used in this situation; you can use EXEC CICS RECEIVE MAP only if you can build a map with a field structure matching that of the screen.

SEND TEXT extensions: SEND TEXT MAPPED and SEND TEXT NOEDIT

BMS provides two special forms of the EXEC CICS SEND TEXT command that allow you to use some of the message delivery facilities of BMS for output that is already formatted. EXEC CICS SEND TEXT MAPPED sends a **page** of device-dependent data **previously built by BMS** and captured with the SET option. You may have used either EXEC CICS SEND MAP or EXEC CICS SEND TEXT commands to build the page originally. See “Using SET” on page 302 for details.

EXEC CICS SEND TEXT NOEDIT is similar, but is used to send a page of device-dependent output built by the program or some method other than BMS.

You can deliver such pages to your own principal facility individually, using a disposition of TERMINAL, or you can include them in a logical message built with the PAGING option. In a logical message, these forms can be mixed with ordinary EXEC CICS SEND TEXT commands or with EXEC CICS SEND MAP⁷ commands, as long as each BMS SEND represents a separate page (that is, the ACCUM option is not used).

You can also use these commands in a routing environment (described in “Message routing: the ROUTE command” on page 294). Whether you are routing or sending to your own terminal, be sure that the data stream is appropriate to the destinations; BMS does not check before transmission, other than to remove 3270 attributes that the destination does not support.

None of the page-formatting options, ACCUM, JUSTIFY, JUSFIRST, JUSLAST, HEADER, and TRAILER, applies to either of these commands, because the page is already formatted and built, by definition.

⁷ The usual restriction against mixing text with mapped output in the same logical method does not apply here, because the page is already formed.

The primary difference between the MAPPED and NOEDIT forms is that EXEC CICS SEND TEXT MAPPED uses the four-byte page control area (PGA) that BMS appends to the end of pages returned via the SET option. This area tells BMS the write command and write control character to be used, which extended attributes were used on the page, and whether the page contains formfeeds, data in SCS format, 14-or 16-bit buffer addresses, structured fields and FMHs. It allows BMS to deliver a page built for one device to one with different hardware characteristics, as might be required for a page copy or a routing operation. With EXEC CICS SEND TEXT NOEDIT, you specify this type of information on the command itself. Use EXEC CICS SEND TEXT MAPPED for output created with BMS, and NOEDIT for output formatted by other means. You cannot include structured fields in the output with either EXEC CICS SEND TEXT MAPPED or EXEC CICS SEND TEXT NOEDIT, incidentally; use a terminal control SEND for such output.

The LENGTH option for an EXEC CICS SEND TEXT MAPPED command should be set from the TIOATDL value returned when the page was built, which does not include the PGA (see “Using SET” on page 302). If you copy the page for later use with EXEC CICS SEND TEXT MAPPED, however, you must be sure to copy the PGA as well as the page itself (TIOATDL + 4 bytes in all).

Message routing: the ROUTE command

The message routing facilities of BMS allow you to send messages to terminals other than the principal facility of your task (your task does not even need to have a principal facility). Routing does not give your task direct control of these terminals, but instead causes the scheduling of a task for each destination to deliver your message. These tasks execute the CICS-supplied transaction CSPG, the same one used for delivery of messages with a disposition of PAGING to your own terminal. Thus the operator at a display terminal who receives a routed message uses CSPG requests to view the message. (See “Terminal operator paging: the CSPG transaction” on page 280 for more information about CSPG.)

Message routing is useful for message-switching and broadcasting applications, and also for printing (see “CICS printers: getting the data to the printer” on page 349). It is the basis for the CICS-supplied transaction CMSG, with which terminal users can send messages to other terminals and users. The *CICS-Supplied Transactions* manual explains how to use CMSG and what you can do with it.

How routing works

To route a message, start by issuing an EXEC CICS ROUTE command. This command tells BMS where to send the message, when to deliver it, what to do about errors, and other details. Then you build your message. It can be a mapped or text message, but it must be a logical message (that is, either ACCUM or PAGING present), and the disposition must be either PAGING or SET, not TERMINAL. PAGING is the more common choice and is assumed in the discussion that follows. We explain SET in a routing context in “Routing with SET” on page 302.

Your EXEC CICS ROUTE command is in effect until you end your message with an EXEC CICS SEND PAGE command, and you must not issue another one until this occurs. (If you issue EXEC CICS ROUTE while building your message you get an invalid request response; if you issue a second EXEC CICS ROUTE before you

start your logical message, it simply replaces the first one.) You can also terminate a message with EXEC CICS PURGE MESSAGE, if you decide you do not want to send it. EXEC CICS PURGE MESSAGE causes the routing environment established with your EXEC CICS ROUTE command to be discarded, along with your logical message.

Specifying destinations for a routed message

You can specify destinations for your routed message in three different ways:

- You can request that certain classes of operators receive the message, by using the OPCLASS option of the EXEC CICS ROUTE command. Classes are associated with an operator in the ESM user definition.
- You can name particular operators who are to receive the message by using a route list, to which you point with the LIST option of the EXEC CICS ROUTE command. Operators are identified by a three-character OPIDENT value, which is also assigned in the ESM user definition.
- You can name particular terminals which are to receive the message; this is also done with a route list. Terminals are identified by their four-character TERMID value, and, for terminal types to which they apply, a two-character logical device code.

Note: If you need to know the identifier or operator class values for the operator signed on at your principal facility to specify the destination of your routed message, you can use the EXEC CICS ASSIGN command with the OPID or OPCLASS options to find out.

Eligible terminals

To format a message properly for a particular destination, BMS needs to know the characteristics of the terminal for which it is formatting. This is true even for destinations that you designate by operator or operator class. The first step in processing a route list, therefore, is to translate your destinations to a list of **terminals** to which the message **may** be delivered. This “eligible terminal” list combines the information in your route list and your OPCLASS specification with the state of the terminal network **at the time of the EXEC CICS ROUTE command**.

Later, when your message is ready for delivery, BMS uses this list to decide which terminals actually get your message. A terminal must be on the list to receive the message, but its presence there does not guarantee delivery. There may be operator restrictions associated with the terminal and, because delivery occurs later in time, the status or even the nature of the terminal may have changed.

Both at the time the list is built and at the time of delivery, BMS is restricted to the terminal definitions installed in its own CICS region (where the routing task is running, or ran) and may not have all of the terminal definitions you expect. First, terminals that are autoinstalled may not be logged on either at the time of the EXEC CICS ROUTE, excluding them from inclusion on the list, or at the times sending is attempted, preventing delivery.

In a multiple-region environment, there is the additional possibility that terminals known to one region may not be known to another. (It depends on how they are defined, as explained in the *CICS Resource Definition Guide*.) In particular, if a terminal definition is shared among regions by designating it as SHIPPABLE in the region that owns it, the terminal is not defined in any other region until something

occurs to cause shipment there. This usually happens the first time the terminal routes a transaction to the region in question. Consequently, an EXEC CICS ROUTE in this region cannot include the terminal before the first such event occurs.

The following sections describe how BMS builds the list of eligible terminals. This occurs at the time of the EXEC CICS ROUTE command:

Destinations specified with OPCLASS only

If you specified operator classes (the OPCLASS option) but no route list, BMS scans all the terminal definitions in the local system. Any terminal that meets all these conditions gets on the eligible terminal list:

- The terminal is of a type supported by BMS
- The terminal can receive routed messages not specifically addressed to it (ROUTEDMSGGS (ALL) in the terminal definition)
- An operator is signed on at the terminal
- The operator belongs to one of the operator classes in your OPCLASS list.

The resulting entry is marked so that delivery occurs only when and if an operator belonging to at least one of the operator classes in your OPCLASS list is signed on. (This operator does not have to be the one that was signed on at ROUTE time.)

OPCLASS and LIST omitted

If you specify neither operator classes nor a route list, BMS puts every terminal that meets the first two tests above on the list, and sets no operator restrictions on delivery. In a network where many terminals are eligible to receive all routed messages, this is a choice you almost certainly want to avoid.

Route list provided

If you provide a route list, BMS builds its list from yours instead of scanning the terminal definitions. Each of your entries is processed as follows. Processing includes setting a status flag in the list entry to tell you whether the entry was used or skipped and why.

- If the entry contains a terminal identifier but no operator identifier, the terminal goes on the eligible list, provided it is defined, of a type supported by BMS, and eligible to receive routed messages. If BMS cannot find the terminal definition, it sets the “entry skipped” and “invalid terminal identifier” bits (X'CO') in the status flag of the route list entry; if the terminal exists but is not supported by BMS or is not allowed to receive any routed messages, the “entry skipped” and “terminal not supported under BMS” bits get set (X'A0').

Note: The eligibility of a terminal to receive routed messages is governed by the ROUTEDMSGGS option in the terminal definition. Three values are possible: a terminal may be allowed to receive all routed messages, only messages routed to it by terminal or operator name, or no routed messages at all. If you specified OPCLASS as well as a route list, BMS checks whether an operator belonging to one of the classes you listed is signed on at the terminal. If not, BMS sets the “operator not signed on” bit (X'10') in the status flag for the entry to inform you, but includes the terminal anyway. There are no operator restrictions associated with the list entry, even when you specify operator classes.

- If the entry contains both a terminal and an operator identifier, the terminal identifier is checked in the same way as it is without an operator identifier, and

the same errors can occur. If the terminal passes these tests, it goes on the eligible list. However, the entry is marked such that the message can be delivered only when the operator named is signed on at the same terminal.

If this operator is not signed on to the terminal at the time of the EXEC CICS ROUTE command, BMS notifies you by turning on the “operator not signed on” bit (X'10') in the status flag, but the terminal goes on the delivery list regardless of sign-on status. (OPCLASS is ignored entirely when an operator identifier is present.)

- If the entry contains only an operator identifier, BMS searches the terminal definitions until it finds one where the operator is signed on. (The operator may be signed on at additional terminals, but BMS ignores these.) If this terminal is of a type not supported by BMS, or if the terminal cannot receive routed messages, BMS sets the “entry skipped” and “operator signed on at unsupported terminal” bits (X'88') in the status flag. It also fills in the terminal identifier in your route list. If the terminal is suitable, BMS treats the entry as if you had specified both that terminal and operator identifier, as described above.

If the operator is not signed on anywhere, BMS sets the “entry skipped” and “operator not signed on” bits (X'90') in the status flag.

Route list format

BMS requires a fixed format for route lists. Each entry in the list is 16 bytes long, as follows:

Bytes	Contents
0-3	Terminal or logical unit identifier (four characters, including trailing blanks), or blanks
4,5	LDC mnemonic (two characters) for logical units with LDC support, or blanks
6-8	Operator identifier, or blanks
9	Status flag for the route entry
10-15	Reserved; must contain blanks

Either a terminal or an operator identifier must be present in each entry. An LDC may accompany either; see “LDCs and routing” on page 313 for more information about LDCs.

The entries in the route list normally follow one another in sequence. However, they do not all have to be adjacent. If you have a discontinuity in your list, you end each group of successive entries except the last group with an 8-byte chain entry that points to the first entry in the next group. This entry looks like this:

Bytes	Contents
0,1	-2 in binary halfword format (X'FFFE')
2,3	Reserved
4-7	Address of the first entry in the next group of contiguous entries

The end of the entire list is signaled by a two-byte entry containing a halfword value of -1 (X'FFFF').

Your list may consist of as many groups as you wish. There is an upper limit on the total number of destinations, but it depends on many variables; if you exceed it, BMS will abend your task with abend code ABMC.

On return from a ROUTE command, BMS raises condition codes to signal errors in your list:

RTESOME means that at least one of the entries in your route list could not be used and was skipped. The default action is to continue the routing operation, using the destinations that were processed successfully.

RTEFAIL means that none of the destinations in your list could be used, and therefore no routing environment has been set up. The default action is to return control to your task. Test for this condition because, with no routing environment, a subsequent BMS SEND command will go to the principal facility, which is probably not your intention.

In addition to the general information reflected by RTESOME and RTEFAIL, BMS tells you what it did with each entry in your list by setting the status flag (byte 9). A null value (X'00') means that the entry was entirely correct. The high-order bit tells you whether the entry was used or skipped, and the other bits tell you exactly what happened. Here are the meanings of each bit being on:

ENTRY SKIPPED (X'80')

The entry was not used. When this bit is on, another bit will also be on to indicate the reason.

INVALID TERMINAL IDENTIFIER (X'40')

There is no terminal definition for the terminal named in the entry. The entry is skipped.

TERMINAL NOT SUPPORTED UNDER BMS (X'20')

The terminal named in the route list entry is of a type not supported by BMS, or it is restricted from receiving routed messages. The entry is skipped.

OPERATOR NOT SIGNED ON (X'10')

The operator named in the entry is not signed on. Any of these conditions causes this flag to be set:

- Both an operator identifier and a terminal identifier were specified, and the operator was not signed on at the terminal. The entry is not skipped.
- An operator identifier was specified without a terminal identifier, and the operator was not signed on at any terminal. The entry is skipped.
- OPCLASS was specified on the ROUTE command, a terminal identifier was specified in the route list entry, and the operator signed on at the terminal did not have any of the specified operator classes. The entry is not skipped.

OPERATOR SIGNED ON AT UNSUPPORTED TERMINAL (X'08')

Only an operator identifier was specified in the route list entry, and that operator was signed on at a terminal not supported by BMS or not eligible to receive routed messages. The entry is skipped. The name of the terminal is returned in the terminal identifier field of the entry.

INVALID LDC MNEMONIC (X'04')

Either of these conditions causes this flag to be set:

- The LDC mnemonic specified in the route list is not defined for this terminal. That is, the terminal supports LDCs but it has no LDC list, or its LDC list is extended but does not contain this entry.
- The device type for this LDC entry is different from that of the first entry in the route list with an LDC (only one LDC device type is allowed, as explained in “LDCs and routing” on page 313).

The entry is skipped.

Note: CICS provides source code which defines a standard route list entry and the values you need to test status flag bit combinations. You can insert this code into your program with a COPY or INCLUDE of the member DFHURLDS, in the same way you can include the BMS attention identifier or attribute byte definitions.

Delivery conditions

We have just explained how BMS determines the terminals eligible to receive your routed message. Actual delivery occurs later in time, much later in some cases, depending on the scheduling options in your ROUTE command (INTERVAL, TIME, AFTER, and AT). You can request delivery immediately, after an interval of time has elapsed, or at a particular time of day.

When the appointed time arrives, BMS attempts to deliver the message to every terminal on the eligible terminal list. All the following conditions must be met for the message to be delivered to any particular terminal:

- The terminal must be defined as a type supported by BMS, and the same type as when the ROUTE command was processed⁸. (Where there is a long delay between creation and delivery of a message, it is possible for the terminal defined with a particular TERMID to change characteristics or disappear, especially in an autoinstall environment.)
- The terminal must be in service and available (that is, there cannot be a task running with the terminal as its principal facility).
- The terminal must be eligible for automatic transaction initiation, or the terminal operator must request delivery of the message with the CSPG transaction.

Note: If several messages accumulate for delivery to a particular terminal, there is no guarantee that the operator will view them in any particular order. In fact, the CSPG transaction allows the operator to control delivery order in some situations. If a specific sequence of pages is required, send them as one message.

- If the delivery list entry restricts delivery to a particular operator or to operators in certain classes, the operator signed on at the terminal must qualify. (See “How routing works” on page 294 for the OPCLASS and LIST specifications that produce these restrictions.)
- The purge delay must not have expired, as explained in the next section.

⁸ A 3270 terminal need not have exactly the same extended attributes that it had at the time the ROUTE command was issued, because BMS removes unsupported attributes from the data stream at the time of delivery.

Undeliverable messages

If BMS cannot deliver a message to an eligible terminal, it continues to try periodically until one of the following conditions occurs:

- A change in terminal status allows the message to be sent.
- The message is deleted by the destination terminal operator.
- The purge delay elapses.

The purge delay is the period of time allowed for delivery of a message once it is scheduled for delivery. After this interval elapses, the message is discarded. The purge delay is a system-wide value, set by the PRGDLY option in the system initialization table. Its use is optional; if the systems programmer sets PRGDLY to zero, messages are kept indefinitely.

When BMS purges a message in this fashion, it sends an error message to the terminal you specify in ERRTERM. (If you use ERRTERM without a specific terminal name, it sends the message to the principal facility of the task that originally created the message. If you omit ERRTERM altogether, no message is sent.)

Temporary storage and routing

Between creation and delivery of a routed message with a disposition of PAGING, BMS stores the message in CICS temporary storage, just as it does in the case of an ordinary PAGING message. Consequently, you can make your routed messages recoverable by your choice of the REQID option value, just as in the case of a nonrouted message. (See “Logical message recovery” on page 282.)

If you are routing to more than one type of terminal, BMS builds a separate logical message for each type, with the appropriate device-dependent data stream, and uses a separate temporary storage queue for each type.

Note: For terminal destinations that have the alternate screen size feature, where two message formats are possible, BMS chooses the default size if the profile under which the task creating the message specifies default size, and alternate size if the profile specifies alternate size.

All of the logical messages use the same REQID value, however, so that you can still choose whether they are recoverable or not.

BMS also uses temporary storage to store the list of terminals eligible to receive your message and to keep track of whether delivery has occurred. When all of the eligible terminals of a particular type have received a message, BMS deletes the associated logical message. When all of the destinations have received delivery, or the purge delay expires, BMS erases all of the information for the message, reporting the number of undeliverable messages by destination to the master terminal operator message queue.

Message identification

You can assign a title to your routed message if you wish. The title is not part of the message itself, but is included with the other information that BMS maintains about your message in CICS temporary storage. Titles are helpful in situations where a number of messages may accumulate for an operator or a terminal, because they allow the operator to control the order in which they are displayed. (See the “query” option of the CSPG command in the *CICS-Supplied Transactions* manual.)

To assign a title, use the TITLE option of the EXEC CICS ROUTE command to point to a data area that consists of the title preceded by a halfword binary length field. The length includes the two-byte length field and has a maximum of 64, so the title itself may be up to 62 characters long. For example:

```
01 MSG-TITLE.
   02 TITLE-LENGTH PIC S9(4) COMP VALUE +19.
   02 TITLE-TEXT   PIC X(17) VALUE 'MONTHLY INVENTORY'.
...
EXEC CICS ROUTE TITLE(MSG-TITLE)....
```

Programming considerations with routing

For the most part, you build a routed message in the same way you do a nonrouted message. However, because BMS builds a separate logical message for each terminal type among your destinations, there are differences. The first involves page overflow.

Routing and page overflow

Because different types of terminals have different page capacities, page overflow may occur at different times for different types. If you are using EXEC CICS SEND MAP commands and intercepting overflows, your program gets control when overflow occurs for each page of each logical message that BMS is creating in response to your EXEC CICS ROUTE command.

If you want to number your pages or do page-dependent processing at overflow time, you may need to keep track of information for each terminal type separately. Data areas kept for this purpose are called **overflow control areas**. You can tell how many such areas you need (that is, how many different terminal types appeared in your EXEC CICS ROUTE command) by issuing an EXEC CICS ASSIGN command with the DESTCOUNT option after your EXEC CICS ROUTE command and before any BMS command that could cause overflow. Issued at this time, EXEC CICS ASSIGN DESTCOUNT returns a count of the logical messages that BMS will build.

When overflow occurs, you can use the same command to determine for which logical message overflow occurred. At this time, EXEC CICS ASSIGN DESTCOUNT returns the relative number of that message among those BMS is building for this EXEC CICS ROUTE command. If you are using overflow control areas, this number tells you which one to use. If you use EXEC CICS ASSIGN PAGENUM at this time, BMS returns the number of the page that overflowed as well.

To handle the complication of different overflow points for different terminal types, the processing you need to do on overflow in a routing environment is:

- Determine which logical message overflowed with EXEC CICS ASSIGN DESTCOUNT (unless you are doing very simple overflow processing).
- Send your trailer maps for the current page, followed by headers for the next page, as you do in a nonrouting environment (see “Page breaks: BMS overflow processing” on page 283). While the OVERFLOW condition is in force, these EXEC CICS SEND MAP commands apply only to the logical message that overflowed (you would not want them in a logical message where you were mid-page, and BMS makes no assumptions about different terminal types that happen to have the same page capacity).

- Reissue the command that caused the overflow, as you do in a nonrouting environment. After you do, however, retest for overflow and repeat the whole process until overflow does not occur. This procedure ensures that you get the trailers and headers and the map that caused the overflow into each of the logical messages that you are building.

Routing with SET

When you specify a disposition of SET in a routing environment, no messages are sent to the destinations in your route list, because the pages are returned to your program as they are completed. However, the EXEC CICS ROUTE command is processed in the usual way to determine these destinations and the terminal types among them. BMS builds a separate logical message for each type, as usual, and returns a page to the program each time one is completed for any of the terminal types. BMS raises the OVERFLOW and RETPAGE conditions as it does with a disposition of PAGING. Consequently, ROUTING with SET allows you to format messages for terminal types other than that of your principal facility.

Interleaving a conversation with message routing

While you are building a message to be routed, you can use BMS SEND commands as well as EXEC CICS RECEIVE MAP and terminal control commands to converse with your principal facility. (Without routing, you cannot use BMS SENDs, as noted in “Rules for logical messages” on page 278.) Such SEND commands must have a disposition option of TERMINAL rather than PAGING or SET, and must not specify ACCUM. The associated inputs and outputs are processed directly and do not interfere with your logical message, even if your terminal is one of the destinations of the message.

Using SET

When you specify a disposition of SET for a BMS message, BMS formats your output and returns it in the form of a device-dependent data stream. No terminal I/O occurs, although the returned data stream usually is sent to a terminal subsequently.

There are several reasons for asking BMS to format a data stream without sending it. You might want to do any of the following:

- Edit the data stream to meet the requirements of a device with special features or restrictions not explicitly supported by CICS.
- Compress the data stream, based on standard 3270 features or special device characteristics.
- Forward the data stream to a terminal not connected directly to CICS. For example, you might want to pass data to a 3270 terminal attached to a system connected to CICS via an APPC link. You can format the data with SET and send the resulting pages to a partner program across the link. If the terminal is of a different type from your principal facility, you can define a dummy terminal of the appropriate type and then ROUTE to that terminal with SET to get the proper formatting, as explained in “Routing with SET.”

BMS returns formatted output by setting the pointer variable named in the SET option to the address of a page **list**. This list consists of one or more four-byte entries in the following format, each corresponding to one page of output.

<i>Table 20. Page list entry format</i>	
Bytes	Contents
0	Terminal type (see Table 17 on page 270)
1-3	Address of TIOA containing the formatted page of output

An entry containing -1 (X'FF') in the terminal type signals the end of the page list. Notice that the addresses in this list are only 24 bits long. If your program uses 31-bit addressing, expand a 24-bit address to a full word by preceding it with binary zeros before using it as an address.

Each TIOA (terminal input-output area) is in the standard format for these areas:

<i>Table 21. TIOA format</i>			
Field name	Position	Length	Contents
TIOASAA	0	8	CICS storage accounting information (8 bytes)
TIOATDL	8	2	Length of field TIOADBA in halfword binary format
(unnamed)	10	2	Reserved field
TIOADBA	12	TIOATDL	Formatted output page
(unnamed)	TIOATDL + 12	4	Page control area, required for the SEND TEXT MAPPED command (if used)

The reason that BMS uses a list to return pages is that some BMS commands produce multiple pages. SEND MAP does not, but SEND TEXT can. Furthermore, if you have established a routing environment, BMS builds a separate logical message for each of the terminal types among your destinations, and you may get pages for several different terminal types from a single BMS command. The terminal type tells you to which message a page belongs. (Pages for a given type are always presented in order.) If you are not routing, the terminal type is always that of your principal facility.

If you are not using the ACCUM option, pages are available on return from the BMS command that creates them. With ACCUM, however, BMS waits until the available space on the page is used. BMS turns on the RETPAGE condition to signal your program that pages are ready. You can detect RETPAGE with an EXEC CICS HANDLE CONDITION command or by testing the response from the BMS command (in EIBRESP or the value returned in the RESP option).

Capture the information in the page list whenever BMS returns one, because BMS reuses the list. You need save only the addresses of the pages, not the contents. BMS does not reuse the pages themselves, and, in fact, moves the storage for them from its control to that of your task. This allows you to free the storage for a page when you are through with it. If you do this, the DATA or DATAPOINTER option in your FREEMAIN command should point to the TIOATDL field, not to TIOASAA.

Partition support

Partitions are the first of several special hardware features that BMS supports. **Standard BMS** is required for partitions.

Some IBM displays allow you to divide the screen into areas that you can write to and read from separately, as if they were independent screens. The areas are called partitions, and features of BMS that allow you to take advantage of the special hardware are collectively called “partition support”.

The IBM 3290 display, which is a member of the 3270 family, and the IBM 8775 are the primary examples of devices that support partitioning. Consult the device manuals⁹ to understand the full capabilities of a partitioned device, but the essential features are these:

- You can divide the physical screen into any arrangement of one to eight nonoverlapping rectangular areas. The areas are independent from one other, in the sense that the operator can clear them separately, the state of the keyboard (locked or unlocked) is maintained separately for each, and you write to and read from them one at a time.
- Only one partition is **active** at any given time. This is the one containing the cursor. The operator is restricted to keying into this partition, and the cursor wraps at partition boundaries. When a key that transmits data is depressed (the ENTER key or one of the program function keys), data is transmitted only from the active partition.
- The operator can change the active partition at any time by using the “jump” key; your program can also, as explained in “Determining the active partition” on page 309.
- BMS also writes to only one partition on a given SEND, but you can issue multiple SENDs and you do not have to write to the active partition.
- The partition configuration is sent to the device as a data stream, so that you can change the partitions for each new task or even within a task. The BMS construct that defines the partitions is called a **partition set** and is described in “How to define partitions” on page 306.
- You also can use the terminal in **base state** (without partitions) and you can switch from partitioned to base state with the same command that you use to change partition arrangements.
- When you specify how to partition the screen area, you also divide up the hardware buffer space from which the screen is driven. In partitioned devices, the capacity of the buffer is generally greater than that of the screen, so that some partitions can be assigned extra buffer space. The screen area allocated to a partition is called its **viewport** and the buffer storage is called its **presentation space**.

BMS uses the presentation space as its page size for the partition, so that you can send as much data as fits there, even though not all of it can be on display at once. Keys on the device allow the operator to scroll the viewport of the

⁹ *IBM 3290 Information Display Panel Description and Reference for the 3290 and IBM 8775 Display Terminal Component Description for the 8775.*

partition vertically to view the entire presentation space. Scrolling occurs without any intervention from the host.

- Some partitioned devices allow you to choose among character sets of different sizes. We talk about this in “3290 character size” on page 307.

In spite of the independence of the partitions, the display is still a single terminal to CICS. You cannot have more than one task at a time with the terminal as its principal facility, although you can use the screen space cooperatively among several pseudoconversational transaction sequences if they use the same partition set (see “Terminal sharing” on page 311).

Note: The 3290 can be configured internally so that it behaves as more than one logical unit (to CICS or any other system); this definition is separate from the partitioning that may occur at any one of those logical terminals.

Uses for partitioned screens

Partitioned screens are particularly useful in certain types of application. For example:

Scrolling

For transactions that produce more output than fits on a single screen, scrolling is an alternative to BMS terminal paging (see “Output disposition options: TERMINAL, SET, and PAGING” on page 277). For example, you can define a partition set that consists of just one partition, where the viewport is the whole screen and the presentation space is the entire buffer. You can write to the entire buffer as a single page, and the operator can scroll through the data using the terminal facilities. Response time to scrolling requests is very short, because there is no interaction with the host. You are limited to the capacity of the buffer, of course.

You may also want to scroll just part of the screen and use some partitions for fixed data.

Data entry

Another good use for a partitioned screen is “heads down” data entry, where the operator’s productivity depends on how fast the application can process an input and reopen the keyboard for the next. With a partitioned screen, you can divide the screen into two identical entry screens. The operator fills one, presses Enter, and then fills the second one while the data entry transaction is processing the first input. If the input is good, the program simply erases it in preparation for the next entry; if not, there is still an opportunity for the operator to make corrections without losing subsequent work. The *CICS Sample Applications Guide* contains an example of such a data entry transaction.

Lookaside

In many online operations, the operator sometimes needs to execute a second transaction in order to finish one in progress. Order entry is an example, where the operator may have to look up codes or prices to complete an entry. Many inquiries are similar. The initial inquiry brings back a summary list of hits. The operator selects one and asks for further detail, then may need to select another for detail, and so on. In such cases, a partitioned screen allows the operator to do the second task while keeping the output of the first, which is needed later, on the

screen. The *CICS Sample Applications Guide* also contains an example of a lookaside transaction.

“Help” text is still another example of “lookaside”. If you allocate one partition of the screen to this text, the operator can get the required tutorial information without losing the main screen.

Data comparison

Applications in which the operator needs to compare two or more sets of data simultaneously are also excellent candidates for a partitioned screen. Partitioning allows a side-by-side comparison, and the scrolling feature makes it possible to compare relatively large documents or records.

Error messages

If you partition a screen and allocate one area to error messages and other explanatory text, usability is enhanced because the operator always knows where to look for messages, and the main screen areas are never overwritten with such information. CICS sends its own messages to such a partition if you designate one in your partition set, as we explain in “How to define partitions.”

How to define partitions

Each partitioning of a screen is defined by a partition set, which is a collection of screen areas (partitions) intended for display together on a screen. You define a partition set with assembler macros, just as you do map sets. There are two of them: DFHPSD and DFHPDI.

The partition set definition begins with a DFHPSD (partition set definition) macro, which defines:

- The name of the partition set
- Screen size (BMS makes sure that the partition viewports do not exceed the total space available)
- Default character cell size (we talk about cell size in “3290 character size” on page 307)
- The partition set suffix, used to associate the partition set with a particular screen size (see “Establishing the partitioning” on page 308)

After the initial DFHPSD macro, you define each partition (screen area) with a DFHPDI macro. DFHPDI specifies:

- The identifier of the partition within the partition set
- Where the partition is located on the screen
- Its viewport size (in lines and columns)
- The presentation space associated with the viewport (that is, the amount of buffer space allocated), also in lines and columns. Because scrolling is strictly vertical, BMS requires that the width of the presentation space match the width of the viewport.
- The character size to be used
- The map set suffix associated with the partition, used to select the map set appropriate for the partition size

- Whether the partition may receive CICS error messages (BMS sends certain error messages that it generates to a partition so designated, if there is one)

You end the partition set with a second DFHPSD macro, containing only the option TYPE=FINAL. See the *CICS Application Programming Reference* manual for full details on DFHPSD and DFHPDI.

Because these are assembler macros, you need to follow assembler format rules in creating them. See “Rules for writing BMS macros” on page 235 if you are not familiar with assembler language. After you do, you need to assemble and link-edit your partition set. The resulting load module can be placed in the same library as your map sets, or in a separate library if your installation prefers. Your systems staff also need to define each partition set to the system with a PARTITION definition.

3290 character size

The 3290 hardware allows you to use up to eight different character sets, of different sizes. Two sets come with the hardware; the others can be loaded with a terminal control SEND command. (See the *IBM 3290 Information Display Panel Description and Reference* manual for further information.)

Each character occupies a rectangular **cell** on the screen. Cell size determines how many lines and columns fit on the screen, or in a particular partition of the screen, because you can specify cell size by partition. Cells are measured in pels (picture elements), both vertically and horizontally. The smallest cell allowed is 12 vertical pels by 6 horizontal pels. The 3290 screen is 750 pels high and 960 pels wide. Using the minimum cell size, therefore, you can fit 62 characters vertically (that is, have 62 lines), and 160 characters horizontally (for 160 columns). (The 3290 always selects the character set that best fits your cell size, and places the character at the top left corner of the cell.)

Partition sizes are expressed in lines and columns, based on the cell size you specify for the partition, which is expressed in pels. (The name of the option is CHARSIZE, but it is really cell size.) To make sure your partitions fit on the screen, you need to work out your allocation in pels, although BMS tells you when you assemble if your partitions overlap or does not fit on the screen. The partition height is the product of the number of rows in the partition and the vertical CHARSIZE dimension; the partition width is the product of the number of columns and the horizontal CHARSIZE value.

If you do not specify a CHARSIZE size in your DFHPDI partition definition, BMS uses the default given in the DFHPSD partition set definition. If DFHPSD does not specify CHARSIZE either, BMS uses the default established for the terminal when it was installed. If you specify cell size for some but not all partitions, specify a default for the partition set too, so that you do not mix your choices with the installation default.

Programming considerations

Partitions affect programming in several areas, as we explain in the sections that follow. These include:

- Partition set loading
- BMS SEND command options
- The active partition

- BMS RECEIVE commands and options
- ASSIGN options
- Logical messages
- Routing
- Attention identifiers and exception conditions

Nonetheless, BMS partition support is designed to have as little impact as possible on existing applications that get executed at a partitioned terminal. We talk about this in “Terminal sharing” on page 311. In addition, options and commands specific to partitions are ignored when executed at a terminal that does not support partitions or that is in base state at the time of the command.

Establishing the partitioning

You can tell BMS which partition set to load for a particular transaction by naming it in the PARTITIONSET option of the RDO TRANSACTION resource definition. If you do this, and the named partition set is not already loaded at the terminal, BMS adds the partition definitions to your data on the first BMS SEND in the task.

You can also direct BMS not to change the partitions from their current state (PARTITIONSET=KEEP in the RDO TRANSACTION resource definition) or indicate that you will load the partitions yourself (PARTITIONSET=OWN). If you do not specify any PARTITIONSET value, BMS sets the terminal to base state (no partitions) at the time it initiates the transaction.

Whatever the PARTITIONSET value associated with the transaction, a task can establish new partitions at almost any time with an EXEC CICS SEND PARTNSET command, except that you cannot issue the command while you are building a logical message.

EXEC CICS SEND PARTNSET does not send anything to the terminal immediately. Instead, BMS remembers to send the partition information along with the next BMS command that sends data or control information, just as it sends a partition set named in the PARTITIONSET option of the TRANSACTION definition on the first BMS SEND. Consequently, issue an EXEC CICS SEND MAP, EXEC CICS SEND TEXT, or EXEC CICS CONTROL command before you issue an EXEC CICS RECEIVE or EXEC CICS RECEIVE MAP that depends on the new partitions. See the *CICS Application Programming Reference* manual for full details on EXEC CICS SEND PARTNSET.

Note: You can get an unexpected change of partitions in the following situation. If CICS needs to send an error message to your terminal, and the current partition set does not include an error partition, CICS will return the terminal to base state, clear the screen, and write the message. For this reason, it is a good idea to designate one partition as eligible for error messages in every partition set.

When BMS loads a partition set, it suffixes the name requested with the letter that represents your terminal type if device-dependent support is in effect, in order to load the one appropriate to your terminal. It takes the suffix from the ALTSUFFIX option value of the RDO TYPETERM resource definition associated with your terminal. Partition set suffixing is analogous to map set suffixing, and the same sequence of steps is taken if there is no partition set with the right suffix (see “Device-dependent maps: map suffixes” on page 269).

Partition options for BMS SEND commands

As noted earlier, when you write to a partitioned screen, you write to only one partition, and the effects of your command are limited to that partition. ERASE and ERASEAUP clear only within the partition, and FREEKB unlocks the keyboard only when the partition becomes active.

You can specify the partition to which you are sending with either the PARTN option in your map definition or with the OUTPARTN option on your SEND MAP. OUTPARTN overrides PARTN. If you don't specify either, BMS chooses the first partition in the set.

The use of partitions affects the suffixing of map set names that we described in "Device-dependent maps: map suffixes" on page 269. The map set suffix is taken from the MAPSFX value for the partition instead of being determined as described in that section.

Determining the active partition

When you send to a partition, you can move the cursor to that partition or another one. A value of ACTIVATE in the PARTN option of the map definition puts the cursor in the partition to which you are writing. If you specify ACTPARTN on your BMS SEND command, you can name any partition (not necessarily the one to which you are writing), and you override the ACTIVATE specification. Both ACTIVATE and ACTPARTN unlock the keyboard for the active partition, as well as placing the cursor there. If neither is present, the cursor does not move and the keyboard is not unlocked.

Although you can make a partition active by placing the cursor there when you send, you do not have the last word on this subject, because the operator can use the jump key on the terminal to move the cursor to another partition. This can complicate receiving data back from the terminal, but BMS provides help, as we are about to explain.

Partition options for BMS RECEIVE commands

When you issue a RECEIVE MAP command, you can tell BMS from which partition you expect data (that is, which partition you expect to be active) with either the PARTN option in the map definition or with the INPARTN option on your RECEIVE MAP. INPARTN overrides PARTN. If you do, and the operator transmits from a different partition than the one you named, BMS repositions the cursor in the partition you named, unlocks the keyboard and repeats the RECEIVE command. It also sends a message to the error partition (the one with ATTRB=ERROR) asking the operator to use the right partition. (No message is sent if there is no error partition.) The input from the wrong partition is discarded, although it is not lost, because it can be reread later. BMS does this up to three times; if the operator persists for a fourth round, BMS raises the PARTNFAIL condition.

You do not have to specify an input partition; sometimes there is only one that allows input, and sometimes the same map applies to all. If you issue RECEIVE MAP without INPARTN and there is no PARTN option in the map, BMS accepts data from any partition and map it with the map named in the command. You also can determine the partition afterward, if you need to, with an ASSIGN command containing the INPARTN option.

INPARTN is not set until after the first BMS operation, however, and so if you need to know which partition is sending to select the right map, you need another approach. In this situation, you can issue a RECEIVE PARTN command, which reads data unmapped and tells you which partition sent it. Then you issue a RECEIVE MAP command using the map that matches the partition with the FROM option, using the map that matches the partition. RECEIVE MAP with FROM maps data already read, as explained in “Formatting other input” on page 266.

ASSIGN options for partitions

In addition to the INPARTN option just described, there are three other ASSIGN options to help you in programming for a partitioned terminal. The PARTNS option tells you whether the terminal associated with your task supports partitions, and the PARTNSET option returns the name of the current partition set (blanks if none has been established). The fourth ASSIGN option, PARTNPAGE applies only to logical messages, which we talk about in “Partitions and logical messages.”

Partitions and logical messages

When you build a BMS logical message for a terminal for which partitions have been established, you can direct the pages of the message to multiple partitions. You can even send text output to some partitions and mapped output to others, provided you do not mix them in the same partition. (This is an exception to the normal rule against mixing text and mapped output in a logical message.)

When the output is displayed, the first page for each partition is displayed initially. The pages are numbered by partition, and CSPG commands that the operator enters into a particular partition apply only to that partition, with the exception of the page purge command. The purge command deletes the entire logical message from all partitions.

On each BMS SEND that contributes to the message, you specify the partition to which the output goes. If you are not using ACCUM, BMS builds a page for that partition. If you are using ACCUM, BMS puts the output on the current page for that partition. Page overflows therefore occur by partition. If you are intercepting overflows and are not sure in which partition the overflow occurred, you can use the PARTNPAGE option of the ASSIGN command to find out.

Note: Because BMS uses both the page size and partition identifiers in building a logical message, you cannot change the partitions mid-message.

The bookkeeping required to handle page overflow when you are distributing pages among partitions is analogous to that required in a routing environment (see “Routing and page overflow” on page 301). In particular, you need to ensure that you finish overflow processing for one partition before doing anything that might cause overflow in another. Failure to do so can cause program loops as well as incorrect output.

Partitions and routing

You cannot route a logical message written to multiple partitions. BMS ignores the OUTPARTN and ACTPARTN options on BMS SEND commands in a routing environment.

You can route an ordinary message to a terminal that supports partitions, but BMS builds the message and the CSPG transaction displays it using the terminal in base (unpartitioned) state.

New attention identifiers and exception conditions

Partitioned terminals have a CLEAR PARTITION key that clears the active partition in the same way that the CLEAR key clears the whole screen (CLEAR still does this on a partitioned terminal). You may need to check for this additional attention identifier in your program logic. The CLEAR PARTITION AID value is included in DFHAID (see “The attention identifier: what caused transmission” on page 260).

There are also some new exception conditions associated with partitions, and new ways to get some of the old ones. The new ones include INVPARTN (naming a partition that does not exist in the partition set), INVPARTNSET (naming a module that is not a partition set), and PARTNFAIL (receiving from a partition other than the one the operator transmitted from). They are all described in the *CICS Application Programming Reference* manual, with the commands to which they apply.

Terminal sharing

With proper planning, you can share a terminal among several processes by assigning each a separate partition. You cannot have more than one task in progress at once at a terminal, of course, but you can interleave the component tasks of several pseudoconversational transaction sequences at a partitioned terminal.

To take a very simple example, suppose you decide to improve response time for an existing pseudoconversational data entry transaction by letting the operator enter data in two partitions (see “Data entry” on page 305). You could modify the application to work on two records at once, or you could simply modify it to send to the same partition from which it got its input. Then you could run it independently from each partition.

You can establish the partitions with the PARTITIONSET option in the TRANSACTION definition (all of the transactions involved, if there are several in the sequence). As noted earlier, BMS does not reload the partitions as long as each transaction has the same PARTITIONSET value. Alternatively, you could establish the partitions with a preliminary transaction (one that displayed the first entry screen in both partitions, say) and use a PARTITIONSET value of KEEP for the data entry transactions. Whenever you share a partitioned screen, whether among like transactions or different ones, you need to ensure that one does not destroy the partition set required by another. Also, if two different CICS systems may share the same screen, they should name partition sets in common, so that BMS does not reload the partitions when it should not.

If the hypothetical data entry transaction sequence uses the TRANSID option on the RETURN command to specify the next transaction identifier, you would need to make another small change, because the option applies to the whole terminal, not the partition. One solution would be to place the next transaction identifier in the first field on the screen (turning on the modified data tag in the field definition) and remove the TRANSID from the RETURN. CICS would then determine the next transaction from the input, as described in “How tasks get started” on page 200.

Restrictions on partitioned screens

We have already noted that you cannot route to a terminal in partitioned state. You also cannot use partitions and logical device codes together (LDCs are described in “Logical device components”). In addition, you cannot use partitions in combination with GDDM, although you can use partitions with outboard formats (see “Outboard formatting” on page 318).

Logical device components

Logical device components (LDCs) are another special hardware feature supported by BMS. Like partitions, LDCs require **standard BMS**.

A terminal that supports LDCs is one that consists of multiple functional components (logical devices) controlled through a single point (the logical unit). The components might be a printer, reader, keyboard and display, representing a remote work station, or they might be multiple like devices, such as word processing stations or passbook printers. The IBM 3601 logical unit, the 3770 batch logical unit, 3770 and 3790 batch data interchange logical units, and LU type 4 logical units all support logical device components.

Because the logical unit is a single entity to CICS, but consists of components that can be written and read independently, the CICS application programming interface for LDC terminals looks similar to that for partitioned terminals, each LDC corresponding to one partition in a partition set. There are many differences, of course. Consult the CICS manual that describes CICS support for your particular terminal type (see “Where to find more information” on page 205 for a list). The sections that follow describe the major differences that affect programming, which are:

- LDC definition
- SEND command options
- RECEIVE command options
- Logical messages
- Routing

Defining logical device components

The logical device components for a terminal are defined by a list called an LDC table. The RDO TYPETERM component of the RDO TERMINAL resource definition points to the table, which may be individual to the logical unit or shared by several logical units that have the same components. The table itself is defined with DFHTCT TYPE=LDC (terminal control) macros. (See the *CICS Resource Definition Guide* for descriptions of both RDO TYPETERM and the DFHTCT macros.)

An LDC table contains the following information for each logical device component of the logical unit:

- A 2-character logical device identifier. These identifiers are usually standard abbreviations, such as CO for console and MS for a magnetic stripe encoder, but they need not be.
- A 1-character device code, indicating the device type (console, card reader, word processing station). Codes are assigned by CICS from the device type and other information provided in the macro.

- A BMS page size. BMS uses this size, rather than one associated with the logical unit, because different logical devices have different page sizes.
- A BMS page status (AUTOPAGE or NOAUTOPAGE); see “The AUTOPAGE option” on page 280.

Sending data to a logical device component

You direct BMS output to a specific logical device component of a terminal by naming it in the LDC option of your SEND MAP, SEND TEXT, or SEND CONTROL command or the LDC option of your mapset. A value in the command overrides one in the map set. If the LDC does not appear in either place, BMS uses a default that varies with the terminal type (see the LDC option discussion of the EXEC CICS SEND MAP command in the *CICS Application Programming Reference* manual for specific information).

LDCs and logical messages

When you build a BMS logical message for your own terminal, you can distribute pages of the message among different logical device components in the same way that you can direct pages to a logical message to different partitions. BMS accumulates pages separately for each logical device component in the same way that it does for partitions (see “Partitions and logical messages” on page 310). You can include both text and mapped output in the message, provided you do not send both to one LDC. Page overflow occurs by LDC, and terminal operator paging commands operate on a logical device component basis.

When retrieving pages, the operator (or user code in the device controller) must indicate the LDC to which the request applies, because not all devices have keyboards. As in the case of partitions, a message purge request deletes the entire message, from all LDCs. See the *CICS-Supplied Transactions* manual for more detail on page retrieval for logical devices.

If you are intercepting page overflows, you can tell which LDC overflowed by issuing an ASSIGN command with either the LDCMNM or LDCNUM option. Both identify the device which overflowed, the first by its 2-character name and the second by the one-byte numeric identifier. You can determine the page number for the overflowing device with ASSIGN PAGENUM, just as with a partitioned device.

There is one restriction associated with LDCs and page overflow that is unique to LDCs. After overflow occurs, send both a trailer map for the current page and a header for the next one to the LDC that overflowed. BMS raises the INVREQ (invalid request) condition if you fail to do this.

LDCs and routing

Routing is supported in an LDC environment, provided the message goes to the same component type for every destination that supports LDCs. (You cannot route a multiple-LDC message.)

You can supply the LDC value in several ways:

- If you use the LDC option on your EXEC CICS ROUTE command, the value supplied overrides all other sources and is used for all eligible destinations to which LDCs apply.

- If you specify an LDC in a route list entry (and not in the EXEC CICS ROUTE command), that value is used for the associated destination. (If you specify both and they do not agree, the ROUTE list value is used and the discrepancy is flagged in the status flag of the entry.)
- If you specify neither, the value is determined from terminal and system LDC tables in the same way as it is in a nonrouting environment when you omit the LDC from the BMS SEND command. (The value on the SEND command is ignored when routing is in effect.)

BMS support for other special hardware

In addition to partitions and LDCs, BMS provides support these other special hardware features:

- 10/63 magnetic slot reader
- Field selection features: cursor select, light pen, trigger fields
- Outboard formatting

The magnetic slot reader and outboard formatting both require **standard** BMS. Support for the cursor select key, light pen and trigger fields is included in **minimum**.

10/63 magnetic slot reader

Some IBM display terminals support a magnetic slot reader (MSR), a device that reads data from small magnetic cards, as an optional feature. The MSR has indicator lights and an audible alarm to prompt operator actions. Some terminals control the MSR themselves, but others, such as the IBM 8775 and the IBM 3643, let you control the functions of the reader by program.

CICS provides an ASSIGN command option, MSR, that tells you whether the principal facility of your task has an MSR or not.

With BMS, you can control the state of such an MSR by using the MSR option of the BMS SEND commands. This option transmits four bytes of control data to the attached MSR, in addition to display data sent to the terminal. BMS provides a copybook, DFHMSRCA, containing most of the control sequences you might need. The *CICS Application Programming Reference* manual describes the supplied constants and explains the structure of the control data, so that you can expand the list if you need to.

The control sequence that you send to an MSR affects the next input from the device; hence it has no effect until a RECEIVE command is issued. Input from MSRs is placed in the device buffer and transmitted in the same way as keyboard input. If the MSR input causes transmission, you can detect this by looking at the attention identifier in EIBAID. A value of X'E6' indicates input from the MSR, and X'E7' signals input from the MSR extended (a second MSR that may be present). See the *IBM 3270 Information Display System Data Stream Programmer's Reference* manual for information on how to format a screen for MSR input and other details on these devices.

Field selection features

BMS supports several special hardware features that allow the operator to enter and transmit input by selecting a field on the screen:

- Trigger fields
- Cursor selectable fields
- Light pen detection

Trigger field support

Trigger fields are a special hardware feature of certain types of terminal, such as the 8775. A field defined as a trigger field causes the terminal to transmit its contents if the operator moves the cursor out of the field when it is **primed**. The field gets primed when the operator moves the cursor into it and enters data or uses either the DELETE or ERASE EOF keys. It becomes unprimed after it causes transmission, or if the operator uses the ERASE INPUT key, or after a send to the terminal (if you are using partitions, the send must be to the partition that contains the trigger field to have this effect).

You define a field as a trigger field by setting the VALIDN extended attribute to a value of TRIGGER, either in the map or by program override.

Only the field itself is sent when a trigger field causes transmission; other fields are not sent, even if they have been modified. You can detect a transmission caused by a trigger field by checking the attention identifier, which has a value of X'7F'.

Terminals that support the validation feature buffer the keyboard, so that the operator can continue to enter data while the host is processing an earlier transmission. The program processing such inputs needs to respond quickly, so that the operator does not exceed the buffer capacity or enter a lot of data before an earlier error is diagnosed.

The customary procedure is for the program receiving the input to check the contents of the trigger field immediately. If correct, the program simply unlocks the keyboard to let the operator continue (a BMS SEND command containing the FREEKB option does this). If the field is in error, you may wish to discard the stored keystrokes, in addition to sending a diagnostic message. Any of the following actions does this:

- A BMS SEND command that contains ERASE, ERASEAUP or ACTPARTN or that lacks FREEKB
- A BMS SEND directed to a partition other than the one containing the trigger field (where partitions are in use)
- A RECEIVE MAP, RECEIVE PARTITION, or terminal control RECEIVE command
- Ending the task

See the *IBM 3270 Information Display System Data Stream Programmer's Reference* manual for more information about trigger fields.

Cursor- and pen-detectable fields

BMS also supports **detectable** fields, another special hardware feature available on some terminals. There are two hardware mechanisms for detectable fields: the “cursor select” key and the light pen. A terminal has either the key or a pen, not both. Both work the same way and, as the key succeeded the pen, we talk about the key.

For a field to be detectable, it must have certain field attributes, and the first character of the data, known as the **designator character**, must contain one of five particular values. You can have other display data after the designator character if you wish.

The bits in the field attributes byte that govern detectability also control brightness. High intensity (ATTRB=BRT) fields are detectable if the designator character is one of the detectable values. Normal intensity fields may or may not be detectable; you have to specify ATTRB=DET to make them so; nondisplay (ATTRB=DRK) fields cannot be detectable.

As usual, you can specify attributes and designator characters either in the map definition or by program override. However, DET has a special effect when it appears in an input-only map, as we explain in a moment.

Note that because high-intensity fields have, by definition, the correct field attributes for detectability, the terminal operator can make an **unprotected** high-intensity field detectable by keying a designator character into the first position of the field.

Selection fields

There are two types of detectable field, **selection** and **attention** fields; the type is governed by the designator character. A selection field is defined by a designator character of either a question mark (?) or a greater-than sign (>). The convention is that (?) means the operator has not selected whatever the field represents, and (>) means he has. The hardware is designed around this convention, but it is not enforced, and you can use another if it suits. You can initialize the designator to either value and initialize the modified data tag off or on with either value.

Every time the operator presses the cursor select key when the cursor is in a selection field, the designator switches from one value to the other (? changes to > and > changes to ?). The MDT is turned **on** when the designator goes from ? to > and **off** when the designator goes from > to ?, regardless of its previous state. This allows the operator to change his mind about a field he has selected (by pressing cursor select under it again) and gives him ultimate control over the status of the MDT. The MDT governs whether the field is included when transmission occurs, as it does for other fields. No transmission occurs at this time, however; selection fields do not of themselves cause transmission; that is the purpose of attention fields.

Attention fields

Attention fields are defined by a designator character of blank, null,¹⁰ or ampersand. In contrast to a selection field, when the cursor select key is pressed with the cursor in an attention field, transmission occurs.

If the designator character is an ampersand, the effect of pressing the cursor select key is the same as depressing the ENTER key. However, if the designator is blank or null, what gets transmitted is the address of every field with the MDT on, the position of the cursor, and an attention identifier of X'7E'. The **contents** of these fields are not transmitted, as they are with the ENTER key (or a cursor select with an ampersand designator). In either case, the fields with the MDT bit on may be selection fields or normal fields which the operator changed or which were sent with the MDT on.

BMS input from detectable fields

After transmission caused by a cursor-select attention field with a blank or null designator, BMS tells you which fields were transmitted (that is, had the MDT on) by setting the first position of the corresponding input (I) subfield to X'FF'. The first position is otherwise set to X'00'. You can tell which attention field caused transmission from this value if it was the only one transmitted, or from the position of the cursor otherwise.

If transmission is caused by a cursor-select attention field with an ampersand designator (or by the ENTER key or a PF key), the I subfield contains the contents of the field if the MDT is on and the L subfield reflects its length, as usual, except in one case: if you specify the DET attribute for a field in an input-only map (that is, MODE=IN, DATA=FIELD), BMS reserves only one byte in the symbolic map for the input subfield, rather than the number indicated by the LENGTH option. After a RECEIVE MAP naming such a map, this I subfield contains X'FF' with a length of 1 if the field is selected (that is, if its MDT was on), and a null (X'00') if not. BMS supplies no other input for the field, even if some was transmitted.

Consequently, if you need to receive data from a detectable field as well as knowing whether it was selected or not, you need to avoid the use of DET in an input-only map. You can define the map as INOUT, even if you do not use it for output, or you can set the DET attribute in the program rather than the map. For high-intensity fields, you do not need to specify DET, because BRT implies DET.

You also need to ensure that the data gets transmitted. When the cause of transmission is the ENTER key, a PF key, or an attention field with an ampersand designator character, field data gets transmitted. It does not when the cause is an attention field with a blank or null designator.

See the *IBM 3270 Information Display System Data Stream Programmer's Reference* manual for more information about detectable fields.

¹⁰ A null in the data stream has the same effect as a blank in this function. In BMS, however, use a blank, because BMS does not transmit nulls in some circumstances, and because you cannot override the first position of a field with a null (see "Where the values come from" on page 250).

Outboard formatting

Outboard formatting is a technique for reducing the amount of line traffic between the host processor and an attached subsystem. The reduction is achieved by sending only variable data across the network. This data is combined with constant data, such as a physical map, by a program within the subsystem. The formatted data can then be displayed.

You can use outboard formatting with a 3650 Host Communication Logical Unit, an 8100 Series processor with DPPX and DPS Version 2, or a terminal attached through a 3174 control unit. Maps used by the 3650 must be redefined using the 3650 transformation definition language before they can be used. For more information, see the section describing BMS in the *CICS/DOS/VS IBM 3650/3680 Guide*. Maps to be used with the 8100 must be generated on the 8100 using either an SDF II utility or the interactive map definition component of the DPS Version 2.

If a program in the host processor sends a lot of mapped data to subsystems, you can reduce line traffic by telling BMS to transmit only the variable data in maps. The subsystem must then perform the mapping operation when it receives the data. BMS prefixes the variable data with information that identifies the subsystem map to be used to format the data.

Terminals that support outboard formatting have OBFORMAT(YES) in their RDO TYPETERM resource definition. When a program issues a SEND MAP command for such a terminal, and the specified map definition contains OBFMT=YES, BMS assumes that the subsystem is going to format the data and generates an appropriate data stream. If you send a map that has OBFMT=YES to a terminal that does not support outboard formatting, BMS ignores the OBFMT operand.

See “Batch data interchange” on page 338 for more information about programming some of the devices that support outboard formatting.

Chapter 19. Terminal control

This chapter describes terminal control commands, the second of two methods that CICS provides for programs to communicate with terminals. The other interface is BMS, described in Chapter 18, “Basic mapping support” on page 225.

This chapter contains the following topics:

- “Access method support”
- “Terminal control commands” on page 320
- “VTAM considerations” on page 333
- “Sequential terminal support” on page 336
- “Batch data interchange” on page 338

Terminal control commands apply to a variety of devices, reducing the sensitivity of programs to the terminals they support and to the access methods controlling the terminals. In addition to the commands themselves, CICS provides the data translation, synchronization of input and output operations, and session control needed to read from or write to a terminal or logical unit. This helps insulate you from the APIs of the individual communications access methods, which are complex and very different from one another.

BMS insulates you even more from the characteristics of particular devices and the mechanics of communication than does terminal control, but at the cost of some flexibility and function. For example, BMS pathlengths are longer, and BMS does not support as many terminal types as does terminal control. See “CICS APIs for terminals” on page 203 for a comparison of BMS and terminal control.

Access method support

CICS Transaction Server for VSE/ESA Release 1 supports terminals directly through interfaces to the following access methods:

- Virtual Telecommunications Access Method (VTAM)
- Sequential Access Method (SAM) for terminals simulated by sequential devices

CICS supports operating system consoles as terminals too, but through operating system services rather than through an access method. The terminal control interface to a console is the same as to other terminals (though certain consoles might have certain restrictions), but BMS is not available. You can find a full list of the terminals supported by CICS Transaction Server for VSE/ESA in the *CICS Resource Definition Guide*.

Basic Telecommunications Access Method (BTAM)

Earlier releases of CICS also supported terminals through Basic Telecommunications Access Method (BTAM).

You can still execute transactions under CICS Transaction Server for VSE/ESA from terminals using these access methods. However, the terminals themselves must be attached to a CICS system at an earlier level that supports the access method. A transaction running under CICS Transaction Server for VSE/ESA Release 1 communicates with a local surrogate for the remote terminal, and the

two CICS systems manage the correspondence between the surrogate and the real terminal. The transaction is invoked either when the CICS that owns the terminal routes the transaction to the CICS region, or by automatic transaction initiation (ATI) in the CICS region. With ATI, this region arranges assignment of the terminal as principal facility for the transaction through the CICS region that owns the terminal.

Terminal control commands

The commands described in this chapter apply only to the principal facility of the task issuing them, where that facility is one of the following:

- A device connected through BTAM or SAM.
- An LU Type 0, 1, 2, 3, or 4 connected through VTAM.

Note: This chapter does not cover program-to-program communication, whether directed to the alternate or principal facility. This is covered in a separate manual, the *CICS Distributed Transaction Programming Guide*.

Terminal control commands fall into four groups:

- Basic data transmission commands: RECEIVE, SEND, and CONVERSE.
- Commands that send device controls, synchronize transmission, end a session, or perform similar control functions.
- Commands to tell you about your terminal: ASSIGN and INQUIRE.
- Special device group commands: the batch data interchange (BDI) commands.

We discuss each of these groups in the sections that follow.

Data transmission commands

There are three commands that transmit data to and from the terminal or logical unit that is the principal facility of your task:

RECEIVE	reads data from the terminal
SEND	Writes data to the terminal
CONVERSE	Writes data to the terminal, waits for input, and reads the input.

CONVERSE is essentially a combination of SEND and RECEIVE and is usually the equivalent of SEND followed by RECEIVE. In certain cases you must use CONVERSE instead of SEND and RECEIVE, for example, sending structured-field data to certain 3270 devices; these are covered in the CICS device guides listed in the bibliography. In other cases you must use SEND and RECEIVE, because CONVERSE is not provided; these are noted in Table 24 on page 329.

The SEND, RECEIVE, and CONVERSE commands are fully described in the *CICS Application Programming Reference* manual. They are broken down by device group, because the options for different devices and access methods vary considerably. “Finding the right commands” on page 325 tells you which device group to use for your particular device.

Send/receive mode

The terminals and logical units covered in this chapter all operate in “half-duplex, flip-flop” mode. This means, essentially, that at any given moment, one partner in a conversation is in send mode (allowed to send data or control commands) and the other is in receive mode (restricted to receiving). This protocol is formally defined and enforced under VTAM. CICS observes the same conventions for terminals attached under other access methods, but both the hardware and the access methods work differently, so that not all operations are identical.

When a terminal is the principal facility of a task, its conversation partner is the task. When it is not associated with a task, its conversation partner is the terminal control component of CICS. Between tasks, under VTAM, the conversation is left in a neutral state where either partner can send first. Ordinarily the terminal goes first, sending the unsolicited input that initiates a task (see “How tasks get started” on page 200).

This transmission also reverses the send/receive roles; thereafter the terminal is in receive mode and CICS, represented by the task that was attached, is in send mode. The task starts and remains in send mode, no matter how many SENDs it executes, until it explicitly changes the direction of the conversation. One way in which you can put the task in receive mode is by specifying the INVITE option on a SEND command. After SEND with INVITE, the task is in receive mode and must issue a RECEIVE before sending to the terminal again. You can also put the task in receive mode simply by issuing a RECEIVE, without a preceding INVITE; INVITE simply optimizes transmissions.

Note that the first RECEIVE command in a task initiated by unsolicited input does not count in terms of send/receive mode, because the input message involved has long since transmitted (it started the task). This RECEIVE just makes the message accessible to the task, and sets the related EIB fields.

ATI tasks—those initiated automatically by CICS—also start out in send mode, just like tasks started by unsolicited input.

Note that if a task is executing normally and performing non-terminal operations when a VTAM/network error occurs, the task is unaware of the error and continues processing until it attempts the next terminal control request. It is at this point that the task receives the TERMERR. If the task does not issue any further terminal control request, it will not receive the TERMERR or ABEND.

Contention for the terminal

CICS satisfies requests for automatic task initiation (ATI) as soon as the terminal required as principal facility is available. When a task ends at a terminal, and CICS has an ATI request for that terminal, there may be contention between CICS, which wants to initiate the ATI task, and the terminal user, who wants to initiate a certain task via unsolicited input. In this situation, CICS always sets itself up as contention **loser**. That is, if the terminal sends unsolicited input quickly enough after the end of the previous transaction, CICS creates a task to process it and delay fulfilling the ATI request. This is intentional—it gives the user priority in contention situations.

EXEC CICS RETURN IMMEDIATE

However, you sometimes need to execute a sequence of particular tasks in succession at a terminal without allowing the user to intervene. CICS provides a way for you to do this, with the IMMEDIATE option on the EXEC CICS RETURN command that ends the task. With EXEC CICS RETURN IMMEDIATE, CICS initiates a task to execute the transaction named in the TRANSID option immediately, before honoring any other waiting requests for tasks at that terminal and without accepting input from the terminal. The old task can even pass data to the new one, using the INPUTMSG option. The new task accesses this data with an EXEC CICS RECEIVE, as if the user had initiated the task with unsolicited input, but no input/output occurs. This RECEIVE, like the first one in a task initiated by unsolicited input, has no effect on send/receive status; it just makes the passed data available to the new task. If the terminal is using bracket protocol (explained in “Preventing interruptions (bracket protocol)” on page 336), CICS does not end the bracket at the end of the first task, as it ordinarily does, but instead continues the bracket to include the following task. Consequently, the automatic opening of the keyboard at the end of bracket between tasks does not occur.

Speaking out of turn

It is usually clear to users when they are supposed to “talk” (key and transmit), and when they are supposed to “listen” (wait for output), because the application makes this clear. On 3270 displays and many other terminals, the keyboard locks after the user has transmitted to reinforce this convention. It remains locked until the task unlocks it, which it usually does on an EXEC CICS SEND before an EXEC CICS RECEIVE, or on the last EXEC CICS SEND in the task. This means the user has to do something particular (press the keyboard reset key) in order to break protocol.

What happens if the user does this? For terminals under BTAM, which is supported by CICS for VSE/ESA 2.3 and earlier, CICS provides two choices for an input it does not expect. CICS can either discard it or save it for use on the next occasion that an input is required (or acceptable) from that terminal. You specify one of these options using the PUNSOL system generation parameter; the option you choose is applied globally to all BTAM terminals.

For terminals connected under VTAM, violating this protocol causes the task to abend (code ATCV) unless read-ahead queueing is in force. Read-ahead queueing allows the logical unit and the task to send and receive at any time; CICS saves input messages in temporary storage until the task needs them. Inputs not read by task end are discarded. Read-ahead queueing is applied at the transaction level (it is specified in the RAQ option of the PROFILE under which the transaction runs). It applies only to LU type 4 devices, and is provided for compatibility reasons, to allow a transaction to support both BTAM-connected and VTAM-connected terminals in the same way. In general, it should not be used except to handle this situation.

Sequential terminals also differ from others in send/receive rules. Because the input is a pre-prepared file, CICS simply provides input messages whenever the task requests them, and it is impossible to break protocol. If the input is improperly prepared, or is not what the task is programmed to handle, it is possible for the task to get out of synchronization with its inputs, to exhaust them prematurely, or to fail to read some of them.

Interrupting

Both VTAM and BTAM provide a mechanism for a terminal in receive mode to tell its partner that it would like to send. This is the “signal” data flow in VTAM, which is detected on the next EXEC CICS SEND, RECEIVE or ISSUE DISCONNECT command from the task. When a signal flow occurs, CICS raises the SIGNAL condition and sets EIBSIG in the EIB. CICS default action for the SIGNAL condition is to ignore it. For the signal to have any effect, the task must first detect the signal and then honor it by changing the direction of the conversation. In BTAM, the corresponding flow is a **reverse interrupt** (RVI), which the terminal sends in place of the usual positive acknowledgment (ACK).

On a 3270 display terminal and some others, the ATTENTION key is the one that generates the interrupt. Not all terminals have this feature, however, and in VTAM, the bind image must indicate support for it as well, or VTAM will ignore the interrupts.

Terminal waits

When a task issues an EXEC CICS SEND command without specifying WAIT, CICS can defer transmission of the output to optimize either its overall terminal handling or the transmissions for your task. When it does this, CICS saves the output message and makes your task dispatchable, so that it can continue executing. The EXEC CICS ISSUE COPY and EXEC CICS ISSUE ERASE commands, which also transmit output, work similarly without WAIT.

If you use the WAIT option, CICS does not return control to your task until the output operation is complete. This wait lengthens the elapsed time of your task, with attendant effects on response time and memory occupancy, but it ensures that your task knows whether there has been an error on the SEND before continuing. You can avoid some of this wait and still check the completion of the operation if you have processing to do after your SEND. You issue the SEND without WAIT, continue processing, and then issue an EXEC CICS WAIT TERMINAL command at the point where you need to know the results of your SEND.

When you issue an EXEC CICS RECEIVE command that requires transmission of input, your task always waits, because the transmission must occur before the RECEIVE can be completed. However, there are cases where a RECEIVE does not correspond to terminal input/output. The first RECEIVE in a task initiated by unsolicited terminal input is the most frequent example of this, but there are others, as explained in the next section.

Also, when you issue any command involving your terminal, CICS ensures that the previous command is complete (this includes any deferred transmissions), before processing the new one.

What you get on an EXEC CICS RECEIVE

We use the terms “input message” and “transmission” to mean both what the terminal sent and what the application received. For the most common types of terminals, these are equivalent. A 3270 display, for example, sends whatever was changed in its buffer as a single entity, and the task associated with the terminal normally gets the entire message in response to a single RECEIVE command.

However, input messages and physical transmissions are not always equivalent, and there are several factors that can affect the one-to-one relationship of either to EXEC CICS RECEIVE commands. These are:

- VTAM chaining
- Logical records
- NOTRUNCATE option
- “Print” PA key

Input chaining

Some SNA devices break up long input messages into multiple physical transmissions, a process called “chaining”. CICS will assemble the component transmissions into a single input message or present them individually, depending on how the terminal associated with the task has been defined. This affects how many RECEIVES you need to read a chained input message. Details on inbound chaining are explained in “Chaining input data” on page 333.

Logical messages

Just as some devices break long inputs into multiple transmissions, others block short inputs and send them in a single transmission. Here again, CICS provides an option about who deblocks, CICS or the receiving program. This choice also affects how much data you get on a single RECEIVE. (See “Handling logical records” on page 334 for more on this subject.)

NOTRUNCATE option

Still another exception to the one-input-message-per-RECEIVE rule occurs when the length of the input data is greater than the program expects. If this occurs and the EXEC CICS RECEIVE command specifies NOTRUNCATE, CICS saves the excess data and uses it to satisfy subsequent EXEC CICS RECEIVE commands from the program with no corresponding read to the terminal. If you are using NOTRUNCATE, issue RECEIVES until the field EIBCOMPL in the EIB is set on (that is, set to X'FF'). CICS turns on EIBCOMPL when no more of the input message is available.

Without NOTRUNCATE, CICS discards the excess data, turns on EIBCOMPL, and raises the LENGERR condition. It reports the true length of the data, before truncation, in the data area named in the LENGTH option, if you provide one.

Print key

If your CICS system has a PA key defined as a “print” key, another exception to the normal send/receive sequence can occur. If the task issues a RECEIVE, and the user presses the “print” key in response, CICS intercepts this input, does the necessary processing to fulfil the request, and puts the terminal in receive mode again. The user must send another input to satisfy the original RECEIVE. (See “CICS print key” on page 355 for more information about the “print” key.)

Control commands

In addition to data transmission commands, the CICS API for terminals includes a series of commands that send instructions or control information, rather than data, to the terminal or to the access method controlling it. These commands are listed in the table below, along with a brief description of their function. Not all of these commands apply to all terminals, and for some, different forms apply to different

terminals. See “Finding the right commands” before going to the descriptions in the *CICS Application Programming Reference* manual.

The terminal in the table below is always the principal facility of the task issuing the command, except where explicitly stated otherwise. It may be a logical unit of a type not ordinarily considered a terminal.

<i>Table 22. Control commands for terminals and logical units</i>	
Command	Action
FREE	Releases the terminal from the task, so that the terminal may be used in another task before the current one ends.
ISSUE COPY	Copies the buffer contents of the terminal named in the TERMID option to the buffer of the terminal owned by the task. Both terminals must be 3270s.
ISSUE DISCONNECT	Schedules termination of the session between CICS and the terminal at the end of the task.
ISSUE ENDFILE	Sends an end-of-file notification to the terminal (for 3740 data entry systems only).
ISSUE ENDOUTPUT	Sends an end-of-output notification to the terminal (for 3740 data entry systems only).
ISSUE EODS	Sends an end-of-data-set function management header (for 3650 interpreter logical units only).
ISSUE ERASEAUP	Erases all the unprotected fields of the terminal (for 3270 devices only).
ISSUE LOAD	Instructs the terminal to load the program named in the PROGRAM option (for 3650 interpreter logical units only).
ISSUE PASS	Schedules disconnection of the terminal from CICS and its transfer to the VTAM application named in the LUNAME option, at the end of the issuing task.
ISSUE PRINT	Copies the terminal buffer to the first printer eligible for a print request (for 3270 displays only).
ISSUE RESET	Disconnects the line to which the terminal is attached, at task end, (for switched lines under BTAM only).
WAIT SIGNAL	Suspends the issuing task until its terminal sends a SIGNAL dataflow command.
WAIT TERMINAL	Suspends the issuing task until the previous terminal operation has completed.

Finding the right commands

Hardware and access method sensitivity is one of the major distinctions between using BMS and using terminal control commands to communicate with a terminal. BMS shields an application from hardware dependencies at the expense of some loss of function, whereas terminal control provides all the function.

The result of providing full function is that not all terminal control commands apply to all devices. Some commands require that you know what type of terminal you have, to determine the options that apply and the exceptional conditions that can occur. For some commands, you will also need to know what access method is in use. The two tables that follow tell you which commands apply to which terminal and access method combinations. If you need to support several types of

terminals, you can find out which type your task has as its principal facility using the commands described in “Finding out about your terminal” on page 330.

To use the tables, look up the terminal type that your program must support in the first column of Table 23. Use the value in the second column to find the corresponding command group in the first column of Table 24 on page 329. The second column of this table tells you the access method, and the third tells you the commands you can use. The commands themselves are described in full in the *CICS Application Programming Reference* manual. Where there is more than one version of a command in that manual, the table tells you which one to use. This information appears in parentheses after the command, just as it does in the manual itself.

<i>Table 23 (Page 1 of 3). Devices supported by CICS</i>	
Device	Use commands for
1050 (1051, 1052, 1053, 1056)	2741
2260, 2265	2260
2740, 2741	2741
2770, 2780	System/3
2980	2980
3101 (supported as TWX 33/35)	3767
3230	
3230 (VTAM)	3767
3230 (non-VTAM)	2741
3270 family (3178, 3179, 3180, 3262, 3275, 3277, 3278, 3279, 3284, 3286, 3287, 3288, 3289, 3290)	
3270 displays, 3270 printers (VTAM SNA)	LU type 2/3
3270 displays, 3270 printers (VTAM non-SNA)	3270 logical
3270 displays, 3270 printers (non-VTAM)	3270 display
SCS printers (VTAM)	SCS
3600	
Pipeline mode (VTAM)	3600 pipeline
3601 (VTAM)	3600-3601
3614 (VTAM)	3600-3614
Non-VTAM	3600 BTAM
3630, attached as 3600 (3631, 3632, 3633, 3643, 3604)	Use 3600 entry
3640	
3641, 3644, 3646, 3647 (VTAM, attached as 3767)	3767
3643 (VTAM, attached as LU type 2)	LU type 2/3
3642, 3645 (VTAM, attached as SCS printer)	SCS
3650	
3650 interpreter LU	3650 interpreter
3650 host conversational LU (3270)	3650-3270
3650 host conversational LU (3653)	3650-3653

<i>Table 23 (Page 2 of 3). Devices supported by CICS</i>	
Device	Use commands for
3650 host command LU (3680, 3684)	3650-3680
3660	System/3
3730	3790 full function or inquiry
3735	3735
3740, 3741	3740
3767	
interactive LU (VTAM)	3767
non-VTAM	2741
3770	
Interactive LU (VTAM)	3767
Full function LU	3790 full function or inquiry
Batch LU (3771, 3773, 3774) (VTAM)	3770
Non-VTAM (3775, 3776, 3777) (supported as 2770)	System/3
3780	System/3
3790	
Full function or inquiry	3790 full function or inquiry
3270 display LU	3790 3270-display
SCS printer	3790 SCS
3270 printer	3790 3270-printer
4700 (supported as 3600)	Use 3600 entry
5100	2741
5110, attached as 2770	System/3
5230, 5231, 5260, 5265 (supported as 3741)	3740
5280	
attached as 3741	3740
attached as 3270	Use 3270 entry
5520	
VTAM, supported as 3790 full-function LU	3790 full function or inquiry
non-VTAM, supported as 2770	System/3
5550 (supported as 3270)	Use 3270 entry
5937 (supported as 3270)	Use 3270 entry
6670	
VTAM	LU type 4
non-VTAM, supported as 2770	System/3
8100	

<i>Table 23 (Page 3 of 3). Devices supported by CICS</i>	
Device	Use commands for
8130, 8140 under DPCX (supported as 3790)	3790 full function or inquiry
DPPX/BASE using Host Presentation Services or Host Transaction Facility (attached as 3790)	3790 full function or inquiry
DPPX/DSC, DPCX/DSC, including 8775 attach (supported as 3270)	LU type 2/3
8775	LU type 2/3
8815	APPC ¹¹
CMCST	2741
Displaywriter	
supported as 3270	Use 3270 entry
supported as APPC	APPC ¹¹
INTLU (interactive LU)	3767
Office System/6 (6640, 6670), attached as 2770	System/3
PC, PS/2, attached as 3270	Use 3270 entry
Scanmaster	APPC ¹¹
Series/1	
supported as 3650 pipeline	3600 pipeline
supported as 3790 full-function LU	3790 full function or inquiry
System/3 (5406, 5408, 5410, 5412, 5415)	System/3
System/7 (5010)	System/7
System/32 (5320)	
VTAM, supported as 3770	Use 3770 entry
non-VTAM, supported as 2770	System/3
System/34 (5340)	
VTAM, supported as 3770	Use 3770 entry
non-VTAM	System/3
System/36 (supported as System/34)	Use System/34 entry
System/38 (5381)	
VTAM, attached as 3770	Use 3770 entry
VTAM, attached as APPC	APPC ¹¹
Non-VTAM	System/3
TWX 33/35	
VTAM NTO	3767
non-VTAM	2741
WTTY	
VTAM NTO	3767
non-VTAM	2741

<i>Table 24 (Page 1 of 2). Terminal control commands by device type</i>		
Device group name	Access methods	Commands applicable
2260	Non-VTAM	RECEIVE (2260), SEND (2260), CONVERSE (2260), ISSUE DISCONNECT (default), ISSUE RESET
2741	Non-VTAM	RECEIVE (2741), SEND (2741), CONVERSE (2741), ISSUE RESET
2980	Non-VTAM	RECEIVE (2980), SEND (2980)
3270 display	Non-VTAM	RECEIVE (3270 display), SEND (3270 display), CONVERSE (3270 display), ISSUE COPY (3270 display), ISSUE DISCONNECT (default), ISSUE ERASEAUP, ISSUE PRINT, ISSUE RESET
LU type 2/3 (3270 SNA)	VTAM	RECEIVE (LU type 2/3), SEND (LU type 2/3), CONVERSE (LU type 2/3), ISSUE COPY (3270 logical), ISSUE DISCONNECT (default), ISSUE ERASEAUP, ISSUE PASS, ISSUE PRINT
3270 logical (3270 non-SNA)	VTAM	RECEIVE (3270 logical), SEND (3270 logical), CONVERSE (3270 logical), ISSUE COPY (3270 logical), ISSUE DISCONNECT (default), ISSUE ERASEAUP, ISSUE PASS, ISSUE PRINT
SCS	VTAM	SEND (SCS), CONVERSE (SCS), ISSUE DISCONNECT (default), ISSUE PASS
3600 pipeline	VTAM	RECEIVE (3600 pipeline), SEND (3600 pipeline), ISSUE DISCONNECT (default), ISSUE PASS
3600-3601	VTAM	RECEIVE (3600-3601), SEND (3600-3601), CONVERSE (3600-3601), ISSUE DISCONNECT (default), ISSUE PASS, WAIT SIGNAL
3600-3614	VTAM	RECEIVE (3600-3614), SEND (3600-3614), CONVERSE (3600-3614), ISSUE DISCONNECT (default), ISSUE PASS
3600 BTAM	Non-VTAM	RECEIVE (3600 BTAM), SEND (3600 BTAM), CONVERSE (3600 BTAM), ISSUE RESET
3650 interpreter	VTAM	RECEIVE (3650), SEND (3650 interpreter), CONVERSE (3650 interpreter), ISSUE DISCONNECT (default), ISSUE EODS, ISSUE LOAD, ISSUE PASS
3650-3270	VTAM	RECEIVE (3650), SEND (3650-3270), CONVERSE (3650-3270), ISSUE DISCONNECT (default), ISSUE ERASEAUP, ISSUE PASS, ISSUE PRINT
3650-3653	VTAM	RECEIVE (3650), SEND (3650-3653), CONVERSE (3650-3653), ISSUE DISCONNECT (default), ISSUE PASS
3650-3680	VTAM	RECEIVE (3650), RECEIVE (3790 full function or inquiry), SEND (3650-3680), SEND (3790 full function or inquiry), CONVERSE(3650-3680), ISSUE DISCONNECT (default), ISSUE PASS

Table 24 (Page 2 of 2). Terminal control commands by device type

Device group name	Access methods	Commands applicable
3735	Non-VTAM	RECEIVE (3735), SEND (3735), CONVERSE (3735), ISSUE DISCONNECT (default), ISSUE RESET
3740	Non-VTAM	RECEIVE (3740), SEND (3740), CONVERSE (3740), ISSUE DISCONNECT (default), ISSUE ENDFILE, ISSUE ENDOUTPUT, ISSUE RESET
3767	VTAM	RECEIVE (3767), SEND (3767), CONVERSE (3767), ISSUE DISCONNECT (default), ISSUE PASS, WAIT SIGNAL
3770	VTAM	RECEIVE (3770), SEND (3770), CONVERSE (3770), ISSUE DISCONNECT (default), ISSUE PASS, WAIT SIGNAL
3790 full function or inquiry	VTAM	RECEIVE (3790 full function or inquiry), SEND (3790 full function or inquiry), CONVERSE (3790 full function or inquiry), ISSUE DISCONNECT (default), ISSUE PASS, WAIT SIGNAL
3790 3270-display	VTAM	RECEIVE (3790 3270-display), SEND (3790 3270-display), CONVERSE (3790 3270-display), ISSUE DISCONNECT (default), ISSUE ERASEAUP, ISSUE PASS, ISSUE PRINT
3790 3270-printer	VTAM	SEND (3790 3270-printer), ISSUE DISCONNECT (default), ISSUE ERASEAUP, ISSUE PASS
3790 SCS	VTAM	SEND (3790 SCS), ISSUE DISCONNECT (default), ISSUE PASS
LU type 4	VTAM	RECEIVE (LU type 4), SEND (LU type 4), CONVERSE (LU type 4), ISSUE DISCONNECT (default), ISSUE PASS, WAIT SIGNAL
Outboard controllers (batch data interchange)	VTAM	ISSUE ABORT, ISSUE ADD, ISSUE END, ISSUE ERASE, ISSUE NOTE, ISSUE QUERY, ISSUE RECEIVE, ISSUE REPLACE, ISSUE SEND, ISSUE WAIT
System/3	Non-VTAM	RECEIVE (System/3), SEND (System/3), CONVERSE (System/3)
System/7	Non-VTAM	RECEIVE (System/7), SEND (System/7), CONVERSE (System/7), ISSUE DISCONNECT (default), ISSUE RESET
All others	VTAM	RECEIVE (VTAM default), SEND (VTAM default), CONVERSE (VTAM default), ISSUE PASS
All others	Non-VTAM	RECEIVE (non-VTAM default), SEND (non-VTAM default), CONVERSE (non-VTAM default)

Finding out about your terminal

Some applications must support more than one type of terminal, and sometimes the types are sufficiently different that they require separate code. If you are writing such a program, and you need to determine what sort of terminal it is currently communicating with, you can use the EXEC CICS ASSIGN command to find out.

The EXEC CICS ASSIGN command returns a variety of information about the executing task, including a number of fields that describe its principal facility. Table 25 lists the ones that relate directly to terminal control operations. There are other ASSIGN options that relate to BMS and to other aspects of the task. You can find details on all ASSIGN options in the *CICS Application Programming Reference* manual. The “terminal” cited in column 2 of the table is always the principal facility of the task.

<i>Table 25 (Page 1 of 2). ASSIGN command options for terminals</i>	
ASSIGN option	Information returned
ALTSCRNHT ALTSCRNWD	The alternate height and width of the terminal screen (from its terminal definition); see also SCRNHT and SCRNWD
APLKYBD	Whether terminal has an APL keyboard
APLTEXT	Whether terminal has the APL text feature
BTRANS	Whether terminal has background transparency capability
COLOR	Whether terminal has extended color capability
DEFSCRNHT DEFSCRNWD	The default height and width of the terminal screen (from its terminal definition); see also SCRNHT and SCRNWD
DELIMITER	The data-link control character for the terminal (for 3600 terminals only)
DESTID DESTIDLENGTH	The identifier of the outboard destination and its length (for BDI operations only)
DSSCS	Whether the terminal is an SCS data stream device
DS3270	Whether the terminal is a 3270 data stream device
EXTDS	Whether the terminal supports “query structured field” orders
EWASUPP	Whether the terminal supports “erase write alternate” orders (i.e. has alternate screen size capability)
FACILITY	The four-character identifier of the terminal
FCI	The type of principal facility associated with the task (terminal, queue, etc.)
GCHARS GCODES	The graphic character set global identifier and the code page global identifier associated with the terminal
HILIGHT	Whether the terminal has extended highlight capability
KATAKANA	Whether the terminal supports Katakana
LANGINUSE	Three-character mnemonic
MAPCOLUMN	The column on the display containing the origin of the most recently positioned map
MAPHEIGHT	The height of the most recently positioned map
MAPLINE	The line on the display containing the origin of the most recently positioned map
MAPWIDTH	The width of the most recently positioned map.
MSRCONTROL	Whether the terminal supports magnetic slot reader control
NATLANGINUSE	The national language in use for the current task
NETNAME	The 8-character identifier of the terminal in the VTAM network

<i>Table 25 (Page 2 of 2). ASSIGN command options for terminals</i>	
ASSIGN option	Information returned
NUMTAB	Number of tabs required to position the print element in the correct passbook area (for 2980s only)
OPID OPCLASS	Operator identifier code and operator class of user signed on at terminal
OUTLINE	Whether the terminal has field outlining capability
PARTNS	Whether the terminal supports screen partitions
PS	Whether the terminal has programmed symbols capability
SCRNHT SCRNWD	Height and width of the terminal screen for the current task
SIGDATA	SIGNAL data received from the terminal
SOSI	Whether the terminal has mixed EBCDIC/double-byte character set capability
STATIONID TELLERID	Station and teller identifier of the terminal (for 2980s only)
TERMCODE	Type and model number of the terminal
TERMPRIORITY	Terminal priority value
TEXTKYBD	Whether the terminal has the TEXTKYBD feature
TEXTPRINT	Whether the terminal has the TEXTPRINT feature
UNATTEND	Whether the terminal is unattended
USERID USERNAME USERPRIORITY	The 8-character identifier, 20-character name and priority of the user signed on at the terminal
VALIDATION	Whether the terminal has validation capability

You can also use the EXEC CICS INQUIRE TERMINAL command to find out about your own terminal or any other terminal. EXEC CICS INQUIRE TERMINAL returns information from the terminal definition, whereas ASSIGN describes the use of that terminal in the current task. For many options, however, particularly the hardware characteristics, the information returned is the same. EXEC CICS INQUIRE TERMINAL is described in the *CICS System Programming Reference* manual.

EIB feedback on terminal control operations

CICS reports the results of processing terminal control commands, including those generated by BMS, in the EIB. Because of the complexity of terminal operations, many EIB fields are specific to terminal commands. Those that apply to the principal facility are listed in Table 26. (Other fields relate only to LU type 6.1, APPC and MRO operations; see the *CICS Application Programming Reference* manual for programming information about these.)

EIB fields are posted when CICS returns control to your task, and they always describe the most recent command to which they apply. This means that if you are conducting program-to-program communication over an alternate facility and using your principal facility, you need to check the EIB before results from one overlay those of the other.

It also means that when a task is initiated by unsolicited input from the terminal, or by an EXEC CICS RETURN IMMEDIATE in the previous task at the same terminal, the EIB fields that describe the input are **not** set at task start. You must issue an EXEC CICS RECEIVE to gain access to the input and post the EIB fields.

Note: If you are interested only in the EIB values and not the data itself, omit both the INTO and SET options from your EXEC CICS RECEIVE command.

Here are the fields that apply to the principal facility:

<i>Table 26. EIB fields that apply to terminal control commands</i>	
Field	Contents
EIBAID	The attention identifier (AID) from the last input operation (3270s only, see "The AID" on page 222)
EIBATT	Whether the input contains attach header data (an attach FMH)
EIBCOMPL	Whether the RECEIVE command just issued used all the input data, or more RECEIVES are required (see "Chaining output data" on page 334)
EIBCPOSN	Cursor position at time of last input operation (3270s only)
EIBEOC	Whether an end-of-chain indicator appeared in the input from the last RECEIVE
EIBFMH	Whether user data just received or retrieved contains an FMH
EIBFREE	Whether the facility just used has been freed
EIBRCODE, EIBRESP, EIBRESP2	CICS response code values from the previously issued command Note: For output commands in which transmission can be deferred, these values reflect only the initial CICS processing of the command, not the eventual transmission (see "Terminal waits" on page 323).
EIBSIG	Whether the terminal has sent a SIGNAL
EIBTRMID	(CICS) identifier of the terminal

VTAM considerations

Under VTAM, communication with logical units is governed by the conventions (protocols) which vary with the type of logical unit. This section describes the options provided by CICS to enable applications to conform to and make best use of these protocols,

The subsystem guides listed in "Where to find more information" on page 205 explain which protocols apply to which logical units.

Chaining input data

As noted earlier, some SNA devices segment long input messages for transmission. Each individual segment is called a **request unit** (RU), and the entire logical message is called a **chain**. CICS provides an option in the RDO TYPETERM resource definition BUILDCHAIN, that governs who assembles the chain. If the BUILDCHAIN value for the terminal is YES, CICS assembles the chain and presents the entire message to the program in response to a single EXEC CICS RECEIVE command. This choice ensures that the whole chain is complete and available before it is presented to the application.

If BUILDCHAIN(NO), the application assembles the chain. CICS provides one RU for each EXEC CICS RECEIVE command. The application can tell when it has received the last RU in the chain, because CICS raises the EOC (end-of-chain) condition at that time. CICS raises this condition even when there is only one RU in the chain, or when it assembles the chain, or when the input is from a terminal that does not support inbound chaining, like a 3270 display. An EOC condition is not considered an error; the CICS default action when it occurs is to ignore the condition.

EOC may occur simultaneously with either the EODS (end-of-data-set) or INBFMH (inbound-FMH) conditions, or both. Either condition will take precedence over EOC in determining where control goes if both it and EOC are the subject of active EXEC CICS HANDLE CONDITION commands.

Chaining output data

VTAM supports the chaining of outbound as well as inbound terminal data. If the length of an output message exceeds the outbound RU size, and the terminal supports outbound chaining, CICS breaks the message into RU-size segments and transmits them separately.

Your application can take advantage of the fact that chaining is permitted by passing a single output message to CICS bit by bit across several EXEC CICS SEND commands. To do this, you specify the CNOTCOMPL (“chain not complete”) option on each SEND except the one that completes the message. (Your message segments do not have to be any particular length; CICS will assemble and transmit as many RUs as are required.) The RDO PROFILE resource definition under which your transaction is running must specify CHAINCONTROL(YES) in order for you to do this.

Note: Options that apply to a complete logical message (that is, the whole chain) must appear only on the first EXEC CICS SEND command for a chain. These include FMH, LAST, and, for the 3601, LDC.

Handling logical records

As noted earlier, some devices block input messages and send multiple inputs in a single transmission. CICS allows you to specify whether CICS or the application should deblock the input. The choice is expressed in the LOGREC option of the PROFILE under which the current transaction is executing.

With LOGREC (NO), CICS provides the entire input message in response to an EXEC CICS RECEIVE (assuming the input is not chained or BUILDCHAIN(YES)). The user is responsible for deblocking the input. If BUILDCHAIN(NO), a RECEIVE will retrieve one RU of the chain at a time. In general, logical records do not span RUs, so that a single RU contains one or more complete logical records. The exception is LU type 4 devices, where a logical record may start in one RU and continue in another; for this reason, BUILDCHAIN(YES) is recommended if you do your own deblocking for these devices.

If the RDO PROFILE resource definition specifies LOGREC (YES), CICS provides one logical record in response to each RECEIVE command (whether or not CICS is assembling input chains).

If an RU contains more than one logical record, the records are separated by new line (NL) characters, X'15', interrecord separators (IRS characters), X'1E', or

transparent (TRN) characters, X'35'. If NL characters are used, they are not removed when the data is passed to the program and appear at the end of the logical record. If IRS characters are used, however, they are removed. If the delimiter is a transparent character, the logical record can contain any characters, including NL and IRS, which are considered normal data in transparent mode. The terminating TRN is removed, however. CICS limits logical records separated by TRNs to 256 characters.

Response protocol

Under VTAM, CICS allows the use of either definite response or exception response protocol for outbound data.

Under exception response, a terminal acknowledges a SEND only if an error occurred. If your task is using exception response, CICS does not wait for the last SEND in the task (which may be the only SEND) to complete before terminating your task. Consequently, if an error does occur, it may not be possible to report it to your task. When this happens, the error is reported to a CICS-supplied task created for the purpose.

Definite response requires that the terminal acknowledge every SEND, and CICS does not terminate your task until it gets a response on the last SEND. Using definite response protocol has some performance disadvantages, but it may be necessary in some applications.

The MSGINTEG option of the PROFILE under which a task is running determines which response mode is used. However, if you select MSGINTEG (NO) (exception response), you can still ask for definite response on any particular SEND by using the DEFRESP option. In this way, you can use definite response selectively, paying the performance penalty only when necessary.

Using function management headers

SNA architecture defines a particular type of header field that accompanies some messages, called a function management header (FMH). It conveys information about the message and how it should be handled. For some logical units, use of an FMH is mandatory, for others it is optional, and in some cases FMHs cannot be used at all. In particular, FMHs do **not** apply to LU type 2 and LU type 3 terminals, which are the most common 3270 devices.

The subsystem guides listed in "Where to find more information" on page 205 tell you which devices accept or require FMHs, and exactly how to format them. FMH data must conform to SNA format specifications; otherwise ATCY abends or unpredictable results can occur.

Inbound FMH

When an FMH is present in an input message, CICS consults the PROFILE definition under which the transaction is executing to decide whether to remove it or pass it on to the application program that issued the RECEIVE. The PROFILE can specify that no FMHs are to be passed, that only the FMH indicating the end of the data set should be passed, or that all FMHs are to be passed. There is also an option that causes the FMH to be passed to the batch data interchange program.

If an FMH is present, it occupies the initial bytes of the input message; its length varies by device type. CICS sets the EIBFMH field in the EIB on (X'FF') to tell

you that one is present, and it also raises the INBFMH condition, which you can detect through an EXEC CICS HANDLE CONDITION command or by testing the RESP value.

Outbound FMH

On output, the FMH can be built by the application program or by CICS. If your program supplies the FMH, you place it at the front of your output data and specify the FMH option on your EXEC CICS SEND command. If CICS is to build the FMH, you reserve the first three bytes of the message for CICS to fill in and omit the FMH option. CICS will build an FMH only for devices that require one; you must supply it for devices for which it is optional.

Preventing interruptions (bracket protocol)

Brackets are an SNA protocol for ensuring that a conversation between two LUs is not interrupted by a request from a third LU. CICS uses bracket protocol to prevent interruption of the conversation between a CICS task and its principal facility for the duration of the task. If the task has an alternate facility, bracket protocol is used there also, for the same reason. The logical unit begins the bracket if it sends unsolicited input to initiate the task, and CICS begins the bracket if it initiates the task automatically. CICS ends the bracket at task end, unless the IMMEDIATE option appears on the final EXEC CICS RETURN command. EXEC CICS RETURN IMMEDIATE lets you initiate another task at your principal facility without allowing it to enter input. CICS does this by **not** ending the bracket between the ending task and its successor when brackets are in use.

CICS requires the use of brackets for many devices under VTAM. For others, the use of brackets is determined by the value of the BRACKET option in the terminal definition. In general, bracket protocol is transparent to an application program, but it is still possible to optimize flows related to bracket protocol using the LAST option on the SEND command. If you know that a particular EXEC CICS SEND will be the last command for the terminal in a task, you can improve performance by adding the LAST option. LAST allows VTAM to send the “end-of-bracket” indicator with the data and saves a separate transmission to send it at task end. If you are sending the last output in a program-built chain (using CNOTCOMPL), LAST must be specified on the first EXEC CICS SEND for the chain in order to be effective.

If your task has significant work to do or may experience a significant delay after its last SEND, you may want to issue an EXEC CICS FREE command. FREE releases the terminal for use in another task.

Sequential terminal support

One of the many types of terminal that CICS supports is not really a terminal at all, but a pair of sequential devices or files simulating a terminal. One of the pair represents the input side of the terminal, and might be a card reader, a spool file or a SAM file on tape or DASD. The other represents the output, and might be a printer, a punch, spool or SAM file. Many device-type combinations are allowed, and either of the pair can be missing; that is, you can have an input-only or output-only sequential terminal.

You read from and write to the devices or files that constitute a sequential terminal with terminal control commands, specifically RECEIVE, SEND, and CONVERSE.

(BMS supports sequential terminals too; see “Special options for non-3270 terminals” on page 268.)

The original purpose of sequential terminal support was to permit application developers to test online code before they had access to real terminals. This requirement rarely occurs any more, but sequential terminals are still useful for:

Printing See “CICS API considerations” on page 352. Sequential terminals are particularly useful for output that is sometimes directed to a low-speed CICS printer, for which BMS or terminal control commands are required, and sometimes directed to a high-speed system printer (spool or transient data commands). If you define the high-speed printer as a sequential terminal, you can use terminal control or BMS commands, and you can use the same code for both types of printers. (If there are differences in the device data streams, you need to use BMS for complete transparency.)

Regression testing

Tests run from sequential terminals leave a permanent record of both input and output. This encourages systematic and verifiable initial testing. Also, it allows you to repeat tests after modifications, to ensure that a given set of inputs produces the same set of outputs after the change as before.

Initialization

Some installations use a sequential terminal to execute one or more initialization transactions, in preference to program list table programs. Transactions initiated from a sequential terminal begin execution as soon as the terminal is in service, and they continue as quickly as CICS can process them until the input is exhausted. Hence the inputs from a sequential terminal can be processed immediately after startup, if the sequential terminal is initially in service, at some later time (when it is put in service) or even as part of a controlled shutdown.

Coding considerations for sequential terminals

The input data submitted from a sequential terminal must be in the form in which it would come from a telecommunication device. For example, the first record usually starts with a transaction code, to tell CICS what transaction to execute. The transaction code must start in the first position of the input, just as it must on a real terminal. Note that this limits the ability to test applications that require input in complex formats. For example, there is no provision for expressing a formatted 3270 input stream as a sequential file, because of all the complex control sequences. However, you can use an unformatted 3270 data stream (or any other similar stream) for input, and you can still use BMS to format your output.

When you build the input file, you place an end-of-data indicator (EODI) character after each of your input records. The EODI character is defined in the system initialization table; the default value is a backslash ('\, X'E0'), but your installation may have defined some other value.

CICS observes only your EODI characters in processing the input stream; it pays no attention to record structure of the input file or device, so that one of your inputs can span records in the input file. Similarly, you do not start each input on a new physical record, but immediately after the EODI terminating the previous one.

The length of an input record (the number of characters between EODIs) should not exceed the size of the input buffer (the INAREAL value in the LINE component of the sequential terminal definition). If it does, the transaction that attempts to RECEIVE the long record abends, and CICS positions the input file after the next EODI before resuming input processing.

An end-of-file marker in the input also acts as an EODI indicator. Any EXEC CICS RECEIVE command issued after end-of-file is detected also causes an abend.

Print formatting

If the definition of a sequential terminal indicates that the output half is a line printer, you can write multiple lines of output with a single EXEC CICS SEND command. For this type of device, CICS breaks your output message into lines after each new line character (X'15') or after the number of characters defined as the line length, whichever occurs first. Line length is defined by the LPLEN value in the terminal definition. Each EXEC CICS SEND command begins a new line.

GOODNIGHT convention

CICS continues to initiate transactions from a sequential terminal until it (or the transactions themselves) have exhausted all the input or until the terminal goes out of service. To prevent CICS from attempting to read beyond the end of the input file (which causes a transaction abend), the last transaction executed can put the terminal out of service after its final output. Alternatively (and this is usually easier), the last input can be a CEF GOODNIGHT transaction, which signs the terminal off and puts it out of service. You cannot normally enter further input from a sequential terminal once CICS has processed its original input, without putting it out of service.

Batch data interchange

Many installations have a host computer and database at a central location, linked to other computers at branch offices. They do not necessarily contain CICS, but they can communicate with CICS in the host system. The CICS batch data interchange program provides for communication between an application program and a named data set (or destination) that is part of a batch data interchange logical unit in an outboard controller, or with a selected medium on a batch logical unit or an LU type 4 logical unit. This medium indicates the required device such as a printer or console.

The term “outboard controller” is a generalized reference to a programmable subsystem, such as the IBM 3770 Data Communication System, the IBM 3790 Data Communication System, or the IBM 8100 System running DPCX, which uses SNA protocols. (Details of SNA protocols and the data sets that can be used are given in *CICS/DOS/VS IBM 3767/3770/6670 Guide* and *CICS/DOS/VS IBM 3790/3780/8100 Guide*.) Figure 60 on page 339 gives an overview of batch data interchange.




Figure 60. CICS batch data interchange

The following batch data interchange commands are provided:

ISSUE QUERY

Initiate transfer of a data set to the CICS application program.

ISSUE RECEIVE

Read a record from a data set or read data from an input medium.

ISSUE SEND

Transmit data to a named data set or to a selected medium.

ISSUE ADD

Add a record to a data set.

ISSUE REPLACE

Update (replace) a record in a data set.

ISSUE ERASE

Delete a record from a data set.

ISSUE END

Terminate processing of a data set.

ISSUE ABORT

Terminate processing of a data set abnormally.

ISSUE NOTE

Request the next record number in a data set.

ISSUE WAIT

Wait for an operation to be completed.

Where the controller is an LU type 4 logical unit, only the ISSUE ABORT, ISSUE END, ISSUE RECEIVE, ISSUE SEND, and ISSUE WAIT commands can be used.

Where the data set is a DPCX/DXAM data set, only the ISSUE ADD, ISSUE ERASE, and ISSUE REPLACE commands can be used.

See Chapter 9, “Dealing with exception conditions” on page 141 for information about how to deal with any exception conditions that occur during execution of a batch data interchange command.

Destination selection and identification

All batch data interchange commands except EXEC CICS ISSUE RECEIVE include options that specify the destination. This is either a named data set in a batch data interchange logical unit, or a selected medium in a batch logical unit or LU type 4 logical unit.

Selection by named data set

The DESTID and DESTIDLENG options must always be specified, to supply the data set name and its length (up to a maximum of eight characters). For destinations having diskettes, the VOLUME and VOLUMELENG options may be specified, to supply a volume name and its length (up to a maximum of six characters); the volume name identifies the diskette that contains the data set to be used in the operation. If the VOLUME option is not specified for a multidiskette destination, all diskettes are searched until the required data set is found.

Selection by medium

As an alternative to naming a data set as the destination, various media can be specified by means of the CONSOLE, PRINT, CARD, or WPMEDIA1–4 options. These media can be specified only in an ISSUE ABORT, ISSUE END, ISSUE SEND, or ISSUE WAIT command.

Definite response

CICS uses terminal control commands to carry out the functions specified in batch data interchange commands. For those commands that cause terminal control output requests to be made, the DEFRESP option can be specified. This option has the same effect as the DEFRESP option of the EXEC CICS SEND terminal control command; that is, to request a definite response from the outboard controller, irrespective of the specification of message integrity for the CICS task (by the system programmer). The DEFRESP option can be specified for the ISSUE ADD, ISSUE ERASE, ISSUE REPLACE, and ISSUE SEND commands.

Waiting for function completion

For those batch data interchange commands that cause terminal control output requests to be made, the NOWAIT option can be specified. This option has the effect of allowing CICS task processing to continue; unless the NOWAIT option is specified, task activity is suspended until the batch data interchange command is completed. The NOWAIT option can be specified only on the ISSUE ADD, ISSUE ERASE, ISSUE REPLACE, and ISSUE SEND commands.

After a batch data interchange command with the NOWAIT option has been issued, task activity can be suspended, by the ISSUE WAIT command, at a suitable point in the program to wait for the command to be completed.

Chapter 20. CICS support for printing

This chapter contains information specific to printing, and contains the following topics:

- “Formatting for CICS printers”
- “CICS printers: getting the data to the printer” on page 349
- “Non-CICS printers” on page 352
- “Printing display screens” on page 354

CICS does not provide special commands for printing, but there are options on BMS and terminal control commands that apply only to printers, and for some printers you use transient data or EXEC CICS SPOOL commands. We cover the factors that determine the API and the choices you have in the sections that follow.

There are two issues associated with printing that do not usually occur in other types of end-user communication:

1. There are additional formatting considerations, especially for 3270 printers
2. The task that needs to print may not have direct access to the printer

In addition, there are two distinct categories of printer, which have different application programming interfaces:

CICS printers

Printers defined as terminals to CICS and managed directly by CICS. They are usually low-speed devices located near the end users, suitable for printing on demand of relatively short documents. The 3289 and 3262 are usually attached as CICS printers.

Non-CICS printers

Printers managed by the operating system or another application. These printers are most often high-speed devices located at the central processing site, appropriate for volume printing that does not have to be available immediately. They may also be advanced function or other printers that require special connections, management or sharing.

Because of the differences in the programming interface, we discuss the two groups separately, although there are formatting and access issues in both cases. We describe CICS printers first, followed by non-CICS printers; the last section of this chapter covers a special category of printing, the copying of display screens to printers.

Formatting for CICS printers

The application programming interface for writing to a printer terminal is essentially the same as for writing to a display. (This section does not discuss the problem of arranging that your task have the printer as its principal facility; this is discussed in “CICS printers: getting the data to the printer” on page 349.)

You can use terminal control commands (EXEC CICS SENDs) for any CICS printer, and most of them are supported by BMS too (EXEC CICS SEND MAP, EXEC CICS SEND TEXT, and EXEC CICS SEND CONTROL). “BMS support levels” on page 226 lists the devices that BMS supports. For printers that are

components of an outboard controller or LU Type 4, you can use batch data interchange (BDI) commands as well as terminal control and BMS. BDI commands are described in “Batch data interchange” on page 338.

The choice between using BMS and terminal control is based on the same considerations as it is for a display terminal. Like displays, printers differ widely from one another, both in function and in the implementation of that function, and the differences are reflected in the data streams and device controls they accept.

When you use terminal control commands, your application code must format the output in the manner required by the printer. For line printers and similar devices, formatting has little programming impact. For high-function printers, however, the data stream often is very complex; formatting requires significant application code and introduces device dependencies into program logic.

For some of these terminals, coding effort is greatly reduced by using BMS, which relieves the programmer of creating or even understanding device data streams. BMS also removes most data stream dependencies from the application code so that the same program can support many types of printers, or a mixture of printers and displays, without change. BMS does not remove all device dependencies and imposes a few restrictions on format. It also involves extra pathlength; the amount depends on how many separate BMS requests you make, the complexity of your requests, and the corresponding pathlength avoided in your own program.

3270 printers

Most of the additional format controls for printers that BMS provides are for a specific type of CICS printer, the 3270 printer. A **3270 printer** is any printer that accepts the 3270 data stream—it is the hardcopy equivalent of a 3270 display. It has a page buffer, corresponding to the display buffer of a 3270 display device. (See “The 3270 buffer” on page 210 for an introductory discussion of the 3270 data stream.) We discuss 3270 printers first and defer the simpler, non-3270 printers, until “Non-3270 CICS printers” on page 346.

A 3270 printer accepts two different types of formatting instructions: **buffer control orders** and **print format orders**. Buffer control orders are executed as they are received by the control unit, and they govern the way in which the buffer is filled. These are same orders that are used to format a 3270 display screen. We have already described some of the important ones in “Orders in the data stream” on page 215, SBA (set buffer address) tells the control unit where in the buffer to place the data that follows, SF (start field), which signals an attributes byte and possibly field data, and so on. You can find a complete list in the *IBM 3270 Information Display System Data Stream Programmer's Reference* manual.

In contrast, print format orders are not executed when they are received, but instead are stored in the buffer along with the data. These orders—NL (new line), FF (form feed), and so on—are interpreted only during a print operation, at which time they control the format of the printed output. (They have no effect on displays, other than to occupy a buffer position; they look like blanks on the screen.)

If you are writing to a 3270 printer, you can format with either buffer control orders or print format orders or a mixture of both. We show an example of formatting with buffer control orders in “Outbound data stream sample” on page 218. If you send this same data stream to a 3270 printer, it will print an image of the screen shown in Figure 48 on page 218. You might choose to format printed output with buffer

control orders so that you can send the same data stream to a display and a printer.

On the other hand, you might choose to format with print format orders so that you can send the same stream to a 3270 printer and a non-3270 printer (print format orders are the same as the format controls on many non-3270 printers). See the discussion of the NLEOM option on page 345 for more details about this choice.

Here is a data stream using print format orders that produces the same **printed** output as the data stream on page 218, which uses buffer control orders.

Table 27. Example of data stream using print control orders

Bytes	Contents	Notes
1	X'FF'	"Formfeed" (FF) order, to cause printer to space to a new page.
2-23	blanks	22 blanks to occupy columns 1-22 on first line.
24-33	Car Record	Text to be printed, which appears in the next available columns (23-32) on line 1.
34	X'1515'	Two successive "new line" (NL) orders, to position printer to beginning of third line.
35-80	Employee No: _____ Tag _____ State: __	Text to be printed, starting at first position of line 3.
81	X'19'	"End-of-message" (EM) print order, which stops the printing.

Notice that the field structure is lost when you use print format orders. This does not matter ordinarily, because you do not use the printer for input. However, even if you format with print control orders, you might need to use buffer control orders as well, to assign attributes like color or underscoring to an area of text.

Options for 3270 printers

For BMS, the special controls that apply to 3270 printers take the form of command options:

- PRINT
- ERASE
- L40, L64, L80 and HONEOM
- NLEOM
- FORMFEED

In terminal control commands, ERASE is also expressed as an option, but the other controls are expressed directly in the data stream. The *IBM CICS/DOS/VS 3270 Data Stream Device Guide* and the *IBM 3270 Information Display System Data Stream Programmer's Reference* manual tell you how to encode them; the discussion that follows explains what they do.

PRINT option and print control bit

Writing to a 3270 display or printer updates the device buffer. On a display, the results are reflected immediately on the screen, which is driven from the buffer. For a printer, however, there might be no visible effect, because printing does not occur until you turn on the appropriate bit in the “write control character”. (The WCC is part of the 3270 data stream; see “Write control character” on page 212.) For BMS, you turn on the print bit by specifying the PRINT option on an EXEC CICS SEND MAP, SEND TEXT or SEND CONTROL command, or in the map used with SEND MAP. If you are using terminal control SENDs, you must turn on the print bit with the CTLCHAR option.

A terminal write occurs on every terminal control SEND, and on every EXEC CICS SEND MAP, EXEC CICS SEND TEXT, and EXEC CICS SEND CONTROL unless you are using the ACCUM or PAGING options. ACCUM delays writing until a page is full or the logical message is ended. When you use ACCUM, use the same print options on every SEND command for the same page. PAGING defers the terminal writes to another task, but they are generated in the same way as without PAGING.

The fact that printing does not occur until the print bit is on allows you to build the print buffer in stages with multiple writes and to change data or attribute bytes already in the buffer. That is, you can use the hardware to achieve some of the effects that you get with the ACCUM option of BMS. The NLEOM option affects this ability, however; see the discussion below.

ERASE

Like the 3270 display buffer, the 3270 printer buffer is cleared only when you use a write command that erases. You do this by specifying the ERASE option, both for BMS and terminal control SENDs. If the printer has the alternate screen size feature, the buffer size is set at the time of the erase, as it is for a display. Consequently, the first terminal write in a transaction should include erasing, to set the buffer to the size required for the transaction and to clear any buffer contents left over from a previous transaction.

Line width options: L40, L64, L80, and HONEOM

In addition to the print bit, the write control character contains a pair of bits that govern line length on printing. If you are using terminal control commands, you use the CTLCHAR option to set these bits. For BMS, the default is the one produced by the HONEOM option, which stands for “honor end-of-message”. With this setting, the printer formats according to the buffer control and print format orders only, stopping printing at the first EM (end-of-message) character in the buffer. Only if you attempt to print beyond the maximum width for the device (the platen width) does the printer move to a new line on its own.

However, you also can specify that the line length is fixed at 40, 64 or 80 characters (the L40, L64 and L80 options, respectively). If you do, the printer ignores certain print format orders, moves to a new line when it reaches the specified line size, and prints the entire buffer. The print format orders that are ignored are NL (new line), CR (carriage return), and EM (end-of-message). Instead they are simply printed, as graphics.

If you use L40, L64 or L80 under BMS, use only the value that corresponds to the page width in your terminal definition (see “Determining the characteristics of a CICS printer” on page 347). The reason is that BMS calculates buffer addresses

based on the page size, and these addresses are wrong if you use a different page width.

NLEOM option

BMS ordinarily uses buffer control orders, rather than print format orders, to format for a 3270 printer, whether you are using EXEC CICS SEND TEXT or SEND MAP. However, you can tell BMS to use print format orders only, by specifying the NLEOM option. If you do, BMS formats the data entirely with blanks and NL (new line) characters, and inserts an EM (end-of-message) character after your data. NLEOM implies HONEOM. (NLEOM support requires **standard** BMS; it is not available in minimum BMS.)

You might want to do this in order to maintain compatibility with an SCS printer (print format orders are compatible with the corresponding SCS control characters). There are also operational differences that might cause you to choose or avoid NLEOM. They are:

Blank lines: The 3270 printer suppresses null lines during printing. That is, a line that has no data fields and appears blank on the display screen is omitted when the same map is sent to a printer. Under BMS, you can force the printed form to look exactly like the displayed form by placing at least one field on every line of the screen; use a field containing a single blank for lines that would otherwise be empty. Specifying NLEOM also has this effect, because BMS uses a new line character for every line, whether or not there is any data on it.

Multiple sends: With NLEOM, data from successive writes is simply stacked in the buffer, since it does not contain positioning information. However, BMS adds an EM (end-of-message) character at the end of data on each SEND with NLEOM, unless you are using the ACCUM option. When printing occurs, the first EM character stops the printing, so that only the data from the first SEND with NLEOM (and any unerased data up to that point in the buffer) gets printed. The net effect is that you cannot print a buffer filled with multiple EXEC CICS SEND commands with NLEOM unless you use the ACCUM option.

Page width: BMS always builds a page of output at a time, using an internal buffer whose size is the number of character positions on the page. (See “Determining the characteristics of a CICS printer” on page 347 for a discussion of how BMS determines the page size.) If you are using buffer control orders to format, the terminal definition must specify a page width of 40, 64, 80 or the maximum for the device (the platen size); otherwise your output might not be formatted correctly. If you are using NLEOM, on the other hand, the terminal definition may specify any page width, up to the platen size.

Total page size: If you are using buffer control orders, the product of the number of lines and the page width must not exceed the buffer size, because the buffer is used as an image of the page. Unused positions to the right on each line are represented by null characters. If you use NLEOM, however, BMS does not restrict page size to the buffer capacity. BMS builds the page according to the page size defined for the terminal and then compresses the stream using new-line characters where possible. If the resulting stream exceeds the buffer capacity, BMS uses multiple writes to the terminal to send it.

FORMFEED

The FORMFEED option causes BMS to put a form feed print format order (X'0C') at the beginning of the buffer, provided that the printer is defined as capable of advancing to the top of the form (with the FORMFEED option in the associated RDO TYPETERM resource definition). CICS ignores a form feed request for a printer defined without this feature.

If you issue an EXEC CICS SEND MAP using a map that uses position (1,1) of the screen, you overwrite the order and lose the form feed. This occurs whether you are using NLEOM or not.

If you use FORMFEED and ERASE together on an EXEC CICS SEND CONTROL command, the results depend on whether NLEOM is present. Without NLEOM, SEND CONTROL FORMFEED ERASE sends the form feed character followed by an entire page of null lines. The printer suppresses these null lines, replacing them with a single blank line. With NLEOM, the same command sends the form feed character followed by one new line character for each line on the page, so that the effect is a full blank page, just as it is on a non-3270 printer.

PRINTERCOMP option

When you EXEC CICS SEND TEXT to a printer, there is one additional option that affects page size. This is the PRINTERCOMP option, which is specified in the RDO PROFILE resource definition associated with the transaction you are executing, rather than on individual SEND TEXT commands. (In the default profile that CICS provides, the PRINTERCOMP value is NO.)

Under PRINTERCOMP(NO), BMS produces printed output consistent with what it would send to a 3270 display. For the display, BMS precedes the text from each SEND TEXT command with an attribute byte, and it also starts each line with an attribute byte. These attribute bytes take space on the screen, and therefore BMS replaces them with blanks for printers if PRINTERCOMP is NO. If PRINTERCOMP is YES, BMS suppresses these blanks, allowing you to use the full width of the printer and every position of the buffer. New line characters that you embed in the text are still honored with PRINTERCOMP(YES), as they are with PRINTERCOMP(NO).

Use PRINTERCOMP(NO) if you can, for compatibility with display devices and to ensure consistent results if the application uses different printer types, even though it reduces the usable line width by one position.

Non-3270 CICS printers

A **non-3270 printer** is any printer that does not accept the 3270 data stream, such as an SNA character set (SCS) printer. The terminology is somewhat confusing, because a non-3270 printer can be a 3270-family device, and many devices, like the 3287 and 3262, can be either 3270 printers or SCS (non-3270) printers, depending on how they are defined at the control unit.

There are special considerations for non-3270 printers, although not so many as for 3270 printers. Non-3270 printers do not have page buffers, and therefore do not understand buffer control orders. Formatting is accomplished entirely with print control orders. For compatibility with 3270 printers, BMS formats for them by constructing an image of a page in memory, and always prints a full page at a time. However, you can define any size page, provided you do not exceed the platen

width, as there is no hardware buffer involved. BMS transmits as many times as required to print the page, just as it does for a 3270 printer using the NLEOM option.

BMS formats for these printers with blanks and NL (new line) characters. It uses form feed (FF) characters as well if the definition of your terminal indicates form feed support.

BMS also uses horizontal tabs to format if the associated RDO TYPETERM resource definition has the HORIZFORM (YES) option and the map contains HTAB specifications. Similarly, it uses vertical tabs if the associated RDO TYPETERM specifies VERTICALFORM and your map includes VTAB. Tab characters can shorten the data stream considerably. If tabs are used, BMS assumes that the current task, or some earlier one, has already set the tabs on the printer. On an SCS printer, you set the tabs with a terminal control SEND command, as explained in the *IBM CICS/DOS/VS 3270 Data Stream Device Guide*. For other non-3270 printers, consult the appropriate device guide.

SCS input

SCS printers also have limited **input** capability, in the form of “program attention” keys. These keys are **not** like the PA keys described in “Attention keys” on page 221, however. Instead they transmit an unformatted data stream consisting of the characters ‘APAK nn’, where “nn” is the two-digit PA key number—‘APAK 01’ for PA key 1, for example.

You can capture such input by defining a transaction named ‘APAK’ (APAK is the transaction identifier, not the TASKREQ attribute value, because SCS inputs do not look like other PA key inputs.) A program invoked by this transaction can determine which PA key was pressed by issuing a RECEIVE and numeric positions of the input.

Determining the characteristics of a CICS printer

If you are writing a program that supports more than one type of CICS printer, you may need to determine the characteristics of a particular printer. As we explained in connection with terminals generally, you can use the ASSIGN and INQUIRE TERMINAL commands for this purpose. Table 25 on page 331 lists the ASSIGN options that apply to terminals, including several that are specific to printers.

The INQUIRE TERMINAL options that apply specifically to printers and the corresponding parameters in the terminal definition are shown in Table 28 on page 348:

<i>Table 28. INQUIRE TERMINAL options for printers</i>		
INQUIRE option	Source in TERMINAL or TYPETERM definition	Description
PAGEHT	x of PAGESIZE(x,y)	Number of lines per page (for alternate screen size terminals, reflects default size).
PAGEWD	y of PAGESIZE(x,y)	Number of characters per line (for alternate screen size terminals, reflects default size).
DEFPAGEHT	x of PAGESIZE(x,y)	Number of lines per page in default mode (alternate screen size terminals only).
DEFPAGEWD	y of PAGESIZE(x,y)	Number of characters per line in default mode (alternate screen size terminals only).
ALTPAGEHT	x of ALTPAGE(x,y)	Number of lines per page in alternate mode (alternate screen size terminals only).
ALTPAGEWD	y of ALTPAGE(x,y)	Number of characters per line in alternate mode (alternate screen size terminals only).
DEVICE	DEVICE	The device type (see the <i>CICS System Programming Reference</i> manual for possible values).
TERMMODEL	TERMMODEL	The model number of the terminal (either 1 or 2).

BMS page size, 3270 printers

BMS uses both the terminal definition and the profile of the transaction that is running to determine the page size of a CICS printer. The profile is used when the terminal has the alternate screen size feature, to determine whether to use default or alternate size. (The default profile in CICS specifies “default” size for the screen.) Table 29 lists the values used.

<i>Table 29. Priority of parameters defining BMS page size. BMS uses the first value in the appropriate column that has been specified in the terminal definition.</i>		
Terminals with alternate screen size, using alternate size	Terminals with alternate screen size, using default size	Terminals without alternate screen size feature
ALTPAGE	PAGESIZE	PAGESIZE
ALTSCREEN	DEFSCREEN	TERMMODEL
DEFSCREEN	TERMMODEL	(12,80)
TERMMODEL	(12,80)	(12,80)

The definition of a “page” is unique to BMS. If you are printing with terminal control SEND commands, you define what constitutes a page, within the physical limits of the device, by your print format. If you need to know the buffer size to determine how much data you can send at once, you can determine this from the SCRNHT and SCRNWD values returned by the EXEC CICS ASSIGN command.

Supporting multiple printer types

When you are writing programs to support printers that have different page sizes, it is not always possible to keep device dependencies like page size out of the program. However, BMS helps with this problem in two ways.

1. You can refer to a map generically and have BMS select the map that was designed for the terminal associated with your task (see the discussion of map suffixes in “Device-dependent maps: map suffixes” on page 269).
2. If you are using EXEC CICS SEND TEXT, BMS breaks the text into lines at word boundaries, based on the page size of the receiving terminal. You can also request header and trailer text on each page.

CICS printers: getting the data to the printer

As we noted at the start of the chapter, the second issue that frequently arises in printing concerns ownership of the printer. Requests for printing often originate from a user at a display terminal. The task that processes the request and generates the printed output is associated with the user's terminal and therefore cannot send output directly to the printer.

If your task does not own the printer it wants to use, it must create another task, which does, to do the work. These are the ways to do this:

1. Create the task with an EXEC CICS START command
2. Write to an intra-partition transient data queue that triggers the task
3. Direct the output to the printer in a BMS ROUTE command
4. Use the ISSUE PRINT command, if you need only a screen copy, and conditions suit.

Details on the first three methods follow. Screen copies are covered in “Printing display screens” on page 354.

Printing with an EXEC CICS START command

The first technique for creating the print task is to issue an EXEC CICS START command in the task that wants to print. The command names the printer as the terminal required by the STARTed task in the TERMID option and passes the data to be printed, or instructions on where to find it, in the FROM option. EXEC CICS START causes CICS to create a task whose principal facility is the designated terminal when that terminal is available.

The program executed by the STARTed task, which you must supply, retrieves the data to be printed (using an EXEC CICS RETRIEVE command), and then writes it to its terminal (the printer) with EXEC CICS SEND, EXEC CICS SEND MAP, or EXEC CICS SEND TEXT commands. For example:

Task that wants to print (on printer PRT1):

```
⋮  
(build output in OUTAREA, formatted as expected by the STARTed task)  
EXEC CICS START TRANSID(PRNT) FROM(OUTAREA) TERMID(PRT1)  
LENGTH(OUTLNG) END-EXEC.  
⋮
```

STARTed task (executing transaction PRNT):

```
⋮  
EXEC CICS RETRIEVE INTO(INAREA) LENGTH(INLNG) END-EXEC.  
⋮  
(do any further data retrieval and any formatting required)  
EXEC CICS SEND TEXT FROM(INAREA) LENGTH(INLNG) ERASE PRINT END-EXEC.  
⋮  
(repeat from the RETRIEVE statement until a NODATA condition arises)
```

The task associated with the printer loops until it exhausts all the data sent to it, in case another task sends data to the same printer before the current printing is done. Doing this saves CICS the overhead of creating new tasks for outputs that arrive while earlier ones are still being printed; it does not change what finally gets printed, as CICS creates new tasks for the printer as long as there are unprocessed EXEC CICS START requests.

Printing with transient data

The second method for creating the print task involves transient data. A CICS intra-partition transient data queue can be defined to have a property called a “trigger”. When the number of items on a queue with a trigger reaches the trigger value, CICS creates a transaction to process the queue. The queue definition tells CICS what transaction this task will execute and what terminal, if any, it requires as its principal facility.

You can use this mechanism to get print data from the task that generates it to a task that owns the printer. A transient data queue is defined for each printer to which you will direct output in this way. A task that wants to print puts its output on the queue associated with the required printer (using EXEC CICS WRITEQ TD commands). When enough items are on the queue and the printer is available, CICS creates a task to process the queue. (For this purpose, the trigger level of “enough” is usually defined as just one item.) The triggered task retrieves the output from the queue (with READQ TD commands) and writes it to its principal facility (the printer), with EXEC CICS SEND, EXEC CICS SEND MAP, or EXEC CICS SEND TEXT commands.

As in the case of a STARTed printer task, you have to provide the program executed by the task that gets triggered. The sample programs distributed with CICS contain a complete example of such a program, called the “order queue print sample program”. The *CICS Sample Applications Guide* describes this program in detail, but the essentials are as follows:

Task that wants to print (on printer PRT1):

```
⋮  
(do any formatting or other processing required)  
EXEC CICS WRITEQ TD QUEUE('PRT1') FROM(OUTAREA)  
LENGTH(OUTLNG) END-EXEC.  
⋮
```

Task that gets triggered:

```
⋮
EXEC CICS ASSIGN QNAME(QID) END-EXEC.
EXEC CICS READQ TD QUEUE(QID) INTO(INAREA) LENGTH(INLNG)
      RESP(RESPONSE) END-EXEC.
IF RESPONSE = DFHRESP(QZERO) GO TO END-TASK.
⋮
(do any error checking, further data retrieval and formatting required)
EXEC CICS SEND FROM(INAREA) LENGTH(INLNG) END-EXEC.
⋮
(repeat from READQ command)
```

The print task determines the name of its queue using an ASSIGN command rather than a hard-coded value so that the same code works for any queue (printer).

Like its START counterpart, this task loops through its read and send sequence until it detects the QZERO condition, indicating that the queue is empty. While this is just an efficiency issue with the STARTed task, it is critical for transient data; otherwise unprocessed queue items can accumulate under certain conditions. (See “Automatic transaction initiation (ATI)” on page 405 for details on the creation of tasks to process transient data queues with triggers.)

If you use this technique, you need to be sure that output to be printed as a single unit appears either as a single item or as consecutive items on the queue. There is no fixed relationship between queue items and printed outputs; packaging arrangements are strictly between the programs writing the queue and the one reading it. However, if a task writes multiple items that need to be printed together, it must ensure that no other task writes to the queue before it finishes. Otherwise the printed outputs from several tasks may be interleaved.

If the TD queue is defined as recoverable, CICS prevents interleaving. Once a task writes to a recoverable queue, CICS delays any other task that wants to write until the first one commits or removes what it has written (by SYNCPOINT or end of task). If the queue is not recoverable, you need to perform this function yourself. One way is to ENQUEUE before writing the first queue item and DEQUEUE after the last. (See Chapter 27, “Transient data control” on page 403 for a discussion of transient data queues.)

Printing with BMS routing

A task also can get output to a printer other than its principal facility with BMS routing. This technique applies only to BMS logical messages (the ACCUM or PAGING options) and thus is most appropriate when you are already building a logical message.

When you complete a routed message, CICS creates a task for each terminal named in a route list. This task has the terminal as its principal facility, and uses CSPG, the CICS-supplied transaction for displaying pages, to deliver the output to the printer. So routing is similar in effect to using START commands, but CICS provides the program that does the printing. (See “Message routing: the ROUTE command” on page 294 for more information about routing.)

Non-CICS printers

Here are the steps to use a printer managed outside CICS:

1. Format your output in the manner required by the application or subsystem that controls the printer you wish to use.
2. Deliver the output to the application or subsystem that controls the printer in the form required by that application.
3. If necessary, notify that application that the output is ready for printing.

Formatting for non-CICS printers

For some printers managed outside CICS, you can format output with BMS, as we explain in “CICS API considerations.” However, for most printers, you need to meet the format requirements of the application that drives the printer. This may be the device format or an intermediate form dictated by the application. For conventional line printers, formatting is simply a matter of producing line images and, sometimes, adding carriage-control characters.

Non-CICS printers: Delivering the data

Print data is usually conveyed to an application outside of CICS by placing the data in an intermediate file, accessible to both CICS and the application. The type of file, as well as the format within the file, is dictated by the receiving application. It is usually one of those listed in the first column of Table 30. The second column of the table shows which groups of CICS commands you can use to create such data.

File type	Methods for writing the data
Spool files	CICS spool commands (EXEC CICS SPOOLOPEN, EXEC CICS SPOOLWRITE, etc.) Transient data commands (WRITEQ TD) Terminal control and BMS commands (SEND, SEND MAP, etc.)
SAM	CICS spool commands (EXEC CICS SPOOLOPEN, EXEC CICS SPOOLWRITE, etc.) Transient data commands (WRITEQ TD)
VSAM	CICS file control commands (EXEC CICS WRITE)
DB2 for VSE/ESA	EXEC SQL commands
DL/I VSE	EXEC DLI commands or CALL DLI statements

CICS API considerations

If you are using VSAM, DB2 for VSE/ESA, or DL/I VSE, the CICS application programming commands you can use are determined by the type of file you are using.

For SAM and spool files, however, you have a choice. The CICS definition of the file (or its absence) determines which commands you use. The file may be:

- An extra-partition transient data queue (see Chapter 27, “Transient data control” on page 403 for information on transient data queues)
- The output half of a sequential terminal (see “Sequential terminal support” on page 336 and “Support for non-3270 terminals” on page 267)

- A spool file (see Chapter 21, “CICS interface to POWER” on page 357)

Both transient data queue definitions and sequential terminal definitions point to an associated data definition (DLBL) statement in the CICS start-up JCL, and it is this DLBL statement that determines whether the file is a SAM file or a spool file. Files created by CICS spool commands do not require definition before use and are spool files by definition.

If the printing application accepts SAM or spool file input, there are several factors to consider in deciding how to define your file to CICS:

System definitions

Files created by the EXEC CICS SPOOLOPEN command do not have to be defined to CICS or the operating system, whereas transient data queues and sequential terminals must be defined to both before use.

Sharing among tasks

A file defined as a transient data queue is shared among all tasks. This allows you to create a print file in multiple tasks, but it also means that if your task writes multiple records to the queue that must be printed together (lines of print for a single report, for example), you must include enqueue logic to prevent other tasks from writing their records between yours. This is the same requirement that was cited for intra-partition queues in “Printing with transient data” on page 350. In the case of extra-partition transient data, however, CICS does not offer the recoverability solution, and your program must prevent the interspersing itself.

In contrast, a file created by an EXEC CICS SPOOLOPEN command can be written only by the task that created it. In the report controller facility, a report can be written to by more than one task (APPENDABLE and LOG report types). This eliminates the danger of interleaving output, but also prevents sharing the file among tasks.

A spool file associated with a sequential terminal can be written by only one task at a time (the task that has the terminal as its principal facility). This also prevents interleaving, but allows tasks to share the file serially.

Release for printing

Both SAM and spool files must be closed in order for the operating system to pass them from CICS to the receiving application, and therefore printing does not begin until the associated file is closed. Files created by EXEC CICS SPOOLOPEN are closed automatically at task end, unless they have already been closed with an EXEC CICS SPOOLCLOSE command. In contrast, an extrapartition transient data queue remains open until some task closes it explicitly, with an EXEC CICS SET command. (It must be reopened with another SET if it is to be used subsequently.) So transient data gives you more control over release of the file for processing, at the cost of additional programming.

A file that represents the output of a sequential terminal does not get closed automatically (and so does not get released for printing) until CICS shutdown, and CICS does not provide facilities to close it earlier. If you use a sequential terminal to pass data to a printer controlled outside of CICS, as you might do in order to use BMS, be aware of this limitation.

Formatting

If you define your file as a sequential terminal, you can use BMS to format your output. This feature allows you to use the same maps for printers managed outside of CICS—for example, line printers managed by VSE POWER—that you use for CICS display and printer terminals.

If you choose this option, remember that BMS always sends a page of output at a time, using the page size in the terminal definition, and that the data set representing the output from a sequential terminal is not released until CICS shutdown.

Spool file limits

Operating systems identify spool files by assigning a sequential number. There is an upper limit to this number, after which numbers are reused. The limit is usually very large, but it is possible for a job that runs a very long time (as CICS can) and creates a huge number of spool files (as an application under CICS can) to exceed the limit. If you are writing an application that generates a very large number of spool files, consult your systems programmer to ensure that you will be within system limits. A new spool file is created at each EXEC CICS SPOOLOPEN statement and each open of a transient data queue defined as a spool file.

Notifying the print application

When you deliver the data to a print application outside CICS, you might need to notify the application that you have data ready to process. You do not need to do this if the application runs automatically and knows to look for your data. For example, to print on a printer owned by VSE POWER, all you need to do is create a spool file with the proper routing information. POWER does the rest.

However, sometimes you need to submit a job to do the processing, or otherwise signal an executing application that you have work for it.

To submit a batch job from a CICS task, you need to create a spool file that contains the JCL for the job, and you need to direct this file to the POWER internal reader. You can create the file in any of the three ways listed for spool files in Table 30 on page 352, although if you use a sequential terminal, the job does not execute until CICS shuts down, as noted earlier. For files written with spool commands, the information that routes the file to VSE POWER is specified in the EXEC CICS SPOOLOPEN command. For transient data queues and sequential terminals, the routing information appears on the first record in the file, the “JOB card”.

The output to be printed can be embedded in the batch job (as its input) or it can be passed separately through any form of data storage that the job accepts.

Printing display screens

If your printing requirement is simply to copy a display screen to a printer, you have choices additional to those already described. Some of these are provided by the terminal hardware itself, and some by CICS. Some of the CICS support also depends on hardware features, and so your options depend on the type of terminals involved and, in some cases, the way in which they are defined to CICS. See the *CICS Resource Definition Guide* for more detail on copying.

CICS print key

The first such option is the **CICS print key** (also called the **local copy key**). This allows a user to request a printed copy of a screen by pressing a program attention key, provided the terminal is a 3270 display or a display in 3270 compatibility mode. Print key support is optional in CICS; the system programmer decides whether to include it and what key is assigned. The default is PA1. (See the PRINT option in the *CICS System Definition Guide*.)

The print key copies the display screen to the first available printer among those defined as eligible. Which printers are eligible depends on the definition of the display terminal from which the request originates, as follows:

- For VTAM 3270 displays defined without the “printer-adapter” feature, the printers named in the PRINTER and ALTPRINTER options of the terminal definition are eligible. PRINTER is used if available; ALTPRINTER is second choice. If both are unavailable, the request is queued for execution when PRINTER becomes available.
- For the 3270 compatibility mode of the 3790 and a 3650 host conversational (3270) logical unit, the same choices apply.
- For a BTAM 3270 display, any printer on the same control unit is eligible, provided it is defined to CICS with the PRINT feature and has a buffer at least the size of the display buffer. CICS chooses the first one in the table that is available. If none are, the request is not queued but discarded.
- For VTAM 3270 displays defined with the printer-adapter feature, copying is limited to printers on the same control unit as the display. The printer authorization matrix within the control unit determines printer eligibility.
- For a 3270 compatibility mode logical unit of the 3790 with the printer-adapter feature, the 3790 determines eligibility and allocates a printer for the copy.
- For a 3275 with the printer-adapter feature, the print key prints the data currently in the 3275 display buffer on the 3284 attached to the display.

Where CICS chooses the printer explicitly, as it does in the first three cases above, the printer has to be in service and not attached to a task to be “available” for a CICS print key request. Where a control unit or subsystem makes the assignment, availability and status are determined by the subsystem. The bracket state of the device usually determines whether it is available or not.

EXEC CICS ISSUE PRINT and EXEC CICS ISSUE COPY

An application can initiate copying a screen to a printer as well as the user, with the EXEC CICS ISSUE PRINT and EXEC CICS ISSUE COPY commands. EXEC CICS ISSUE PRINT simulates the user pressing the CICS print key, and printer eligibility and availability are the same as for CICS print key requests.

There is also a command you can use to copy a screen in a task that owns the printer, as opposed to the task that owns the terminal that is to be copied. This is the EXEC CICS ISSUE COPY command. It copies the buffer of the terminal named in the TERMID option to the buffer of the principal facility of the issuing task. The method of copying and the initiation of printing once the copy has occurred is controlled by the “copy control character” defined in the CTLCHAR option of the EXEC CICS ISSUE COPY command; see the *IBM CICS/DOS/VS 3270 Data Stream Device Guide* for the bit settings in this control character. The

terminal whose buffer is copied and the printer must both be either 3270 logical units or BTAM 3270s, and they must be on the same control unit.

Hardware print key

Some 3270 terminals also have a **hardware print key**. Pressing this key copies the screen to the first available and eligible printer on the same control unit as the display. This function is performed entirely by the control unit, whose configuration and terminal status information determine eligibility and availability. If no printer is available, the request fails; the user is notified by a symbol in the lower left corner of the screen and must retry the request later.

BMS screen copy

Both the CICS and hardware print keys limit screen copies to a predefined set of eligible printers, and if more than one printer is eligible, the choice depends on printer use by other tasks. For screens created as part of a BMS logical message, a more general screen copy facility is available. Users can print any such screen with the “page copy” option of the CICS-supplied transaction for displaying logical messages, CSPG. With page copy, you name the specific printer to receive the output, and it does not have to be on the same control unit as the display. CSPG is described in the *CICS-Supplied Transactions* manual.

Chapter 21. CICS interface to POWER

CICS interface to VSE/POWER

The CICS interface to VSE/POWER provides a programming interface to functions that access the system spool files maintained by VSE/POWER. Using VSE/POWER facilities, you can support the requirements of other products to exchange files with other systems connected through VSE/POWER.

To access VSE/POWER from CICS, you must code SPOOL=(YES, identifier, class) at system initialization time.

With VSE/POWER, you can exchange documents with other systems, through an EXEC interface to:

1. Retrieve data for a specific user from the local POWER spool (see Figure 61 on page 358).
2. Write data to the POWER spool, using VSE/POWER facilities (see Figure 62 on page 359).
3. Send a VSE/POWER spool file to a specific remote destination (see Figure 63 on page 359).

The interface described here provides an opportunity for you to use VSE/POWER facilities to access VSE/POWER spool files. The CICS report controller, gives much greater scope in the handling of VSE/POWER spool files. For a description of the commands supported by the report controller, see the Chapter 22, "Report Controller" on page 363, and the *CICS Report Controller Planning Guide*.

Input to CICS from VSE/POWER

A remote application must route any files intended for a CICS transaction to a specific userid at the system where CICS resides; see Figure 61 on page 358 for an example of a CP command used by a VM system to do this.

The CICS transaction issues the EXEC CICS SPOOLOPEN INPUT command. The normal response is that the file is allocated to you for retrieval, and is identified by the "token" returned by CICS. This token must be included on every spool command for retrieving the data set, and for closing the data set.

If the response is normal, the transaction should then retrieve the file with EXEC CICS SPOOLREAD commands. When there is more than one transaction using the interface, their files may be differentiated by using different userids or different classes within a single userid.

When the file is exhausted, or you have read as much as required, the transaction should issue on EXEC CICS SPOOLCLOSE command.

The input spool file cannot be shared with other tasks-until the task that opened the spool file closes it. The file must be read in the same logical unit of work that opened it, and CICS will close it automatically at an EXEC CICS SYNCPOINT command or at task end if the task has not done so. However, you can close the

file in such a way that your task (or another one) can read it again from the beginning.




Figure 61. Retrieve data from the VSE/POWER spool

Output from CICS using VSE/POWER

The transaction program issues EXEC CICS SPOOLOPEN OUTPUT to cause creation of an output report, specifying a remote NODE and USERID.

- If the node is a remote VSE system, then the report is queued on the VSE/POWER spool against the userid specified on the SPOOLOPEN OUTPUT USERID command.
- If the node is a remote MVS system, then the report is queued on the JES spool against the external writename. The identifier of this writename was specified on the SPOOLOPEN OUTPUT USERID parameter.
- If the node is a remote VM system, then the data is queued in the VM RDR queue for the identifier that was specified on the same USERID parameter.

EXEC CICS SPOOLOPEN OUTPUT returns a unique token to the transaction, which must be used in all subsequent EXEC CICS SPOOLWRITE and EXEC CICS SPOOLCLOSE commands to identify the relevant file. Finally, the transaction issues EXEC CICS SPOOLCLOSE to close the report, and permit its immediate printing or onward routing by the system spooler.

A task can create multiple output spool files, and it can have more than one open at a time; operations are kept separate via the token returned at SPOOLOPEN time. However, a spool file cannot be shared among tasks (if you want to be able to do this, see Chapter 22, "Report Controller" on page 363 for information.) nor across logical units of work in the same task. If the task fails to close the file before an EXEC CICS SYNCPOINT command or task end, CICS closes it automatically at these points.




Figure 62. Create a file and write directly to the VSE/POWER spool. See also Figure 63 on page 359




Figure 63. Send the written file to a remote destination

Typical use of EXEC CICS SPOOL commands

Figure 64 on page 360 gives an example of the use of the EXEC CICS spool commands where:

1. The spool file is opened. If there is no data, or the single-thread input path is busy, the program sends an error message and restarts in five seconds.
2. If there is data and the input file is available, the program opens the output file. Then it reads and writes the logical records back to the output file.
3. At the comment PROCESS DATA, a report could be generated. For clarity, only comments have been inserted here.
4. The input and output files are closed.

```

DFHEISTG DSECT
CICSAMPL CSECT
OPENIN EQU *
* OBTAIN AND CLEAR GETMAIN AREA
EXEC CICS GETMAIN SET(PTR) LENGTH(500) INITIMG(BLANK)
USING RENTAREA,PTR
* OPEN INPUT SPOOLFILE
EXEC CICS SPOOLOPEN INPUT USERID(XWTRNAME) *
TOKEN(INTOKEN) *
RESP(RESPLD) RESP2(RESPLD)
CLC RESPLD,DFHRESP(NORMAL) OPEN OK?
BE OPENOUT YES
CLC RESPLD,DFHRESP(SPOLBUSY) NO, IS SPOOL IN USE?
BE RESTART YES, TRY AGAIN LATER
CLC RESPLD,DFHRESP(NOTFND) NO, IS ANY DATA ON SPOOL?
BE RESTART NO, TRY AGAIN LATER
OPENERR EQU * ERROR TERMINATE TRANSACTION AND RETURN MESSAGE
MVC MSGAREA(L'ERRMSG1),ERRMSG1
B SAMPLEOJ
RESTART EQU * SPOOL BUSY OR NO DATA. TRY AGAIN LATER.
EXEC CICS START TRANSID('LADM') INTERVAL(000005) *
TERMID('BLA3') NOCHECK PROTECT
MVC MSGAREA(L'OKMSG02),OKMSG02
B SAMPLEOJ
OPENOUT EQU * EVERYTHING OK. DATA ON INPUT FILE. OPEN OUTPUT. *
EXEC CICS SPOOLOPEN OUTPUT NODE(VM1) USERID(USER1) *
TOKEN(OUTTOKEN) NOCC *
RESP(RESPLD) RESP2(RESPLD)
CLC RESPLD,DFHRESP(NORMAL) OPEN OK?
BE CONTINU0 YES
MVC MSGAREA('ERRMSG2'),ERRMSG2 NO! GIVE ERROR MSG
B SAMPLEOJ
CONTINU0 EQU * READ INPUT DATA AND PROCESS
EXEC CICS SPOOLREAD INTO(INREC) *
TOKEN(INTOKEN) *
TOFLENGTH(L80) MAXFLENGTH(L132) *
RESP(RESPLD) RESP2(RESPLD)
CLC RESPLD,DFHRESP(NORMAL) READ OK?
BE CONTINU1 YES
CLC RESPLD,DFHRESP(ENDFILE) ANY MORE DATA?
BE ENDATA NO
MVC MSGAREA(L'ERRMSG3),ERRMSG3 MOVE IN ERROR MSG
B SAMPLEOJ
CONTINU1 EQU * YES, PROCESS DATA
* PROCESS DATA. CREATE REPORT AND PREPARE TO OUTPUT REPORT TO SPOOLER OUTPUT FILE
MVI CTLCHAR,C' '
MVC OUTDATA,CTLCHAR
MVC OUTDATA,INREC
EXEC CICS SPOOLWRITE FROM(OUTREC) *
TOKEN(OUTTOKEN) *
FLENGTH(L80) *
RESP(RESPLD) RESP2(RESPLD)
CLC RESPLD,DFHRESP(NORMAL) WRITE OK?
BE CONTINU2 YES
MVC MSGAREA('ERRMSG4'),ERRMSG4 NO MOVE ERROR MSG
B SAMPLEOJ

```

Figure 64 (Part 1 of 2). An example using EXEC CICS SPOOL commands

```

CONTINU2 EQU *
          B      CONTINU0      READ NEXT INPUT RECORD FROM SPOOL FILE
ENDATA   EQU *
* CREATE TOTAL SALE PAGE AND OUTPUT TO SPOOLER FILE.
* CLOSE INPUT AND OUTPUT SPOOL FILES.
CLOSEIN  EQU *
          EXEC CICS SPOOLCLOSE TOKEN(INTOKEN) *
              RESP(RESPLD) RESP2(RESPLD) DELETE
          CLC  RESPLD,DFHRESP(NORMAL)      CLOSE INPUT OK?
          BE   CLOSEOUT      YES
          MVC  MSGAREA(L'ERRMSG5A),ERRMSG5A      NO MOVE ERROR MSG
          MVC  MSGAREA+23(L'ERRMSG5B),ERRMSG5B
          MVC  MSGAREA+36(L'ERRMSG5D),ERRMSG5D
CLOSEOUT EQU *
          EXEC CICS SPOOLCLOSE TOKEN(OUTTOKEN) *
              RESP(RESPLD) RESP2(RESPLD) KEEP
          CLC  RESPLD,DFHRESP(NORMAL)      CLOSE OUTPUT OK?
          BE   SAMPLE0J      YES
          MVC  MSGAREA(L'ERRMSG5A),ERRMSG5A      NO MOVE ERROR MSG
          MVC  MSGAREA+28(L'ERRMSG5C),ERRMSG5C
          MVC  MSGAREA+36(L'ERRMSG5D),ERRMSG5D
SAMPLE0J EQU *
          MVC  MSGAREA+50(L'OKMSG01),OKMSG01      MOVE EOJ MESSAGE
          EXEC CICS SEND TEXT FROM(MSGAREA) LENGTH(80)
          EXEC CICS RETURN
R1       EQU 1      REGISTER 1
PTR      EQU 5      REGISTER 5
OKMSG01 DC  CL20'  TRAN LADM EOJ.      '
OKMSG02 DC  CL45'  TRAN LADM RESTARTED, NO DATA OR SPOOL BUSY. '
ERRMSG1  DC  CL45'  TRAN LADM TERMINATED. SPOOLOPEN INPUT ERROR. '
ERRMSG2  DC  CL46'  TRAN LADM TERMINATED. SPOOLOPEN OUTPUT ERROR. '
ERRMSG3  DC  CL39'  TRAN LADM TERMINATED. SPOOLREAD ERROR. '
ERRMSG4  DC  CL40'  TRAN LADM TERMINATED. SPOOLWRITE ERROR. '
ERRMSG5A DC  CL23'  TRAN LADM, SPOOLCLOSE '
ERRMSG5B DC  CL5'  INPUT '
ERRMSG5C DC  CL7'  OUTPUT '
ERRMSG5D DC  CL6'  ERROR. '
L80      DC  F'80'
L132     DC  F'132'
BLANK    DC  X'00'
XWTRNAME DC  CL8' CICSAMPL '
USER1    DC  CL8' USER1 '
VM1      DC  CL8' VM1 '
RENTAREA DSECT
INTOKEN  DS  CL8
OUTTOKEN DS  CL8
RESPLD   DS  CL4
RESP2FLD DS CL4
INREC    DS  CL80
OUTREC   DS  0CL133
CTLCHAR  DS  CL1
OUTDATA  DS  CL132
MSGAREA  DC  CL80 '
          END

```

Figure 64 (Part 2 of 2). An example using EXEC CICS SPOOL commands

Programming note for spool commands

You must specify either RESP or NOHANDLE on all spool commands. If you do not, your program abends. These options are described in the *CICS Application Programming Reference* manual, but they do not appear explicitly in the syntax boxes or the option lists for the spool commands. The list of conditions that can occur for each command includes RESP values.

Chapter 22. Report Controller

The report controller provides an end-user interface to the POWER LST and POWER RDR queues. This makes it possible for end-users (who may not be DP professionals) to control the printing of their own reports, without needing to know anything about either CICS or POWER.

For more detailed information about using the report controller transactions, see the *CICS Report Controller User's Guide*, which has been designed especially for non-DP end-users.

The report controller transactions (CEMS and CEOS) allow the end-user to control when the report is printed, which printer is used, how many copies to print, and so on. They do this by **changing the characteristics of the report** (this is the phrase used to include all these operations: you see it used in the *CICS Report Controller Planning Guide*).

The characteristics of the report are set up initially by the program you write to create the report. The basic commands for creating reports (EXEC CICS SPOOLOPEN, EXEC CICS SPOOLWRITE, EXEC CICS SPOOLCLOSE) are provided by CICS itself, but you need the report controller feature to be able to use the options that set many of these characteristics. In addition to options related to controlling printing, the report controller provides several other options that make report creation easier than existing methods.

Some of these command options can be specified in the JCL for reports created by batch programs. To handle the nonstandard stream of batch-generated reports, you can specify the name of an ESCAPE program in the JCL for the report. The named ESCAPE program is invoked when the report is printed.

If you are designing application programs to create reports that you want end-users to control and print using the report controller, first read the *CICS Report Controller Planning Guide*. This gives an overview of report creation using the commands and options that are described in detail here. It also gives an overview of how end-users control and print their reports: how they use the various characteristics that you set up. This helps you to understand how your choice of options influences their work.

The Low Balance Report sample program shows how to create a report using the report controller. The source code is supplied in the VSE/ESA sublibrary PRD1.BASE in four languages as follows:

COBOL	DFH0CRPS
C	DFH\$DRPS
PL/I	DFH\$PRPS
Assembler	DFH\$ARPS

In addition, the supplied transaction (SREP) creates a selection of reports so that end-users can try out the report controller. This also shows you how to create reports, and the source program is also supplied in each of the languages. For more details about running both the sample program and SREP, see the *CICS*

System Definition Guide. The reports that SREP creates form the basis of the tutorial in the *CICS Report Controller User's Guide*.

You can use the command interpreter (CECI) to try out the report controller commands. (You can see the commands, even when you do not have the report controller installed.) However, we recommend that you do not use CECI to create reports.

Reports

In the *CICS Report Controller Planning Guide*, a **report** is a file of data to be printed. In one special case, it can be a file of job control language (JCL) statements and, optionally, data to be processed by a batch program. Reports are held as POWER spool files during and after creation, and before, during, and sometimes after, printing. Because a report is equivalent to a spool file, the commands you use to create them begin with EXEC CICS SPOOL, to distinguish them from commands you use for creating ordinary files. These are sometimes referred to collectively as the "EXEC CICS SPOOL" commands. They are:

- EXEC CICS SPOOLOPEN REPORT ...
- EXEC CICS SPOOLWRITE REPORT ...
- EXEC CICS SPOOLCLOSE REPORT ...

The REPORT option identifies the command as a report controller command, as well as specifying a name for the report. There are other options that:

- Identify the report (to users and to CICS)
- Specify how the report is to be created
- Affect its appearance
- Influence how, when, and where it is to be printed
- Provide for security and recovery

Your choice depends on the way in which your transactions are going to add data to the report. There are three basic **report types** that you can create, as follows:

Standard report

You can have one transaction create the whole report in one go.

Resumable report

You can have several transactions or a pseudoconversational transaction create a report in separate stages.

Log report

You can have many independent transactions interleaving their data to create a continuous report.

You can also format your reports according to several different methods. The **report format** that you choose depends on how you want to control such things as page breaks, line lengths, new lines, headings, footings, and page numbering. Your choice depends largely on whether you are converting an existing program to use the EXEC CICS SPOOL commands, or writing a new one.

If your program uses ASA carriage control characters, machine code carriage control characters (MCC), the SNA character string (SCS), a 3270 data stream, or BMS mapping, there is an output format for you. Not only that, but all these formats (with the exception of BMS) are converted to MCC, so that you can print

any report on any of your printers, with the same results. However, if you want to use special features of your SCS or 3270 printers, you can avoid the conversion.

It is partly because of this convertibility and partly because of the options for control over page breaks, line lengths, headings, footings, and page numbering, that you need not worry about the format for new programs. For simplicity, choose NOCC (no carriage control) or ASA; you then have maximum flexibility with minimum fuss.

If you have special requirements for graphics or for an output format not mentioned, you can write your own **escape** program that is invoked when the report is ready to be printed.

Finally, there is a special format (JCL) for creating input to batch jobs. See the *CICS Report Controller Planning Guide* for further details about the report controller.

Part 5. CICS management functions

This part describes several groups of operations that are not specifically database or data communication operations, but control the processing of tasks within a CICS system. These are:

- Interval control—time-dependent functions
- Task control—for temporarily giving control back to CICS or for synchronizing resource usage
- Program control—for affecting the flow of control between application programs
- Storage control—to get and release areas of main storage
- Transient data control—for transferring data between CICS tasks and between the CICS region and other regions
- Temporary storage control—for the temporary storage of data

Each group of operations has its own chapter.

Chapter 23. Interval control	369	Defining the execution key	393
Expiration times	369	Selecting and defining the storage key for applications	394
Request identifiers	371	Deciding what execution and storage key to specify	397
Chapter 24. Task control	373	User-key applications	397
Controlling sequence of access to resources	374	CICS-key applications	398
Chapter 25. Program control	377	Storage protection exception conditions	400
Application program logical levels	378	Protecting CICS from being passed invalid addresses	400
Link to another program expecting return	378	C programs and CICS storage	401
Passing data to other programs	379	Chapter 27. Transient data control	403
COMMAREA	379	Introduction	403
INPUTMSG	381	Intrapartition queues	403
Using the INPUTMSG option on the RETURN command	383	Extrapartition destinations	404
Other ways of passing data	383	Indirect destinations	404
Mixed addressing mode transactions	383	Automatic transaction initiation (ATI)	405
Examples of passing data with the LINK command	384	Chapter 28. Temporary storage control	407
Examples of passing data with the RETURN command	387	Temporary storage queues	407
Chapter 26. Storage control and storage protection	391	Typical uses of temporary storage control	408
CICS storage control	391	Chapter 29. Security control	409
CICS storage protection	392	QUERY SECURITY command	409
CICS-key and user-key storage	392	Using QUERY SECURITY	409
Selecting the execution key for applications	393	Non-Terminal transaction security	410

Chapter 23. Interval control

The CICS interval control facility provides functions that are related to time.

Using interval control commands, you can:

- Start a task at a specified time or after a specified interval, and pass data to it (EXEC CICS START command)¹².
- Retrieve data passed on a START command (EXEC CICS RETRIEVE command)¹².
- Delay the processing of a task (EXEC CICS DELAY command).
- Request notification when a specified time has expired (EXEC CICS POST command).
- Wait for an event to occur (EXEC CICS WAIT EVENT command).
- Cancel the effect of previous interval control commands (EXEC CICS CANCEL command).
- Request the current date and time of day (EXEC CICS ASKTIME command).
- Select the format of date and time (EXEC CICS FORMATTIME command). Options are available that help you to handle dates in the twenty-first century. Programming information about these is in the *CICS Application Programming Reference* manual.

Note: On a lightly used system, the interval time specified can be exceeded by as much as a quarter of a second.

If you use EXEC CICS WAIT EVENT, EXEC CICS START, EXEC CICS RETRIEVE with the WAIT option, EXEC CICS CANCEL, EXEC CICS DELAY, or EXEC CICS POST commands, you could create intertransaction affinities that adversely affect your ability to perform dynamic transaction routing.

See Chapter 5, “Affinity among transactions” on page 91 for more information about transaction affinity.

Expiration times

The time at which a time-controlled function is to be started is known as the **expiration time**. You can specify expiration times absolutely, as a time of day (using the TIME option), or as an interval that is to elapse before the function is to be performed (using the INTERVAL option). For the EXEC CICS DELAY command, you can use the FOR and UNTIL options; and for the POST and START commands, you can use the AFTER and AT options. See the *CICS Application Programming Reference* manual for programming information about these commands.

Note: The C language does not provide the support for the packed decimal types used by the TIME and INTERVAL options.

¹² Do not use EXEC CICS START TRANSID() TERMID(EIBTRMID) to start a remote transaction. Use EXEC CICS RETURN TRANSID() IMMEDIATE instead. START, used in this way, ties up communications resources unnecessarily and can lead to performance degradation across the connected regions.

You use an **interval** to tell CICS when to start a transaction in a specified number of hours, minutes, and seconds from the current time. A nonzero INTERVAL value always indicates a time in the future—the current time plus the interval you specify. The hours may be 0–99, but the minutes and seconds must not be greater than 59. For example, to start a task in 40 hours and 10 minutes, you would code:

```
EXEC CICS START INTERVAL(401000)
```

You can use an **absolute time** to tell CICS to start a transaction at a specific time, again using hhmmss. For example, to start a transaction at 3:30 in the afternoon, you would code:

```
EXEC CICS START TIME(153000)
```

An absolute time is always relative to the midnight before the current time and may therefore be earlier than the current time. TIME may be in the future or the past relative to the time at which the command is executed. CICS uses the following rules:

- If you specify a task to start at any time within the previous six hours, it starts immediately. This happens regardless of whether the previous six hours includes a midnight. For example:

```
EXEC CICS START TIME(123000)
```

This command, issued at 05:00 or 07:00 on Monday, expires at 12:30 on the same day.

```
EXEC CICS START TIME(020000)
```

This command, issued at 05:00 or 07:00 on Monday expires immediately because the specified time is within the preceding six hours.

```
EXEC CICS START TIME(003000)
```

This command, issued at 05:00 on Monday, expires immediately because the specified time is within the preceding six hours. However, if it is issued at 07:00 on Monday, it expires at 00:30 on Tuesday, because the specified time is not within the preceding six hours.

```
EXEC CICS START TIME(230000)
```

This command, issued at 02:00 on Monday, expires immediately because the specified time is within the preceding six hours.

- If you specify a time with an hours component that is greater than 23, you are specifying a time on a day following the current one. For example, a time of 250000 means 1 a.m. on the day following the current one, and 490000 means 1 a.m. on the day after that.

If you do not specify an expiration time or interval option on the EXEC CICS DELAY, EXEC CICS POST, or EXEC CICS START command, CICS responds using the default of INTERVAL(0), which means immediately.

Because each end of an intersystem link may be in a different time zone, use the INTERVAL form of expiration time when the transaction to be started is in a remote system.

If the system fails, the times associated with unexpired START commands are remembered across the restart.

Note: If your expiration time falls within a possible CICS shutdown, consider whether your task should test the status of CICS before attempting to run. You can do this using the EXEC CICS INQUIRE SYSTEM CICSSTATUS command described in the *CICS System Programming Reference* manual. During a normal shutdown, your task could run at the same time as the PLT programs with consequences known only to you.

Request identifiers

As a means of identifying the request and any data associated with it, a unique request identifier is assigned by CICS to each EXEC CICS DELAY, EXEC CICS POST, and EXEC CICS START command. You can specify your own request identifier by means of the REQID option. If you do not, CICS assigns (for EXEC CICS POST and EXEC CICS START commands only) a unique request identifier and places it in field EIBREQID in the EXEC interface block (EIB). Specify a request identifier if you want the request to be canceled at some later time by an EXEC CICS CANCEL command.

Chapter 24. Task control

The CICS task control facility provides functions that synchronize task activity, or that control the use of resources.

CICS assigns priorities based on the value set by the CICS system programmer. Control of the processor is given to the highest-priority task that is ready to be processed, and is returned to the operating system when no further work can be done by CICS or by your application programs.

You can:

- Suspend a task (EXEC CICS SUSPEND command) to enable tasks of higher priority to proceed. This can prevent processor-intensive tasks from monopolizing the processor. When other eligible tasks have proceeded and terminated or suspended processing, control is returned to the issuing task; that is, the task remains dispatchable.
- Schedule the use of a resource by a task (EXEC CICS ENQ and DEQ commands). This is sometimes useful in protecting a resource from concurrent use by more than one task; that is, by making that resource serially reusable. Each task that is to use the resource issues an enqueue command (ENQ). The first task to do so has the use of the resource immediately but, if a HANDLE CONDITION ENQBUSY command has not been issued, subsequent ENQ commands for the resource, issued by other tasks, result in those tasks being suspended until the resource is available.

If the NOSUSPEND option is coded on the ENQ command, control is always returned to the next instruction in the program. By inspecting the contents of the EIBRESP field, you can see whether the ENQ command was successful or not.

Each task using a resource should issue a dequeue command (DEQ) when it has finished with it. However, when using the enqueue/dequeue mechanism, there is no way to guarantee that two or more tasks issuing ENQ and DEQ commands issue these commands in a given sequence relative to each other. For a way to control the sequence of access, see “Controlling sequence of access to resources” on page 374.

- Change the priority assigned to a task (EXEC CICS CHANGE TASK PRIORITY command).
- Wait for events that post VSE format ECBs when they complete.

Two commands are available: EXEC CICS WAITCICS and EXEC CICS WAIT EXTERNAL. These commands cause the issuing task to be suspended until one of the ECBs has been posted; that is, until one of the events has occurred. The task can wait on one or more ECBs. If it waits on more than one, it is dispatchable as soon as one of the ECBs is posted. You must ensure that each ECB is cleared (set to binary zeros) no later than the earliest time it could be posted. CICS cannot do this for you. If you wait on an ECB that has been previously posted and is not subsequently cleared, your task is not suspended and continues to run as though the WAITCICS or WAIT EXTERNAL command had not been issued.

In CICS Transaction Server for VSE/ESA Release 1, EXEC CICS WAITCICS and EXEC CICS WAIT EXTERNAL perform the same function and have the same restrictions.

However, care should be taken when choosing which command to use to avoid any future problems if the application is migrated to CICS Transaction Server for OS/390®.

The extra restrictions imposed by CICS Transaction Server for OS/390 are:

- No ECB in the ECBLIST for an EXEC CICS WAIT EXTERNAL should be hand-posted. If you are in any doubt about the method of posting, use the EXEC CICS WAITCICS command.
- A given ECB must not be waited on by more than one task at once (or appear more than once in one task's ECBLIST). Failure to follow this rule will lead to an INVREQ response to such an EXEC CICS WAIT EXTERNAL command.

WAITCICS should be used if ECBs are to be posted by any method other than the VSE POST facility or by an optimized post. For example, if your application posts the ECB by moving a value into it, WAITCICS should be used. (The WAITCICS command can also be used for ECBs that are posted using the VSE POST facility or optimized post.) A given ECB must not be waited on by more than one task. Whenever CICS goes into an VSE WAIT, it passes a list to VSE of all the ECBs being waited on by tasks that have issued a WAITCICS or WAIT EXTERNAL command. If the ECBLIST passed by CICS on the VSE WAIT contains duplicate addresses, VSE abends CICS.

If you use VSE POST, WAIT EXTERNAL, WAITCICS, ENQ, or DEQ commands, you could create intertransaction affinities that adversely affect your ability to perform dynamic transaction routing.

See Chapter 5, "Affinity among transactions" on page 91 for more information about transaction affinity.

CICS checks that the timer-event control area and the ECBs are not in read-only storage.

Controlling sequence of access to resources

If you want a resource to be accessed by two or more tasks in a specific order, instead of the ENQ and DEQ commands, use one or more WAITCICS commands in conjunction with one or more hand-posted ECBs.

To hand-post an ECB, a CICS task sets a 4-byte field to either the cleared state of binary zeros, or the posted state of X'40008000'. The task can use a START command to start another task and pass the address of the ECB. The started task receives the address through a RETRIEVE command.

Either task can set the ECB or wait on it. Use the ECB to control the sequence in which the tasks access resources. Two tasks can share more than one ECB if necessary. You can extend this technique to control as many tasks as you wish.

Note: Only one task can wait on a given ECB at any one time.

The example in Figure 65 on page 375 shows how two tasks can sequentially access a temporary storage queue by using hand-posted ECBs and the WAITCICS command.




Figure 65. Two tasks sequentially accessing a temporary storage queue

The example uses two ECBs, (ECB1 and ECB2), addressed by the pointers illustrated in Table 31.

In theory, these tasks could exchange data through the temporary storage queue for ever. In practice, some code would be included to close the process down in an orderly way.

Table 31. Example of task control	
Task A	Task B
Delete temporary storage queue Clear ECB1 (set to X'00000000') Clear ECB2 EXEC CICS START TASK B and pass the addresses of PTR_ECB1_ADDR_LIST and PTR_ECB2_ADDR_LIST.	EXEC CICS RETRIEVE the addresses passed.
START OF LOOP: EXEC CICS WAITCICS ECBLIST(PTR_ECB1_ADDR_LIST NUMEVENTS(1) Clear ECB1 Read T/S queue < act on data from queue > Delete T/S queue Write to T/S queue Post ECB2 Go to START OF LOOP	START OF LOOP: Write to T/S queue Post ECB1 (set to X'40008000') EXEC CICS WAITCICS ECBLIST(PTR_ECB2_ADDR_LIST NUMEVENTS(1) Clear ECB2 Read T/S queue < act on data from queue > Delete T/S queue Go to START OF LOOP

Chapter 9, “Dealing with exception conditions” on page 141 describes how the exception conditions occur during processing of a task control command are handled.

Chapter 25. Program control

The CICS program control facility governs the flow of control between application programs in a CICS system. The name of the application referred to in a program control command must have been defined as a program to CICS. You can use program control commands to:

- Link one of your application programs to another, anticipating subsequent return to the requesting program (LINK command). The COMMAREA, INPUTMSG, and INPUTMSGLEN options of this command allow data to be passed to the requested application program. For further information, see “COMMAREA” on page 379 and “INPUTMSG” on page 381.
- Link one of your application programs to another program in a separate CICS region, anticipating subsequent return to the requesting program (LINK command). The COMMAREA option of this command allows data to be passed to the requested application program. This is referred to as distributed program link (DPL). (You cannot use the INPUTMSG and INPUTMSGLEN options of the LINK command when using DPL. See the *CICS Application Programming Reference* manual for programming information, including details about this restriction.) For more information about DPL, see Chapter 6, “Intercommunication considerations” on page 119.
- Transfer control from one of your application programs to another, with no return to the requesting program (XCTL command). The COMMAREA, INPUTMSG, and INPUTMSGLEN options of this command allow data to be passed to the requested application program. (You cannot use the INPUTMSG and INPUTMSGLEN options of the XCTL command when using DPL. (See the *CICS Application Programming Reference* manual for programming information, including details about this restriction.)
- Return control from one of your application programs to another, or to CICS (RETURN command). The COMMAREA and INPUTMSG options of this command allow data to be passed to a newly initiated transaction. (You cannot use the INPUTMSG and INPUTMSGLEN options of the RETURN command when using DPL. (See the *CICS Application Programming Reference* manual for programming information, including details about this restriction.)
- Load a designated application program, table, or map into main storage (LOAD command).

If you use the HOLD option with the LOAD and RELEASE command to load a program, table or map that is not read-only, you could create intertransaction affinities that could adversely affect your ability to perform dynamic transaction routing.

See Chapter 5, “Affinity among transactions” on page 91 for more information about transaction affinity.

- Delete a previously loaded application program, table, or map from main storage (RELEASE command).

You can use the RESP option to deal with abnormal terminations.

Application program logical levels

Application programs running under CICS are executed at various logical levels. The first program to receive control within a task is at the highest logical level. When an application program is linked to another, expecting an eventual return of control, the linked-to program is considered to reside at the next lower logical level. When control is simply transferred from one application program to another, without expecting return of control, the two programs are considered to reside at the same logical level.

Link to another program expecting return

The LINK command is used to pass control from an application program at one logical level to an application program at the next lower logical level. If the program receiving control is not already in main storage, it is loaded. When the RETURN command is processed in the linked program, control is returned to the program initiating the link at the next sequential process instruction.

The linked program operates independently of the program that issues the LINK command with regard to handling exception conditions, attention identifiers, and abends. For example, the effects of HANDLE commands in the linking program are not inherited by the linked-to program, but the original HANDLE commands are restored on return to the linking program. You can use the HANDLE ABEND command to deal with abnormal terminations in other link levels. See the *CICS Application Programming Reference* manual for programming information about this command. Figure 66 on page 379 shows the concept of logical levels.




Figure 66. Application program logical levels

Passing data to other programs

You can pass data to another program when control is passed to that other program using a program control command.

COMMAREA

The COMMAREA option of the LINK and XCTL commands specifies the name of a data area (known as a **communication area**) in which data is passed to the program being invoked.

In a similar manner, the COMMAREA option of the RETURN command specifies the name of a communication area in which data is passed to the transaction identified in the TRANSID option. (The TRANSID option specifies a transaction that is initiated when the next input is received from the terminal associated with the task.) For programming information about the length of the communication area, see the *CICS Application Programming Reference* manual.

The invoked program receives the data as a parameter. The program must contain a definition of a data area to allow access to the passed data.

In a receiving COBOL program, you must give the data area the name DFHCOMMAREA. In this COBOL program, if a program passes a COMMAREA as part of a LINK, XCTL, or RETURN command, either the WORKING_STORAGE or the LINKAGE SECTION can contain the data area. A program receiving a COMMAREA should specify the data in the LINKAGE SECTION. This applies when the program is either of the following:

- The receiving program during a LINK or XCTL command where a COMMAREA is passed
- The initial program, where the RETURN command of a previously called task specified a COMMAREA and TRANSID

In a C program that is receiving a COMMAREA, the COMMAREA must be defined as a pointer to a structure. The program then must issue the ADDRESS COMMAREA command to gain addressability to the passed data.

In a PL/I program, the data area can have any name, but it must be declared as a based variable, based on the parameter passed to the program. The pointer to this based variable should be declared explicitly as a pointer rather than contextually by its appearance in the declaration for the area. This prevents the generation of a PL/I error message. No ALLOCATE statement can be processed within the receiving program for any variable based on this pointer. This pointer must not be updated by the application program.

In an assembler language program, the data area should be a DSECT. The register used to address this DSECT must be loaded from DFHEICAP, which is in the DFHEISTG DSECT.

The receiving data area need not be of the same length as the original communication area; if access is required only to the first part of the data, the new data area can be shorter. However, it must not be longer than the length of the communication area being passed. If it is, your transaction may inadvertently attempt to read data outside the area that has been passed. It may also overwrite data outside the area, which could cause CICS to abend.

To avoid this happening, your program should check whether the length of any communication area that has been passed to it is as expected, by accessing the EIBCALEN field in the EIB of the task. If no communication area has been passed, the value of EIBCALEN is zero; otherwise, EIBCALEN always contains the value specified in the LENGTH option of the LINK, XCTL, or RETURN command, regardless of the size of the data area in the invoked program. Ensure that the value in EIBCALEN matches the value in the DSECT for your program, and make sure that your transaction is accessing data within that area.

You may also add an identifier to COMMAREA as an additional check on the data that is being passed. This identifier is sent with the sending transaction and is checked for by the receiving transaction.

When a communication area is passed using an EXEC CICS LINK command, the invoked program is passed a pointer to the communication area itself. Any changes made to the contents of the data area in the invoked program are available to the invoking program, when control returns to it. To access any such

changes, the program names the data area specified in the original COMMAREA option.

When a communication area is passed using an XCTL command, a copy of that area is made unless the area to be passed has the same address and length as the area that was passed to the program issuing the command. For example, if program A issues a LINK command to program B, which in turn issues an XCTL command to program C, and if B passes to C the same communication area that A passed to B, program C will be passed addressability to the communication area that belongs to A (not a copy of it), and any changes made by C will be available to A when control returns to it.

When a lower-level program, which has been accessed by a LINK command, issues the RETURN command, control passes back one logical level higher than the program returning control. If the task is associated with a terminal, the TRANSID option can be used at the lower level to specify the transaction identifier for the next transaction to be associated with that terminal. The transaction identifier comes into play only after the highest logical level has relinquished control to CICS using the RETURN command and input is received from the terminal. Any input entered from the terminal, apart from an attention key, is interpreted wholly as data. You may use the TRANSID option without COMMAREA when returning from any link level, but it can be overridden on a later RETURN command. If a RETURN command fails at the top level because of an invalid COMMAREA, the TRANSID becomes null. Also, you can specify COMMAREA or IMMEDIATE only at the highest level, otherwise you will get an INVREQ with EIBRESP2=2.

In addition, the COMMAREA option can be used to pass data to the new task that is to be started.

The invoked program can determine which type of command invoked it by accessing field EIBFN in the EIB. This field must be tested before any CICS commands are issued. If the program was invoked by a LINK or XCTL command, the appropriate code is found in the EIBFN field. If it was invoked by a RETURN command, no CICS commands have been issued in the task, and the field contains zeros.

INPUTMSG

The INPUTMSG option of the EXEC CICS LINK, XCTL, and RETURN commands is another way of specifying the name of a data area to be passed to the program being invoked. For programming information about the use of these commands, see the *CICS Application Programming Reference* manual. In this case, the invoked program gets the data by processing a RECEIVE command. This option enables you to invoke (“front-end”) application programs that were written to be invoked directly from a terminal, and which contain RECEIVE commands, to obtain initial terminal input.

If a program that has been accessed by means of a LINK command issues a RECEIVE command to obtain initial input from a terminal, but the initial RECEIVE request has already been issued by a higher-level program, there is no data for the program to receive. In this case, the application waits on input from the terminal. You can ensure that the original terminal input continues to be available to a linked program by invoking it with the INPUTMSG option.

When an application program invokes another program, specifying INPUTMSG on the LINK (or XCTL or RETURN) command, the data specified on the INPUTMSG continues to be available even if the linked program itself does not issue an RECEIVE command, but instead invokes yet another application program. See Figure 67 on page 382 for an illustration of INPUTMSG.




Figure 67. Use of INPUTMSG in a linked chain

Notes:

1. In this example, the “real” first RECEIVE command is issued by program A. By linking to program B with the INPUTMSG option, it ensures that the next program to issue a RECEIVE request can also receive the terminal input. This can be either program B or program C.
2. If program A simply wants to pass on the unmodified terminal input that it received, it can name the same data area for the INPUTMSG option that it used for the RECEIVE command. For example:

```

EXEC CICS RECEIVE
      INTO(TERMINAL-INPUT)
      :
EXEC CICS LINK
      PROGRAM(PROGRAMB)
      INPUTMSG(TERMINAL-INPUT)
      :
  
```

3. As soon as one program in a LINK chain issues a RECEIVE command, the INPUTMSG data ceases to be available to any subsequent RECEIVE command. In other words, in the example shown, if B issues a RECEIVE request before linking to C, the INPUTMSG data area is not available for C.
4. This method of communicating data from one program to another can be used for any kind of data—it does not have to originate from a user terminal. In our

example, program A could move any data into the named data area, and invoke program B with INPUTMSG referencing the data.

5. The “data” passed on INPUTMSG also ceases to be available when control is eventually returned to the program that issued the link with INPUTMSG. In our example, if C returns to B, and B returns to A, and neither B nor C issues a RECEIVE command, the data is assumed by A to have been received. If A then invokes another program (for example, D), the original INPUTMSG data is no longer available to D, unless the INPUTMSG option is specified.
6. The INPUTMSG data ceases to be available when a SEND or CONVERSE command is issued.

Using the INPUTMSG option on the RETURN command

You can specify INPUTMSG to pass data to the next transaction specified on a RETURN command with the TRANSID option. To do this, the RETURN command must be issued at the highest logical level to return control to CICS, and the command must also specify the IMMEDIATE option.

EXEC CICS RETURN with INPUTMSG can be specified at any link-level, but be aware that it might be overridden by a later EXEC CICS RETURN command at a higher level.

If you specify INPUTMSG with TRANSID, and do not also specify IMMEDIATE, the next real input from the terminal overrides the INPUTMSG data, which is therefore lost. See the *CICS Application Programming Reference* manual for programming information about the RETURN command.

If you specify INPUTMSG with TRANSID some time after a SEND command, the SEND message is immediately flushed out to the terminal.

The other use for INPUTMSG, on a RETURN command without the TRANSID option, is intended for use with a dynamic transaction routing program. See the *CICS Customization Guide* for programming information about the user-replaceable dynamic transaction routing program.

Other ways of passing data

Data can also be passed between application programs and transactions in other ways. For example, the data can be stored in a CICS storage area outside the local environment of the application program, such as the transaction work area (TWA). Another way is to store the data in temporary storage; see Chapter 28, “Temporary storage control” on page 407 for details.

Mixed addressing mode transactions

CICS supports the use of the LINK, XCTL, and RETURN commands between programs with different addressing modes and between programs with the same addressing mode.

The following restrictions apply to programs passing data using a communication area named by the COMMAREA option:

- Addresses passed within a communication area to an AMODE(31) program must be 31 bits long. Do not use 3-byte addresses with flag data packed into

the top byte, unless the called program is specifically designed to ignore the top byte.

- Addresses passed as data to an AMODE(24) program must be below the 16MB line if they are to be interpreted correctly by the called program.

These restrictions apply to the address of the communication area itself, and also to addresses within it. However, a COMMAREA above the 16MB line can be passed to an AMODE(24) subprogram. CICS copies the COMMAREA into an area below the 16MB line for processing. It copies it back again when control returns to the linking program. See Chapter 26, “Storage control and storage protection” on page 391 for information about copying CICS-key storage.

CICS does not validate any data addresses passed within a communication area between programs with different addressing modes.

Examples of passing data with the LINK command

Figures 68 to 71 show how, in COBOL, C, PL/I, and assembler language, the LINK command causes data to be passed to the program being linked to; the XCTL command is coded in a similar way.

```

                                Invoking program

IDENTIFICATION DIVISION.
PROGRAM ID. 'PROG1'.
:
WORKING-STORAGE SECTION.
01 COM-REGION.
    02 FIELD PICTURE X(3).
:
PROCEDURE DIVISION.
    MOVE 'ABC' TO FIELD.
    EXEC CICS LINK PROGRAM('PROG2')
        COMMAREA(COM-REGION)
        LENGTH(3) END-EXEC.
:

                                Invoked program

IDENTIFICATION DIVISION.
PROGRAM-ID. 'PROG2'.
:
LINKAGE SECTION.
01 DFHCOMMAREA.
    02 FIELD PICTURE X(3).
:
PROCEDURE DIVISION.
    IF EIBCALEN GREATER ZERO
    THEN
        IF FIELD EQUALS 'ABC' ...
```

Figure 68. COBOL example—LINK command

```

                                Invoking program

main()
{
    unsigned char field[3];
    memcpy(field, "ABC", 3);
    EXEC CICS LINK PROGRAM("PROG2") COMMAREA(field)
                LENGTH(sizeof(field));
}

                                Invoked program

main()
{
    unsigned char *commarea;
    EXEC CICS ADDRESS COMMAREA(commarea) EIB(dfheiptr);
    if (dfheiptr->eibcalen > 0)
    {
        if (memcmp(commarea, "ABC", 3) == 0)
        {
            :
        }
    }
}

```

Figure 69. C example—LINK command. In this example, the COMMAREA contains a character string. For an example of a COMMAREA that contains a structure, see Figure 73 on page 388.

```

                                Invoking program

PROG1: PROC OPTIONS(MAIN);
DCL 1 COM_REGION AUTOMATIC,
      2 FIELD CHAR(3),
:
FIELD='ABC';
EXEC CICS LINK PROGRAM('PROG2') COMMAREA(COM_REGION) LENGTH(3)
END;

                                Invoked program

PROG2:
PROC(COMM_REG_PTR) OPTIONS(MAIN);
DCL COMM_REG_PTR PTR;
DCL 1 COM_REGION BASED(COMM_REG_PTR),
      2 FIELD CHAR(3),
:
IF EIBCALEN>0 THEN DO;
    IF FIELD='ABC' THEN ...
:
    END;
:
END;

```

Figure 70. PL/I example—LINK command

```

                                Invoking program
DFHEISTG DSECT
COMREG  DS 0CL20
FIELD   DS CL3
:
:
PROG1   CSECT
:
:
:           MVC FIELD,=C'XYZ'
:           EXEC CICS LINK
:           PROGRAM('PROG2')
:           COMMAREA(COMREG) LENGTH(3)
:
:           END
                                Invoked program
COMREG  DSECT
FIELD   DS CL3
:
:
PROG2   CSECT
:
:
:           L COMPTR,DFHEICAP
:           USING COMREG,COMPTR
:           CLC FIELD,=C'ABC'
:
:           END

```

Figure 71. ASM example—LINK command

Examples of passing data with the RETURN command

Figures 72 to 75 show how, in COBOL, C, PL/I, and assembler language, the RETURN command is used to pass data to a new transaction.

```

                                Invoking program

IDENTIFICATION DIVISION.
PROGRAM-ID. 'PROG1'.
:
WORKING-STORAGE SECTION.
01  TERMINAL-STORAGE.
    02  FIELD PICTURE X(3).
    02  DATAFLD PICTURE X(17).
:
PROCEDURE DIVISION.
    MOVE 'XYZ' TO FIELD.
    EXEC CICS RETURN TRANSID('TRN2')
        COMMAREA(TERMINAL-STORAGE)
        LENGTH(20) END-EXEC.
:
                                Invoked program

IDENTIFICATION DIVISION.
PROGRAM-ID. 'PROG2'
:
LINKAGE SECTION.
01  DFHCOMMAREA.
    02  FIELD PICTURE X(3).
    02  DATAFLD PICTURE X(17).
:
PROCEDURE DIVISION.
    IF EIBCALEN GREATER ZERO
    THEN
        IF FIELD EQUALS 'XYZ'
        MOVE 'ABC' TO FIELD.
    EXEC CICS RETURN END-EXEC.
```

Figure 72. COBOL example—RETURN command

```

                                Invoking program

struct ter_struct
{
    unsigned char field[3];
    unsigned char datafld[17];
};

main()
{
    struct ter_struct ter_stor;

    memcpy(ter_stor.field,"XYZ",3);
    EXEC CICS RETURN TRANSID("TRN2")
           COMMAREA(&ter_stor)
           LENGTH(sizeof(ter_stor));
}

                                Invoked program

struct term_struct
{
    unsigned char field[3];
    unsigned char datafld[17];
};

main()
{
    struct term_struct *commarea;

    EXEC CICS ADDRESS COMMAREA(commarea) EIB(dfheiptr);
    if (dfheiptr->eibcalen > 0)
    {
        if (memcmp(commarea->field, "XYZ", 3) == 0)
            memcpy(commarea->field, "ABC", 3);
    }
    EXEC CICS RETURN;
}

```

Figure 73. C example—RETURN command

```

                                Invoking program

PROG1: PROC OPTIONS(MAIN);
DCL 1 TERM_STORAGE,
      2 FIELD CHAR(3),
:
FIELD='XYZ';
EXEC CICS RETURN TRANSID('TRN2')
      COMMAREA(TERM_STORAGE);
END;

                                Invoked program

PROG2:
PROC(TERM_STG_PTR) OPTIONS(MAIN);
DCL TERM_STG_PTR PTR;
DCL 1 TERM_STORAGE
      BASED(TERM_STG_PTR),
      2 FIELD CHAR(3),
:
IF EIBCALEN>0 THEN DO;
  IF FIELD='XYZ' THEN FIELD='ABC';
END;
EXEC CICS RETURN;
END;

```

Figure 74. PL/I example—RETURN command

```

                                Invoking program

DFHEISTG DSECT
TERMSTG  DS 0CL20
FIELD    DS CL3
DATAFLD  DS CL17
:
PROG1    CSECT
:
          MVC FIELD,=C'ABC'
          EXEC CICS RETURN
          TRANSID('TRN2')
          COMMAREA(TERMSTG)
:
          END

                                Invoked program

TERMSTG  DSECT
FIELD    DS CL3
DATAFLD  DS CL17
:
PROG2    CSECT
:
          CLC EIBCALEN,=H'0'
          BNH LABEL2
          L  COMPTR,DFHEICAP
          USING TERMSTG,COMPTR
          CLC FIELD,=C'XYZ'
          BNE LABEL1
          MVC FIELD,=C'ABC'
LABEL1   DS 0H
:
LABEL2   DS 0H
:
          END

```

Figure 75. ASM example—RETURN command

Chapter 26. Storage control and storage protection

This chapter describes storage control and storage protection, and contains the following topics:

- “CICS storage control”
- “CICS storage protection” on page 392
- “CICS-key and user-key storage” on page 392
- “Deciding what execution and storage key to specify” on page 397
- “C programs and CICS storage” on page 401

CICS storage control

The CICS storage control facility controls requests for main storage to provide intermediate work areas and other main storage not provided automatically by CICS but needed to process a transaction.

CICS makes working storage available within each command-level program automatically, without any specific request from the application program, and provides other facilities for intermediate storage, both within and among tasks. Chapter 3, “Designing efficient applications” on page 51 describes storage within individual programs. If you need working storage in addition to the working storage provided automatically by CICS, however, you can use the following commands:

- EXEC CICS GETMAIN to get and initialize main storage
- EXEC CICS FREEMAIN to release main storage

You can initialize the acquired main storage to any bit configuration by supplying the INITIMG option on the GETMAIN command; for example, zeros or EBCDIC blanks.

CICS releases all main storage associated with a task when the task is ended normally or abnormally. This includes any storage acquired, and not subsequently released, by your application program, except for areas obtained with the SHARED option. This option of the GETMAIN command prevents storage being released automatically when a task completes.

If you use the GETMAIN command with the SHARED option, and the FREEMAIN command, you could create intertransaction affinities that adversely affect the ability to perform dynamic transaction routing.

See Chapter 5, “Affinity among transactions” on page 91 for information about transaction affinity.

If there is no storage available when you issue your request, CICS suspends your task until space is available, unless you specify the NOSUSPEND option. While the task is suspended, it may be canceled (timed out) if the RDO TRANSACTION resource definition specifies SPURGE(YES) and DTIMOUT(mmss). NOSUSPEND returns control to your program if storage is not available, allowing you to do alternative processing, as appropriate.

CICS storage protection

Storage protection protects CICS code and control blocks from applications.

CICS Transaction Server for VSE/ESA Release 1 exploits the ESA/390™ subsystem storage protection facility in a way that enables you to prevent CICS code and control blocks from being overwritten accidentally by your application programs. It does **not** provide protection against deliberate overwriting of CICS code or control blocks. CICS cannot prevent an application obtaining the necessary access (execution key) to modify CICS storage.

The use of storage protection is optional. You choose whether you want to use storage protection facilities by means of CICS system initialization parameters described in the *CICS System Definition Guide*.

CICS-key and user-key storage

When you are running with the storage protection facility active, CICS Transaction Server for VSE/ESA Release 1 separates storage into two categories:

CICS-key storage

is used for CICS system code and control blocks and, at the discretion of the installation, other code and data areas that require protection from overwriting.

User-key storage

is where application programs and their data areas normally reside.

There are two associated execution modes:

- CICS system programs run in **CICS key**. CICS-key execution allows a program read-write access to both CICS-key and user-key storage.
- Application programs normally execute in **user key**. User-key execution allows a program read-write access to user-key storage, but only read access to CICS-key storage.

The terms “user key” and “CICS key” thus apply both to storage and to the execution of programs with respect to that storage. They are reflected in the resource definition keywords on the RDO TRANSACTION and PROGRAM resource definitions. These keywords are described in the *CICS Resource Definition Guide*.

CICS allows you to run your application programs in either user-key or CICS-key storage. CICS storage is automatically protected from being overwritten by application programs that execute in user-key storage (the default). The concept of isolating CICS code and control blocks (CICS internal data areas) from user application programs is illustrated in Figure 76 on page 393.




Figure 76. Protecting CICS code and control blocks from user application programs

Selecting the execution key for applications

The execution key controls the type of access your application programs have to CICS-key storage. The default is that application programs are given control in user key. Define CICS key only for those programs where it is essential that they execute in CICS key. The programs you might select to run in CICS key are typically those that are written by system programmers, and are usually designed to provide special function in support of user applications. Such programs are generally considered to be an extension of CICS rather than part of an application. Some examples of such programs are described in “CICS-key applications” on page 398.

The storage protection facility does not protect CICS code and control blocks from being overwritten by this type of program, or by ordinary application programs that you choose to execute in CICS key.

Defining the execution key

To run your programs in CICS key, use the execution key keyword (EXECKEY) on the RDO PROGRAM resource definition. See “Deciding what execution and storage key to specify” on page 397 for an explanation of EXECKEY. The EXECKEY parameter determines the key in which CICS passes control to an application program.

Selecting and defining the storage key for applications

CICS enables you to choose between user-key storage and CICS-key storage for a number of CICS data areas and application program data areas that your applications can use. Depending on the data area, you select the storage key by:

- System initialization parameters
- TASKDATAKEY keyword on RDO TRANSACTION resource definition.
- Selecting an option on the EXEC CICS GETMAIN command

Defining the storage key for storage areas that your applications need to access is described in the following sections.

System-wide storage areas

For each CICS region, your installation can choose between user-key and CICS-key storage for the common work area (CWA) and for the terminal control table user areas (TCTUAs). If these areas are in user-key storage, all programs have read-write access to them; if they are in CICS-key storage, user-key application programs are restricted to read-only access. The storage keys for the CWA and the TCTUAs are set by the system initialization parameters CWAKEY and TCTUAKEY, respectively. In both cases the default option is that CICS obtains user-key storage.

See the *CICS System Definition Guide* for information about how to specify these and other storage-protection related system initialization parameters.

Task lifetime storage

You can also specify whether user-key or CICS-key storage is used for the storage that CICS acquires at transaction attach time, and for those elements of storage directly related to the individual application programs in a transaction. You do this by means of the TASKDATAKEY option on the RDO TRANSACTION resource definition. See “Deciding what execution and storage key to specify” on page 397 for an explanation about the TASKDATAKEY option. This governs the type of storage allocated for the following storage areas:

- The transaction work area (TWA) and the EXEC interface block (EIB)
- The copies of working storage that CICS obtains for each execution of an application program
- Any storage obtained for an application program in response to:
 - Explicit storage requests by means of an EXEC CICS GETMAIN command which do not specify either USERDATAKEY or CICSDATAKEY options.
 - Implicit storage requests as a result of an EXEC CICS command that uses the SET option

For information about how to specify the TASKDATAKEY parameter, see the *CICS Resource Definition Guide*.

Figure 77 on page 396 shows the TASKDATAKEY controls for both task lifetime storage and program working storage.

See the *CICS Application Programming Reference* manual for programming information about EXEC CICS commands; see the *CICS Resource Definition Guide* for information about specifying the TASKDATAKEY option on the transaction resource definition.

Program working storage specifically for exit and PLT programs

CICS uses the TASKDATAKEY option of the calling transaction to determine the storage key for the storage acquired for global user exits, task-related user exits, user-replaceable modules, and PLT programs. For programming information about storage key, including details of how this affects the different types of program, see the *CICS Customization Guide*.

Passing data via a COMMAREA

In a pseudoconversational application, CICS ensures that a COMMAREA you specify on a RETURN command is always accessible in read-write mode to the next program in the conversation. The same is true when passing a COMMAREA within a transaction that comprises more than one program (using a LINK or XCTL command). CICS ensures that the target program has read-write access to the COMMAREA.

The GETMAIN command

The GETMAIN command provides USERDATAKEY and CICS DATAKEY options to enable the application program to explicitly request user-key or CICS-key storage, regardless of the TASKDATAKEY option specified on the associated transaction resource definition. For example, this option allows application programs, which are executing with TASKDATAKEY(CICS) specified, to obtain user-key storage for passing to, or returning to, a program executing in user key.

CICS-key storage obtained by GETMAIN commands issued in a program defined with EXECKEY(CICS) can be freed explicitly only if the FREEMAIN command is issued by a program defined with EXECKEY(CICS). If an application program defined with EXECKEY(USER) attempts to free CICS-key storage using FREEMAIN commands, CICS returns the INVREQ condition. However, an application can free user-key storage with FREEMAIN commands regardless of the EXECKEY setting.

All task lifetime storage acquired by an application, whether in CICS key or user key, is freed by CICS at task termination. You can also specify STORAGECLEAR(YES) on the associated transaction resource definition. This clears the storage and so prevents another task accidentally viewing sensitive data.

For programming information about commands, see the *CICS Application Programming Reference* manual; for information about defining commands, see the *CICS Resource Definition Guide*.




Figure 77. Illustration of the use of the TASKDATAKEY and EXECCKEY options

Notes:

1. The TASKDATAKEY option ensures the TWA and EIB are allocated from user-key storage, required for PROGRAM1, which executes in user key—specified by EXECCKEY(USER).
2. PROGRAM1 executes in user key (controlled by EXECCKEY), and has its working storage obtained in user-key storage (controlled by the TASKDATAKEY option). Any other storage the program obtains by means of GETMAIN commands or by using the SET option on a CICS command is also obtained in user-key storage.
3. PROGRAM2 executes in CICS key (controlled by EXECCKEY), but has its working storage obtained in user-key storage, which again is controlled by the TASKDATAKEY option.
4. PROGRAM2 issues an explicit GETMAIN command using the CICS DATAKEY option and, because it executes in CICS key, can store data into the CICS-key protected storage before returning control to PROGRAM1.
5. PROGRAM1 cannot write to the CICS-key protected storage that PROGRAM2 acquired, but can read what PROGRAM2 wrote there.

When deciding whether you need to specify EXECCKEY(CICS) and TASKDATAKEY(CICS), you must consider all the reasons that make these options necessary.

Programs that modify their storage protection key should ensure they are running in the correct key when attempting to access storage. CICS can only use the EXECCKEY defined in the RDO PROGRAM resource definition when invoking a program.

Deciding what execution and storage key to specify

When you are running CICS with storage protection, the majority of your application programs should execute in user key, with all their storage obtained in user key. You only need to define EXECKEY(CICS) on RDO PROGRAM resource definitions, and TASKDATAKEY(CICS) on the associated RDO TRANSACTION resource definitions, for those programs that use facilities that are not permitted in user key, or for any special “system-type” transactions or vendor packages.

Specify TASKDATAKEY(CICS) for only those transactions where all the component programs have EXECKEY(CICS), and for which you want to protect their task lifetime and working storage from being overwritten by user-key applications. For example, the CICS-supplied transactions such as CEDF are defined with TASKDATAKEY(CICS).

Note that you cannot specify EXECKEY(USER) on any programs that form part of a transaction defined with TASKDATAKEY(CICS) because, in this situation, a user-key program would not be able to write to its own working storage. Transactions abend with an AEZD abend if any program is defined with EXECKEY(USER) within a transaction defined with TASKDATAKEY(CICS), regardless of whether storage protection is active.

You cannot define a program so that it inherits its caller’s execution key. The execution key and data storage keys are derived for each program from its program and associated transaction resource definitions respectively, which you either specify explicitly or allow to default; the default is always user key. Table 32 summarizes the various combinations of options.

EXECKEY	TASKDATAKEY	Recommended usage and comments
USER	USER	For normal applications using the CICS API
USER	CICS	Not permitted. CICS abends any program defined with EXECKEY(USER) invoked under a transaction defined with TASKDATAKEY(CICS).
CICS	USER	For programs that need to issue restricted VSE requests or modify CICS-key storage.
CICS	CICS	For transactions (and component programs) that function as extensions to CICS, such as the CICS-supplied transactions, or which require the same protection.

User-key applications

For most applications, define your programs with EXECKEY(USER), and the related transactions with TASKDATAKEY(USER). To obtain the maximum benefits from the CICS storage protection facility, you are recommended to run your application programs in user key storage. Specifying USER on these options has the following effect:

EXECKEY(USER)

This specifies that CICS is to give control to the program in user key when it is invoked. Programs defined with EXECKEY(USER) are restricted to read-only access to CICS-key storage. These include:

- Storage belonging to CICS itself
- CICS-key storage belonging to user transactions defined with TASKDATAKEY(CICS)
- Application programs defined with EXECKEY(CICS) and thus loaded into CICS-key storage

TASKDATAKEY(USER)

This specifies that all task lifetime storage, such as the transaction work area (TWA) and the EXEC interface block (EIB), is obtained from the user-key storage.

It also means that all storage directly related to the programs within the transaction is obtained from user-key storage.

USER is the default for both the EXECKEY and TASKDATAKEY options, therefore you do not need to make any changes to resource definitions for existing application programs.

CICS-key applications

Most application programs can be defined with EXECKEY(USER), which is the default value, and this is the option you are recommended to use in the majority of cases. These include programs that use DL/I or DB2 and programs that access vendor products through the resource manager interface (RMI) or a LINK command.

However, some application programs need to be defined with EXECKEY(CICS) because they need to use certain facilities that are listed later. Widespread use of EXECKEY(CICS) diminishes the protection offered by the storage protection facility because there is no protection of CICS code and control blocks from being overwritten by application programs that execute in CICS key.

Any application program causing a protection exception when defined with EXECKEY(USER) must be examined to determine why it is attempting to modify storage it is not allowed to modify. Change a program's definition to EXECKEY(CICS) only if you are satisfied that the application program legitimately requires the facilities described below.

- The program uses VSE macros or services directly, rather than through the CICS API. The only VSE macros that are supported in user-key programs are STXIT, POST, WAIT, WTO, and WTOR. If a program uses any other VSE macro or service, it must be defined with EXECKEY(CICS). Some particular examples are:
 - Use of VSE GETVIS and FREEVIS requests
 - Use of VSE OPEN, CLOSE, or other file access requests

Direct use of some VSE macros and services is undesirable, even in a CICS application defined with EXECKEY(CICS). This is because they may cause VSE to suspend the whole CICS region until the request is satisfied.

Some COBOL, PL/I, and C language statements, and compiler options, cause operating system functions to be invoked. See Chapter 2, “Language considerations” on page 19 for information about which of these should not be used in CICS application programs. It is possible that some of these functions may have worked in previous releases of CICS, or at least may not have caused the application to fail. They **do not work** when the program is defined with EXECKEY(USER). When the use of prohibited options or statements is the cause of a protection exception, remove these from the program rather than simply redefine the program with EXECKEY(CICS). The use of prohibited statements and options can have other side effects on the overall execution of CICS, and these should be removed.

- The program needs to modify the CWA, and the CWA is in CICS-key storage (CWAKEY=CICS system initialization parameter).

If you decide to protect the CWA by specifying CWAKEY(CICS), restrict the programs that are permitted to modify the CWA to as few as possible, perhaps only one. See “Common work area (CWA)” on page 83 for information about how you can control access to a protected CWA.

- The program needs to modify the TCTUA, and the TCTUAs are in CICS-key storage (TCTUAKEY=CICS system initialization parameter).

See “TCTTE user area (TCTUA)” on page 87 for information about using TCTUAs in a storage protection environment.

- The program can be invoked from PLT programs, from transactions defined with TASKDATAKEY(CICS), from task-related or global user exits programs, or from user-replaceable programs.
- The program modifies CICS control blocks—for example, some vendor products that do need to manipulate CICS control blocks. These must be defined with EXECKEY(CICS).
- The program provides user extensions to CICS and requires protection and data access similar to CICS system code. For example, you may consider that such programs are a vital part of your CICS installation, and that their associated storage, like CICS storage, should be protected from ordinary application programs.
- CICS always gives control in CICS key to the following types of user-written program, regardless of the option specified on their program resource definitions:
 - Global user exits (GLUEs)
 - Task-related user exits (TRUEs)
 - User-replaceable modules (URMs)
 - Program list table (PLT) programs.

CICS ensures that when control is passed to a PLT program, a global or task-related user exit, or a user-replaceable program, the first program so invoked executes in CICS key, regardless of the EXECKEY specified on its program resource definition. However, if this first program LINKs or XCTLs to other programs, these programs execute under the key specified in their program definitions. If these subsequent programs are required to write to CICS-key data areas, as often occurs in this type of situation, they must be defined as EXECKEY(CICS).

For programming information about the execution of GLUEs, TRUEs, URM, and PLT programs in a CICS region running with storage protection, see the *CICS Customization Guide*.

Tables

In addition to executable programs, you can define tables, map sets, and partition sets as program resources. EXECCKEY has less relevance to these objects, because they are not actually executed. However, EXECCKEY does control where nonexecutable objects are loaded, and thus affects whether other programs can store into them.

Map sets and partition sets

Map sets are not reentrant (BMS itself updates fields in maps when calculating absolute screen positions). However, map sets should not be modified by application programs; they must be modified only by CICS, which always executes in CICS key. CICS always loads map sets and partition sets into CICS-key storage.

Storage protection exception conditions

If an application program executing in user key attempts to modify CICS-key storage, a protection exception occurs. The protection exception is processed by normal CICS program error handling, and the offending transaction abends with an ASRA abend. The exception condition appears to the transaction just as if it had attempted to reference any other protected storage. CICS error handling checks whether the reference is to a CICS-key dynamic storage area (DSA), and sends a message (DFHSR0622) to the console. Otherwise, CICS does not treat the failure any differently from any other ASRA abend. See the *CICS Problem Determination Guide* for more information about the storage protection exception conditions.

Protecting CICS from being passed invalid addresses

CICS protects itself against applications that pass invalid addresses that would result in CICS causing storage violations. This occurs when an application program issues an EXEC CICS command that causes CICS to modify storage on the program's behalf, but the program does not own the storage. In earlier releases, CICS did not check ownership of the storage referenced by the passed address, and executed such commands with consequent storage violations.

CICS validates the start address of the storage, and establishes that the application program has write access to the storage that begins with that address, before executing the command.

This address checking is controlled using the CMDPROT system initialization parameter. If a program passes an invalid address to CICS as an output field on an EXEC CICS command, an AEYD abend occurs. It is completely independent of storage protection.

C programs and CICS storage

A C program running under CICS acquires all storage from the dynamic storage area (DSA).

A C program can acquire or release storage from or to CICS, either explicitly or implicitly.

Storage is acquired and released explicitly:

- By invoking the C library functions `calloc`, `malloc`, or `realloc` to acquire storage for your program.
- By invoking the C library function `free`, to release storage for your program.
- By issuing an EXEC CICS GETMAIN command to acquire storage directly from CICS.
- By issuing an EXEC CICS FREEMAIN command to return storage directly to CICS.

Storage that is explicitly acquired by C is placed in the heap area and exists until either it is explicitly freed or the program is terminated.

Storage that is acquired using the EXEC CICS GETMAIN command exists for the duration of the CICS task.

When you use the C library functions to acquire storage, the C library returns a pointer to storage acquired from CICS. If you request more storage than is currently available, the task is suspended until adequate storage becomes available. If you request less than one byte of storage, or more than the maximum allowed by CICS for the area requested (that is, DSA) the `calloc`, `malloc`, or `realloc` function returns a null pointer.

Acquiring and releasing storage through C library functions follows normal C rules.

Storage is acquired and released implicitly by the C run-time environment. When your program invokes a C function, C might require storage to perform the requested function. In such cases, the storage is acquired and released by the run-time environment on behalf of your program.

Chapter 27. Transient data control

This chapter describes the CICS transient data control facility, and contains the following topics:

- “Introduction”
- “Intrapartition queues”
- “Extrapartition destinations” on page 404
- “Indirect destinations” on page 404
- “Automatic transaction initiation (ATI)” on page 405

Introduction

The CICS transient data control facility provides a generalized queuing facility. Data can be queued (stored) for subsequent internal or external processing. Selected data, specified in the application program, can be routed to or from predefined symbolic destinations: either **intrapartition** or **extrapartition**.

Transient data queues are intrapartition if they are associated with a facility allocated to the CICS partition, and extrapartition if the data is directed to a destination that is external to the CICS partition. Transient data queues must be defined and installed before first reference by an application program. The destinations must be defined in the destination control table (DCT) by the system programmer when the CICS system is generated.

You can:

- Write data to a transient data queue (EXEC CICS WRITEQ TD command)
- Read data from a transient data queue (EXEC CICS READQ TD command)
- Delete an intrapartition transient data queue (EXEC CICS DELETEQ TD command)

If the TD keyword is omitted, the command is assumed to be for temporary storage. (See Chapter 28, “Temporary storage control” on page 407.)

Intrapartition queues

“Intrapartition” refers to data on direct-access storage devices for use with one or more programs running as separate tasks. Data directed to or from these internal destinations is referred to as intrapartition data; it must consist of variable-length records. Intrapartition destinations can be associated with either a terminal or an output data set. Intrapartition data may ultimately be transmitted upon request to the destination terminal or retrieved sequentially from the output data set.

Typical uses of intrapartition data include:

- Message switching
- Broadcasting
- Database access
- Routing of output to several terminals (for example, for order distribution)

- Queuing of data (for example, for assignment of order numbers or priority by arrival)
- Data collection (for example, for batched input from 2780 Data Transmission Terminals)

There are three types of intrapartition transient data queue:

- **Nonrecoverable**

Nonrecoverable intrapartition transient data queues are recovered only on a warm start of CICS. If a logical unit of work (LUW) updates a nonrecoverable intrapartition queue and subsequently backs out the updates, the updates made to the queue are **not** backed out.

- **Physically recoverable**

Physically recoverable intrapartition transient data queues are recovered on warm and emergency restarts. If a LUW updates a physically recoverable intrapartition queue and subsequently backs out the updates, the updates made to the queue are **not** backed out.

- **Logically recoverable**

Logically recoverable intrapartition transient data queues are recovered after individual transaction failures and on warm and emergency restarts, to the status they had at the end of the last completed LUW.

Extrapartition destinations

Extrapartition queues (data sets) reside on any sequential device (DASD, tape, printer, and so on) that are accessible by programs outside (or within) the CICS region. In general, sequential extrapartition destinations are used for storing and retrieving data outside the CICS region. For example, one task may read data from a remote terminal, edit the data, and write the results to a data set for subsequent processing in another region. Logging data, statistics, and transaction error messages are examples of data that can be written to extrapartition destinations. In general, extrapartition data created by CICS is intended for subsequent batched input to non-CICS programs. Data can also be routed to an output device such as a printer.

Data directed to or from an external destination is referred to as extrapartition data and consists of sequential records that are fixed-length or variable-length, blocked or unblocked. The record format for an extrapartition queue must be defined in the DCT by the system programmer. (See the *CICS Resource Definition Guide* for details.)

Note that you cannot delete an extrapartition queue. If you try to do so, you cause an INVREQ condition.

Indirect destinations

Intrapartition and extrapartition queues can be used as indirect queues. Indirect queues provide some flexibility in program maintenance in that data can be routed to one of several destinations with only the DCT, not the program itself, having to be changed.

When the DCT has been changed, application programs continue to route data to the queue using the original symbolic name; however, this name is now an indirect queue that refers to the original symbolic name. Because indirect queues are established by means of DCT table entries, the application programmer does not usually have to be concerned with how this is done. Further information about transient data resource definition is in the *CICS Resource Definition Guide*.

Automatic transaction initiation (ATI)

For intrapartition queues, CICS provides the option of automatic transaction initiation (ATI).

A basis for ATI is established by the system programmer by specifying a nonzero trigger level for a particular intrapartition queue in the DCT. See the discussion of the DFHDCT TYPE=INTRA macro in the *CICS Resource Definition Guide* for more information about trigger levels. When the number of entries (created by WRITEQ TD commands issued by one or more programs) in the queue reaches the specified trigger level, a task specified in the definition of the queue is automatically initiated. Control is passed to a program that processes the data in the queue; the program must issue repetitive READQ TD commands to deplete the queue.

When the queue has been emptied, a new ATI cycle begins. That is, a new task is scheduled for initiation when the specified trigger level is again reached, whether or not execution of the earlier task has ended.

If an automatically initiated task does not empty the queue, access to the queue is not inhibited.

The task may be normally or abnormally ended before the queue is emptied (that is, before a QZERO condition occurs in response to a READQ TD command).

If the contents of the queue are to be sent to a terminal (DESTFAC=TERMINAL on a DFHDCT TYPE=INTRA macro), the fact that QZERO has not been reached means that trigger processing has not been reset and the same task will be reinitiated later. Without the reset of trigger processing, a subsequent WRITEQ TD command will not trigger a new task.

If the contents of the queue are to be sent to a file (DESTFAC=FILE on a DFHDCT TYPE=INTRA macro), the task termination has the same effect as QZERO (that is, trigger processing is reset) and the next WRITEQ TD command will initiate the trigger transaction (if the trigger level has been reached).

If the trigger level of a queue is zero, no task is automatically initiated.

If a queue is logically recoverable, initiation of the trigger transaction is deferred until the next syncpoint.

The trigger level may already have been exceeded because the last triggered transaction abended before clearing the queue, or because the transaction was never started because the MXT limit was reached. For DESTFAC=TERMINAL, as described above, another task will not be scheduled. This is because QZERO has not been raised to reset trigger processing. The task that has already been scheduled is reinitiated as soon as possible.

For DESTFAC=FILE, the termination of the task will have reset trigger processing and so the next WRITEQ TD command will trigger a new task.

To ensure that an automatically initiated task completes when the queue is empty, the application program should test for a QZERO condition in preference to some other application-dependent factor (such as an anticipated number of records). Only the QZERO condition indicates an emptied queue.

DESTFAC=(SYSTEM, sysid) is used for starting up the front end of a DTP transaction. The transaction is started in the local system on an automatically allocated session on the connection given by the sysid. the transaction is started in the local system on an automatically allocated session on the connection given by the system identifier. The session name is then held in EIBTERMID. If a transaction (started with a queue of system) abends, a new transaction is started in the same way as a terminal.

As with a terminal destination, the QZERO condition must be reached by an automatically initiated task for trigger processing to be reset. If the trigger processing is not reset, the same task is reinitiated regardless of trigger level.

If you use ATI with a transient data trigger mechanism, it could create intertransaction affinities that adversely affect your ability to perform dynamic transaction routing. See Chapter 5, "Affinity among transactions" on page 91 for more information about transaction affinity.

Chapter 28. Temporary storage control

The CICS temporary storage control facility provides the application programmer with the ability to store data in temporary storage queues, either in main storage, or in auxiliary storage on a direct-access storage device. Data stored in a temporary storage queue is known as temporary data.

You can:

- Write data to a temporary storage queue (EXEC CICS WRITEQ TS command)
- Update data in a temporary storage queue (EXEC CICS WRITEQ TS REWRITE command)
- Read data from a temporary storage queue (EXEC CICS READQ TS command)
- Read the next data from a temporary storage queue (EXEC CICS READQ TS NEXT command)
- Delete a temporary storage queue (EXEC CICS DELETEQ TS command)

The TS keyword may be omitted; temporary storage is assumed if it is not specified.

Exception conditions that occur during execution of a temporary storage control command are handled as described in Chapter 9, “Dealing with exception conditions” on page 141.

If you use these commands, you could create intertransaction affinities that adversely affect your ability to perform dynamic transaction routing.

See Chapter 5, “Affinity among transactions” on page 91 for more information about transaction affinity.

Temporary storage queues

Temporary storage queues are identified by symbolic names that may be up to eight characters, assigned by the originating task. Temporary data can be retrieved by the originating task or by any other task using the symbolic name assigned to it. To avoid conflicts caused by duplicate names, a naming convention should be established; for example, the operator identifier or terminal identifier could be used as a suffix to each programmer-supplied symbolic name. Specific items (logical records) within a queue are referred to by relative position numbers.

Temporary storage queues remain intact until they are deleted by the originating task, by any other task, or by a cold start; before deletion, they can be accessed any number of times. Even after the originating task is terminated, temporary data can be accessed by other tasks through references to the symbolic name under which it is stored.

Temporary data can be stored either in main storage or in auxiliary storage. Generally, main storage should be used if the data is needed for short periods of time; auxiliary storage should be used if the data is to be kept for long periods of time. Data stored in auxiliary storage is retained after CICS termination and can be

recovered in a subsequent restart, but data in main storage cannot be recovered. Main storage might be used to pass data from task to task, or for unique storage that allows programs to meet the requirement of CICS that they be quasi-reentrant (that is, serially reusable between entry and exit points of the program).

Typical uses of temporary storage control

A temporary storage queue that has only one record can be treated as a single unit of data that can be accessed using its symbolic name. Using temporary storage control in this way provides a typical scratch-pad capability. This type of storage should be accessed using the READQ TS command with the ITEM option; not doing so may cause the ITEMERR condition to be raised.

In general, temporary storage queues of more than one record should be used only when direct access or repeated access to records is necessary; transient data control provides facilities for efficient handling of sequential data sets.

Some uses of temporary storage queues are:

Terminal paging

A task could retrieve a large master record from a direct-access data set, format it into several screen images (using BMS), store the screen images temporarily in auxiliary storage, and then ask the terminal operator which “page” (screen image) is desired. The application programmer can provide a program (as a generalized routine or unique to a single application) to advance page by page, advance or back up a relative number of pages, and so on.

A suspend data set

Suppose a data collection task is in progress at a terminal. The task reads one or more units of input and then allows the terminal operator to interrupt the process by some kind of coded input. If not interrupted, the task repeats the data collection process. If interrupted, the task writes its incomplete data to temporary storage and terminates. The terminal is now free to process a different transaction (perhaps a high-priority inquiry). When the terminal is available to continue data collection, the operator initiates the task in a “resume” mode, causing the task to recall its suspended data from temporary storage and continue as though it had not been interrupted.

Preprinted forms

An application program can accept data to be written as output on a preprinted form. This data can be stored in temporary storage as it arrives. When all the data has been stored, it can first be validated and then transmitted in the order required by the format of the preprinted form.

Chapter 29. Security control

To avoid attempting accesses that would cause security violations, it can be useful for an application to determine the security authorization of the terminal user. The EXEC CICS QUERY SECURITY command provides a way of doing this.

This chapter gives some guidance on the use of the QUERY SECURITY command. For additional guidance, see the *CICS Security Guide*. For programming information about EXEC CICS commands, see the *CICS Application Programming Reference* manual.

QUERY SECURITY command

QUERY SECURITY is effective with any RACROUTE compliant external security manager (ESM). You can use this command to query whether the terminal user has access to resources that are defined to the external security manager. These can be:

- Resources in CICS resource classes
- Resources in user-defined resource classes.

The terminal user in this context is the user invoking the transaction that contains the QUERY SECURITY command.

In response to a QUERY SECURITY command, CICS returns information about the terminal user's security authorizations. CICS obtains this information from the external security manager. You can code the application to proceed in different ways depending on the user's permitted accesses.

You specify the type of resource that you are querying by the CICS resource type name. For example, if you want to query a user's authorization to access a file, you can specify RESTYPE('FILE'). To identify a particular file within the type, you specify the RESID parameter.

Using QUERY SECURITY

A typical use of the QUERY SECURITY command is to check whether a user is authorized to use a particular transaction **before** displaying the transaction code in a menu.

Security protection at the record or field level

Another use for QUERY SECURITY is to enable you to control access to data at the record or field level. The normal CICS resource security checking for file resources, for example, works only at the file level. To control access to individual records, or even fields within records, you can use QUERY SECURITY. For this purpose, your security administrator must define resource profile names, with appropriate access authorizations, for the records or fields that you want to protect. These profiles are defined in user resource classes defined by the administrator, **not** in CICS resource classes.

To query these classes and resources, the QUERY SECURITY command uses the RESCLASS and RESID options (RESCLASS and RESTYPE are mutually exclusive

options). You can use the CVDA values returned by QUERY SECURITY to determine whether to access the record or field.

CICS-defined resource identifiers

In all cases except for the SPCOMMAND resource type, the resource identifiers are user-defined. However, for the SPCOMMAND type, the identifiers are fixed by CICS. The *CICS Security Guide* details the possible RESID values for the SPCOMMAND resource type.

SEC system initialization parameter

The setting of the SEC system initialization parameter affects the CVDA values returned by the QUERY SECURITY command. The SEC system initialization parameters are described in more detail in the *CICS Security Guide*.

Programming hints

- A transaction can use the QUERY SECURITY command to query a number of resources in order to prepare a list of resources to which the terminal user has access. The use of this technique could generate up to four resource violation messages for each query on a resource that the transaction is not authorized to access. These messages appear on the system console, and the CICS TD queue. If you want to suppress these messages, code NOLOG in the QUERY SECURITY command.
- If a transaction accesses the same resource many times in one execution, you can probably improve performance by defining the transaction with RESSEC(NO) in the transaction resource definition. You can then code the transaction to issue a single QUERY SECURITY command, and to permit access to the resource according to the CVDA values returned. For detailed guidance, see the *CICS Security Guide*.

Non-Terminal transaction security

CICS can now protect, against unauthorized use, resources used in transactions that are not associated with a terminal. These transactions are of three types:

- Transactions that are started by a START command and that do not specify a terminal ID.
- Transactions that are started, without a terminal, as a result of the trigger level being reached for an intrapartition transient data queue.
- The CICS internal transaction (CPLT), which runs during CICS startup, to execute programs specified in the program list table (PLT). This transaction executes both first and second phases of PLTs.

Also, resource security checking can now be carried out for PLT programs that are run during CICS shutdown. PLT shutdown programs execute as part of the transaction that requests the shutdown, and therefore run under the authorization of the user issuing the shutdown command.

The START command handles security for non-terminal transactions started by the START command.

A surrogate user who is authorized to attach a transaction for another user, or cause it to be attached, or who inherits all the resource access authorizations for that transaction, can act for the user.

CICS can issue up to three surrogate user security checks on a single START command, depending on the circumstances:

1. The userid of the transaction that issues the START command, if USERID is specified
2. The userid of the CEDF transaction, if the transaction that issues the START command is being run in CEDF dual-screen mode
3. The CICS region userid of the remote system, if the START command is function shipped to another CICS system and link security is in effect.

A separate surrogate user security check is done for each of these userids, as required, before the transaction is attached.

For programming information about the USERID option, USERIDERR condition, and INVREQ, and NOTAUTH conditions, see the *CICS Application Programming Reference* manual.

Part 6. Testing applications

Chapter 30. Testing applications: the process	415
Preparing the application and system table entries	415
Preparing the system for debugging	415
Single-thread testing	416
Multithread testing	416
Regression testing	416
Year 2000 testing	417

Chapter 31. CICS Execution diagnostic facility (EDF)	419
Getting started	419
Where does EDF intercept the program?	419
What does EDF display?	420
The header	421
The body	421
How you can intervene in program execution	427
EDF menu functions	427
How to use EDF	433
Using EDF in single-screen mode	434
Using EDF in dual-screen mode	435
EDF and remote transactions	435
EDF and DTP programs	436
EDF and DPL programs	436
Stopping EDF	437
Overtyping to make changes	437
Restrictions when using EDF	439
Security considerations	440

Chapter 32. Temporary storage browse (CEBR)	441
How to use the CEBR transaction	441

What does the CEBR transaction display?	442
The header	442
The command area	443
The body	443
The message line	443
The CEBR options on function keys	443
The CEBR commands	444
Using the CEBR transaction with transient data	446
Security considerations	447

Chapter 33. Command-level interpreter (CECI)	449
How to use CECI	449
What does CECI display?	450
The command line	450
The status line	451
The body	453
The message line	454
CECI options on function keys	454
Additional displays	455
Expanded area	455
Variables	455
The EXEC interface block (EIB)	457
Error messages display	458
Making changes	459
How CECI runs	460
CECI sessions	460
Abends	460
Exception conditions	460
Program control commands	461
Terminal sharing	461
Saving commands	461
Security considerations	462

Chapter 30. Testing applications: the process

You have to do two main tasks before you can test and debug an application:

1. Prepare the application and the system table entries
2. Prepare the system for debugging

Preparing the application and system table entries

To prepare the application and system table entries, do the following:

1. Translate, assemble or compile, and link-edit each program. Make sure that there are no error messages on any of these three steps for any program before you begin testing.
2. Use the DEBUG and EDF options on your translator step, so that you can use translator statement numbers with CICS execution diagnostic facility (EDF) displays.
3. Use the COBOL option LIST so that you can relate storage locations in dumps and CEDF displays to the original COBOL source statements, and find your variables in working storage.
4. Use the RDO DEFINE PROFILE command to generate a profile for your transactions to use, and make sure the definitions are INSTALLED.
5. Use the RDO DEFINE TRANSACTION command for each transaction in your application, and make sure the definitions are INSTALLED.
6. If your system does not use program autoinstall, use the RDO DEFINE PROGRAM command for each program used in the application, and make sure the definitions are INSTALLED.
7. If your system does not use program autoinstall, use the RDO DEFINE MAPSET command for each map set in the application, and make sure each definition is INSTALLED.
8. Use the RDO DEFINE FILE command, or put an entry in the FCT, for each file used. If you use RDO, make sure the definitions are INSTALLED.
9. Build at least a test version of each of the files required.
10. Use the DFHDCT macro instruction to define in the DCT each of the transient data destinations to be used by the application.
11. Put job control DLBL statements in the startup job stream for each file used in the application.
12. Prepare some test data.

Preparing the system for debugging

To prepare the system for debugging, do the following:

1. Make sure that the CEDF transaction is available in your system, by including group DFHEDF in the list you specify in the GRPLIST system initialization parameter.
2. Set up appropriate tracing options for your application. For details about setting up tracing options, see the *CICS Problem Determination Guide*.

3. Make sure that transaction dumping is enabled for all transactions under test, by ensuring that DUMP(YES) is specified on the RDO TRANSACTION definitions.
4. Also, ensure that transaction dumping is enabled for all transaction dump codes, and that system dumping is enabled for all system dump codes. These are, anyway, the default settings. For information about setting up dump options, see the *CICS Problem Determination Guide*.
5. Be prepared to print the transaction dumps. Have a DFHDU410 job stream or procedure ready, and have the CICS dump data sets defined in your startup procedure. See the *CICS Operations and Utilities Guide* for information on DFHDU410.
6. Contact your system programmer for information about VSE SYSDUMP data sets available on your system and access to JCL for processing them.
7. Enable CICS to detect loops, by setting the ICVR system initialization parameter in to a number greater than zero. Something between five and ten seconds (ICVR=5000 to ICVR=10000) is usually a workable value.
8. Generate statistics. For more information about using statistics, see the *CICS Performance Guide*.

Single-thread testing

A **single-thread** test takes one application transaction at a time, in an otherwise “empty” CICS system, and sees how it behaves. This enables you to test the program logic, and also shows whether or not the basic CICS information (such as resource definition) is correct. It is quite feasible to test this single application in one CICS region while your normal, online production CICS system is active in another.

Multithread testing

A **multithread** test involves several concurrently active transactions. Naturally, all the transactions are in the same CICS region, so you can readily test the ability of a new transaction to coexist with them.

You may find that a transaction that works perfectly in its single-thread testing still fails in the multithread test. Or it may cause other transactions to fail, or even terminate CICS.

Regression testing

A **regression** test is used to make sure that all the transactions in a system continue to do their processing in the same way both before and after changes are applied to the system. This is to ensure that fixes applied to solve one problem do not cause further problems. It is a good idea to build a set of miniature files to perform your tests on, because it is much easier to examine a small data file for changes.

A good regression test exercises all the code in every program; that is, it explores all tests and possible conditions. As your system develops to include more transactions, more possible conditions, and so on, add these to your test system to

keep it in step. The results of each test should match those from the previous round of testing. Any discrepancies are grounds for suspicion. You can compare terminal output, file changes, and log entries for validity.

Sequential terminal support can be useful for regression testing. When you have a module that has worked for some time and is now being modified, you will need to rerun your old tests to ensure that the function still works. Sequential terminal support makes it easy to maintain a “library” of old test cases and to rerun them when needed.

Sequential terminal support allows you to test programs without having to use a telecommunication device. System programmers can specify that sequential devices be used as terminals (using the terminal control table (TCT)). These sequential devices may be card readers, line printers, disk units, or magnetic tape units. They can also be combinations of sequential devices such as:

- A card reader and line printer (CRLP)
- One or more disk or tape data sets as input
- One or more disk or tape data sets as output

You can prepare a stream of transaction test cases to do the basic testing of a program module. As the testing progresses, you can generate additional transaction streams to validate the multiprogramming capabilities of the programs or to allow transaction test cases to run concurrently.

See “Sequential terminal support” on page 336 for more information about sequential terminal support.

Year 2000 testing

Year 2000 testing should be performed by running VSE with a system date beyond 1999. If you are running VSE as a VM guest you can set up a system date for a specific guest only. If native, you should test VSE under an LPAR. Ensure the testing system does not share resources with any other system. The use of the // DATE card in the CICS startup JCL to set the date for year 2000 testing is not recommended.

Chapter 31. CICS Execution diagnostic facility (EDF)

You can use the execution diagnostic facility (EDF) to test an application program online, without modifying the program or the program-preparation procedure. The names of your programs should not begin with the letters “DFH” because this prefix is used for CICS system modules and samples. Attempting to use the CEDF transaction on a CICS-supplied transaction has no effect. However, you can use the CEDF transaction with CICS sample programs. (For example, you can use CEDF to debug DFHPEP.)

EDF intercepts the execution of CICS commands in the application program at various points, allowing you to see what is happening. Each command is displayed before execution, and most are displayed after execution is complete. Screens sent by the application program are preserved, so you can converse with the application program during testing, just as a user would on a production system.

Each time EDF interrupts the execution of the application program a new CEDF task is started. Each CEDF task is short lived, lasting only long enough for the appropriate display to be processed.

Getting started

The terminal that you are using for the EDF interaction must be in transceive (ATI/TTI) status and be able to send and receive data. This is the most common status for display terminals, but you can find out by asking your system programmer to check its status, or you can use CEMT.

You can use EDF on the same terminal as the transaction you are testing or on a different one. On the same terminal, you **must** start by clearing the screen and entering the transaction code CEDF, otherwise you may get unpredictable results. The message **THIS TERMINAL: EDF MODE ON** is displayed at the top of an empty screen. You clear the screen again and run your transaction in the normal way.

If you are using two terminals, you enter CEDF tttt at one, naming the second in tttt. Then you run your transaction on the second terminal. “How to use EDF” on page 433 gives all the details.

Where does EDF intercept the program?

When a transaction runs under EDF control, EDF intercepts it at the following points, allowing you to interact with it:

- At **program initiation**, after the EXEC interface block (EIB) has been updated, but before the program is given control.
- At the **start of the execution of each CICS command**. This interrupt happens after the initial trace entry has been made, but before the command has been performed. Both standard CICS commands and the Front End Programming Interface (FEPI) commands are intercepted. EXEC DLI and EXEC SQL commands and any requests processed through the resource manager interface are also intercepted at this point.

- At the **end of the execution of every command** except for ABEND, XCTL, and RETURN commands (although these commands could raise an error condition that EDF displays). EDF intercepts the transaction when it finishes processing the command, but before the HANDLE CONDITION mechanism is invoked, and before the response trace entry is made.
- At **program termination**.
- At **normal task termination**.
- When an **ABEND** occurs and after **abnormal task termination**.

Note: For a program translated with the option NOEDF, this still applies, apart from before and after the execution of each command. For a program with CEDF defined as NO by resource definition online (RDO) or program autoinstall exit, the program initiation and termination screens are suppressed as well.

What does EDF display?

All EDF displays have the same general format, but the contents depend on the point at which the task was interrupted. The display indicates which of these interception points has been reached and shows information relevant to that point. Figure 78 shows a typical display; it occurred after execution of a SEND MAP command.

```

TRANSACTION: AC20 PROGRAM: DFH0VT1 TASK: 00032 APPLID: 1234567 DISPLAY:00
STATUS:  COMMAND EXECUTION COMPLETE
EXEC CICS SEND MAP
MAP ('T1 ')
FROM ('.....'.....)
LENGTH (154)
MAPSET ('DFH0T1 ')
CURSOR
TERMINAL
ERASE
NOFLUSH
NOHANDLE

OFFSET:X'002522'   LINE:00673   EIBFN=X'1804'
RESPONSE: NORMAL   EIBRESP=0

ENTER: CONTINUE
PF1 : UNDEFINED      PF2 : SWITCH HEX/CHAR  PF3 : END EDF SESSION
PF4 : SUPPRESS DISPLAYS  PF5 : WORKING STORAGE  PF6 : USER DISPLAY
PF7 : SCROLL BACK      PF8 : SCROLL FORWARD  PF9 : STOP CONDITIONS
PF10: PREVIOUS DISPLAY  PF11: UNDEFINED       PF12: ABEND USER TASK

```

Note: **1** Header **2** Body **3** Message line **4** Menu of functions

Figure 78. Typical EDF display

The display consists of a header, a body (the primary display area), a message line, and a menu of functions you can select at this point. If the body does not fit on one screen, EDF creates multiple screens, which you can scroll through using PF7 and PF8. The header, menu and message areas are repeated on each screen.

The header

The header shows:

- The identifier of the transaction being executed
- The name of the program being executed
- The internal task number assigned by CICS to the transaction
- The applid of the system where the transaction is being executed
- A display number
- STATUS, that is, the reason for the interception by EDF

The body

The body or main part of the display contains the information that varies with the point of intercept.

- At **program initiation**, as shown in Figure 79, EDF displays the COMMAREA and the contents of the principal fields in the EIB. For programming information about these EIB fields, see the *CICS Application Programming Reference* manual. If no COMMAREA value is supplied, line 4 on the screen is not displayed, and EIBCALEN has a value of zero.

```
TRANSACTION: AC20 PROGRAM: DFH0VT1 TASK: 00032 APPLID: 1234567 DISPLAY:00
STATUS: PROGRAM INITIATION

COMMAREA      = '3476559873'
EIBTIME       = 184225
EIBDATE       = 0098105
EIBTRNID      = 'AC20'
EIBTASKN      = 32
EIBTRMID      = 'S246'

EIBCPOSN      = 4
EIBCALEN      = 10
EIBAID        = X'7D'                AT X'032F059A'
EIBFN         = X'0000'              AT X'032F059B'
EIBRCODE      = X'000000000000'     AT X'032F059D'
EIBDS         = '.....'
+ EIBREQID    = '.....'

ENTER: CONTINUE
PF1 : UNDEFINED      PF2 : SWITCH HEX/CHAR  PF3 : END EDF SESSION
PF4 : SUPPRESS DISPLAYS  PF5 : WORKING STORAGE  PF6 : USER DISPLAY
PF7 : SCROLL BACK      PF8 : SCROLL FORWARD  PF9 : STOP CONDITIONS
PF10: PREVIOUS DISPLAY  PF11: UNDEFINED       PF12: UNDEFINED
```

Figure 79. Typical EDF display at program initiation

- At the **start of execution of a CICS command**, EDF displays the command, including keywords, options, and argument values, as shown in Figure 80 on page 422. You can display the information in hexadecimal or character form (and switch from one to the other) by pressing PF2. If character format is requested, numeric arguments are shown in signed numeric character format.

```

TRANSACTION: AC20 PROGRAM: DFH0VT1 TASK: 00032 APPLID: 1234567 DISPLAY:00
STATUS: ABOUT TO EXECUTE COMMAND
EXEC CICS SEND MAP
MAP ('T1      ')
FROM ('.....'..)
LENGTH (154)
MAPSET ('DFH0T1 ')
CURSOR
TERMINAL
ERASE
NOFLUSH
NOHANDLE

OFFSET:X'002522'   LINE:00673   EIBFN=X'1804'

ENTER: CONTINUE
PF1 : UNDEFINED          PF2 : SWITCH HEX/CHAR      PF3 : UNDEFINED
PF4 : SUPPRESS DISPLAYS  PF5 : WORKING STORAGE     PF6 : USER DISPLAY
PF7 : SCROLL BACK       PF8 : SCROLL FORWARD     PF9 : STOP CONDITIONS
PF10: PREVIOUS DISPLAY   PF11: UNDEFINED          PF12: ABEND USER TASK

```

Figure 80. Typical EDF display at start of execution of a CICS command

Figure 81 shows a similar screen for the **start of execution** of an EXEC SQL command running with DB2 version 2.3.

```

TRANSACTION: LOKO PROGRAM: TLOKO TASK: 00082 APPLID: 1234567 DISPLAY:00
STATUS: ABOUT TO EXECUTE COMMAND
CALL TO RESOURCE MANAGER DSNCSQL
EXEC SQL UPDATE
DBRM=TLOK0, STMT=00242, SECT=00001
IVAR 001: TYPE=CHAR, LEN=00010 AT X'001E5A99'
DATA=X'F0F0F0F0F0F0F0F0F0F0F0F0F0F0F0F0'

OFFSET:X'000298'   LINE: UNKNOWN EIBFN= X'0A02'

ENTER: CONTINUE
PF1 : UNDEFINED          PF2 : UNDEFINED          PF3 : UNDEFINED
PF4 : SUPPRESS DISPLAYS  PF5 : WORKING STORAGE     PF6 : USER DISPLAY
PF7 : SCROLL BACK       PF8 : SCROLL FORWARD     PF9 : STOP CONDITIONS
PF10: PREVIOUS DISPLAY   PF11: UNDEFINED          PF12: ABEND USER TASK

```

Figure 81. Typical SQL display at start of execution of a SQL command

In addition to options and values, the command is identified by its hexadecimal offset within the program. If the program was translated with the DEBUG translator option, the line number also appears, as shown in Figure 80. (See “Translator options” on page 7 for information about this option.)

At the start of an EXEC SQL or EXEC DLI command, the body of the EDF display shows you the parameter list of the CALL to which your command translates. If a DLI command generates multiple CALL statements, you see only the last CALL statement.

Note: If you use a PL/I return rather than an EXEC CICS RETURN command to terminate a PL/I program, as is permissible, EDF does not intercept the program, and you do not get the associated panels.

- At the **end of execution of a command**, EDF provides a display in the same format as at the start of the command. At this point, you can see the effects of executing the command, in the values of the variables returned or changed and in the response code. EDF does not provide this display for the ABEND, XCTL, and RETURN commands (although these commands could raise an error condition that EDF displays). The completion screen corresponding to the **about to execute** screen in Figure 80 on page 422 is shown in Figure 82.

```

TRANSACTION: AC20 PROGRAM: DFH0VT1 TASK: 00054 APPLID: 1234567 DISPLAY:00
STATUS: COMMAND EXECUTION COMPLETE
EXEC CICS SEND MAP
MAP ('T1 ')
FROM ('.....')
LENGTH (154)
MAPSET ('DFH0T1 ')
CURSOR
TERMINAL/
ERASE
NOFLUSH
NOHANDLE

OFFSET:X'002522' LINE:00673 EIBFN=X'1804'
RESPONSE: NORMAL EIBRESP=0

ENTER: CONTINUE
PF1 : UNDEFINED PF2 : SWITCH HEX/CHAR PF3 : END EDF SESSION
PF4 : SUPPRESS DISPLAYS PF5 : WORKING STORAGE PF6 : USER DISPLAY
PF7 : SCROLL BACK PF8 : SCROLL FORWARD PF9 : STOP CONDITIONS
PF10: PREVIOUS DISPLAY PF11: UNDEFINED PF12: ABEND USER TASK

```

Figure 82. Typical EDF display at completion of a CICS command

For CICS commands, response codes are described both by name (for example, NORMAL or NOTFND) and by the corresponding EIBRESP value in decimal form. For DL/I, the response code is a 2-character DL/I status code, and there is no EIBRESP value. Programming information, including a list of EIBRESP codes, is in the *CICS Application Programming Reference* manual, and DL/I codes are documented in the *DL/I VSE Application Programming: High Level Programming Interface*.

Figure 83 and Figure 84 show typical screens for an EXEC DLI command.

```

TRANSACTION: XDLI PROGRAM: UPDATE TASK: 00111 APPLID: 1234567 DISPLAY: 00
STATUS: COMMAND EXECUTION COMPLETE
EXEC DLI GET NEXT
USING PCB (+00003)
FIRST
SEGMENT ('A ')
INTO (' ')
SEGLENGTH (+00012)
FIRST
VARIABLE
+SEGMENT ('B ')

OFFSET:X'000246' LINE: 00000510 EIBFN:X'000C'
RESPONSE: 'AD'

ENTER: CONTINUE
PF1 : UNDEFINED PF2 : SWITCH HEX/CHAR PF3 : END EDF SESSION
PF4 : SUPPRESS DISPLAYS PF5 : WORKING STORAGE PF6 : USER DISPLAY
PF7 : SCROLL BACK PF8 : SCROLL FORWARD PF9 : STOP CONDITIONS
PF10: PREVIOUS DISPLAY PF11: UNDEFINED PF12: ABEND USER TASK

```

Figure 83. Typical EDF display at completion of a DLI command (screen one)

```

TRANSACTION: XDLI PROGRAM: UPDATE TASK: 00111 APPLID: 1234567 DISPLAY: 00
STATUS: COMMAND EXECUTION COMPLETE
EXEC DLI GET NEXT
+
FIRST
SEGMENT ('C      ')
SEGLLENGTH (+00010)
LOCKED
INTO ('SMITH  ')
WHERE (ACCOUNT = '12345')
FIELDLENGTH (+00005)

OFFSET:X'000246'  LINE: 00000510      EIBFN:X'000C'
RESPONSE: 'AD'

ENTER: CONTINUE
PF1 : UNDEFINED          PF2 : SWITCH HEX/CHAR      PF3 : END EDF SESSION
PF4 : SUPPRESS DISPLAYS  PF5 : WORKING STORAGE      PF6 : USER DISPLAY
PF7 : SCROLL BACK        PF8 : SCROLL FORWARD      PF9 : STOP CONDITIONS
PF10: PREVIOUS DISPLAY   PF11: UNDEFINED           PF12: ABEND USER TASK

```

Figure 84. Typical EDF display at completion of a DLI command (screen two)

Figure 85 shows a typical screen for an EXEC SQL command at completion.

```

TRANSACTION: LOKO PROGRAM: TLOK0 TASK: 00111 APPLID: 1234567 DISPLAY: 00
STATUS: COMMAND EXECUTION COMPLETE
CALL TO RESOURCE MANAGER DSNCSQL
EXEC SQL UPDATE
PLAN=TLOK0, DBRM=TLOK0, STMT=00242, SECT=00001
SQL COMMUNICATION AREA:
SQLCABC = 136 AT X'001E5A18'
SQLCODE = 000 AT X'001E5A1C'
SQLERRML = 000 AT X'001E5A20'
SQLERRMC = '' AT X'001E5A22'
SQLERRP = 'DSN' AT X'001E5A68'
SQLERRD(1-6) = 000, 000, 00001, -1, 00000, 000 AT X'001E5A70'
SQLWARN(0-A) = ' AT X'001E5A88'
SQLSTATE = 00000----- AT X'001E5A93'

OFFSET:X'000298'  LINE: UNKNOWN EIBFN= X'0A02'
RESPONSE:

ENTER: CONTINUE
PF1 : UNDEFINED          PF2 : UNDEFINED          PF3 : END EDF SESSION
PF4 : SUPPRESS DISPLAYS  PF5 : WORKING STORAGE    PF6 : USER DISPLAY
PF7 : SCROLL BACK        PF8 : SCROLL FORWARD     PF9 : STOP CONDITIONS
PF10: PREVIOUS DISPLAY   PF11: UNDEFINED          PF12: ABEND USER TASK

```

Figure 85. Typical EDF display at completion of a SQL command

- At **program termination** and **normal task termination**, there is no body information; all the pertinent information is in the header. Figure 86 on page 425 and Figure 87 on page 425 show typical screens for program and task termination.

```

TRANSACTION: AC20 PROGRAM: DFH0VT1 TASK: 00054 APPLID: 1234567 DISPLAY:00
STATUS: PROGRAM TERMINATION

ENTER: CONTINUE
PF1 : UNDEFINED          PF2 : SWITCH HEX/CHAR    PF3 : UNDEFINED
PF4 : SUPPRESS DISPLAYS PF5 : WORKING STORAGE    PF6 : USER DISPLAY
PF7 : SCROLL BACK       PF8 : SCROLL FORWARD    PF9 : STOP CONDITIONS
PF10: PREVIOUS DISPLAY  PF11: UNDEFINED         PF12: ABEND USER TASK

```

Figure 86. Typical EDF display at program termination

```

TRANSACTION: AC20          TASK: 00054 APPLID: 1234567 DISPLAY: 00
STATUS: TASK TERMINATION

CONTINUE EDF? (ENTER YES OR NO)
ENTER: CONTINUE
PF1 : UNDEFINED          PF2 : SWITCH HEX/CHAR    PF3 : END EDF SESSION
PF4 : SUPPRESS DISPLAYS PF5 : WORKING STORAGE    PF6 : USER DISPLAY
PF7 : SCROLL BACK       PF8 : SCROLL FORWARD    PF9 : STOP CONDITIONS
PF10: PREVIOUS DISPLAY  PF11: UNDEFINED         PF12: UNDEFINED

```

Figure 87. Typical EDF display at task termination

- When an **abend** or **abnormal task termination** occurs, EDF displays the screens shown in Figure 88 on page 426 and Figure 89 on page 426.

```

TRANSACTION: AC20 PROGRAM: DFH0VT1 TASK:00054 APPLID: 1234567 DISPLAY: 00
STATUS: AN ABEND HAS OCCURRED
COMMAREA = '1287656678'
EIBTIME = 135510
EIBDATE = 91163
EIBTRNID = 'AC20'
EIBTASKN = 76
EIBTRMID = 'S232'
EIBCPOSN = 4
EIBCALEN = 10
EIBAID = X'7D' AT X'032F059A'
EIBFN = X'1804' SEND AT X'032F059B'
EIBRCODE = X'000000000000' AT X'032F059D'
EIBDS = '.....'
+ EIBREQID = '.....'

ABEND : ABCD

ENTER: CONTINUE
PF1 : UNDEFINED PF2 : SWITCH HEX/CHAR PF3 : END EDF SESSION
PF4 : SUPPRESS DISPLAYS PF5 : WORKING STORAGE PF6 : USER DISPLAY
PF7 : SCROLL BACK PF8 : SCROLL FORWARD PF9 : STOP CONDITIONS
PF10: PREVIOUS DISPLAY PF11: UNDEFINED PF12: UNDEFINED

```

Figure 88. Typical EDF display when an abend occurs

```

TRANSACTION: AC20 TASK: 00054 APPLID: 1234567 DISPLAY: 00
STATUS: ABNORMAL TASK TERMINATION
COMMAREA = '2934564671'
EIBTIME = 135510
EIBDATE = 0098115
EIBTRNID = 'AC20'
EIBTASKN = 76
EIBTRMID = 'S232'
EIBCPOSN = 4
EIBCALEN = 10
EIBAID = X'7D' AT X'032F059A'
EIBFN = X'1804' SEND AT X'032F059B'
EIBRCODE = X'000000000000' AT X'032F059D'
EIBDS = '.....'
+ EIBREQID = '.....'

ABEND : ABCD
CONTINUE EDF? (ENTER YES OR NO) REPLY: YES
ENTER: CONTINUE
PF1 : UNDEFINED PF2 : SWITCH HEX/CHAR PF3 : END EDF SESSION
PF4 : SUPPRESS DISPLAYS PF5 : WORKING STORAGE PF6 : USER DISPLAY
PF7 : SCROLL BACK PF8 : SCROLL FORWARD PF9 : STOP CONDITIONS
PF10: PREVIOUS DISPLAY PF11: UNDEFINED PF12: UNDEFINED

```

Figure 89. Typical EDF display at abnormal task termination

The body displays the COMMAREA and the values of the fields in the EIB as well as the following items:

- The abend code
- If the abend code is ASRA (that is, a program interrupt has occurred), the program status word (PSW) at the time of interrupt, and the source of the interrupt as indicated by the PSW
- If the PSW indicates that the instruction giving rise to the interrupt is within the application program, the offset of that instruction

How you can intervene in program execution

The power of EDF lies in what you can do at each of the intercept points. For example, you can:

- Change the argument values before a command is executed. For CICS commands, you cannot change the actual command, or add or delete options, but you can change the value associated with any option. You can also suppress execution of the command entirely using NOOP. See page 438 for further details.
- Change the results of a command, either by changing the argument values returned by execution or by modifying the response code. This allows you to test branches of the program that are hard to reach using ordinary test data (for example, what happens on an input/output error). It also allows you to bypass the effects of an error to check whether this eliminates a problem.
- Display the working storage of the program, the EIB, and for DL/I programs, the DIB.
- Invoke the command interpreter (CECI). Under CECI you can execute commands that are not present in the program to gain additional information or change the execution environment.
- Display any other location in the CICS region.
- Change the working storage of the program and most fields in the EIB and the DIB. EDF stops your task from interfering with other tasks by preventing you from changing other areas of storage.
- Display the contents of temporary storage and transient data queues.
- Suppress EDF displays until one or more of a set of specific conditions is fulfilled. This speeds up testing.
- Retrieve up to 10 previous EDF displays or saved screens.
- Switch off EDF mode and run the application normally.
- Abend the task.

The first two types of changes are made by overtyping values in the body of the command displays. “Overtyping to make changes” on page 437 tells you how to do this. You use the function keys in the menu for the others; “EDF menu functions” tells you exactly what you can do and how to go about it.

EDF menu functions

The function keys that you can use at any given time are displayed in a menu at the bottom of every EDF display (see Figure 78 on page 420). The function of the ENTER key for that display is also shown. Functions that apply to all displays are always assigned to the same key, but definitions of some keys depend on the display and the intercept point. To select an option, press the indicated function key. Where a terminal has 24 function keys, EDF treats PF13 through PF24 as duplicates of PF1 through PF12 respectively. If your terminal has no PF keys, place the cursor under the option you want and press the ENTER key.

ABEND USER TASK

terminates the task being monitored. EDF asks you to confirm this action by displaying the message “ENTER ABEND CODE AND REQUEST ABEND AGAIN.” After entering the code at the position indicated by the cursor, the user

must request this function again to abend the task with a transaction dump identified by the specified code. If you enter "NO," the task is abended without a dump and with the 4-character default abend code of ????.

Abend codes beginning with the character A are reserved for use by CICS. Using a CICS abend code may cause unpredictable results.

You cannot use this function if an abend is already in progress or the task is terminating.

BROWSE TEMP STORAGE

produces a display of the temporary storage queue CEBRxxxx, where xxxx is the terminal identifier of the terminal running EDF. This function is only available from the working storage (PF5) screen. You can then use CEBR commands, discussed in "The CEBR commands" on page 444, to display or modify temporary storage queues and to read or write transient data queues.

CONTINUE

redispays the current screen if you have made any changes, incorporating the changes. If you had not made changes, CONTINUE causes the transaction under test to resume execution up to the next intercept point. To continue, press ENTER.

CURRENT DISPLAY

redispays the current screen if you have made any changes, with the changes incorporated. If you have not made changes, it causes EDF to display the command screen for the last intercept point. To execute this function, press ENTER from the appropriate screen.

DIB DISPLAY

shows the contents of the DL/I interface block (DIB). This function is only available from the working-storage screen (PF5). See the *DL/I VSE Application Programming: EXEC DLI Commands* manual for information on DIB fields.

EIB DISPLAY

displays the contents of the EIB. This function is only available from the working-storage screen (PF5). See Figure 79 on page 421 for an example of an EIB display. For programming information about the EIB, see the *CICS Application Programming Reference* manual. If COMMAREA exists, EDF also displays its address and one line of data in the dump format.

INVOKE CECI

accesses CECI. This function is only available from the working storage (PF5) screen. See Figure 91 on page 432 for an example of the screen from which CECI is invoked. You can then use CECI commands, discussed in Chapter 33, "Command-level interpreter (CECI)" on page 449. These CECI commands include INQUIRE and SET commands against the resources referenced by the original command before and after command execution. See page 439 for restrictions when running CECI in dual-screen mode. The use of CECI from this panel is similar to the use of CEBR within CEDF.

END EDF SESSION

ends the EDF control of the transaction. The transaction continues running from that point but no longer runs in EDF mode.

NEXT DISPLAY

is the reverse of PREVIOUS DISPLAY. When you have returned to a previous display, this option causes the next one forward to be displayed and the display number to increase by one.

PREVIOUS DISPLAY

causes the previous display to be sent to the screen. This is the previous command display, unless you saved other displays. The number of the display from the current intercept point is always 00. As you request previous displays, the display number decreases by 1 to -01 for the first previous display, -02 for the one before that, and so on, down to the oldest display, -10. When no more previous screens are available, the PREVIOUS option disappears from the menu, and the corresponding function key becomes inoperative.

REGISTERS AT ABEND

displays storage containing the values of the registers should a local ASRA abend occur. The layout of the storage is:

- Register values (0 through 15)
- PSW at abend (8 bytes)

In some cases, when a second program check occurs in the system before EDF has captured the values of the registers, this function does not appear on the menu of the abend display. If this happens, a second test run generally proves to be more informative.

REMEMBER DISPLAY

places a display that would not usually be kept in memory, such as an EIB display, in the EDF memory. (EDF automatically saves the displays at the start and completion of each command.) The memory can hold up to 10 displays. The displays are numbered in reverse chronological order (that is, -10 is the oldest display, and -01 is the newest). All pages associated with the display are kept in memory and can be scrolled when recalled. Note, however, that if you save a working-storage display, only the screen on view is saved.

SCROLL BACK

applies to an EIB, DIB, or command display that does not all fit on one screen. When the screen on view is not the first one of the display, and there is a plus sign (+) before the first option or field, you can view previous screens in the display by selecting SCROLL BACK. See Figure 79 on page 421 for an example.

SCROLL FORWARD

applies to an EIB, DIB, or command display that does not all fit on one screen. When this happens, a plus sign (+) appears after the last option or field in the display, to show that there are more screens. Using SCROLL FORWARD brings up the next screen in the display.

SCROLL BACK FULL

has the same function for displays of working storage as the SCROLL BACK option for EIB and DIB displays. SCROLL BACK FULL gives a working-storage display one full screen backward, showing addresses lower in storage than those on the current screen.

SCROLL FORWARD FULL

has the same function for displays of working storage as the SCROLL FORWARD option for EIB and DIB displays. SCROLL FORWARD FULL gives

a working-storage display one full screen forward, showing addresses higher in storage than those on the current screen.

SCROLL BACK HALF

is similar to SCROLL BACK FULL, except that the display of working storage is reversed by only half a screen.

SCROLL FORWARD HALF

is similar to SCROLL FORWARD FULL, except that the display of working storage is advanced by only half a screen.

STOP CONDITIONS

produces the menu screen shown in Figure 90. You use this screen to tell EDF when to resume its displays after you have pressed the SUPPRESS DISPLAYS key. You can use STOP CONDITIONS and SUPPRESS DISPLAYS together to cut down on the interaction between you and EDF when you are checking a program that you know is partly working.

```
TRANSACTION: AC20 PROGRAM: DFH0VT1 TASK: 0086 APPLID: 1234567 DISPLAY: 00
DISPLAY ON CONDITION:-

COMMAND:                                EXEC CICS
OFFSET:                                  X'.....'
LINE NUMBER:
CICS EXCEPTION CONDITION:                ERROR
ANY CICS CONDITION                       NO
TRANSACTION ABEND                        YES
NORMAL TASK TERMINATION                  YES
ABNORMAL TASK TERMINATION                 YES

DLI ERROR STATUS:
ANY DLI ERROR STATUS

ENTER: CURRENT DISPLAY
PF1 : UNDEFINED                          PF2 : UNDEFINED                          PF3 : UNDEFINED
PF4 : SUPPRESS DISPLAYS                   PF5 : WORKING STORAGE                   PF6 : USER DISPLAY
PF7 : UNDEFINED                          PF8 : UNDEFINED                          PF9 : UNDEFINED
PF10: UNDEFINED                          PF11: UNDEFINED                         PF12: REMEMBER DISPLAY
```

Figure 90. Typical EDF display for STOP CONDITIONS

You can specify any or all of these events as STOP CONDITIONS:

- A specific type of function and option, such as EXEC CICS READNEXT file or EXEC CICS ENQ resource, is encountered.
- The command at a specific offset or on a specific line number (assuming the program has been translated with the DEBUG option) is encountered.
- Any DLI error status occurs, or a particular DLI error status occurs.
- A specific exception condition occurs. If CICS exception condition is specified as ERROR (the default), CEDF redisplay a screen in response to any ERROR condition (for example, NOTOPEN, EOF, or INVREQ). If you specify a specific condition such as EOF, CEDF redisplay the screen only when an EOF condition arises, provided that ANY CICS CONDITION is left as the default NO.

If this field is changed to YES, CEDF overrides the CICS exception conditions and redisplay a screen whenever any command results in a nonzero EIBRESP value such as NOTOPEN, EOF, or QBUSY.

- Any exception condition occurs for which the system action is to raise ERROR; for example, INVREQ or NOTFND.
- An abend occurs.
- The task ends normally.
- The task ends abnormally.

You do not always have to set STOP CONDITIONS in order to use the SUPPRESS DISPLAYS function, because EDF sets a default in the following fields on the assumption that you usually want to resume displays if any of them occurs:

- CICS exception condition
- Transaction abend
- Normal task termination
- Abnormal task termination

These are the options described in Figure 90 on page 430. You can turn off any of the defaults that do not apply when you bring up the STOP CONDITIONS menu, as well as adding conditions specific to your program.

When you use an offset for STOP CONDITIONS, you must specify the offset of the BALR instruction corresponding to a command. The offset can be determined from the code listing produced by the compiler or assembler. In COBOL, C, or PL/I, you must use the compiler option that produces the assembler listing to determine the relevant BALR instruction.

When you use a line number, you must specify it exactly as it appears on the listing, including leading zeros, and it must be the line on which a command starts. If you have used the NUM or the SEQUENCE translator options, the translator uses your line numbers as they appear in the source. Otherwise, the translator assigns line numbers.

Line numbers can be found in the translator listing (SYSLST in the translator step) if you have used either the SOURCE or VBREF translator options. If you have used the DEBUG translator option, as you must to use line numbers for STOP CONDITIONS, the line number also appears in your compilation (assembly) listing, embedded in the translated form of the command, as a parameter in the CALL statement.

You can tell EDF to stop suppressing displays at DL/I commands as well as at CICS commands. You do this by overtyping the qualifier "CICS" on the command line with "DLI" and entering the type of DL/I command at which you want suppression to stop. You must be executing a DL/I program or have executed one earlier in the same task. You can suppress DL/I commands as early as the program initiation panel.

You can also stop suppression when a particular DL/I status code occurs. However, you can specify the code in STOP CONDITIONS only when a DL/I command (not a CICS command) is about to be processed. For information about the status codes that you can use, see the list of codes in the DL/I interface block (DIB) in the *DL/I VSE Application Programming: EXEC DLI Commands* manual.

SUPPRESS DISPLAYS

suppresses all EDF displays until one of the specified STOP CONDITIONS occurs. When the condition occurs, however, you still have access to the 10

previous command displays, even though they were not actually sent to the screen when they were originally created.

SWITCH HEX/CHAR

switches displays between character and hexadecimal form. The switch applies only to the command display, and has no effect on previously remembered displays, STOP CONDITIONS displays, or working-storage displays.

In DL/I command displays that contain the WHERE option, only the key values (the expressions following each comparison operator) can be converted to hexadecimal.

UNDEFINED

means that the indicated function key is not defined for the current display at the current intercept point.

USER DISPLAY

causes EDF to display what would be on the screen if the transaction was not running in EDF mode. (You can use it only for single terminal checkout.) To return to EDF after using this key, press the ENTER key.

WORKING STORAGE

allows you to see the contents of the working-storage area in your program, or of any other address in the CICS region. Figure 91 shows a typical working-storage screen.

```

TRANSACTION: AC20 PROGRAM: DFH0VT1 TASK: 00030 APPLID: 1234567 DISPLAY:00
ADDRESS: 035493F0 WORKING STORAGE
035493F0 000000 E3F14040 00000000 00010000 00000000 T1 .....
03549400 000010 00000000 00000000 F1000000 00000000 .....1.....
03549410 000020 F0000000 00000000 F0000000 00000000 0.....0.....
03549420 000030 F0000000 00000000 F0000000 00000000 0.....0.....
03549430 000040 00000000 00000000 00000000 00000000 .....
03549440 000050 D7C1D5D3 00000000 D9C5C3C4 00000000 PANL...RECD...
03549450 000060 D3C9E2E3 00000000 C8C5D3D7 00000000 LIST...HELP...
03549460 000070 84000000 00000000 A4000000 00000000 d.....u.....
03549470 000080 82000000 00000000 C4000000 00000000 b.....D.....
03549480 000090 E4000000 00000000 C2000000 00000000 U.....B.....
03549490 0000A0 D5000000 00000000 E2000000 00000000 N.....S.....
035494A0 0000B0 7B000000 00000000 6C000000 00000000 #.....%.....
035494B0 0000C0 4A000000 00000000 F1000000 00000000 ¢.....1.....
035494C0 0000D0 F2000000 00000000 F3000000 00000000 2.....3.....
035494D0 0000E0 00000000 00000000 C4C6C8F0 E5C1C240 .....DFH0VAB
035494E0 0000F0 C4C6C8F0 E5E3C2D3 C4C6C8F0 E5D3C9D6 DFH0VTBLDFH0VLIO

ENTER: CURRENT DISPLAY
PF1 : UNDEFINED PF2 : BROWSE TEMP STORAGE PF3 : UNDEFINED
PF4 : EIB DISPLAY PF5 : INVOKE CECI PF6 : USER DISPLAY
PF7 : SCROLL BACK HALF PF8 : SCROLL FORWARD HALF PF9 : UNDEFINED
PF10: SCROLL BACK FULL PF11: SCROLL FORWARD FULL PF12: REMEMBER DISPLAY

```

Figure 91. Typical EDF display for working-storage

The working-storage contents are displayed in a form similar to that of a dump listing, that is, in both hexadecimal and character representation. The address of working storage is displayed at the top of the screen. You can browse through the entire area using the scroll commands, or you can simply enter a new address at the top of the screen. This address can be anywhere within the CICS region. The **working-storage** display provides two additional scrolling keys, and a key to display the EIB (the DIB if the command is a DL/I command).

The meaning of “working storage” depends on the programming language of the application program, as follows:

COBOL

All data storage defined in the WORKING-STORAGE section of the program

C and PL/I

The dynamic storage area (DSA) of the current procedure

Assembler language

The storage defined in the current DFHEISTG DSECT

Assembler language programs do not always acquire working storage; it may not be necessary, for example, if the program does not issue CICS commands. You may get the message "Register 13 does not address DFHEISTG" when you LINK to such a program. The message does not necessarily mean an error, but there is no working storage to look at.

Except for COBOL programs, working storage starts with a standard format save area; that is, registers 14 to 12 begin at offset 12 and register 13 is stored at offset 4.

Working storage can be changed at the screen; either the hexadecimal section or the character section can be used. Also, the ADDRESS field at the head of the display can be overtyped with a hexadecimal address; storage starting at that address is then displayed when ENTER is pressed. This allows any location in the address space to be examined. Further information on the use of overtyping is given in "Overtyping to make changes" on page 437.

If the program storage examined is not part of the working storage of the program currently executing (which is unique to the particular transaction under test), the corresponding field on the screen is protected to prevent the user from overwriting storage that might belong to or affect another task.

If the initial part of a working-storage display line is blank, the blank portion is not part of working storage. This can occur because the display is doubleword aligned.

At the beginning and end of a task, working storage is not available. In these circumstances, EDF generates a blank storage display so that the user can still examine any storage area in the region by overtyping the address field.

Note that if you terminate an LE/VSE program with an ordinary non-CICS return), EDF does not intercept the return, and you are not able to see working storage. If you use an EXEC CICS RETURN command instead, you get an EDF display before execution and at program termination. This is normally the last point at which working storage is available to you.

If you are using an LE/VSE-enabled program, working-storage is freed at program termination if the program is terminated using a non-CICS return. In this case, working storage is not available for display.

How to use EDF

You can run EDF on the same terminal as the transaction to be tested (this is called "single-screen mode"), or on a different terminal ("dual-screen mode"). Generally, you can use whichever method you prefer, but there are a few situations in which one or the other is required. You must use single-screen mode for remote

transactions. See “Restrictions when using EDF” on page 439 for other conditions that affect your

Using EDF in single-screen mode

When you use EDF with just one terminal, the EDF inputs and outputs are interleaved with those from the transaction. This sounds complicated, but works quite easily in practice. The only noticeable peculiarity is that when a SEND command is followed by a RECEIVE command, the display sent by the SEND command appears twice: once when the SEND is executed, and again when the RECEIVE is executed. It is not necessary to respond to the first display, but if you do, EDF preserves anything that was entered from the first display to the second.

You start EDF by:

- Entering transaction code CEDF from a cleared screen, or
- Pressing the appropriate PF key (if one has been defined for EDF)

Next, you start the transaction to be tested by:

1. Pressing the CLEAR key to clear the screen
2. Entering the transaction code of the transaction you want to test

When both EDF and the user transaction are sharing the same terminal, EDF restores the user transaction display at the following times:

- When the transaction requires input from the operator
- When you change the transaction display
- At the end of the transaction
- When you suppress the EDF displays
- When you request USER DISPLAY

To enable restoration, user displays are remembered at the following times:

1. At start of task, before the first EDF screen for the task is displayed
2. Before the next EDF screen is displayed, if the user display has been changed
3. On leaving SCREEN SUPPRESS mode

If a program has been translated with option NOEDF, or has NO specified for CEDF in its RDO resource definition, it is not possible for EDF to ascertain when the user display is being changed. This means that, unless either situation 1 or 3 also apply, the next EDF screen to be displayed overwrites any user display sent by this program without saving it first, so that it cannot be later restored.

When EDF restores the transaction display, it does not sound the alarm or affect the keyboard in the same way as the user transaction. The effect of the user transaction options is seen when the SEND command is processed, but not when the screen is restored. When you have NOEDF specified in single-screen mode, take care that your program does not send and receive data because you will not see it.

When EDF restores the transaction display on a device that uses color, programmed symbols, or extended highlighting, these attributes are no longer present and the display is monochrome without the programmed symbols or extended highlighting. Also, if the inbound reply mode in the application program is set to “character” to enable the attribute-setting keys, EDF resets this mode,

causing these keys to be disabled. If these changes will prevent your transaction from executing properly, test in a dual-screen mode.

If you end your EDF session part way through the transaction, EDF restores the screen with the keyboard locked if the most recent RECEIVE command has not been followed by a SEND command; otherwise, the keyboard is unlocked.

Checking pseudoconversational programs

EDF makes a special provision for testing pseudoconversational transactions from a single terminal. If the terminal came out of EDF mode between the several tasks that make up a pseudoconversational transaction, it would be very hard to do any debugging after the first task. So, when a task terminates, EDF asks the operator whether EDF mode is to continue to the next task. If you are debugging a pseudoconversational task, press enter, as the default is "yes." If you have finished, reply "no."

Using EDF in dual-screen mode

In dual-screen mode, you use one terminal for EDF interaction and another for sending input to, and receiving output from, the transaction under test.

You start by entering, at the EDF terminal, the transaction CEDF tttt, where tttt is the name of the terminal on which the transaction is to be tested.

The message that CEDF gives in response to this depends on whether there is already a transaction running on the second terminal. If the second terminal is not busy, the message displayed at the first terminal is:

```
TERMINAL tttt: EDF MODE ON
```

and nothing further happens until a transaction is started on the second terminal, when the PROGRAM INITIATION display appears.

You can also use EDF in dual-screen mode to monitor a transaction that is already running on the second terminal. If, for example, you believe a transaction at a specific terminal to be looping, you can go to another terminal and enter a CEDF transaction naming the terminal at which this transaction is running. The message displayed at the first terminal is:

```
TERMINAL tttt: TRANSACTION RUNNING: EDF MODE ON
```

EDF picks up control at the next EXEC CICS command executed, and you can then observe the sequence of commands that are causing the loop, assuming that at least one EXEC CICS command is executed.

EDF and remote transactions

You cannot use EDF in dual-screen mode if the transaction under test, or the terminal that invokes it, is owned by a different system. Furthermore, if the release of CICS on the remote system is earlier than CICS Transaction Server for VSE/ESA Release 1, you cannot run the transaction directly under EDF. In this situation, you must use the routing transaction CRTE. You enter CEDF at the terminal, clear the screen, and then enter CRTE followed by the system identifier of the remote system. This action causes CICS to route subsequent inputs to the remote system, and you can then enter the transaction identifier of the transaction

you want to test. The *CICS-Supplied Transactions* manual explains how to use CRTE.

If a remote application abends while under CEDF using CRTE, the CEDF abnormal task termination screen is displayed after message DFHAC2206 for the user transaction and before message DFHAC2206 for CRTE.

There is a difference in execution as well. For remote transactions, EDF purges its memory of your session at the termination of each transaction, whether EDF is to be continued or not. This means that any options you have set and any saved screens are lost between the individual tasks in a pseudoconversational sequence.

EDF and DTP programs

You can also test a transaction that is using distributed transaction processing across a remote link by telling EDF to monitor the session on the link. You can do this on either (or both) of the participating systems that are running under CICS and has EDF installed. (You cannot do this if the transaction has been routed from another CICS system because you must use single-screen mode for remote transactions.)

For APPC and MRO links, you can name the system identifier (sysid) of the remote system:

```
CEDF sysid
```

This causes EDF to associate itself with any transaction attached across any session belonging to the specified system.

For APPC, MRO, and LU6.1 links, you can use the session identifier (sessionid) that the transaction is using:

```
CEDF sessionid
```

You can determine the session identifier with the CEMT transaction, but this means that the transaction must be running and have reached the point of establishing a session before you start EDF.

If a transaction using distributed transaction processing also has a terminal associated with it, or if you can invoke it from a terminal (even though it does not use one), you can use EDF to test it in the ordinary way from that terminal.

When you have finished testing the transaction on the remote system, turn off EDF on that SYSID or sessionid before logging off from CICS with CESF. For example:

```
CEDF sysid,OFF
```

Failure to do this could cause another transaction using a link to that system to be suspended.

EDF and DPL programs

You can use EDF to test a transaction that includes a distributed program link using single-or dual-terminal mode. However, EDF only shows command invocation and response (which could include a remote abend and the message: **a remote abend has occurred** is sent to your local terminal along with the SYSID of the system from which the abend was received). After control is returned to your local

program, EDF continues to test as normal, but the PSW is not displayed if the abend is in a remote program.

Stopping EDF

If you want to end EDF control of a terminal, the method you use depends on where you are in the testing. If the transaction under test is still executing and you want it to continue, but without EDF, press the END EDF SESSION function key. If you have reached the task termination intercept, EDF asks if you want to continue. If you do not, overwrite the reply as NO (YES is the default). If no transaction is executing at the terminal, clear the screen and enter:

```
CEDF ,OFF
```

(The space and comma are required.)

If you are logging off from dual-screen mode, clear the screen and enter `CEDF tttt,OFF`.

In all these cases, the message **THIS TERMINAL: EDF MODE OFF** is displayed at the top of an empty screen.

Overtyping to make changes

Most of the changes you make with EDF involve changing information in memory. You do this simply by typing over the information shown on the screen with the information you want used instead. You can change any area where the cursor stops when you use the tab keys, except for the menu area at the bottom.

When you change the screen, you must observe the following rules:

- On CICS command screens, any argument value can be overtyped, but not the keyword of the argument. An optional argument cannot be removed, nor can an option be added or deleted.
- When you change an argument in the command display (as opposed to the working storage screen), you can change only the part shown on the display. If you attempt to overwrite beyond the value displayed, the changes are not made and no diagnostic message is generated. If the argument is so long that only part of it appears on the screen, change the area in working storage to which the argument points. (To determine the address, display the argument in hexadecimal format; the address of the argument location also appears.)
- In character format, numeric values always have a sign field, which can be overtyped with a minus or a blank only.
- When an argument is to be displayed in character format, some of the characters may not be displayable (including lowercase characters). EDF replaces each nondisplayable character with a period. When overtyping a period, you must be aware that the storage may in fact contain a character other than a period. Do not overwrite any character with a period; if you do, the change is ignored and no diagnostic message is issued. If you need to overwrite a character with a period, you can do so by switching the display to hexadecimal format, using PF2, and overtyping with X'4B'.
- When storage is displayed in both character and hexadecimal format and changes are made to both, the value of the hexadecimal field takes precedence should the changes conflict; no diagnostic message is issued.

- The arguments for some commands, such as HANDLE CONDITION, are program labels rather than numeric or character data. The form in which EDF displays (and accepts modifications to) these arguments depends on the programming language in use:
 - For COBOL, a null argument is displayed: for example, ERROR (), and because of this, you cannot modify it.
 - For C, labels are not valid.
 - For PL/I, the address of the label constant is used; for example, ERROR (X'001D0016').
 - For assembler language, the address of the program label is used; for example, ERROR (X'00030C').

If no label value is specified on an EXEC CICS HANDLE CONDITION command, EDF displays the condition name alone without the parentheses.

- The response field can be overtyped with the name of any exception condition, including ERROR, that can occur for the current function, or with the word NORMAL. The effect when EDF continues is that the program takes whatever action has been prescribed for the specified response. You can get the same effect by changing the EIBRESP field in the EIB display to the corresponding values. If you change the EIBRESP value or the response field on the **command execution complete** screen, EIBRCODE is updated. EIBRESP appears on second EIB screen and is the only one you can change (EIBRCODE protected). You can get the same effect by changing the EIBRESP value on the EIB display; EDF changes related values in the EIB and command screens accordingly if you do this.
- If uppercase translation is not specified for the terminal you are using you must take care to always enter uppercase characters.
- Any command can be overtyped with NOOP or NOP before processing; this suppresses processing of the command. Use of the ERASE EOF key, or overtyping with blanks, gives the same effect. When the screen is redisplayed with NOOP, the original verb line can be restored by erasing the whole verb line with the ERASE EOF key and pressing the ENTER key.

When you overtype a field representing a data area in your program, the change is made directly in application program storage and is permanent. However, if you change a field that represents a constant (a program literal), program storage is not changed, because this may affect other parts of the program that use the same constant or other tasks using the program. The command is executed with the changed data, but when the command is displayed after processing, the original argument values reappear. For example, suppose you are testing a program containing a command coded:

```
EXEC CICS SEND MAP('MENU') END-EXEC.
```

If you change the name MENU to MENU2 under EDF before executing the command, the map actually used is MENU2, but the map displayed on the response is MENU. (You can use the “previous display” key to verify the map name you used.) If you process the same command more than once, you must enter this type of change each time.

EDF responses

The response of EDF to any keyboard entry follows the rules listed below, in the order shown:

1. If the CLEAR key is used, EDF redisplay the screen with any changes ignored.
2. If invalid changes are made, EDF accepts any valid changes and redisplay the screen with a diagnostic message.
3. If the display number is changed, EDF accepts any other changes and shows the requested display.
4. If a PF key is used, EDF accepts any changes and performs the action requested by the PF key. Pressing ENTER with the cursor under a PF key definition in the menu at the bottom of the screen is the same as pressing a PF key.
5. If the ENTER key is pressed and the screen has been modified (other than the REPLY field), EDF redisplay the screen with changes included.
6. If the ENTER key is pressed and the screen has not been modified (other than the REPLY field), the effect differs according to the meaning of the ENTER key. If the ENTER key means CONTINUE, the user transaction continues to execute. If it means CURRENT DISPLAY, EDF redisplay the current status display.

Restrictions when using EDF

There are some restrictions on the use of EDF that make it preferable or even necessary to use one particular screen mode:

- EDF can be used only in single-screen mode when running a remote transaction.
- VM PASSTHRU is not supported by EDF when testing in single-screen mode.
- In single-screen mode, neither the user transaction nor CEDF should specify message journaling, because the messages interfere with the EDF displays. Message journaling is controlled by the profile definition for each transaction.
- In single screen mode, the CEDF transaction should not specify PROTECT(YES) in its RDO PROFILE resource definition. If this option is specified, message protection for the CEDF transaction is ignored. The user transaction can still specify the PROTECT(YES) option even when running under CEDF. This restriction does not apply to dual-screen mode.
- If a SEND LAST command is issued, EDF is ended before the command is processed if you are using single-screen mode.
- If you want to test an application program that uses screen partitions, or that does its own request unit (RU) chaining, you must run in dual-screen mode.
- In single-screen mode, if the RDO PROFILE for the user transaction specifies INBFMH(ALL) or INBFMH(DIP), the profile for CEDF must have the same INBFMH value. Otherwise the user transaction abends ADIR. Dual-screen mode does not require the profiles to match in this respect.
- If the inbound reply mode is set to "character" to enable the attribute setting keys, EDF disables the keys in single-screen mode.

- When using CECI under EDF in dual-screen mode, be aware that certain commands (for example, ASSIGN and ADDRESS) are issued against the EDF terminal and not the transaction terminal. See page 428 for information about how to invoke CECI from CEDF.
- When using EDF in dual-screen mode, avoid starting a second task at the EDF terminal, for example by issuing a START command. Because EDF is a pseudoconversational transaction, it does not prevent a second task from starting at the terminal it is using. This may lead to a deadlock in certain circumstances.
- When using EDF screen suppression in dual screen mode, commands like long wait, or a second RECEIVE, may cause EDF to appear as if it had finished. If the task is then ABENDED, EDF is reactivated at the monitoring terminal.

Other restrictions apply to both screen modes:

- If a transaction issues the FREE command, EDF is switched off without warning.
- To test a user transaction executing on a remote CICS at a release level earlier than CICS Transaction Server for VSE/ESA Release 1, you must run the transaction under control of CRTE, as explained in “EDF and remote transactions” on page 435.
- EDF does not intercept calls to the CPI Communications interface (CPI-C) or the SAA Resource Recovery interface (CPI-RR). You can test transactions that use CPI calls under EDF, but you cannot see EDF displays at the call points.
- User application programs that are to be debugged using EDF must be assembled (or compiled) with the translator option EDF, which is the default. If you specify NOEDF, the program cannot be debugged using EDF. There is no performance advantage in specifying NOEDF, but the option can be useful to prevent commands in well debugged subprograms appearing on EDF displays.
- When processing an EXEC CICS SIGNON command, CEDF suppresses display of the password value to reduce the risk of accidental disclosure.

Security considerations

EDF is such a powerful tool that your installation may restrict its use with attach-time security. (The external security manager used by your installation defines the security attributes for the EDF transaction.) If this has been done, and you are not authorized to use CEDF, you cannot initiate the transaction.

For guidance on using security, see your system programmer or the *CICS Security Guide*.

Chapter 32. Temporary storage browse (CEBR)

You can use the browse transaction (CEBR) to browse temporary storage queues and delete them. You can also use the CEBR transaction to transfer the contents of a transient data queue to temporary storage in order to look at them, and to reestablish the transient data queue when you have finished. The CEBR commands that perform these transfers allow you to add records to a transient data queue and remove all records from a transient data queue. See “The CEBR commands” on page 444 and “The CEBR options on function keys” on page 443 for more information about their use.

How to use the CEBR transaction

You start the CEBR transaction by entering the transaction identifier CEBR, followed by the name of the queue you want to browse. For example, to display the temporary storage queue named CEBRS209, you type CEBR CEBRS209 and press ENTER. CICS responds with a display of the queue, for example, as shown in Figure 92:

```
CEBR TS QUEUE CEBRS209 SYSID ISC1 RECORD 1 OF 3 COL 1 OF 22
ENTER COMMAND ==>
***** TOP OF QUEUE *****
00001 00005500123400001S209
00002 00005600345600002S209
00003 00010200056400001S209
***** BOTTOM OF QUEUE *****

PF1 : HELP          PF2 : SWITCH HEX/CHAR    PF3 : TERMINATE BROWSE
PF4 : VIEW TOP      PF5 : VIEW BOTTOM          PF6 : REPEAT LAST FIND
PF7 : SCROLL BACK HALF PF8 : SCROLL FORWARD HALF PF9 : UNDEFINED
PF10: SCROLL BACK FULL PF11: SCROLL FORWARD FULL PF12: UNDEFINED
```

Figure 92. Typical CEBR display of temporary storage queue contents

Alternatively, you can start the CEBR transaction from the CEDF transaction. You do this by pressing PF5 from the initial CEDF screen (see Figure 78 on page 420), which takes you to the working-storage screen, and then pressing PF2 from that screen to browse temporary storage (that is, invoke the CEBR transaction). The CEBR transaction responds by displaying the temporary storage queue whose name consists of the four letters CEBR followed by the four letters of your terminal identifier. (CICS uses this same default queue name if you invoke the CEBR transaction directly and do not supply a queue name.) The result of invoking the CEBR transaction without a queue name or from an EDF session at terminal S21A is shown in Figure 93. If you enter the CEBR transaction from the CEDF transaction, you return to the EDF panel when you press PF3 from the CEBR screen.

```

CEBR TS QUEUE CEBRS21A SYSID ISC1 RECORD 1 OF 0 COL 1 OF 0 1
ENTER COMMAND ==> 2
***** TOP OF QUEUE *****
***** BOTTOM OF QUEUE *****

3

TEMPORARY STORAGE QUEUE CEBRS21A IS EMPTY 4
PF1 : HELP PF2 : SWITCH HEX/CHAR PF3 : TERMINATE BROWSE 5
PF4 : VIEW TOP PF5 : VIEW BOTTOM PF6 : REPEAT LAST FIND
PF7 : SCROLL BACK HALF PF8 : SCROLL FORWARD HALF PF9 : UNDEFINED
PF10: SCROLL BACK FULL PF11: SCROLL FORWARD FULL PF12: UNDEFINED

```

Note: 1 Header 2 Command area 3 Body 4 Message line 5 Menu of options

Figure 93. Typical CEBR display of default temporary storage queue

What does the CEBR transaction display?

As shown in Figure 93, a CEBR transaction display consists of a header, a command area, a body (the primary display area), a message line, and a menu of functions you can select at this point.

The header

The header shows:

- The transaction being run, that is, CEBR.
- The identifier of the temporary storage queue (CEBRS209 in Figure 92 on page 441 and CEBRS21A in Figure 93). You can overwrite this field in the header if you want to switch the screen to another queue.
- The system name that corresponds to a temporary storage pool name or to a remote system. If you have not specified one, the name of the local system is displayed. You can overwrite this field in the header if you want to browse a shared or remote queue.
- The number of the highlighted record.
- The number of records in the queue (three in CEBRS209 and none in CEBRS21A).
- The position in each record at which the screen starts (position 1 in both cases) and the length of the longest record (22 for queue CEBRS209 and zero for queue CEBRS21A).

The command area

The command area is where you enter commands that control what is to be displayed and what function is to be performed. These commands are described in “The CEBR commands” on page 444. You can also modify the screen with function keys shown in the menu of options at the bottom of the screen. The function keys are explained in “The CEBR options on function keys.”

The body

The body is where the queue records are shown. Each line of the screen corresponds to one queue record. If a record is too long for the line, it is truncated. You can change the portion of the record that is displayed, however, so that you can see an entire record on successive screens. If the queue contains more records than will fit on the screen, you can page forward and backward through them, or specify at what record to start the display, so that you can see all the records you want.

The message line

CEBR uses the message line between the body and menu to display messages to the user, such as the “Queue...empty” message in Figure 93 on page 442.

The CEBR options on function keys

The function keys that you can use at any time are displayed at the bottom of every CEBR transaction screen. The keys have the same meaning on all screens. If your terminal does not have PF keys, you can simulate their use by placing the cursor under the description and pressing ENTER. Where a terminal has 24 function keys, The CEBR transaction treats PF13 through PF24 as duplicates of PF1 through PF12 respectively.

PF1 HELP

Displays a help screen that lists all the commands you can use when the CEBR transaction is running. You can return to the main screen by pressing ENTER.

PF2 SWITCH HEX/CHAR

Switches the screen from character to hexadecimal format, and back again.

PF3 TERMINATE BROWSE

Terminates the CEBR transaction. If you entered the CEBR transaction directly, it frees up your terminal for the next transaction. If you entered from an CEDF session, it returns you to the working-storage screen from which you entered.

PF4 VIEW TOP

Displays the first records in the queue and has the same effect as the TOP command.

PF5 VIEW BOTTOM

Displays the last records in the queue and has the same effect as the BOTTOM command.

PF6 REPEAT LAST FIND

Repeats the previous FIND command.

PF7 SCROLL BACK HALF

Moves the display backward by one-half the number of records that fit on the screen, so that the records on the top half of the screen move to the bottom half.

PF8 SCROLL FORWARD HALF

Advances the display by one-half the number of records that fit on the screen, so that the records on the bottom half of the screen move to the top half.

PF9 VIEW RIGHT (or VIEW LEFT)

Changes the screen to show the columns immediately after (to the right of) or before (to the left of) the columns currently on display. The key is not defined if the entire record fits on one line of the screen. It moves you to the right until the end of the record is reached, and then reverses to move left back to the beginning of the record. You can also use the COLUMN command to change the column at which the display begins.

PF10 SCROLL BACK FULL

Moves the screen backward by the number of records that fit on the screen, to show the records immediately before those currently on display.

PF11 SCROLL FORWARD FULL

Advances the screen by the number of records that will fit on the screen, to show the records immediately after those currently on display.

The CEBR commands

Here is a list of the CEBR commands that you can use to view and manipulate the records in the temporary storage queue.

BOTTOM

(Abbreviation: B)

Shows the last records in the temporary storage queue (as many as fill up the body of the screen, with the last record on the last line).

COLUMN nnnn

(Abbreviation: C nnnn)

Displays the records starting at character position (column) nnnn of each record. The default starting position, assumed when you initiate the CEBR transaction, is the first character in the record.

FIND /string

(Abbreviation: F /string)

Finds the next occurrence of the specified string. The search starts in the record **after** the **current record**. The current record is the one that is highlighted. In the initial display of a queue, the current record is set to one, and therefore the search begins at record two.

If the string is found, the record containing the string becomes the highlighted line, and the display is changed to show this record on the second line. If you cannot see the search string after a successful FIND, it is in columns of the record beyond those on display; use the scroll key or the COLUMN command to shift the display right or left to show the string.

For example:

```
FIND /05-02-93
```

will locate the next occurrence of the string "05-02-93" The / character is a delimiter. It does not have to be /, but it must not be a character that appears in the search argument. For example, if the string you were looking for was "05/02/93" instead of "05-02-93," you could not use the following:

```
FIND /05/02/93
```

There is a slash in the search string. The following examples would work:

```
FIND X05/02/93    or    FIND S05/07/93
```

Any delimiter except a / or one of the digits in the string works. If there are any spaces in the search string, you must repeat the delimiter at the end of the string. For example:

```
FIND /CLARE JACKSON/
```

The search string is not case-sensitive. When you have entered a FIND command, you can repeat it (that is, find the next occurrence of the string) by pressing PF6.

GET xxxx

(Abbreviation: G xxxx)

Transfers the named transient data queue to the end of the temporary storage queue currently on display. This enables you to browse the contents of the queue. xxxx must be either the name of an intrapartition transient data queue, or the name of an extrapartition transient data queue that has been opened for input. See "Using the CEBR transaction with transient data" on page 446 for more information about browsing transient data queues.

LINE nnnn

(Abbreviation: L nnnn)

Starts the body of the screen at the queue record one prior to nnnn, and sets the current line to nnnn. (This arrangement causes a subsequent FIND command to start the search after record nnnn.)

PURGE

Deletes the queue being browsed.

Do not use PURGE to delete the contents of an internally generated queue, such as a BMS logical message.

Note: If you purge a recoverable temporary storage queue, no other task can update that queue (add a record, change a record, or purge) until your task ends.

PUT xxxx

(Abbreviation: P xxxx)

Copies the temporary storage queue that is being browsed to the named transient data queue. xxxx must be either the name of an intrapartition transient data queue, or the name of an extrapartition transient data queue that has been opened for output. See "Using the CEBR transaction with transient data" on page 446 for more information about creating or restoring a transient data queue.

QUEUE xxxxxxxx**(Abbreviation: Q xxxxxxxx)**

Changes the name of the queue you are browsing. The value that you specify can be in character format (for example, QUEUE ABCD) or in hexadecimal format (for example, QUEUE X'C1C2C3C4'). The CEBR transaction responds by displaying the data that is in the named queue.

You can also change the queue name by overtyping the current value in the header.

SYSID xxxx**(Abbreviation: S xxxx)**

Changes the name of the temporary storage pool or remote system where the queue is to be found.

You can also change this name by overtyping the current SYSID value in the header.

TERMINAL xxxx**(Abbreviation: TERM xxxx)**

Changes the name of the queue you are browsing, but is tailored to applications that use the convention of naming temporary storage queues that are associated with a terminal by a constant in the first four characters and the terminal name in the last four. The new queue name is formed from the first four characters of the current queue name, followed by xxxx.

TOP**(Abbreviation: T)**

Causes the CEBR transaction to start the display at the first record in the queue.

Using the CEBR transaction with transient data

The GET command reads each record in the transient data queue that you specify and writes it at the end of the temporary storage queue you are browsing, until the transient data queue is empty. You can then view the records that were in the transient data queue. When you have finished your inspection, you can copy the temporary storage queue back to the transient data queue (using the PUT command). This usually leaves the transient data queue as you found it, but not always. Here are some points you need to be aware of when using the GET and PUT commands:

- If you want to restore the transient data queue unchanged after you have browsed it, make sure that the temporary storage queue on display at the time of the GET command is empty. Otherwise, the existing temporary storage records will be copied to the transient data queue when the subsequent PUT command is issued.
- After you get a transient data queue and before you put it back, other tasks may write to that transient data queue. When you issue your PUT command, the records in the temporary storage queue will be copied **after** the new records, so that the records in the queue are no longer in the order in which they were originally created. Some applications depend on sequential processing of the records in a queue.

- After you get a **recoverable** transient data queue, no other task can access that queue until your transaction ends. If you entered the CEBR transaction from the CEDF transaction, the CEDF transaction must end, although you can respond “yes” to the “continue” question if you are debugging a pseudoconversational sequence of transactions. If you invoked the CEBR transaction directly, you must end it.
- Likewise, after you issue a PUT command to a recoverable transient data queue, no other task can access that queue until your transaction ends.

The GET and PUT commands do not need to be used as a pair. You can add to a transient data queue from a temporary storage queue with a PUT command at any time. If you are debugging code that reads a transient data queue, you can create a queue in temporary storage (with the CECI transaction, or the CEBR GET command, or by program) and then refresh the transient data queue as many times as you like from temporary storage. Similarly, you can empty a transient data queue by using a GET command without a corresponding PUT command.

Security considerations

Some installations restrict the use of the CEBR transaction, particularly in production systems, to prevent modifications that were not intended or not authorized. Installations also may protect individual resources, including temporary storage and transient data queues. If you are using the CEBR transaction and experience an abend described as a security failure, you probably have attempted to access a queue to which your user ID is not authorized.

Chapter 33. Command-level interpreter (CECI)

You can use the command-level interpreter (CECI) transaction to check the syntax of CICS commands and process these commands interactively on a 3270 screen. CECI allows you to follow through most of the commands to execution and display the results. It also provides you with a reference to the syntax of the whole of the CICS command-level application programming and system programming interface.

CECI interacts with your test system to allow you to create or delete test data, temporary storage queues, or to deliberately introduce wrong data to test out error logic. You can also use CECI to repair corrupted database records on your production system.

How to use CECI

You start the command-level interpreter by entering either of two transaction identifiers, CECS or CECI, followed by the name of the command you want to test. You can list command options too, although you can also do this later. For example:

```
CECS READ FILE('FILEA')
```

or

```
CECI READ FILE('FILEA')
```

CICS responds with a display of the command and its associated functions, options, and arguments, as shown in Figure 94. If you leave out the command, CECI provides a list of possible commands to get you started. You can use any of the commands described for programming purposes in the *CICS Application Programming Reference* manual and the *CICS System Programming Reference* manual. CECI also supports the FEPI commands provided for the CICS Front End Programming Interface.

```
READ FILE('FILEA')
STATUS:  COMMAND SYNTAX CHECK          NAME= 1 2
EXEC CICS READ
  File( 'FILEA  ' )
  < SYsid() >                            3
  SEt() | Into()
  < Length() >
  RIDfld()
  < Keylength() < GGeneric > >
  < RBa | RRn | DEBRec | DEBKey >
  < GTeq | Equal >
  < Update < Token() > >

S Option RIDFLD has been omitted or specified with an invalid value,
the command cannot be executed.        4
PF 1 HELP 2 HEX 3 END 4 EIB 5 VAR 6 USER 7 SBH 8 SFH 9 MSG 10 SB 11 SF 5
```

Note: 1 Command line 2 Status line 3 Body 4 Message line 5 Menu of functions

Figure 94. Typical CECI display for command syntax check

If you use the transaction code CECS, the interpreter simply checks your command for correct syntax. If you use CECI, you have the option of executing your command once the syntax is correct. (CICS uses two transaction identifiers to allow different security to be assigned to syntax checking and execution.)

What does CECI display?

All CECI screens have the same basic layout. As shown in Figure 94 on page 449, CECI displays consist of a command input line, a status line, the body or main part of the screen, a message line, and a menu of functions you can select at this point.

The command line

The command line is the first line of the screen. You enter the command you want to process or whose syntax you want to check here. This can be the full or abbreviated syntax. The rules for entering and abbreviating the command are:

- The keywords EXEC CICS are optional.
- The options of a command can be abbreviated to the number of characters sufficient to make them unique. Valid abbreviations are shown in uppercase characters in syntax displays in the body of the screen.
- The quotes around character strings are optional, and all strings of characters are treated as character-string constants unless they are preceded by an ampersand (&), in which case they are treated as variables.
- Options of a command that receive a value from CICS when the command is processed are called **receivers**, and need not be specified. The value received from CICS is included in the syntax display, and stored in the variable if one has been specified, after the command has been processed.
- If you issue a CECI command with two of the keywords in conflict, CECI ignores the first keyword and issues an error message, such as this one, from a READ command:

```
E INTO option conflicts with SET option and is ignored
```
- If you put a question mark in front of your command, the interpreter will stop after the syntax check, even if you have used the transaction code CECI. If you want to proceed with execution, remove the question mark.

The following example shows the abbreviated form of a command. The file control command:

```
EXEC CICS READ FILE('FILEA') RIDFLD('009000') INTO(&REC)
```

can be entered on the command input line, as:

```
READ FIL(FILEA) RID(009000)
```

or at a minimum, as:

```
READ F(FILEA) RI(009000)
```

In the first form, the INTO specification creates a variable, &REC, into which the data is to be read. However, INTO is a receiver (as defined above) and you can omit it. When you do, CICS creates a variable for you automatically.

The status line

As you go through the process of interpreting a command, CECI presents a sequence of displays. The format of the body of the screen is essentially the same for all; it shows the syntax of the command and the option values selected. The status line on these screens tells you where you are in the processing of the command, and is one of:

- COMMAND SYNTAX CHECK
- ABOUT TO EXECUTE COMMAND
- COMMAND EXECUTION COMPLETE
- COMMAND NOT EXECUTED

From any of these screens, you can select additional displays. When you do, the body of the screen shows the information requested, and the status line identifies the display, which may be any of:

- EXPANDED AREA
- VARIABLES
- EXEC INTERFACE BLOCK
- SYNTAX MESSAGES

These screens are described in “Additional displays” on page 455. You can request them at any time during processing and then return to the command interpretation sequence.

There is also one input field in the status line called NAME=. This field is used to create and name variables, as explained in “Variables” on page 455 and “Saving commands” on page 461.

Command syntax check

When the status line shows **command syntax check** (as shown in Figure 94 on page 449), it indicates that the command entered on the command input line has been syntax checked but is not about to be processed. This is always the status if you enter CECS or if you precede your command with a question mark. It is also the status when the syntax check of the command gives severe error messages.

In addition, you will get this status if you attempt to execute one of the commands that the interpreter cannot execute. Although any command can be syntax-checked, using either CECS or CECI, the interpreter cannot process the following commands any further:

- EXEC CICS commands that depend upon an environment that the interpreter does not provide:

- HANDLE ABEND
- HANDLE AID
- HANDLE CONDITION
- IGNORE CONDITION
- PUSH HANDLE
- POP HANDLE
- GETMAIN
- FREEMAIN
- WAIT EVENT
- WAIT EXTERNAL
- WAITCICS
- SEND PARTNSET

```
FREE
SEND LAST
```

- BMS commands that refer to partitions (because the display cannot be restored after the screen is partitioned)
- EXEC DLI
- CPI Communication (CPI-C) commands
- SAA Resource Recovery interface (CPI-RR) commands

About to execute command

This display (as shown in Figure 95) appears when none of the reasons for stopping at **command syntax check** applies.

```
READ FILE('FILEA') RIDFLD('009000')
STATUS: ABOUT TO EXECUTE COMMAND                                NAME=
EXEC CICS READ
  File( 'FILEA  ' )
  < SYsid() >
  SEt() | Into()
  < Length() >
  RIDfld( '009000' )
  < Keylength() < GGeneric > >
  < RBa | RRn | DEBRec | DEBKey >
  < GTeq | Equal >
  < Update < Token() > >

PF 1 HELP 2 HEX 3 END 4 EIB 5 VAR 6 USER 7 SBH 8 SFH 9 MSG 10 SB 11 SF
```

Figure 95. Typical CECI display for about to execute command

If you press the ENTER key at this point without changing the screen, CECI will execute the command. You can still modify it at this point, however. If you do, CECI ignores the previous command and processes the new one from scratch. This means that the next screen displayed will be **command syntax check** if the command cannot be executed or else **about to execute command** if the command is correct.

Command execution complete

This display (as shown in Figure 96 on page 453) appears after the interpreter has executed a command, in response to the ENTER key from an unmodified **about to execute command** screen.

```

INQUIRE FILE NEXT
STATUS:  COMMAND EXECUTION COMPLETE                                NAME=
EXEC CICS INquire File( 'FILEA  ' )
  < STArt | END | Next >
  < ACcessmethod( +0000000003 ) >
  < ADd( +00000000041 ) >
  < BAseDsname( '                                ' ) >
  < BLOCKFormat( +0000000016 ) >
  < BLOCKKeylen( -0000000001 ) >
  < BLOCKSize( -0000000001 ) >
  < BRowse( +00000000039 ) >
  < Catname( 'UCATC1C' ) >
  < DElete( +00000000043 ) >
  < DSname( 'CICS410.FILEA                        ' ) >
  < EMptystatus( +00000000032 ) >
  < ENAblestatus( +00000000033 ) >
  < EXclusivE( +00000000001 ) >
  < Fwdrecstatus( +00000000361 ) >
  < Journalnum( +000000 ) >
+ < KEYLength( +00000000000 ) >

RESPONSE: NORMAL                                EIBRESP=+0000000000 EIBRESP2=+0000000000
PF 1 HELP 2 HEX 3 END 4 EIB 5 VAR 6 USER 7 SBH 8 SFH 9 MSG 10 SB 11 SF

```

Figure 96. Typical CECI display for command execution complete

The command has been processed and the results are displayed on the screen.

Any **receivers**, whether specified or not, together with their CICS-supplied values, are displayed intensified.

The body

The body of **command syntax check**, **about to execute command**, and **command execution complete** screens contains information common to all three displays.

The full syntax of the command is displayed. Options specified in the command line or assumed by default are intensified, to show that they will be used in executing the command, as are any **receivers**. The < > brackets indicate that you can select an option from within these brackets. If you make an error in your syntax, CECI diagnoses it in the message area that follows the body, described in “The message line” on page 454. If there are too many diagnostic messages, the rest of the messages can be displayed using PF9.

Arguments can be displayed in either character or hexadecimal format. You can use PF2 to switch between formats. In character format, some characters are not displayable (including lowercase characters on some terminals); CECI shows them as periods. You need to switch to hexadecimal to show the real values, and you need to use caution when modifying them, as explained in “Making changes” on page 459.

If the value of an option is too long for the line, only the first part is displayed followed by “...” to indicate there is more. You can display the full value by

positioning the cursor at the start of the option value and pressing ENTER. This produces an expanded display described in “Expanded area” on page 455.

If the command has more options than will fit on one screen, a plus sign (+) appears at the left-hand side of the last option of the current display to indicate that there are more. An example of this is shown in Figure 96 on page 453. You can display additional pages by scrolling with the PF keys.

The message line

CECI uses the message line to display error messages. After execution of a command, the message line shows the response code. Figure 94 on page 449 shows an error message, where the user has omitted a required field. The **S** that precedes the message indicates that it is severe (bad enough to prevent execution). There are also warning messages (flagged by **W**) and error messages (flagged by **E**), which provide information without preventing execution. **E** messages indicate option combinations unusual enough that they may not be intended and warrant a review of the command before you proceed with execution.

Where there are multiple error messages, CECI creates a separate display containing all of them, and uses the message line to tell you how many there are, and of what severity. You can get the message display with PF9, as explained in “Additional displays” on page 455.

Figure 96 on page 453 shows the second use of the message line, to

show the result of executing a command. CECI provides the information in both text (NORMAL in the example in Figure 96 on page 453) and in decimal form (the EIBRESP and EIBRESP2 value).

CECI options on function keys

The single line at the foot of the screen provides a menu indicating the effect of the PF keys for the display. See Figure 94 on page 449.

The PF keys are described below. If the terminal has no PF keys, the same effect can be obtained by positioning the cursor under the required item in the menu and pressing ENTER.

PF1 HELP

displays a HELP panel giving more information on how to use the command interpreter and on the meanings of the PF keys.

PF2 HEX

(SWITCH HEX/CHAR) switches the display between hexadecimal and character format. This is a mode switch; all subsequent screens stay in the chosen mode until the next time this key is pressed.

PF3 END

(END SESSION) ends the current session of the interpreter.

PF4 EIB

(EIB DISPLAY) shows the contents of the EXEC interface block (EIB). An example of this screen is shown in Figure 98 on page 458.

PF5 VAR

(VARIABLES) shows all the variables associated with the current command interpreter session, giving the name, length, and value of each. See “Variables” for more information about the use of this PF key.

PF6 USER

(USER DISPLAY) shows the current contents of the user display panel (that is, what would appear on the terminal if the commands processed thus far had been executed by an ordinary program rather than the interpreter). This key is not meaningful until a terminal command is executed, such as EXEC CICS SEND MAP.

PF7 SBH

(SCROLL BACK HALF) scrolls the body half a screen backward.

PF8 SFH

(SCROLL FORWARD HALF) scrolls the body half a screen forward.

PF9 MSG

(DISPLAY MESSAGES) shows all the messages generated during the syntax check of a command.

PF10 SB

(SCROLL BACK) scrolls the body one full screen backward.

PF11 SF

(SCROLL FORWARD) scrolls the body one full screen forward.

Additional displays

Additional screens of information are available when you press the relevant PF key. You can get back to your original screen by pressing ENTER from an unmodified screen.

Expanded area

This display uses the whole of the body of the screen to display a variable selected with the cursor. The cursor can be positioned at the start of the value of an option on a syntax display, or under the ampersand of a variable in a variables display. Pressing ENTER then gives the expanded area display. The scrolling keys can be used to display all the information if it exceeds a full screen.

Variables

Figure 97 on page 456 shows the result of requesting a **variables** display, obtained by pressing PF5. For each variable associated with the current interpreter session, it shows the name, length, and value.

```

READ FILE('FILEA') RIDFLD('009000')INTO(&REC)
VARIABLES  LENGTH  DATA
&DFHC      +00016  THIS IS A SAMPLE
&DFHW      +00046  EXEC CICS WRITEQ QUEUE('  CIS200') FROM(&DFHC)
&DFHR      +00045  EXEC CICS READQ  QUEUE('  CIS200') INTO(&DFHC)
&REC       +00080  482554 D694 72 WIDGET, .007 TEST 100

PF 1 HELP 2 HEX 3 END 4 EIB 5 VAR 6 USER          9 MSG

```

Figure 97. Typical CECI display of variables associated with a CECI session

The first three variables displayed are created for you by CECI and always appear unless you explicitly delete them. They are designed to help you create command lists, as described in “Saving commands” on page 461, as well as to serve as examples.

After these three, you will see any variables that you have created. The fourth one in Figure 97, &REC, is the result of executing:

```
READ FILE('FILEA') RID('009000') INTO(&REC)
```

Normally, the value supplied for an option in the command line is taken as a character string constant. However, sometimes you need to specify a variable to represent this value, as when you want to connect two commands through option values.

For example, to change a record with CECI, you might first enter:

```
EXEC CICS READ UPDATE INTO(&REC)
      FILE('FILEA') RID('009000')
```

You would then modify the record as required by changing the variable &REC, and then enter:

```
EXEC CICS REWRITE FROM(&REC) FILE('FILEA')
```

The ampersand (&) in the first position tells CECI that you are specifying a variable.

A variable is also useful when the values of the options cause the command to exceed the line length of the command input area. Creating variables with the required values and specifying the variable names in the command overcomes the line length limitation.

Defining variables

Variables can have a data type of character, fullword, halfword, or packed decimal, and you can create them in any of the following ways:

- By naming the variable in a receiver (&REC in Figure 97 on page 456, for example). The variable is created when the command is processed. The data type and length are implied by the option.
- By adding new entries to the list of variables already defined. To create a new variable, simply type its name and length in the appropriate columns on the first unused line of the variables display, and then press ENTER. For character variables, use the length with which the variable has been defined. For fullwords or halfwords, type **F** or **H**. For packed variables, use the length in bytes, preceded by a **P**.

Character variables are initialized to blanks. The others are initialized to zero in the appropriate form. After a variable is created, you can change the value by modifying the data field on the **variables** display.

- By using the NAME field on the status line when you have produced an expanded area display of a particular option. You do this by positioning the cursor under the option on a syntax display and pressing ENTER. Then you assign the variable name you want associated with the displayed option value by typing it into the NAME field and pressing ENTER again.
- By copying an existing variable. You do this by obtaining an expanded area display of the variable to be copied, overkeying the name displayed with the name of the new variable, and pressing ENTER.
- By using the NAME field directly on a syntax display. This creates a character variable whose contents are the character string on the command line, for use in command lists as explained in “Saving commands” on page 461.

You can also delete a variable, although you do not usually need to, as CECI discards all variables at session end. To delete one before session end, position the cursor under the ampersand that starts the name, press ERASE EOF, and then press ENTER.

The EXEC interface block (EIB)

You can display the EIB associated with the CECI transaction by pressing PF4. Figure 98 on page 458 shows an example of the contents of the EXEC interface block (EIB).

```
READ FILE('FILEA') RIDFLD('009000')
EXEC INTERFACE BLOCK
  EIBTIME      = +0124613
  EIBDATE      = +0098175
  EIBTRNID     = 'CECI'
  EIBTASKN     = +0000046
  EIBTRMID     = 'S200'
  EIBCPOSN     = +00004
  EIBCALEN     = +00000
  EIBAID       = X'7D'
  EIBFN        = X'0000'
  EIBRCODE     = X'000000000000'
  EIBDS        = '.....'
  EIBREQID     = '.....'
  EIBRSRCE     = ' '
  EIBSYNC      = X'00'
  EIBFREE      = X'00'
  EIBRECV      = X'00'
  EIBATT       = X'00'
  EIBEOC       = X'00'
+ EIBFMH      = X'00'

PF 1 HELP 2 HEX 3 END 4 EIB 5 VAR 6 USER 7 SBH 8 SFH 9 MSG 10 SB 11 SF
```

Figure 98. Typical CECI display of the EIB

The fields in the EIB are described for programming purposes in the *CICS Application Programming Reference* manual.

Error messages display

When there are more messages than CECI can display on the message line, you can display all of them by pressing PF9.

```
READ
SYNTAX MESSAGES
S Option FILE has been omitted or specified with an invalid value.
The command cannot be executed.
S Option RIDFLD has been omitted or specified with an invalid value.
The command cannot be executed.

PF 1 HELP 2 HEX 3 END 4 EIB 5 VAR 6 USER 7 SBH 8 SFH      10 SB 11 SF
```

Figure 99. Typical CECI display of the message display

Making changes

Until CICS executes a command, you can change it by changing the contents of the command line, by changing the option values shown in the syntax display in the body, or by changing the values of variables on the **variables** screen. (You can still make changes after a command is executed, but, unless they are in preparation for another command, they have no effect.)

When you make your changes in the command line or on the **variables** screen, they last for the duration of the CECI transaction. If you make them in the body of the syntax screen, however, they are temporary. They last only until the command is executed and are not reflected in the command line.

As noted earlier, not all characters are displayable on all terminals. When the display is in character rather than hexadecimal format, CECI shows these characters as periods (X'4B'). When you overwrite a period, be aware that the current value may not be a period but an undisplayable character.

Furthermore, you cannot change a character to a period when the display is in character mode. If you attempt this, CECI ignores your change, and does not issue a diagnostic message. To make such a change, you have to switch the display to hexadecimal and enter the value (X'4B') that represents a period.

There is a restriction on changes in hexadecimal format as well. If you need to change a character to a blank, you cannot enter the code (X'40') from a

hexadecimal display. Again, your change is ignored and CECI does not issue a message. Instead, you must switch to character mode and blank out the character.

After every modification, CECI rechecks your syntax to ensure that no errors have appeared. It restarts processing at the **command syntax check** if there are any execution-stoppers, and at **about to execute command** if not. Only after you press ENTER on an unmodified **about to execute command** screen does CECI execute your command.

How CECI runs

There are several things to know about how the interpreter works, in order to use it properly. These include:

- CECI sessions
- Abends
- Exception conditions
- Program control commands
- Terminal sharing
- Saving commands

CECI sessions

The interpreter runs as a transaction, using programs supplied by CICS. It is conversational, which means that everything you do between the start of a session (entering CECI) and the end (PF3) is a single logical unit of work in a single task. This means that locks and enqueues produced by commands you execute remain for the duration of your session. If you read a record for update from a recoverable file, for example, that record is not available to any other task until you end CECI.

Abends

CECI executes all commands with the NOHANDLE option, so that execution errors do not ordinarily cause abends.

CECI also issues a HANDLE ABEND command at the beginning of the session, so that it will not lose control even if an abend occurs. Consequently, when you get one, CECI handles it and there is no resource backout. If you are doing a series of related updates to protected resources, you need to be sure that you do not do some of them and then find you cannot complete the others. If you find yourself in this situation, you can execute a SYNCPOINT ROLLBACK command or an ABEND command with the CANCEL option to remove the effects of your earlier commands on recoverable resources.

Exception conditions

For some commands, CECI may return exception conditions even when all specified options are correct. This occurs because, on some commands, CECI uses options that you do not specify explicitly. For example, the ASSIGN command always returns the exception condition INVREQ under CECI. Even though it may return the information you requested correctly, it will have attempted to get information from other options, some of which are invalid.

Program control commands

Because the interpreter is itself an application program, the interpretation of some program control commands may produce different results from an application program executing those commands. For example, ABEND command is intercepted, as noted above, unless you use the CANCEL option.

If you execute a LINK command, the target program executes, but in the environment of the interpreter, which may not be the one expected. In particular, if you modify a user display during a linked-to program, the interpreter will not be aware of the changes.

Similarly, if you interpret an XCTL command, CECI passes control to the named program and never gets control back, so that the CECI session is ended.

Terminal sharing

When the command being interpreted is one that uses the same screen as the interpreter, the command interpreter manages the sharing of the screen between the interpreter display and the user display.

The user display is restored:

- When the command being processed requires input from the operator
- When the command being processed is about to modify the user display
- When USER DISPLAY is requested

Thus, when a SEND command is followed by a RECEIVE command, the display sent by the SEND command appears twice, once when the SEND command is processed, and again when the RECEIVE command is processed. It is not necessary to respond to the SEND command, but, if a response is made, the interpreter stores it and redisplay it when the screen is restored for the RECEIVE command.

When the interpreter restores the user display, it does not sound the alarm or affect the keyboard in the same way as when a SEND command is processed. This is similar to EDF (see “Using EDF in single-screen mode” on page 434 for more information).

Saving commands

Sometimes you may want to execute a command, or a series of commands, under CECI, repeatedly. One technique for doing this is to create a temporary storage queue containing the commands. You then alternate reading the command from the queue and executing it.

CECI provides shortcuts both for creating the queue and for executing commands from it. To create the queue:

1. Start a CECI session.
2. Enter the first (or next) command you want to save on the command line, put &DFHC in the NAME field in the status line, and press ENTER. This action causes the usual syntax check, and it also stores your command as the value of &DFHC, which is the first of those three variables that CECI always defines

for you. (See Figure 97 on page 456.) If you select the variables display at this point, you will see that &DFHC is the value of your command.

3. After the syntax is correct but before execution (on the **about to execute command** screen), change the command line to &DFHW and press ENTER. This causes CECI to use the value of &DFHW for the command to be executed. &DFHW is the second of the variables CECI supplies, and it contains a command to write the contents of variable &DFHC (that is, your command) to the temporary storage queue named "CI^{ttt}," where "ttt" is the name of your terminal and two blanks precede the letters "CI."
4. Execute this WRITEQ command (through the **command execution complete** screen). This stores your command on the queue.
5. If you want to save more than one command, repeat steps 2 on page 461 to 4 for each.

When you want to execute the saved commands from the list, do the following:

1. Enter &DFHR on the command line and press ENTER. &DFHR is the last of the CECI-supplied variables, and it contains a command to read the queue that was written earlier. Execute this command; it will bring the first (next) of the commands you saved into the variable &DFHC.
2. Then enter &DFHC on the command line and press ENTER. CECI replaces the command line with the value of &DFHC, which is your command. Press ENTER to execute your command.
3. Repeat these two steps, alternating &DFHR and &DFHC on the command line, until you have executed all of the commands you saved.

You can vary this procedure to suit your needs. For example, you can skip commands in the sequence by simply skipping step (2). You can change the options of the saved command before executing it in the same way as a command entered normally.

If you want to repeat execution of the saved sequence, you need to specify the option ITEM(1) on the first execution of the READQ command, in order to reposition your read to the beginning of the queue.

Security considerations

The interpreter is such a powerful tool that your installation may restrict its use with attach-time security. (The external security manager used by your installation defines the security attributes for the CECI and CECS transactions.) If this has been done, and you are not authorized to use the interpreter transaction you select, you will not be able to initiate the transaction.

Part 7. Appendix

Appendix. CICS commands and their equivalent obsolete macros 465

Appendix. CICS commands and their equivalent obsolete macros

This appendix provides a list of the CICS commands available to the application programmer. These commands replace the CICS macro instructions, which are now obsolete. This list gives a command that performs the same function as the obsolete macro instruction. Command options may have different defaults or functions from macro-level operands having similar names. Some CICS macros do not have an equivalent command; for example, there is only one CICS built-in function that can be invoked by a command.

Macro	Command		
		TYPE=RECEIVE	ISSUE RECEIVE
		TYPE=REPLACE	ISSUE REPLACE
DFHBFTCA	-	TYPE=SEND	ISSUE SEND
		TYPE=WAIT	ISSUE WAIT
DFHBIF			
TYPE=DEEDIT	BIF DEEDIT	DFHEMP	
TYPE=PHONETIC ¹		TYPE=ENTRY	ENTER TRACEID MONITOR
DFHBMS		DFHFC	
TYPE=CHECK	RESP option	TYPE=CHECK	RESP option
TYPE=IN	RECEIVE MAP	TYPE=DELETE	DELETE RIDFLD
TYPE=MAP	RECEIVE MAP FROM	(DL/I types)	-
TYPE=OUT	SEND TEXT	TYPE=ESETL	ENDBR
TYPE=OUT,MAP=	SEND MAP	TYPE=GET	READ
TYPE=PAGEBLD	SEND MAP ACCUM	TYPE=GET,	
TYPE=PAGEOUT	SEND PAGE	TYPOPER=UPDATE	READ UPDATE
TYPE=PURGE	PURGE MESSAGE	TYPE=GETAREA	-
TYPE=RETURN	SEND{MAP TEXT} SET	TYPE=GETNEXT	READNEXT
TYPE=ROUTE	ROUTE	TYPE=GETPREV	READPREV
TYPE=STORE	SEND{MAP TEXT} PAGING	TYPE=PUT,	
TYPE=TEXTBLD	SEND TEXT ACCUM	TYPOPER=DELETE	DELETE
		TYPE=PUT,	
DFHDC		TYPOPER=NEWREC	WRITE
TYPE=CICS	DUMP TABLES	TYPE=PUT,	
TYPE=COMPLETE	DUMP COMPLETE	TYPOPER=UPDATE	REWRITE
TYPE=PARTIAL		TYPE=RELEASE	UNLOCK
LIST=PROGRAM	DUMP PROGRAM	TYPE=RESETL	RESETBR
LIST=TERMINAL	DUMP TERMINAL	TYPE=SETL	STARTBR
LIST=TRANSACTION	DUMP STORAGE	DFHIC	
LIST=SEGMENT	DUMP FROM	TYPE=CANCEL	CANCEL
TYPE=TRANSACTION	DUMP[TASK]	TYPE=CHECK	RESP option
DFHDI		TYPE=GET	RETRIEVE
TYPE=ABORT	ISSUE ABORT	TYPE=GETIME	ASKTIME
TYPE=ADD	ISSUE ADD	TYPE=INITIATE	START
TYPE=CHECK	RESP option	TYPE=POST	POST
TYPE=END	ISSUE END	TYPE=PUT	START FROM
TYPE=ERASE	ISSUE ERASE	TYPE=RETRY	RETRIEVE
TYPE=NOTE	ISSUE NOTE	TYPE=WAIT	DELAY
TYPE=QUERY	ISSUE QUERY		

DFHJC

TYPE=CHECK	RESP option
TYPE=GETJCA	-
TYPE=PUT	JOURNALNUM WAIT
TYPE=WAIT	WAIT JOURNALNUM
TYPE=WRITE	WRITE JOURNALNUM
TYPE=GETB	-
TYPE=GETF	-
TYPE=OPEN,INPUT	-
TYPE=OPEN,OUTPUT	SET JOURNALNUM OPENOUT
TYPE=CLOSE	SET JOURNALNUM CLOSE
TYPE=NOTE	-
TYPE=POINT	-

DFHKC

TYPE=ATTACH	-
TYPE=CHAP	CHANGE TASK PRIORITY
TYPE=DEQ	DEQ
TYPE=ENQ	ENQ
TYPE=NOPURGE	-
TYPE=PURGE	-
TYPE=WAIT	SUSPEND
TYPE=WAIT,DCI=LIST	WAITCICS or WAIT EXTERNAL ⁴
TYPE=WAIT,ECADDR	WAIT EVENT

DFHKP

TYPE=CHECK	-
TYPE=RTBOCTL	-
TYPE=RTBODATA	-
TYPE=RTBOEND	-

DFHMDF**DFHMDI****DFHMDS****DFHPC**

TYPE=ABEND	ABEND
TYPE=CHECK	RESP option
TYPE=COBADDR	-
TYPE=DELETE	RELEASE
TYPE=LINK	LINK
TYPE=LOAD	LOAD
TYPE=RESETXIT	HANDLE ABEND RESET
TYPE=RETURN	RETURN
TYPE=SETXIT	HANDLE ABEND
TYPE=XCTL	XCTL

DFHSC

TYPE=FREE MAIN	FREEMAIN
TYPE=GETMAIN	GETMAIN

DFHSP

TYPE=USER	SYNCPOINT
TYPE=ROLLBACK	SYNCPOINT ROLLBACK

DFHTC

CTYPE=CHECK	RESP option
CTYPE=COMMAND	-
TYPE=CBUFF	SEND CBUFF
TYPE=CONVERSE	CONVERSE
TYPE=COPY	ISSUE COPY
TYPE=DISCONNECT	ISSUE DISCONNECT
TYPE=EODS	ISSUE EODS
TYPE=ERASEAUP	ISSUE ERASEAUP
TYPE=GET	RECEIVE
TYPE=PAGE	-
TYPE=PASSBK	SEND PASSBK
TYPE=PRINT	ISSUE PRINT
TYPE=PROGRAM	ISSUE LOAD
TYPE=PUT	SEND WAIT
TYPE=READ	RECEIVE(WAIT assumed)
TYPE=READB	RECEIVE BUFFER
TYPE=READL	RECEIVE LEAVEKB
TYPE=RESET	ISSUE RESET
TYPE=SIGNAL	WAIT SIGNAL
TYPE=WAIT	WAIT TERMINAL
TYPE=WRITE	SEND
TYPE=WRITEL	SEND LEAVEKB

DFHTD

TYPE=CHECK	RESP option
TYPE=FEOV	-
TYPE=GET	READQ TD
TYPE=PURGE	DELETEQ TD
TYPE=PUT	WRITEQ TD

DFHTR

TYPE=ENTRY	ENTER
TYPE=OFF ²	SET TRACEDEST
	SET TRACEFLAG
TYPE=ON	SET TRACETYPE

DFHTS

TYPE=CHECK	RESP option
TYPE=GET ³	READQ TS
TYPE=GETQ	READQ TS
TYPE=PURGE	DELETEQ TS
TYPE=PUT ³	WRITEQ TS
TYPE=PUTQ	WRITEQ TS
TYPE=RELEASE	DELETEQ TS

Language	Command
COBOL	CALL 'DFHPHN' USING lang name phon.
C	
	In the program prior to the main()[statement, code:
	#pragma linkage(DFHPHN,OS) void DFHPHN();
	In the program AFTER the main(){ statement, code:
PL/I	DFHPHN(lang,name,phon);
Assembler	CALL DFHPHN (lang,name,phon); CALL DFHPHN,(lang,name,phon)

Notes

¹ You can no longer use this macro to perform phonetic conversion of 16-character names. For online code, you must supply the name to be converted as input to DFHPHN. DFHPHN is supplied in the VSE sublibrary PRD1.BASE. The result is a 4-byte phonetic equivalent. The general form of this routine is as follows: The “lang” is a symbolic reference to a 1-byte argument indicating the programming language being used. This is X'F0' for COBOL, C, and assembler, and X'F1' for PL/I. If an error occurs during the processing of this request, then X'50' is returned in this argument; if no error occurs, X'00' is returned. This argument must be reset to indicate the programming language, before it can be reused. The “name” is a symbolic reference to the field that contains the 16-character name you want to convert. The “phon” is a symbolic reference to the field that contains the 4-byte phonetic equivalent of the name passed from the subroutine to the calling program.

The phonetic code conversion subroutine (DFHPHN) also assists you to load and access data sets offline. The steps in creating such a data set would typically be:

1. Create the keys
 - a. Read a record from the source data set
 - b. Generate the “phon” using a call to the DFHPHN subroutine
 - c. Write the record to a temporary sequential data set
2. Sort the temporary data set on the 4-byte phonetic code
3. Load the key-sequenced VSAM data set
 - a. Read the sorted temporary data set
 - b. Write to the keyed data set

² TYPE=OFF and TYPE=ON have been replaced by SET TRACEDEST, SET TRACEFLAG, and SET TRACETYPE.

³ Because single units of information cannot be handled by the command-level interface, data stored by a DFHTS TYPE=PUT macro cannot be retrieved by a READQ TS command or be deleted by a DELETEQ TS command. Conversely, data stored by a WRITEQ TS command cannot be retrieved by a DFHTS TYPE=GET macro.

⁴ The choice between WAIT EXTERNAL and WAITCICS depends on whether any of the ECBs in the ECBLIST are hand-posted. If any ECBs are hand-posted, WAITCICS must be used, otherwise WAIT EXTERNAL should be used.

Bibliography

CICS Transaction Server for VSE/ESA Release 1 library

Evaluation and planning	
<i>Release Guide</i>	GC33-1645
<i>Migration Guide</i>	GC33-1646
<i>Report Controller Planning Guide</i>	GC33-1941
General	
<i>Master Index</i>	SC33-1648
<i>Trace Entries</i>	SC34-5556
<i>User's Handbook</i>	SC34-5555
<i>Glossary (softcopy only)</i>	GC33-1649
Administration	
<i>System Definition Guide</i>	SC33-1651
<i>Customization Guide</i>	SC33-1652
<i>Resource Definition Guide</i>	SC33-1653
<i>Operations and Utilities Guide</i>	SC33-1654
<i>CICS-Supplied Transactions</i>	SC33-1655
Programming	
<i>Application Programming Guide</i>	SC33-1657
<i>Application Programming Reference</i>	SC33-1658
<i>Sample Applications Guide</i>	SC33-1713
<i>Application Migration Aid Guide</i>	SC33-1943
<i>System Programming Reference</i>	SC33-1659
<i>Distributed Transaction Programming Guide</i>	SC33-1661
<i>Front End Programming Interface User's Guide</i>	SC33-1662
Diagnosis	
<i>Problem Determination Guide</i>	GC33-1663
<i>Messages and Codes Vol 3 (softcopy only)</i>	SC33-6799
<i>Diagnosis Reference</i>	LY33-6085
<i>Data Areas</i>	LY33-6086
<i>Supplementary Data Areas</i>	LY33-6087
Communication	
<i>Intercommunication Guide</i>	SC33-1665
<i>CICS Family: Interproduct Communication</i>	SC33-0824
<i>CICS Family: Communicating from CICS on System/390</i>	SC33-1697
Special topics	
<i>Recovery and Restart Guide</i>	SC33-1666
<i>Performance Guide</i>	SC33-1667
<i>Shared Data Tables Guide</i>	SC33-1668
<i>Security Guide</i>	SC33-1942
<i>External CICS Interface</i>	SC33-1669
<i>XRF Guide</i>	SC33-1671
<i>Report Controller User's Guide</i>	GC33-1940
CICS Clients	
<i>CICS Clients: Administration</i>	SC33-1792
<i>CICS Universal Clients Version 3 for OS/2: Administration</i>	SC34-5450
<i>CICS Universal Clients Version 3 for Windows: Administration</i>	SC34-5449
<i>CICS Universal Clients Version 3 for AIX: Administration</i>	SC34-5348
<i>CICS Universal Clients Version 3 for Solaris: Administration</i>	SC34-5451
<i>CICS Family: OO programming in C++ for CICS Clients</i>	SC33-1923
<i>CICS Family: OO programming in BASIC for CICS Clients</i>	SC33-1671
<i>CICS Family: Client/Server Programming</i>	SC33-1435
<i>CICS Transaction Gateway Version 3: Administration</i>	SC34-5448

Books from VSE/ESA 2.4 base program libraries

VSE/ESA Version 2 Release 4

Book title	Order number
Administration	SC33-6705
Diagnosis Tools	SC33-6614
Extended Addressability	SC33-6621
Guide for Solving Problems	SC33-6710
Guide to System Functions	SC33-6711
Installation	SC33-6704
Licensed Program Specification	GC33-6700
Messages and Codes Volume 1	SC33-6796
Messages and Codes Volume 2	SC33-6798
Messages and Codes Volume 3	SC33-6799
Networking Support	SC33-6708
Operation	SC33-6706
Planning	SC33-6703
Programming and Workstation Guide	SC33-6709
System Control Statements	SC33-6713
System Macro Reference	SC33-6716
System Macro User's Guide	SC33-6715
System Upgrade and Service	SC33-6702
System Utilities	SC33-6717
TCP/IP User's Guide	SC33-6601
Turbo Dispatcher Guide and Reference	SC33-6797
Unattended Node Support	SC33-6712

High-Level Assembler Language (HLASM)

Book title	Order number
General Information	GC26-8261
Installation and Customization Guide	SC26-8263
Language Reference	SC26-8265
Programmer's Guide	SC26-8264

Language Environment for VSE/ESA (LE/VSE)

Book title	Order number
C Run-Time Library Reference	SC33-6689
C Run-Time Programming Guide	SC33-6688
Concepts Guide	GC33-6680
Debug Tool for VSE/ESA Fact Sheet	GC26-8925
Debug Tool for VSE/ESA Installation and Customization Guide	SC26-8798
Debug Tool for VSE/ESA User's Guide and Reference	SC26-8797
Debugging Guide and Run-Time Messages	SC33-6681
Diagnosis Guide	SC26-8060
Fact Sheet	GC33-6679
Installation and Customization Guide	SC33-6682
LE/VSE Enhancements	SC33-6778
Licensed Program Specification	GC33-6683
Programming Guide	SC33-6684
Programming Reference	SC33-6685
Run-Time Migration Guide	SC33-6687
Writing Interlanguage Communication Applications	SC33-6686

VSE/ICCF

Book title	Order number
Administration and Operations	SC33-6738
User's Guide	SC33-6739

VSE/POWER

Book title	Order number
Administration and Operation	SC33-6733
Application Programming	SC33-6736
Networking Guide	SC33-6735
Remote Job Entry User's Guide	SC33-6734

VSE/VSAM

Book title	Order number
Commands	SC33-6731
User's Guide and Application Programming	SC33-6732

VTAM for VSE/ESA

Book title	Order number
Customization	LY43-0063
Diagnosis	LY43-0065
Data Areas	LY43-0104
Messages and Codes	SC31-6493
Migration Guide	GC31-8072
Network Implementation Guide	SC31-6494
Operation	SC31-6495
Overview	GC31-8114
Programming	SC31-6496
Programming for LU6.2	SC31-6497
Release Guide	GC31-8090
Resource Definition Reference	SC31-6498

Books from VSE/ESA 2.4 optional program libraries

C for VSE/ESA (C/VSE)

Book title	Order number
C Run-Time Library Reference	SC33-6689
C Run-Time Programming Guide	SC33-6688
Diagnosis Guide	GC09-2426
Installation and Customization Guide	GC09-2422
Language Reference	SC09-2425
Licensed Program Specification	GC09-2421
Migration Guide	SC09-2423
User's Guide	SC09-2424

COBOL for VSE/ESA (COBOL/VSE)

Book title	Order number
Debug Tool for VSE/ESA Fact Sheet	GC26-8925
Debug Tool for VSE/ESA Installation and Customization Guide	SC26-8798
Debug Tool for VSE/ESA User's Guide and Reference	SC26-8797
Diagnosis Guide	SC26-8528
General Information	GC26-8068
Installation and Customization Guide	SC26-8071
Language Reference	SC26-8073
Licensed Program Specifications	GC26-8069
Migration Guide	GC26-8070
Migrating VSE Applications To Advanced COBOL	GC26-8349
Programming Guide	SC26-8072

DB2 Server for VSE

Book title	Order number
Application Programming	SC09-2393
Database Administration	GC09-2389
Installation	GC09-2391
Interactive SQL Guide and Reference	SC09-2410
Operation	SC09-2401
Overview	GC08-2386
System Administration	GC09-2406

DL/I VSE

Book title	Order number
Application and Database Design	SH24-5022
Application Programming: CALL and RQDLI Interface	SH12-5411
Application Programming: High-Level Programming Interface	SH24-5009
Database Administration	SH24-5011
Diagnostic Guide	SH24-5002
General Information	GH20-1246
Guide for New Users	SH24-5001
Interactive Resource Definition and Utilities	SH24-5029
Library Guide and Master Index	GH24-5008
Licensed Program Specifications	GH24-5031
Low-level Code and Continuity Check Feature	SH20-9046
Library Guide and Master Index	GH24-5008
Messages and Codes	SH12-5414
Recovery and Restart Guide	SH24-5030
Reference Summary: CALL Program Interface	SX24-5103
Reference Summary: System Programming	SX24-5104
Reference Summary: HLPI Interface	SX24-5120
Release Guide	SC33-6211

PL/I for VSE/ESA (PL/I VSE)

Book title	Order number
Compile Time Messages and Codes	SC26-8059
Debug Tool For VSE/ESA User's Guide and Reference	SC26-8797
Diagnosis Guide	SC26-8058
Installation and Customization Guide	SC26-8057
Language Reference	SC26-8054
Licensed Program Specifications	GC26-8055
Migration Guide	SC26-8056
Programming Guide	SC26-8053
Reference Summary	SX26-3836

Screen Definition Facility II (SDF II)

Book title	Order number
VSE Administrator's Guide	SH12-6311
VSE General Introduction	SH12-6315
VSE Primer for CICS/BMS Programs	SH12-6313
VSE Run-Time Services	SH12-6312

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:

IBM World Trade Asia Corporation
Licensing
2-31 Roppongi 3-chome, Minato-ku
Tokyo 106, Japan

The following paragraph does not apply in 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 may not apply to you.

This publication 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.

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 United Kingdom Laboratories, MP151, Hursley Park, Winchester, Hampshire, England, SO21 2JN. 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 Programming License Agreement, or any equivalent agreement between us.

Trademarks and service marks

The following terms, used in this publication, are trademarks or service marks of IBM Corporation in the United States or other countries:

AIX	DATABASE 2	OS/400
CICS	DB2	OS/390
CICS/ESA	ESA/390	System 36
CICS/OS2	GDDM	System/38
CICS/VSE	IBM	VSE/ESA
CICS/400	IMS	VTAM
CICS/6000	OS/2	

Other company, product, and service names may be the trademarks or service marks of others.

Index

Special Characters

/string, FIND 444

Numerics

10/63 magnetic slot reader 314

31-bit

addressing

restrictions for 22

mode transaction 383

3262 printer 341

3270 family 207—224

attention keys 221

attributes, extended 215

base color 214

buffer 210

color, base 214

data stream 207

data stream orders 215

data stream, outbound 218

display 323

display characteristics 212

display data 212

emulating 209

extended attributes 215

field attributes 213

field format, inbound 223

fields 212

inbound field format 223

input from 221

intensity 214

MDT 214

modified data tag 214

of terminals 207

orders in data stream 215

outbound data stream 218

protection 213

reading from 222

screen fields 208, 259

terminal, writing to 211

unformatted mode 224

write commands 211

write control character 212

writing to terminal 211

3270 printers 342

BMS page size, 348

options for 343

3270 terminal

PC as 210

3289 printer 341

3290 display 304, 307

character size 307

3601 logical unit 312

3770 batch data interchange logical unit 312

3790 batch data interchange logical unit 312

A

ABEND command 156

abend exit

creating program-level 156

facility 155

program 156

routine 156

ABEND USER TASK 427

ABEND, REGISTERS AT 429

abends 141, 146, 460

abnormal termination recovery 155

about to execute command 452

access

method support 319

register support 46

sequential data set 74

to resources, controlling sequence of 374

to system information 153

EXEC interface block (EIB) 153

accessing

data sets from CICS application programs 172

user-related information with EXEC CICS ASSIGN

command 129

ACCUM option 242, 277, 282, 344

ACK 323

positive acknowledgment 323

acquiring storage for maps 242

active partition

determining 309

ACTPARTN option 278, 309

adding

records 178

(WRITE MASSINSERT command),

sequential 179

to DAM data sets 191

to ESDS 178

to KSDS 178

to RRDS 178

additional displays 455

ADDRESS

command 153

COMMAREA command 44, 380

register, COBOL/VSE 19

addresses

protecting CICS from being passed invalid 400

- affinity 91
 - among transactions 91
 - intertransaction 91
 - transaction groups 113
 - transaction-system 92
- AFTER option 369
- AID 222
- AID command, HANDLE 261
- ALARM option 247
- ALLOCATE command 63
 - inhibit wait, NOSUSPEND option 63
- alternate
 - facility 200
 - index 169
 - indexes, VSAM 169
 - key 169
- ALTPAGE value 282
- ANSI85
 - option 9
 - standards, COBOL/VSE 30, 37
- APAK transaction 347
- APCG 52
- API
 - application programming interface 153
 - considerations, CICS 352
 - for terminals, CICS 203
 - restrictions for user-maintained data tables 171
 - subset for DPL 129
- APOST option 9
- application
 - and system table entries, preparing 415
 - CICS-key 398
 - errors
 - protecting CICS from 139
 - generators, intertransaction affinities caused by 113
 - performance
 - monitoring 162
 - programming interface (API) 153
 - selecting and defining storage key for 394
 - selecting execution key for 393
 - testing 415
 - user-key 397
 - using VSE macros from within 18
- application programs
 - accessing data sets from CICS 172
 - assembler language considerations 46
 - asynchronous processing 132
 - C considerations 38
 - design considerations 51
 - designing efficient 51
 - distributed program link 121
 - distributed transaction processing 132
 - function shipping 120
 - intercommunication considerations 119
 - logical levels 378
 - PL/I considerations 44
 - application programs (*continued*)
 - preparing to run 3
 - program size 52
 - program structure 51
 - testing 139, 415
 - transaction routing 120
 - translation 3—16
 - writing 3
 - area
 - dynamic storage 52
 - expanded 455
 - passed to CICS commands, lengths of 67
 - arguments
 - passing values as 41
 - ASIS option 260
 - ASKTIME command 369
 - assembler 157
 - use of DFHRESP in 143
 - assembler language 5, 68, 433
 - applications 57
 - CALL statement 47
 - DFHECALL macro 68
 - programming techniques 46
 - restrictions 46
 - assembling map 237
 - assembly 3
 - TYPE=DSECT 238
 - ASSIGN command 153, 286, 301, 330, 331, 347
 - accessing user-related information with EXEC CICS 129
 - DESTCOUNT option 301
 - MAPCOLUMN option 286
 - MAPHEIGHT option 286
 - MAPLINE option 286
 - MAPWIDTH option 286
 - MSR option 314
 - options 331
 - PAGENUM option 301
 - ASSIGN options
 - for cumulative processing 286
 - for partitions 310
 - asynchronous
 - journal output 136
 - processing 119, 132
 - AT option 369
 - ATI 200, 321, 405
 - ATNI 121
 - attention
 - fields 317
 - identifiers 260
 - and exception conditions, new 311
 - keys 221
 - ATTENTION key 323
 - attribute
 - (SA) order, set 218
 - value definitions 246

- attributes
 - changing 246
 - extended 215
- automatic
 - task initiation 200, 321
 - transaction initiation 405
- AUTOPAGE option 280
- auxiliary
 - storage
 - temporary data 407
 - temporary storage 69
 - trace 63
- avoid
 - transaction deadlocks 181
 - WAIT option on BTAM 81
- AZI6 121

B

- BACK
 - FULL, PF10 SCROLL 444
 - HALF, PF7 SCROLL 444
- backout of resources 137
- base color 214
- base locator linkage (BLL)
 - See BLL (base locator linkage)
- BASE option 243
- based addressing 19
- basic
 - mapping support 225
 - telecommunications access method (BTAM) 319
- basic mapping support
 - See BMS
- batch compilation, COBOL/VSE (ANSI85) 31
 - ANSI85 standards
 - nested programs 32
- batch data interchange 338
 - definite response 340
 - DEFRESP option 340
 - destination identification 340
 - ISSUE WAIT command 340
 - NOWAIT option 340
- BDI 342
- blank
 - ANSI85 standards
 - sequence numbers 30
 - fields 76
 - do not send to screen 76
 - lines 345
 - lines, COBOL/VSE (ANSI85) 30
- blanks, use nulls instead of 78
- BLL (base locator linkage)
- block
 - data 276
 - reference subfield 189
 - references 189

- blocked-data set 170
- BMS 225—318, 342
 - assembling map 237
 - assembly, TYPE=DSECT 238
 - BMS support levels 226
 - complex fields 239
 - composite fields 239
 - considerations 75
 - copy facility 80
 - creating map 231
 - cursor
 - finding 261
 - position 252
 - data streams 75
 - data, moving to map 244
 - DFHMDI macro 231
 - DFHMDI macro 231, 234
 - DFHMSD macro 231, 234
 - display, receiving data from 253
 - EOC condition 266
 - fields 230
 - complex 239
 - composite 239
 - group 239
 - repeated 240
 - finding cursor 261
 - full 227
 - group field 239
 - GRPNAME option 239
 - initializing output map 244
 - input from detectable fields 317
 - introduction to 225
 - invalid data 253
 - link-edit 238
 - logical messages 277
 - macro 231
 - macros, rules for writing 235
 - map 228
 - assembling 237
 - creating 231
 - moving data to 244
 - physical 238
 - sets 239
 - symbolic 238
 - mapping input data 257
 - maps 75, 76, 77, 80
 - preparing 16
 - storage 242
 - MAPSET resource definition 238
 - MDT 259
 - message lengths, reducing 288
 - minimizing pathlength 287
 - minimum 226
 - modified data tag (MDT) 75, 259
 - moving data to map 244
 - multimap screens 80

BMS (continued)

- OCCURS option 240
- output example 228
- output map, initializing 244
- page building operations 78
- page routing operations 78
- page size, 3270 printers 348
- page-copy facility, using 80
- pathlength, minimizing 287
- performance considerations 287
- physical map 238
- preparing 16
- PROGRAM resource definition 238
- RECEIVE commands, partition options for 309
- receiving data from display 253
- reducing message lengths 288
- repeated fields 240
- routing 351
 - printing with 351
- rules for writing macros 235
- screen copy 356
- SEND commands, partition options for 309
- SEND MAP command 241
- SEND options, other 248
- standard 227
- storage for maps 242
- support across platforms 227
- support for other special hardware 314
- support levels 226
- symbolic map 238
- TERMINAL NOT SUPPORTED UNDER 298
- terminals supported 227
- TYPE=DSECT assembly 238
- upper case translation 260
- using GDDM and 252

body 421, 443, 453

BOTTOM

- command, CEBR transaction 444
- PF5 VIEW 443

BRACKET option 336

bracket protocol

- LAST option 336
- preventing interruptions 336

breaks, page 283

brightness 316

browse

- ending 175
- operations
 - DAM 190
 - simultaneous 176
- PF3 TERMINATE 443
- TEMP STORAGE option, CEDF 428
- transaction (CEBR) 441

browsing 74

- data sets 74
- DELAY 74

browsing (continued)

- records 173
- records from DAM data sets 190
- sequential reading 173
- SUSPEND 74
- through ESDS 175
- through KSDS 174
- through RRDS 175
- using path 175

BTAM 319

- avoid WAIT option 81
- basic telecommunications access method 319

buffer

- 3270 210
- address order, set 217

BUILDCHAIN 333

C

C

- and PL/I 433
- compilers supported 38
- considerations 38
- data declarations 39
- example of exception handling in 143
- for using storage under CICS 401
- LENGTH option default 67
- naming EIB fields 39
- programming techniques 38
- programs and CICS storage 401
- restrictions 38
- use of RESP and DFHRESP in 143

CALL statement

- assembler language 47
- in COBOL/VSE 25—29

calling subprograms from COBOL/VSE 25

CALLOC (C function) 401

CANCEL command 369

CANCEL REQID commands

- using 111
- using DELAY and 109
- using START and 108

CARD option 340

CBLCARD option 9

CEBR transaction 441—447

- body 443
- BOTTOM command 444
- browse transaction 441
- CEBR initiation 441
- COLUMN command 444
- command area 443
- commands 444
- displays 442
- FIND command 444
- GET command 445
- header 442

- CEBR transaction (*continued*)
 - how to use 441
 - initiation 441
 - LINE command 445
 - message line 443
 - options on function keys 443
 - PURGE command 445
 - PUT command 445
 - QUEUE command 446
 - security considerations 447
 - SYSID command 446
 - temporary storage browse 441
 - TERMINAL command 446
 - TOP command 446
 - transient data 446
 - with transient data, using 446
- CECI transaction
 - about to execute command 452
 - ampersand (&) 456
 - body 453
 - command execution complete 452
 - command input 450
 - command input line 450
 - command line 450
 - command-level interpreter 449
 - EIB 457
 - ENTER key 454
 - expanded area 455
 - how to use 449
 - information area 453
 - introduction 449
 - INVOKE 428
 - invoking 449
 - making changes 459
 - message line 454
 - messages display 458
 - options on function keys 454
 - PF key values area 454
 - runs, how 460
 - screen layout 450
 - security considerations 462
 - sessions 460
 - status area 451
 - terminal sharing 461
 - variables 455
- CECS transaction 449
- CEDF transaction 419—440
 - body 421
 - header 421
- chaining 324
 - input 324
 - input data 333
 - output data 334
- changes
 - making 459
 - overtyping to make 437
- changing attributes 246
- character size, 3290 307
- characteristics
 - display 212
 - of CICS printer, determining 347
 - setting display 244
- checking
 - command syntax 451
 - pseudoconversational programs 435
- checkout, program 419
- choosing between pseudoconversational and conversational design 53
- CICS
 - API considerations 352
 - APIs for terminals 203
 - application programs
 - accessing data sets from 172
 - techniques used by to pass data 92
 - commands
 - and equivalent obsolete macros 465
 - lengths of areas passed to 67
 - data tables 170
 - dump utility program 163
 - exception handling
 - default 141
 - modifying default 146
 - for VSE interface to POWER 357
 - interface (EXCI), external 133
 - option 9
 - print key 355
 - printers 341, 349
 - determining characteristics of 347
 - formatting for 341
 - non-3270 346
 - programs
 - writing 3
 - protecting from application errors 139
 - RETURN IMMEDIATE, EXEC 322
 - storage control 391
 - storage protection 392
 - storage, C programs and 401
 - support for printing 341
 - system
 - mixing DPL and DTP to same 128
 - mixing DPL and function shipping to same 128
 - terms, basic 199
- CICS terms 199
- CICS-defined resource identifiers 410
- CICS-key
 - applications 398
 - storage 392
- CICS-maintained data tables
 - application programming 171
- CICS-value data area
 - See CVDA

- CICSDATAKEY option 65, 395
- CLEAR key 89, 311
- CLEAR PARTITION AID value 311
- CLEAR PARTITION key 311
- client
 - program
 - and server program, sharing resources
 - between 128
 - exception conditions returned to 130
 - region 121
 - task, issuing multiple distributed program links from
 - same 128
- CNOTCOMPL option 334
- COBOL 19
 - and PL/I, use of RESP and DFHRESP in 143
 - compilers supported 19
 - DL/I CALL interface 23
 - example of exception handling in 145
 - programming techniques 19
 - restrictions 19, 57
- COBOL/VSE 38
 - addressing CICS data areas 19
 - ANSI85 standards 30—38
 - CALL statement 25—29
 - calling subprograms 25—29
 - compiler options not used under CICS 21
 - considerations 19
 - programming restrictions summary 37
 - run unit 24, 29
 - USE FOR DEBUGGING facility 20
- COBOL2 option 9
- COBOL3 option 9
- coding
 - considerations for sequential terminals 337
 - server program (DPL) 127
- color, base 214
- COLUMN
 - command, CEBR transaction 444
 - nnnn 444
- comma as delimiter
 - COBOL/VSE 36
- command
 - about to execute 452
 - area 443
 - control options on SEND MAP 247
 - execution complete 452
 - GETMAIN 395
 - GETMAIN SHARED 71
 - language translator 5
 - line 450
 - per screen, use only one physical EXEC CICS
 - SEND 81
 - SYNCPOINT 137
 - syntax check 451
 - use EXEC CICS CONVERSE 82
 - use of EXEC CICS HANDLE ABEND 150
- command language translator
 - See translation
- command-level interpreter
 - See CECI transaction
- commands
 - and equivalent obsolete macros, CICS 465
 - CEBR 444
 - contained within macros and COPY code 47
 - control 324
 - data transmission 320
 - lengths of areas passed to CICS 67
 - limit use of DEFRESP option on EXEC CICS
 - SEND 82
 - program control 461
 - saving 461
 - supported in C 38
 - system programming 153
- COMMAREA 53, 56, 66, 94, 379
 - in EXEC CICS LINK and EXEC CICS XCTL
 - commands 66
 - in EXEC CICS RETURN commands 88
 - option 88, 377, 381
 - passing data via a 395
- common programming interface communications
 - See CPI-C
- common work area (CWA) 83
 - protecting 84
 - using 97
- communications
 - common programming interface (CPI) 132
- comparison, data 306
- compilation 3
- compiler 5
- compilers supported
 - C 38
 - COBOL 19
 - PL/I 44
- complex fields 239
- composite
 - fields 239
 - screen, input from 286
- compress data sent to screen 78
- computers, personal 209
- condition
 - EOC 266
 - exception 141, 460
 - storage protection exception 400
- CONDITIONS, STOP 430
- CONNECT PROCESS command 132
- considerations
 - recovery 135
 - security 440, 462
- console
 - writing message to system 62
- CONSOLE option 340

- contention
 - for resources 53
 - for terminal 321
- CONTINUE 428
- control
 - character, write 212
 - CICS storage 391
 - commands 324
 - exclusive of DAM 192
 - of resources, exclusive 59
 - of VSAM blocks 186
 - operational 61
 - options on SEND MAP command 247
 - security 409
 - storage 391
 - typical uses of temporary storage 408
- controlling sequence of access to resources 374
- conversation
 - partner 321
 - with message routing, interleaving 302
- conversational
 - programming 53, 82
- conversational design
 - choosing between pseudoconversational and 53
- CONVERSE command 82, 320, 336
 - use EXEC CICS 82
- COPY code
 - commands contained within 47
- copy facility
 - BMS 80
- CPI communications
 - See CPI-C
- CPI-C 132
 - references 3
- CQRY transaction 201
- creating
 - map 231
- CSPG transaction 279, 280, 351, 356
- CSPP transaction 202
- CTLCHAR option 344
- cumulative processing, ASSIGN options for 286
- CURRENT DISPLAY 428
- cursor
 - finding 261
 - position 252
 - positioning 252
 - positioning, symbolic 253
- CURSOR option 242, 252, 278
- cursor-detectable field 316
- CVDA 410
- CWA
 - common work area 83
 - protecting 84
- CWAKEY parameter 84, 394

D

- DAM 73
 - browsing operations 190
 - considerations, file control 189
 - data set operations, VSAM and 72
 - data sets 73, 169, 189
 - data sets, adding records to 191
 - data sets, browsing records from 190
 - data sets, updating records from 190
 - exclusive control 192
 - updating operations 190
- data
 - across transactions, sharing 83
 - block 276
 - chaining 333
 - communication, introduction to 199
 - comparison 306
 - definition 56
 - entry 305
 - from display, receiving 253
 - initialization 56
 - interchange block (DIB) 5
 - intrapartition transient 70
 - modified 259
 - moving to map 244
 - operations 71
 - other ways of passing 383
 - passing to other program 379
 - reading from display 257
 - records 68
 - sent to screen, compress 78
 - storing within transaction 64
 - stream
 - orders in 215
 - sample, outbound 218
 - sent to terminal, length of 75
 - tables, CICS 170
 - techniques used by CICS application programs to pass 92
 - to logical device component, sending 313
 - transmission commands 320
 - types in EIB fields 40
 - via a COMMAREA, passing 395
 - within a transaction, storing 64
- data entry operations
 - design to reduce line traffic 78
- data sets 72
 - access, sequential 74
 - adding records to DAM 191
 - batch data interchange 338
 - blocked 170
 - browsing 74
 - browsing records from DAM 190
 - DAM 73, 169, 189
 - definitions 68

- data sets (*continued*)
 - empty 168
 - entry-sequenced (ESDS) 168
 - from CICS application programs, accessing 172
 - key-sequenced (KSDS) 167
 - KEYLENGTH option for remote 183
 - logging 74
 - operations, VSAM and DAM 72
 - relative record (RRDS) 168
 - selection by named 340
 - sequential 74
 - user 71
 - VSAM 72, 167, 185
 - your own 71
- data streams
 - compressing 78
 - inbound 76
 - RA order 78
 - repeat-to-address orders (SBA) 78
 - SBA order 78
 - set buffer address order 78
- DATA(24) 22
- DATA(31) 22
- DATABASE 2 (DB2) 193
- databases
 - control 193
 - DATABASE 2 (DB2) 193
 - DB2 193
 - DL/I 193
 - operations, DL/I 71
- DATAONLY option 77, 242, 248
- date field of EIB 153
- DB2 193
- DBCS option 9, 11
- DDS
 - device-dependent support 270
- deadlocks 59
 - avoiding transaction 181
 - prevention 186
 - rules for avoiding 183
- DEBKEY option 191
- deblocking argument subfield 189
- DEBREC option 189, 191
- DEBUG option 10
- debugging 419
 - interactive 413, 463
 - preparing system for 415
- deciding what execution and storage key to specify 397
- default
 - action for condition 141
 - CICS exception handling 141
 - modifying 146
- deferred journal output 136
- defining
 - execution key 393
- defining (*continued*)
 - logical device components 312
 - map 234
 - map set 234
 - storage for maps 242
 - storage key for applications, selecting and 394
 - variables 457
- definite response 340
 - protocol
 - terminal control 335
- definitions, attribute value 246
- DEFRESP option 82, 340
 - on EXEC CICS SEND commands, limit use of 82
 - terminal control 335
- DELAY command 369
 - using 109
- DELETEQ TD command 403
- DELETEQ TS command 407
- deleting
 - groups of records (generic delete) 178
 - records 177
 - single records 177
- delimiters
 - ANSI85 standards
 - symbolic characters defined in program 36
 - comma 36
 - semicolon 36
- delivery conditions 299
- DEQ command 373
 - using 102
- DEQUEUE command 351
- design considerations 119
 - of applications
 - exclusive control of resources 59, 60
 - other suggestions 62
- design implications, recovery 60
- designator character 316
- designing efficient applications 51
- DESTCOUNT option 301
- DESTID option 340
- DESTIDLENG option 340
- destinations
 - extrapartition 404
 - for routed message, specifying 295
 - identification 340
 - indirect 404
 - other than input terminal, sending messages to 79
 - selection and identification 340
 - specified with OPCLASS only 296
- detectable fields 316
 - BMS input from 317
- determining
 - active partition 309
 - characteristics of CICS printer 347
 - how program was invoked 129

- device-dependent
 - maps 269
 - support (DDS)
- devices
 - output considerations for non-3270 267
- DFHBMSCA 246
 - definitions 262
- DFHCOMMAREA 380
- DFHCPLC 133
- DFHDU41 163
- DFHEDF group 415
- DFHEIEND macro 14, 15
- DFHEIENT macro 14, 15
- DFHEIRET macro 10, 13
- DFHEISTG macro 14, 15
- DFHMDF macro 231
 - display characteristics 244
 - DSATTS option 244
 - MAPATTS option 244
- DFHMDI macro 231, 234
- DFHMIRS program 125
- DFHMSD macro 231, 234
 - BASE option 243
 - STORAGE option 243
- DFHMSRCA 314
- DFHPDI macro 306
- DFHPEP program 156
- DFHPSD macro 306
- DFHRESP
 - translator action 31
 - translator function 5, 142
 - use of in assembler 143
 - use of in C 143
 - use of in COBOL and PL/I 143
- DFHSIT, SPOOL operand 357
- DFHURLDS 299
- DFHVALUE 5
 - ANSI85 standards
 - batch compilation 31
 - translator action 31
- DIB
 - data interchange block 5
 - DISPLAY 428
 - DL/I interface block 428
- differences on input 268
- direct reading
 - by way of path 173
 - from ESDS 173
 - from KSDS 173
 - from RRDS 173
 - using READ command 172
- direct terminal 302
- display
 - CECI 450
 - characteristics 212, 244
 - setting 244
- display (*continued*)
 - CURRENT 428
 - data, 3270 212
 - DIB 428
 - EIB 428
 - NEXT 429
 - PREVIOUS 429
 - reading from 257
 - receiving data from 253
 - register, EDF 429
 - REMEMBER 429
 - screens 89
 - printing 354
 - USER 432
 - what does EDF 420
- DISPLAYS
 - additional 455
 - SUPPRESS 431
- disposition options, output 277
- distributed program link 121, 377
 - See also* DPL
 - coding server program 127
 - examples of 124
 - from same client task, issuing multiple 128
 - function, using the 122
 - programming considerations for 128
 - subset, restricting program to 129
- distributed transaction processing 132
 - See also* DTP
- DL/I 60
 - CALL interface 23
 - database operations 71
 - databases 193
 - EXEC DLI interface 193
 - interface block (DIB) 428
 - references 3
 - segment-search area (SSA) 71
 - syncpoints 138
- DLI 10
 - option 10
- DPL 119, 138, 377, 436
 - and DTP to same CICS system, mixing 128
 - and function shipping to same CICS system, mixing 128
 - client region 121
 - coding server program 127
 - COMMAREA option 124
 - distributed program link 121
 - DPL API subset 129
 - exception conditions 130
 - independent syncpoints 125
 - options 123
 - programming considerations 128
 - programs, EDF and 436
 - REMTENAME option 124
 - REMOTESYSTEM option 125

- DPL (*continued*)
 - server program 124
 - server region 121, 125
 - SYSID option 125
 - TRANSID option 125
 - DSA 52
 - DSATTS option 244
 - DTP 119
 - distributed transaction processing 132
 - programs, EDF and 436
 - to same CICS system, mixing DPL and 128
 - dual-screen mode, using EDF in 435
 - DUMP TRANSACTION command 163
 - dumping 163
 - DUPKEY condition 175
 - duration of intertransaction affinities 113
 - dynamic
 - COBOL call 25
 - log 60
 - program 58
 - storage
 - PL/I and 46
 - storage area 52
 - transaction backout program 156
 - transaction routing 383
- E**
- ECBLIST 374
 - EDF 6, 10, 419
 - and DPL programs 436
 - and DTP programs 436
 - and remote transactions 435
 - display 420
 - EDF option 10
 - execution diagnostic facility 419
 - how to use 433
 - in dual-screen mode, using 435
 - in single-screen mode, using 434
 - menu functions 427
 - option 10
 - responses 439
 - restrictions when using 439
 - SESSION, END 428
 - stopping 437
 - where it intercepts program 419
 - efficient applications
 - designing 51
 - EIB 5, 141, 321
 - DISPLAY 428
 - EXEC interface block 153, 457
 - feedback on terminal control operations 332
 - fields
 - data types in 40
 - PF4 454
 - terminal control feedback 332
 - EIBRESP option 150
 - eligible terminals 295
 - empty data sets 168
 - END EDF SESSION 428
 - end-of-data indicator character 337
 - END, PF3 454
 - ENDBR command 175
 - ending
 - browse 175
 - logical message 279
 - ENQ command 63, 373
 - using 102
 - ENQBUSY condition 63
 - ENQUEUE command 351
 - enqueuing
 - in a VSAM file 59
 - VSAM internal 59
 - entries, preparing application and system table 415
 - entry point, trace 161
 - ENTRY SKIPPED 298
 - entry-sequenced data set (ESDS) 168
 - EOC condition 266, 334
 - EODI character 337
 - EODS condition 334
 - EPILOG option 10
 - EQUAL option 173
 - equivalent obsolete macros, CICS commands and 465
 - ERASE option 247, 278, 344
 - ERASEAUP option 247, 278, 289
 - ERRATT option 61
 - errors
 - flagging 263
 - handling 158
 - handling input 263
 - messages 306
 - messages display 458
 - minimizing 139
 - protecting CICS from application 139
 - sending invalid data and other 253
 - ESDS
 - adding to 178
 - browsing through 175
 - direct reading from 173
 - entry-sequenced data set 168
 - establishing partitioning 308
 - events
 - monitoring point 162
 - example
 - BMS output 228
 - input 223
 - input-output 254
 - of exception handling in C 143
 - of exception handling in COBOL 145
 - of passing data with LINK command 384
 - of passing data with RETURN command 387

- exception
 - conditions 460
 - handling 158
 - default CICS 141
 - in C, example of 143
 - in COBOL, example of 145
 - modifying default CICS 146
 - using only C and CICS functions 158
 - trace entry points 161, 162
 - user 162
- exception conditions
 - by in-line code, handling 142
 - dealing with 141
 - description 141
 - for EXEC CICS LINK command 130
 - for globally accessible queues 105, 106
 - HANDLE CONDITION command 146, 147
 - MAPFAIL and other 265
 - new attention identifiers and 311
 - returned to client program 130
 - returned to server program 130
 - storage protection 400
- EXCI
 - CALL 133
 - communications 133
 - external call interface 119
 - external CICS interface 133
 - option 10
- exclusive
 - control of records
 - DAM 192
 - VSAM 186
 - control of resources 59
 - resources 60
- EXEC CICS
 - ADDRESS COMMAREA 44
 - ASSIGN command, accessing user-related information with 129
 - CONVERSE command, use 82
 - HANDLE ABEND command 155
 - HANDLE ABEND command, use of 150
 - HANDLE CONDITION command, use of 147
 - HANDLE CONDITION ERROR command, use of 148
 - IGNORE CONDITION command, use of 149
 - ISSUE COPY 355
 - ISSUE PRINT 355
 - LINK command 25
 - exception conditions for 130
 - LINK commands
 - COMMAREA in 66
 - POP HANDLE command 152
 - PUSH HANDLE command 152
 - RECEIVE, what you get on 323
 - RETURN commands, COMMAREA in 88
 - RETURN IMMEDIATE 322
- EXEC CICS (*continued*)
 - SEND command per screen, use only one physical 81
 - SEND command, on BTAM, avoid WAIT option on 81
 - SEND commands, limit use of DEFRESP option on 82
 - START command, printing with 349
 - XCTL commands, COMMAREA in 66
- EXEC DLI interface 193
- EXEC interface block (EIB) 5, 153, 457
- EXEC interface stubs 68
- EXEC SQL
 - interface 193
- EXECKEY 66, 84
 - parameter 393
- execution
 - complete, command 452
 - key
 - deciding what to specify 397
 - defining 393
 - for applications, selecting 393
- execution diagnostic facility (EDF) 6, 10, 419
- existing screen, send only changed fields to 77
- exit
 - creating program-level abend 156
 - programs
 - program working storage specifically for 395
- expanded area 455
- expiration time
 - specifying 369
- extended addressing 19
 - assembler language considerations 46
 - C considerations 38
 - COBOL considerations 19
 - PL/I considerations 44
 - program size 52
- extended attributes 215
- extensions, SEND TEXT 293
- external call interface (EXCI) 119
- external CICS interface (EXCI) 133
- extrapartition
 - destinations 404
 - queues 404
 - transient data 71, 74

F

- FE option 10
- features, field selection 315
- FEPI
 - Front End Programming Interface 119
 - option 10
 - references 3
- field
 - format, inbound 223

field (*continued*)
 group 239
 level, security protection at record or 409
 selection features 315

fields
 attention 317
 attributes, 3270 213
 blank 76
 BMS 230
 complex 239
 composite 239
 cursor- and pen-detectable 316
 repeated 240
 screen 208
 selection 316
 to an existing screen, send only changed 77
 to screen, do not send blank 76

file control
 command options, review of 179
 DAM considerations 189
 DAM data sets 189
 overview 167
 VSAM considerations 185
 VSAM data sets 185

FIND /string 444
 FIND command, CEBR transaction 444
 FIND, PF6 REPEAT LAST 443

finding
 cursor 261
 out about your terminal 272

flag byte, route list 298

FLAG option 10
 FLAG(i|w|e|s) 10
 flagging errors 263

flip-flop mode 321
 floating maps 282

FMH
 inbound 335
 option 336
 outbound 336

FMHPARM option 278

FOR option 369

format
 inbound field 223
 route list 297

formatted screen, reading from 259

FORMATTIME command 369

formatting
 for CICS printers 341
 for non-CICS printers 352
 other input 266
 outboard 318
 print 338
 text output 288

FORMFEED 346
 option 278, 346

FORWARD
 FULL, PF11 SCROLL 444
 HALF, PF8 SCROLL 444

FREE
 C function 401
 command 336

FREEKB option 247

FREEMAIN command 391

FROM option 180, 242

Front End Programming Interface (FEPI) 119

FRSET
 option 247
 use to reduce inbound traffic 76

full BMS 227

function
 completion, waiting for 340
 keys
 CEBR options on 443
 CECI options on 454
 management headers (FMH)
 description 335
 using 335
 shipping 119, 120
 and TOKEN option 181
 to same CICS system, mixing DPL and 128

function (PF) keys, CEBR transaction 443

functions
 EDF 419
 EDF menu 427

G

GDDM and BMS, using 252

GDS option 11

general programming techniques 55

generic
 delete (deleting groups of records) 178
 key 173

GENERIC option 173, 185

GET
 command, CEBR transaction 445
 xxxx 445

GETMAIN command 57, 65, 71, 391, 395
 CICS DATAKEY option 65, 395
 INITIMG option 57, 391
 NOSUSPEND option 391
 SHARED option 65, 71, 391
 TASK DATAKEY option 65
 USER DATAKEY option 65, 395

GETMAIN requests (MVS)
 GETMAIN command 63

GETMAIN SHARED
 command 71
 storage, using 98

getting storage for mapped input 258

- global
 - relation 114
 - user exits 395
 - using 92
 - variables
 - COBOL/VSE (ANSI85) 36
- globally accessible queues, exception conditions
 - for 105, 106
- GOODNIGHT
 - convention 338
 - transaction 338
- GRAPHIC option 11
- group field 239
- grouping maps into map sets 239
- groups
 - affinity transaction 113
 - of records (generic delete), deleting 178
- GRPNAME option 239
- GTEQ option 173, 185

H

- half-duplex mode 321
- HANDLE ABEND
 - command 142, 146, 155
 - EXEC CICS 155
 - use of EXEC CICS 150
 - PROGRAM 159
- HANDLE AID command 261
- HANDLE CONDITION
 - command 142, 146, 152
 - use of EXEC CICS 147
 - ERROR command 148
 - use of EXEC CICS 148
- handling
 - exception 158
 - exception conditions by in-line code 142
 - input errors 263
 - logical records 334
- hardware
 - BMS support for other special 314
 - print key 356
- header 421, 442
 - and trailer format for text messages 292
- HELP, PF1 443, 454
- HEX, PF2 454
- HEX/CHAR
 - PF2 SWITCH 443
 - SWITCH 432
- HOLD option 377
- HONEOM option 344
- horizontal tabs 347
- how
 - BMS breaks text into lines 291
 - CECI runs 460
 - CICS keeps track of what to do 151

- how (*continued*)
 - tasks get started 200
- how to
 - define partitions 306
 - intervene in program execution 427
 - use CEBR transaction 441
 - use CECI 449
 - use EDF 433

I

- ICVR parameter 416
- identification
 - DAM record 189
 - destination selection and 340
 - message 300
 - record 185, 189
 - VSAM record 185
- identifiers
 - attention 260
 - request 371
- IGNORE CONDITION command 142, 146, 149
 - use of EXEC CICS 149
- IGREQID condition 278
- IMMEDIATE
 - EXEC CICS RETURN 322
 - option 322, 336, 381
- implications, recovery design 60
- improving performance
 - other suggestions for 62
- in-line code, handling exception conditions by 142
- INBFMH condition 334
- inbound
 - data streams 76
 - field format 223
 - FMH 335
 - traffic, use FRSET to reduce 76
- index, alternate 169
- indirect destinations 404
- INFOANA 163
- initializing output map 244
- INITIMG option 57, 391
- input
 - chaining 324
 - data sets 5
 - data, chaining 333
 - differences on 268
 - errors, handling 263
 - example 223
 - formatting other 266
 - from 3270 terminal 221
 - from composite screen 286
 - getting storage for mapped 258
 - map, symbolic 256
 - mapped output after mapped 264
 - processing mapped 262

- input (*continued*)
 - programming simple mapped 257
 - saving good 263
 - SCS 347
 - terminal, sending messages to destinations other than 79
 - unsolicited 201
- input-output example 254
- INPUTMSG 381
 - option 377, 383
 - on RETURN command, using 383
- INQUIRE
 - command 153
 - using 92
 - TERMINAL command 332, 347
- INRTN option 309
- integrity
 - options, limit use of message 82
 - read 178
- intensity 214
- interactive debugging
 - CECI transaction 449
 - CECS transaction 449
 - CEDF transaction 419
- interchange, batch data 338
- intercommunication considerations 119
- interface
 - block (EIB), EXEC 457
 - external CICS (EXCI) 133
 - stubs, EXEC 68
 - to POWER, CICS 357
- interface to VSE/POWER
 - See VSE/POWER
- interleaving conversation with message routing 302
- interregion communication 133
- interruptibility, terminal 61
- interrupting
 - terminal control 323
- interruptions (bracket protocol), preventing 336
- intertransaction affinity 91
 - affinity life times 113
 - affinity transaction groups 113
 - caused by application generators 113
 - detecting 112
 - duration of 113
 - programming techniques 92
 - recommendations 117
 - relations and lifetimes 114
 - global relation 114
 - terminal relation 115
 - userid relation 116
 - safe programming techniques 94
 - COMMAREA 94
 - TCTUA 95
 - scope of 113
 - suspect programming techniques 92—111
 - DELAY and CANCEL REQID commands 109
- intertransaction affinity (*continued*)
 - suspect programming techniques (*continued*)
 - global user exits 92
 - INQUIRE and SET commands 92
 - POST command 111
 - RETRIEVE WAIT and START commands 106
 - START and CANCEL REQID commands 108
 - temporary storage 103
 - transient data 105
 - unsafe programming techniques 97
 - CWA 97
 - using DEQ 102
 - using ENQ 102
 - using LOAD PROGRAM HOLD 98
 - using shared storage 98
 - using task lifetime storage 100
 - using WAIT EVENT 101
- interval control 369
 - cancel interval control command 369
 - DELAY command 369
 - delay processing of task 369
 - expiration time 369
 - POST command 369
 - specifying request identifier 371
 - START command 369
 - starting task 369
- INTERVAL option 369
- INTO option 180
- intrapartition
 - queues 403
 - transient data 70
- introduction to data communication 199
- invalid
 - addresses
 - protecting CICS from being passed 400
 - data and other errors, sending 253
 - LDC MNEMONIC 299
 - TERMINAL IDENTIFIER 298
- INVITE option 321
- INVOKE CECI 428
- invoked, determining how program was 129
- invoking EDF 419
- INVPARTN condition 311
- INVPARTNSET condition 311
- INVREQ condition 278
- IRC 133
- ISCINVREQ 121
- ISSUE ABORT command 339, 340
 - CARD option 340
 - CONSOLE option 340
 - PRINT option 340
 - WPMEDIA1–4 option 340
- ISSUE ADD command 339
- ISSUE COPY command 323, 355
- ISSUE DISCONNECT command 323

ISSUE END command 339, 340
 CARD option 340
 CONSOLE option 340
 PRINT option 340
 WPMEDIA1–4 option 340
 ISSUE ERASE command 323, 339
 ISSUE NOTE command 339
 ISSUE PRINT command 355
 ISSUE QUERY command 339
 ISSUE RECEIVE command 339
 ISSUE REPLACE command 339
 ISSUE SEND command 339, 340
 CARD option 340
 CONSOLE option 340
 PRINT option 340
 WPMEDIA1–4 option 340
 ISSUE WAIT command 339, 340
 CARD option 340
 CONSOLE option 340
 PRINT option 340
 WPMEDIA1–4 option 340

J

journal
 control
 output synchronization 136
 identifier 137
 output synchronization 136
 records 68, 135
 type identifier 137
 journaling 135–137
 JCT 135
 journal control table 135
 JOURNALNUM 137
 JTYPEID 137
 JUSFIRST option 290
 JUSLAST option 290
 JUSTIFY option 290

K

key 185
 key-sequenced data set (KSDS) 167
 KEYLENGTH option for remote data sets 183
 keys
 alternate (secondary) 169
 attention 221
 CEBR options on function 443
 CICS print 355
 generic 173
 hardware print 356
 physical 189
 print 324
 KSDS
 adding to 178

KSDS (*continued*)
 browsing through 174
 direct reading from 173
 key-sequenced data set 167

L

language
 assembler 433
 programming 120
 language considerations
 assembler 46
 assembler language 46
 C 38
 COBOL 19
 PL/I 44
 programming techniques 44
 LAST FIND, PF6 REPEAT 443
 LAST option 248, 336
 bracket protocol 336
 LDC
 logical device component 312
 MNEMONIC, INVALID 299
 option 313
 LDCMNEM option 313
 LDCNUM option 313
 LDCs
 and logical messages 313
 and routing 313
 LEFT, PF9 VIEW 444
 LENGERR condition 324
 LENGTH
 of areas passed to CICS commands 67
 of data stream sent to terminal 75
 option 67, 324
 register, COBOL/VSE 19
 LENGTH option 11
 LENGTHLIST option
 multiple dumps of task-related storage areas 163
 lifetimes, relations and 114
 light pen-detectable field 316
 line
 length on printing 344
 message 443
 nnnn 445
 status 451
 traffic reduction 77
 traffic, design data entry operations to reduce 78
 width options 344
 LINE command
 CEBR transaction 445
 LINECOUNT option 11
 linecount(n) 11
 lines
 blank 345
 how BMS breaks text into 291

- LINK command 25, 52, 65, 377, 378, 381
 - COMMAREA in EXEC CICS 66
 - COMMAREA option 377, 379
 - examples of passing data with 384
 - exception conditions for EXEC CICS 130
 - in COBOL/VSE 25—29
 - INPUTMSG option 377
- LINK PROGRAM 133
- link to program, expecting return 378
- link-edit 3, 5
 - of map 238
- LINKAGE option 11
- LIST omitted, OPCLASS and 296
- LIST option 295
- load
 - module size 52
- LOAD command
 - HOLD option 377
- LOAD PROGRAM HOLD command, using 98
- local copy key 355
- locality of reference 56
- locking of VSAM records in recoverable files, CICS 186
- logging data sets 74
- logical
 - device component
 - sending data to 313
 - device components
 - defining 312
 - device components (LDCs) 312
 - levels, application program 25, 29, 378
 - record presentation 334
 - records, handling 334
 - unit 199
- logical messages 324
 - BMS 277
 - ending 279
 - LDCs and 313
 - partitions and 310
 - recovery 282
 - rules for 278
- logical unit of work (LUW)
 - database operations, control of PSB 72
 - description 54
 - recoverable resources 54
 - syncpoints used 137
- logical units (LUs)
 - facilities for 333
- lookaside 305
- lowercase characters, COBOL/VSE (ANSI85) 31
 - ANSI85 standards
 - REPLACE statement 31
- LU 199
 - facilities for 333
 - type 4
 - batch data interchange 339
 - logical record presentation 334

- LU type 4 device 322
- luname (terminal) relation 115
- LUW (logical unit of work)
 - See logical unit of work (LUW)

M

- macros
 - CICS commands and equivalent obsolete 465
 - commands contained within 47
 - instructions 465
- magnetic slot reader (MSR) 314
 - 10/63 314
- main storage 52
 - temporary 69
 - temporary data 407
- making changes 459
- MALLOC (C function) 401
- map
 - assembling 237
 - BMS 228
 - creating 231
 - defining 234
 - initializing output 244
 - link-edit 238
 - map
 - initializing output 244
 - moving variable data to 244
 - outside 252
 - placement rules 284
 - set
 - defining 234
 - sets
 - grouping maps into 239
 - physical and symbolic 238
 - sets and partition sets 400
 - symbolic input 256
 - symbolic output 256
- MAPATTS option 244
- MAPCOLUMN option 286
- MAPFAIL
 - and other exceptional conditions 265
 - condition 259, 265
- MAPHEIGHT option 286
- MAPLINE option 286
- MAPONLY option 77, 242, 248
- mapped
 - input
 - getting storage for 258
 - mapped output after 264
 - processing 262
 - programming simple 257
 - output
 - after mapped input 264
 - sending 241, 276

- MAPPED option 293
- mapping
 - input data 257
- maps
 - acquiring and defining storage for 242
 - BMS 76, 77, 80
 - device-dependent 269
 - floating 282
 - into map sets, grouping 239
 - options for merging symbolic and physical 248
 - preparing BMS 16
 - send unformatted data without 82
 - sending pages built from multiple 80
 - sets 58
- MAPSET
 - option 242
 - resource definition 238
- MAPWIDTH option 286
- MARGINS option 12
- MASSINSERT option 73, 179
- MDTs 259
 - modified data tags
 - avoid turning on unnecessarily 75
- medium, selection by 340
- menu functions, EDF 427
- message
 - error 306
 - identification 300
 - integrity options, limit use of 82
 - lengths, reducing 288
 - line 443, 454
 - logical 324
 - routing 294
 - routing, interleaving conversation with 302
 - sending to destinations other than input terminal 79
 - text logical 289
 - title 300
 - undeliverable 300
 - writing to system console 62
- MF
 - modify field 216
- minimizing
 - errors 139
 - pathlength 287
- minimum BMS 226
- mixed addressing mode transactions 383
- mixing
 - DPL and DTP to same CICS system 128
 - languages 24
- mode, unformatted 224
- modification 214
- modified
 - data 259
 - data tag 75, 259
 - avoid turning on unnecessarily 75
- modify field (MF) order 216
- modifying
 - default CICS exception handling 146
 - execution, EDF 437
- modular program 57
- MONITOR
 - command 162
 - POINT command 162
- monitoring application performance 162
- moving variable data to map 244
- MSG, PF9 455
- MSGINTEG option 82
- MSR
 - magnetic slot reader 314
 - option 278, 314
- MSR option 314
- multimap screens 80
- multipage output, sending 79
- multiple
 - distributed program links from same client task, issuing 128
 - maps, sending pages built from 80
 - printer types, supporting 349
 - sends 345
- multithreading 64
 - testing 416
- MXT parameter 61

N

- naming conventions for remote queues 104
- NATLANG option 12
- natlang(en|ch|ge|ka) 12
- nested programs, COBOL/VSE (ANSI85) 32
 - ANSI85 standards
 - reference modification 36
- new attention identifiers and exception conditions 311
- NEXT DISPLAY 429
- NLEOM option 278, 343, 344, 345
- NOAUTOPAGE option 280
- NOCBLCARD option 12
- NODEBUG option 13
- NOEDF option 13
- NOEDIT option 293
- NOEPILOG option 13
- NOFE option 13
- NOFEPI option 13
- NOFLUSH option 283, 289
- NOHANDLE option 142, 150
- NOJBUFSP condition 63
- NOLENGTH option 13
- NOLINKAGE option 13
- non-3270
 - CICS printers 346
 - devices
 - output considerations for 267

- non-3270 (*continued*)
 - terminals
 - special options for 268
 - support for 267
- non-CICS
 - printer 341
 - printers 352
 - formatting for 352
- non-terminal transaction security 410
- nonconversational programming 53
- NONUM option 13
- NOOPSEQUENCE option 13
- NOOPTIONS option 14
- NOPROLOG option 14
- NOQUEUE option 63
- NOSEQ option 14
- NOSEQUENCE option 14
- NOSOURCE option 14
- NOSPACE condition 151
- NOSPIE option 14
- NOSUSPEND option 63
 - GETMAIN command 391
 - READNEXT command 178
 - READPREV command 178
- NOTAUTH condition
- notifying print application 354
- NOTRUNCATE option 324
- NOVBREF option 14
- NOWAIT option 340
- NOXREF option 14
- null values, use of 78
- nulls instead of blanks, use 78
- NUM option 14
- NUMREC option 178
- NUMSEGMENTS option
 - multiple dumps of task-related storage areas 163

O

- obsolete macros, CICS commands and equivalent 465
- OCCURS option 240
- OPCLASS
 - and LIST omitted 296
 - only, destinations specified with 296
 - option 295
- operating system waits 62
- operational control 61
- operations
 - data 71
 - page-building and routing 78
 - restrictions on retrying 157
 - terminal 75
 - VSAM and DAM data set 72
- OPERATOR NOT SIGNED ON 298
- OPERATOR SIGNED ON AT UNSUPPORTED TERMINAL 298

- OPID option 295
- OPIDENT value 295
- OPMARGINS option 14
- OPSEQUENCE option 15
- opsequence(m,n) 15
- option
 - ACCUM 282
 - BASE 243
 - DATAONLY 248
 - LENGTH 67
 - MAPONLY 248
 - PAGING 279
 - RIDFLD 179
 - STORAGE 243
- options
 - (main) specification 45
 - for 3270 printers 343
 - for merging symbolic and physical maps 248
 - HANDLE CONDITION command 147
 - on function keys, EDF 427
 - specifying translator 6
 - translator 7
- OPTIONS option 15
- orders in data stream 215
- outboard
 - controller 338
 - formatting 318
- outbound
 - data stream sample 218
 - FMH 336
- output
 - considerations for non-3270 devices 267
 - data, chaining 334
 - disposition options 277
 - example, BMS 228
 - map, initializing 244
 - map, symbolic 256
 - requests for printed 80
 - sending multipage 79
- outside map 252
- OVERFLOW condition 283
- overlays 58
- overtyping
 - EDF displays 437
 - to make changes 437
- overview of file control 167

P

- PA key 347
- page
 - breaks 283
 - building operations 78
 - fault 56
 - format for text messages 290
 - formation 282

page (*continued*)
 overflow 301
 overflow, routing and 301
 routing operations 78
 size, 3270 printers, BMS 348
 size, total 345
 width 345
 page-copy facility, using BMS 80
 PAGENUM option 301
 pages
 built from multiple maps, sending 80
 PAGESIZE value 282
 paging
 reducing effects 56
 terminal operator 280
 PAGING option 242, 277, 279, 344
 partitioned screens
 restrictions on 312
 uses for 305
 partitioning
 establishing 308
 partitions 304
 active 309
 and logical messages 310
 and routing 310
 ASSIGN options for 310
 defining 306
 how to define 306
 options for BMS RECEIVE commands 309
 options for BMS SEND commands 309
 sets, map sets and 400
 support 304
 PARTITIONSET option 308
 PARTN option 309
 partner, conversation 321
 partners, range of 132
 PARTNFAIL condition 311
 PARTNPAGE option 310
 passing
 control, anticipating return (LINK) 378
 data
 other ways of 383
 to other programs 379
 via a COMMAREA 395
 values as arguments 41
 path
 browsing using 175
 direct reading by way of 173
 pathlength, minimizing 287
 PCs as 3270s 210
 pen-detectable field 316
 PERFORM command 153
 PERFORM DUMP command 163
 performance
 monitoring application 162
 other suggestions for improving 62
 personal computers 209
 PF (program function) key 420, 442, 454
 PF1 HELP 443, 454
 PF10 SB 455
 PF10 SCROLL BACK FULL 444
 PF11 SCROLL FORWARD FULL 444
 PF11 SF 455
 PF2 HEX 454
 PF2 SWITCH HEX/CHAR 443
 PF3 END 454
 PF3 TERMINATE BROWSE 443
 PF4 EIB 454
 PF4 VIEW TOP 443
 PF5 VAR 455
 PF5 VIEW BOTTOM 443
 PF6 REPEAT LAST FIND 443
 PF6 USER 455
 PF7 SBH 455
 PF7 SCROLL BACK HALF 444
 PF8 SCROLL FORWARD HALF 444
 PF8 SFH 455
 PF9 MSG 455
 PF9 VIEW LEFT 444
 PF9 VIEW RIGHT 444
 phonetic command equivalent 465
 macro equivalent 465
 physical
 and symbolic map sets 238
 EXEC CICS SEND command per screen, use only
 one 81
 key subfield 189
 keys 189
 maps, options for merging symbolic and 248
 PL/I
 and dynamic storage 46
 C and 433
 programming language 44
 restrictions 44
 use of RESP and DFHRESP in COBOL and 143
 PL/I language considerations
 compilers supported 44
 OPTIONS(MAIN) specification 45
 program segments 46
 restrictions 44
 PLT programs
 program working storage specifically for 395
 points
 exception trace entry 162
 system trace entry 161
 trace entry 161
 user exception trace entry 162
 user trace entry 162
 POP HANDLE command 146, 152, 156
 positioning cursor 252
 positive acknowledgment (ACK) 323

- POST command 369
 - using 111
- POWER 357
- POWER, typical use 359
- preparing
 - application table entries 415
 - BMS maps 16
 - system for debugging 415
 - system table entries 415
 - your application to run 3
 - your program 3
- preprinted form 408
- presentation space 304
- preventing interruptions (bracket protocol) 336
- PREVIOUS DISPLAY 429
- PRGDLY option 300
- principal facility 200
- print
 - application, notifying 354
 - control bit, PRINT option and 344
 - formatting 338
 - key 324, 355
 - hardware 356
 - keys
 - CICS 355
 - on printer PRT1, task that wants to 349, 350
 - option 340
- PRINT option
 - and print control bit 344
- printed output, requests for 80
- PRINTERCOMP option 346
- printers
 - 3270 342, 343
 - options 343
 - BMS page size, 3270 348
 - CICS 341, 349
 - determining characteristics of 347
 - formatting for CICS 341
 - formatting for non-CICS 352
 - non-3270 CICS 346
 - non-CICS 341, 352
 - options for 3270 343
 - PRT1, task that wants to print on 349, 350
 - SCS 345
 - spooling subsystem 364
 - types, supporting multiple 349
- printing 341—356
 - CICS API considerations 352
 - CICS support for 341
 - display screens 354
 - line length on 344
 - START command 349
 - with BMS routing 351
 - with EXEC CICS START command 349
 - with transient data 350
- PRNT, started task executing transaction 350
- process, translation 5
- processing
 - asynchronous 132
 - mapped input 262
 - skip-sequential 176
- processor usage 60
- program
 - control 377
 - commands 461
 - linking to another program 378
 - passing data to another program 379
 - program logical levels 378
 - design
 - conversational 53, 82
 - nonconversational 53
 - pseudoconversational 53
 - determining how invoked 129
 - execution, how you can intervene in 427
 - expecting return, link to another 378
 - labels in EDF 437
 - logical levels 25, 29
 - option 155
 - resource definition 238
 - restricting to distributed program link subset 129
 - segments
 - PL/I 46
 - size 52
 - source 5
 - storage 67
 - structure 51
 - testing 419
 - working storage specifically for exit and PLT programs 395
- program-level abend exit, creating 156
- programming
 - considerations 307
 - for distributed program link 128
 - with routing 301
 - hints 410
 - language 120
 - note for spool commands 362
 - restrictions
 - assembler language 46
 - C 38
 - COBOL 19
 - PL/I 44
 - safe techniques 94
 - simple mapped input 257
 - suspect techniques 103
 - unsafe techniques 97
- programming techniques
 - general 52, 55
 - PL/I 44
 - structure 51

- programs
 - DPL 436
 - DTP 436
 - EDF 436
 - passing data to other 379
 - preparing your 3
 - writing CICS 3
- PROLOG option 15
- PROTECT option 82
- protecting
 - CICS from application errors 139
 - CICS from being passed invalid addresses 400
- resources
 - server program in DPL 127
- protection 213
 - CICS storage 392
 - storage 391
- protocol, response 335
- PRT1, task that wants to print on printer 349, 350
- pseudoconversational
 - and conversational design
 - choosing between 53
 - programming 53
 - programs, checking 435
- PURGE 445
 - command, CEBR transaction 445
 - MESSAGE command 281, 295
- purge delay 300
- PUSH HANDLE command 146, 152, 156
- PUT
 - command, CEBR transaction 445
 - xxxx 445

Q

- QBUSY condition 63
- QUERY
 - NOLOG option 410
 - SECURITY command 409, 410
 - RESCLASS option 409
 - RESID option 409
 - RESTYPE option 409
 - using 409
- query transaction 201
- QUEUE
 - command, CEBR transaction 446
 - xxxxxxx 446
- queues
 - exception conditions for globally accessible 105, 106
 - extrapartition 404
 - intrapartition 403
 - naming conventions for remote 104
 - temporary storage 407
 - transient data 403

- QUOTE option 15
- QZERO condition 405

R

- range of partners 132
- RBA (relative byte address) 168, 185
- RDF 192
- READ command
 - direct reading using 172
 - updating records with 176
- read integrity 178
- read-ahead queueing 322
- reading
 - by way of path, direct 173
 - data from display 257
 - direct (using READ command) 172
 - from 3270 terminal 222
 - from ESDS, direct 173
 - from formatted screen 259
 - from KSDS, direct 173
 - from RRDS, direct 173
 - records 172
 - sequential browsing 173
- READNEXT command 173, 174
 - NOSUSPEND option 178
- READPREV command 174
 - NOSUSPEND option 178
- READQ TD command 63, 403
- READQ TS command 407, 408
 - ITEM option 408
- REALLOC (C function) 401
- RECEIVE command 82, 265, 320, 321, 323, 336
- RECEIVE MAP command 257, 260
 - ASIS option 260
- RECEIVE MAP, other information from 260
- RECEIVE PARTN command 310
- RECEIVE, what you get on EXEC CICS 323
- receiving data from display 253
- rechecking 264
- recommendation 68
- record
 - data set (RRDS), relative 168
 - description field 192
 - identification 185, 189
 - length, specifying 179
 - level, security protection at 409
 - locking 186
- records
 - adding 178
 - adding to DAM data set 191
 - browsing 172
 - deleting 177
 - deleting groups of (generic delete) 178
 - deleting single 177
 - from DAM data sets, browsing 190

- records (*continued*)
 - from DAM data sets, updating 190
 - journal 68, 135
 - length of 68
 - reading 172
 - sequential adding of (WRITE MASSINSERT command) 179
 - to DAM data sets, adding 191
 - updating with READ command 176
 - writing 178
- recoverable
 - resources 53
 - exclusive use 53
- recoverable files, CICS locking of VSAM records in 186
- recovery
 - abnormal termination 155
 - considerations 135
 - design implications 60
 - logical message 282
 - of resources 59, 60
 - problem avoidance 139
 - sequential terminal support 336
 - syncpoint 137
- reducing
 - inbound traffic, use of FRSET 76
 - line traffic 77
 - design data entry operations for 78
 - message lengths 288
 - paging effects 56
- reentrancy 64
- reference modification, COBOL/VSE (ANSI85) 36
 - ANSI85 standards
 - global variables 36
- reference set 58
- reference, locality of 56
- REGISTERS AT ABEND 429
- regression testing 416
- relation
 - and lifetime 114
 - global 114
 - LUname (terminal) 115
 - userid 116
- relative
 - byte address (RBA) 168, 185
 - record data set (RRDS) 168
 - record number (RRN) 168, 185
- RELEASE command
 - HOLD option 377
- RELEASE option 279
- RELTYPE keyword 170
- REMEMBER DISPLAY 429
- remote
 - data sets, KEYLENGTH option for 183
 - queues, naming conventions for 104
 - transactions, EDF and 435
 - remote-linked programs
 - DPL 436
 - EDF 436
 - REMOTENAME option 124
 - REMOTESYSTEM option 125
 - REPEAT LAST FIND, PF6 443
 - repeated fields 240
 - REPLACE statement, COBOL/VSE (ANSI85) 31
 - report controller 363
 - reports 364
 - reporty controller 364
 - REQID option 176, 278, 300, 371
 - request identifiers 371
 - request/response unit (RU) 333
 - requests
 - for printed output 80
 - RESCLASS option 409
 - RESETBR command 174
 - RESID option 409
 - resource identifiers, CICS-defined 410
 - resources
 - contention for 53
 - control of 53
 - controlling sequence of access to 374
 - exclusive control of 59
 - exclusive use 53
 - recoverable 53, 59, 60
 - sharing between client program and server program 128
 - RESP
 - option 142
 - deactivating NOHANDLE 148
 - how to use 143
 - use of in C 143
 - use of in COBOL and PL/I 143
 - value 141
 - RESP2
 - option
 - how to use 143
 - RESP2 option 142
 - RESP2 value 141, 142
 - response
 - definite 340
 - EDF 439
 - protocol 335
 - restricting program to distributed program link subset 129
 - restrictions
 - C 40
 - on partitioned screens 312
 - on retrying operations 157
 - when using EDF 439
 - RESTYPE option 409
 - RETPAGE condition 290, 303
 - RETRIEVE command 369, 374

- RETRIEVE WAIT and START commands, using 106
- retrying operations, restrictions on 157
- RETURN CODE register, COBOL/VSE 19
- RETURN command 66, 381
 - COMMAREA in EXEC CICS 88
 - COMMAREA option 88, 377
 - ERRATT option 61
 - examples of passing data with 387
 - IMMEDIATE option 322, 336
 - INPUTMSG option 377, 383
 - TRANSID option 61
 - using INPUTMSG option on 383
- RETURN IMMEDIATE, EXEC CICS 322
- return, link to another program expecting 378
- reverse interrupt 323
- REWRITE command 176
- RIDFLD option 170, 173, 179
- RIGHT, PF9 VIEW 444
- ROUTE command 294
 - LIST option 295
 - page overflow 301
 - TITLE option 300
- route list 295
 - format 297
 - LIST option 297, 298
 - provided by BMS 296
 - segment chain entry format 297
 - standard entry format 297
- routed message, specifying destinations for 295
- ROUTEDMSGS option 296
- routing
 - and page overflow 301
 - how it works 294
 - LDCs and 313
 - message 294
 - operations 78
 - partitions and 310
 - printing with BMS 351
 - programming considerations with 301
 - temporary storage and 300
 - terminals 302
 - transaction 119, 120
 - with SET 302
- RRDS
 - adding to 178
 - browsing through 175
 - direct reading from 173
 - relative record data set 168
- RRN 186
 - relative record number 168
- RRN (relative record number) 185
- RTEFAIL condition 298
- RTESOME condition 298
- RU (request/response unit) 333
- rules
 - for logical messages 278
 - rules (*continued*)
 - for writing BMS macros 235
 - map placement 284
 - run unit in COBOL/VSE 24, 29
 - runaway tasks 62
 - RVI 323

S

- SA
 - set attribute 218
- SAA Resource Recovery 138
- safe
 - programming techniques 94
 - techniques 92
- SAM 319
- saving
 - commands 461
 - good input 263
- SB, PF10 455
- SBA
 - set buffer address 217
- SBH, PF7 455
- scope of intertransaction affinities 113
- screen
 - compress data sent to 78
 - copy, BMS 356
 - display 89
 - do not send blank fields to 76
 - fields 208, 259
 - printing display 354
 - reading from formatted 259
 - use only one physical EXEC CICS SEND command per 81
- Screen Definition Facility II for VSE
 - See SDF/CICS
- SCROLL
 - BACK 429
 - FULL 429
 - FULL, PF10 444
 - HALF 430
 - HALF, PF7 444
 - FORWARD 429
 - FULL 429
 - FULL, PF11 444
 - HALF 430
 - HALF, PF8 444
- scrolling 305
- SCS
 - input 347
 - printer 345
 - printers 347
- SDF
 - alternative 239
- SDF/CICS 17, 231

- SEC
 - system initialization option 410
 - system initialization parameter 410
- security 409—411
 - CICS-defined resource identifiers 410
 - considerations 440, 447, 462
 - control 409
 - EDF 440
 - non-terminal transaction 410
 - programming hints 410
 - protection at record or field level 409
 - record or field level 409
 - SEC system initialization option 410
 - SPCOMMAND resource type 410
- SEGMENTLIST option
 - multiple dumps of task-related storage areas 163
- segments, program
 - PL/I 46
- selection
 - and definition of storage key for applications 394
 - and identification, destination 340
 - by medium 340
 - by named data set 340
 - fields 316
 - of execution key for applications 393
- semicolon as delimiter
 - COBOL/VSE 36
- SEND command 82, 320, 321, 323, 336, 344
 - avoid WAIT option on BTAM 81
 - CNOTCOMPL option 334
 - CTLCHAR option 344
 - FMH option 336
 - INVITE option 321
 - LAST option 336
 - MSR option 314
- SEND command per screen, use only one physical EXEC CICS 81
- SEND commands, limit use of DEFRESP option on EXEC CICS 82
- SEND CONTROL command 249, 341
- SEND MAP command 241, 242, 277, 341
 - ACCUM option 242, 277
 - ALARM option 247
 - control options on 247
 - CURSOR option 242, 252
 - DATAONLY option 242
 - ERASE option 247
 - ERASEAUP option 247, 289
 - FREEKB option 247
 - FROM option 242
 - FRSET option 247
 - LAST option 248
 - MAPONLY option 242
 - MAPSET option 242
 - NOFLUSH option 283, 289
 - PAGING option 242
 - SEND MAP command (*continued*)
 - SET option 242
 - TERMINAL option 242, 277
 - WAIT option 248
- SEND PAGE command 138, 279, 280, 294
 - AUTOPAGE option 280
 - NOAUTOPAGE option 280
 - RELEASE option 279
- SEND PARTNSET command 308
- SEND TEXT command 288, 289, 341
 - extensions 293
 - MAPPED option 293
 - NOEDIT option 293
- send/receive mode 321
- sending
 - data to logical device component 313
 - invalid data and other errors 253
 - mapped output 241, 276
- sends, multiple 345
- SEQ option 15
- sequence numbers, COBOL/VSE (ANSI85) 30
 - ANSI85 standards
 - lowercase characters 31
- sequence of access to resources, controlling 374
- SEQUENCE option 15
- sequence(m,n) 15
- sequential
 - adding of records (WRITE MASSINSERT command) 179
 - data set access 74
 - reading (browsing) 173
 - terminal support 336, 417
 - terminals, coding considerations for 337
- server
 - program 124
 - coding considerations 127
 - distributed program link (DPL) 127
 - exception conditions returned to 130
 - sharing resources between client program and 128
 - syncpointing 127
 - region 121, 125
- SERVICE RELOAD
 - elimination, COBOL/VSE 19
- SESSBUSY condition 63
- SESSION, END EDF 428
- sessions, CECI 460
- set
 - attribute (SA) order 218
 - buffer address (SBA) order 217
 - command 153, 277
 - commands and global user exits, using INQUIRE and 92
 - option 180, 242, 277
 - reference 58
 - routing with 302

- set (*continued*)
 - using 302
 - working 57
- setting display characteristics 244
- SF
 - PF11 455
 - start field 216
- SFE
 - start field extended 216
- SFH, PF8 455
- SHARED
 - command, GETMAIN 71
 - GETMAIN command 391
 - option 65, 71, 391
 - storage 71
 - storage, using GETMAIN 98
- sharing
 - data across transactions 83
 - resources between client and server programs 128
 - terminal 311, 461
- shipping, function 120
- SIGNAL condition 323
- simple mapped input, programming 257
- simultaneous browse operations 176
- single records, deleting 177
- single-screen mode, EDF 434
- single-screen mode, using EDF in 434
- single-thread testing 416
- single-threading 64
- size
 - program 52
 - total page 345
- skip-sequential processing 176
- SOURCE option 16
- source program 5
- SP option 16
- SPACE option 16
- space, presentation 304
- SPCOMMAND resource type 410
- speaking out of turn 322
- special
 - hardware, BMS support for other 314
 - options for non-3270 terminals 268
- specifying
 - destinations for routed message 295
 - record length 179
 - translator options 6
- SPIE option 16
- SPOOL
 - system initialization parameter 357
- spool commands
 - programming note for 362
- spooler
 - See* VSE/POWER
- SPOOLOPEN command 353
- SPURGE parameter 101
- SQL 3
 - interface, EXEC 193
- SSA (segment-search area)
 - See* DL/I, segment-search area (SSA)
- standard BMS 227
- START command 349, 369, 374
 - and CANCEL REQID command, using 108
 - printing with EXEC CICS 349
 - using RETRIEVE WAIT and 106
- start field
 - (SF) order 216
 - extended (SFE) order 216
- STARTBR command 173
- started task executing transaction PRNT 350
- STARTIO option 136
- static
 - COBOL call 25
 - storage 58
- status
 - flag byte, route list 298
 - line 451
- STOP CONDITIONS 430
- stopping EDF 437
- storage
 - area
 - dynamic 52
 - areas, system-wide 394
 - BROWSE TEMP 428
 - C programs and CICS 401
 - CICS-key 392
 - CICS-key and user-key 392
 - control 391
 - control, CICS 391
 - for mapped input, getting 258
 - for maps, acquiring and defining 242
 - key for applications, selecting and defining 394
 - key, deciding what to specify 397
 - main 52
 - option 243
 - program 67
 - program working
 - specifically for exit and PLT programs 395
 - protection 391, 392
 - protection exception conditions 400
 - protection, CICS 392
 - shareable 391
 - sharing task-lifetime 100
 - static 58
 - task lifetime 394
 - temporary 69
 - under CICS 401
 - user 65
 - user-key 392
 - using GETMAIN SHARED 98
 - using temporary 103

- storage (*continued*)
 - virtual 55
 - WORKING 432
- storing data within a transaction 64
- stream, orders in data 215
- structure
 - program 51
- stubs, EXEC interface 68
- subfield
 - block reference 189
 - deblocking argument 189
 - physical key 189
- subprogram, calling from COBOL/VSE 25—29
- subroutines 56
- support
 - device-dependent 270
 - for multiple printer types 349
 - for non-3270 terminals 267
 - levels, BMS 226
- SUPPRESS DISPLAYS 431
- suspect
 - programming techniques 103
 - techniques 93
- SUSPEND command 373
- suspend data set 408
- SWITCH HEX/CHAR 432
 - PF2 443
- symbolic
 - and physical maps, options for merging 248
 - ANSI85 standards
 - programming restrictions 37
 - characters defined in program, COBOL/VSE (ANSI85) 36
 - cursor positioning 253
 - input map 256
 - map sets, physical and 238
 - output map 256
- synchronization, journal output 136
- synchronize action
 - journal output 136
- SYNCONRETURN option 125, 131
- SYNCPOINT command 137, 138, 156, 279
 - ROLLBACK option 156
- syncpointing 137—138
 - DPL 125
- syntax check, command 451
- SYSEIB option 16
- SYSID
 - command, CEBR transaction 446
 - option 125
 - xxxx 446
- SYSIDERR 121
- SYSIPT 5
- SYSLST 5
- SYSPCH 5

- system
 - console, writing message to 62
 - for debugging, preparing 415
 - information, access to 153
 - initialization parameter, SEC 410
 - mixing DPL and function shipping to same CICS 128
 - programming commands 153
 - table entries, preparing application and 415
 - trace entry points 161
 - waits, operating 62
- system-wide storage areas 394

T

- table entries, preparing application and system 415
- tables 400
- tabs
 - horizontal 347
 - vertical 347
- task 200
 - control 373
 - executing transaction PRNT, started 350
 - lifetime storage 394
 - that gets triggered 351
 - that wants to print on printer PRT1 349, 350
- task control
 - sequence of access to resources 374
- task-lifetime storage, sharing 100
- task-related user exit 395
- TASKDATAKEY option 65, 394
- TASKDATALOC option 16, 65
- tasks
 - runaway 62
- tasks get started, how 200
- TCTTE user area (TCTUA) 87
- TCTUA 95, 394
 - in an unsafe way, using 96
 - TCTTE user area 87
- TCTUAKEY 87, 394
- TCTUALOC 87
- techniques
 - general programming 55
 - programming 52
 - safe 92
 - safe programming 94
 - suspect 93
 - suspect programming 103
 - unsafe 93
 - unsafe programming 97
 - used by CICS application programs to pass data 92
- TEMP STORAGE, BROWSE 428
- temporary data 407
- temporary storage 69
 - and routing 300
 - auxiliary 69, 407

- temporary storage (*continued*)
 - browse (CEBR) 441
 - browse transaction, CEBR 441
 - control 407
 - typical uses of 408
 - data 407
 - main 69, 407
 - queues 407
 - using 103
- TERM option 270
- TERMIN value 295
- terminal
 - command, CEBR transaction 446
 - contention for 321
 - finding out about your 272, 330
 - input from 3270 221
 - interruptibility 61
 - length of data stream sent to 75
 - NOT SUPPORTED UNDER BMS 298
 - operations 75
 - operator paging 280
 - option 242, 277
 - reading from 3270 222
 - sending messages to destinations other than
 - input 79
 - sharing 311, 461
 - support, sequential 336
 - waits 323
 - xxxx 446
- terminal control 319—340
 - bracket protocol, LAST option 336
 - break protocol 322
 - chaining of input data 333
 - chaining of output data 334
 - commands 320
 - considerations, additional 81
 - conversation partner 321
 - definite response 335
 - facilities for logical units 333
 - flip-flop mode 321
 - FMH, inbound 335
 - FMH, outbound 336
 - function management header (FMH) 335
 - half-duplex mode 321
 - interrupting 323
 - logical record presentation 334
 - map input data 257
 - operations, EIB feedback on 332
 - partner, conversation 321
 - print formatting 338
 - protocol, break 322
 - read-ahead queueing 322
 - table user area (TCTUA) 87
- terminal control table user area
 - See TCTUA
- terminals
 - 3270 family of 207
 - CICS APIs for 203
 - coding considerations for sequential 337
 - eligible 295
 - special options for non-3270 268
 - support for non-3270 267
- TERMINATE BROWSE, PF3 443
- termination recovery, abnormal 155
- testing applications 139, 415
 - multithread testing 416
 - preparing application table entries 415
 - preparing system 415
 - preparing system table entries 415
 - regression testing 416
 - sequential terminal support 417
 - single-thread testing 416
 - using sequential devices 336, 417
- text
 - command, SEND 289
 - extensions, SEND 293
 - into lines, how BMS breaks 291
 - logical messages 289
 - messages, header and trailer format for 292
 - messages, page format for 290
 - output, formatting 288
- time field of EIB 153
- times, expiration 369
- TIOATDL value 294
- TITLE option 300
- title, message 300
- TOKEN option 181
- TOP 446
 - command, CEBR transaction 446
 - PF4 VIEW 443
- total page size 345
- trace 161
 - auxiliary 63
 - description 161
 - entry points 161
 - exception 162
 - system 161
 - user 162
 - user exception 162
- trailer format for text messages, header and 292
- transaction 199
 - affinity 120, 369, 374, 391, 407
 - affinity, what is 91
 - deadlocks, avoiding 181
 - display, CEBR 442
 - how to use CEBR 441
 - identifier
 - CEBR 441
 - CECI 449
 - CEDF transaction 419
 - PRNT, started task executing 350

- transaction (*continued*)
 - routing 119, 120
 - dynamic 383
 - security, non-terminal 410
 - storing data within a 64
 - system affinity 92
 - which 201
 - with transient data, using CEBR 446
 - work area (TWA) 65
- transaction affinity 91—118
 - intertransaction affinity 91
 - transaction-system affinity 92
- transaction work area 65
- transactions
 - affinity among 91
 - avoid using unnecessary 82
 - conversational 53
 - EDF and remote 435
 - mixed addressing mode 383
 - nonconversational 53
 - pseudoconversational 53
 - sharing data across 83
- transferring control to a subprogram 25
- TRANSID option 61, 125, 381
- transient data
 - control 403
 - automatic transaction initiation (ATI) 405
 - queues 403, 404
 - extrapartition 71, 74
 - intrapartition 70
 - printing with 350
 - queue 70, 105
 - using 105
 - using CEBR transaction with 446
- translation 3—16
 - ANSI85 option 9
 - APOST option 9
 - CBLCARD option 9
 - CICS option 9
 - COBOL2 option 9
 - COBOL3 option 9
 - DBCS option 9
 - DEBUG option 10
 - DLI option 10
 - EDF option 10
 - EPILOG option 10
 - EXCI option 10
 - FE option 10
 - FEPI option 10
 - FLAG option 10
 - GDS option 11
 - GRAPHIC option 11
 - LENGTH option 11
 - line numbers 6
 - LINECOUNT option 11
 - LINKAGE option 11
- translation (*continued*)
 - MARGINS option 12
 - NATLANG option 12
 - NOCBLCARD option 12
 - NODEBUG option 13
 - NOEDF option 13
 - NOEPILOG option 13
 - NOFE option 13
 - NOFEPI option 13
 - NOLENGTH option 13
 - NOLINKAGE option 13
 - NONUM option 13
 - NOOPSEQUENCE option 13
 - NOOPTIONS option 14
 - NOPROLOG option 14
 - NOSEQ option 14
 - NOSEQUENCE option 7, 14
 - NOSOURCE option 6, 14
 - NOSPIE option 14
 - NOVBREF option 14
 - NOXREF option 14
 - NUM option 14
 - OPMARGINS option 14
 - OPSEQUENCE option 15
 - options 6, 7
 - OPTIONS option 15
 - process 5
 - PROLOG option 15
 - QUOTE option 15
 - SEQ option 15
 - SEQUENCE option 15
 - SOURCE option 6, 16
 - SP option 16
 - SPACE option 16
 - SPIE option 16
 - SYSEIB option 16
 - upper case 260
 - VBREF option 6, 16
 - XOPTS keyword 6
 - XREF option 16
- translator
 - data sets 5
 - options 7
- translator options
 - specifying 6
- trigger field support 315
- triggered, task that gets 351
- TS queue 104
- TWA 65
- TWASIZE option 65
- TYPE=DSECT assembly 238
- typical uses of temporary storage control 408

U

UNDEFINED 432
undeliverable messages 300
unformatted
 data without maps, send 82
 mode 224
unit of compilation, COBOL
 description 30
 start of, nested programs 33
 submitting to translator 33
 translator options for 32
unit of work 137
UNLOCK command 179
unnecessary transactions, avoid using 82
unsafe
 programming techniques 97
 techniques 93
 way, using TCTUA in an 96
unsolicited input 201
UNSUPPORTED TERMINAL, OPERATOR SIGNED ON
 AT 298
UNTIL option 369
UOW 137
update operation, DAM 190
updating records
 from DAM data sets 190
 with READ command 176
upgrade set 169
upper case translation in BMS 260
usage, processor 60
USE FOR DEBUGGING facility 20
use FRSET to reduce inbound traffic 76
user
 data sets 71
 DISPLAY 432
 exception trace entry points 162
 PF6 455
 storage 65
 trace entry point 161
 trace entry points 162
user-key
 applications 397
 storage 392
user-maintained data tables
 application programming 171
user-related information with EXEC CICS ASSIGN
 command, accessing 129
user-replaceable module 395
USERDATAKEY option 65, 395
userid relation 116
using
 CEBR transaction with transient data 446
 EDF
 in dual-screen mode 435
 in single-screen mode 434

using (*continued*)
 function management headers 335
 VSE macros from within an application 18

V

validity of reference 56
values as arguments
 passing 41
VAR, PF5 455
variable data to map, moving 244
variables
 CECI/CECS 455
 defining 457
VBREF option 16
vertical tabs 347
VIEW BOTTOM, PF5 443
VIEW LEFT, PF9 444
VIEW RIGHT, PF9 444
VIEW TOP, PF4 443
viewport 304
virtual storage 55
 environment 52
virtual storage constraint relief (VSCR)
 assembler language considerations 46
 C considerations 38
 COBOL considerations 19
 PL/I considerations 44
 program size 52
virtual storage environment 55
VOLUME option 340
VOLUMELENG option 340
VS COBOL II
VSAM
 alternate indexes 169
 and DAM data set operations 72
 considerations, file control 185
 data sets 72, 167, 185
 enqueues 59
 MASSINSERT option 73
 processor overhead 73
 records in recoverable files, CICS locking of 186
VSCR (virtual storage constraint relief)
 assembler language considerations 46
 C considerations 38
 COBOL considerations 19
 program size 52
VSE macros 18
 from within an application, using
VSE/ESA Info/Analysis program 163
VSE/ESA transaction 383
VSE/POWER 357
 CICS interface to 357
 input 357
 output 358
 retrieve data from VSE/POWER spool 358

VSE/POWER (*continued*)
 send file to destination 359
 write directly to VSE/POWER spool 359
VTAM 319, 321
 considerations 333

W

WAIT
 EVENT command 101, 369, 374
 using 101
 EXTERNAL command 102, 373
 JOURNALNUM command
 synchronize with journal output 135
 option 248, 323
 avoid on BTAM 81
 TERMINAL command 323
WAITCICS command 102, 373
waiting for function completion 340
waits
 operating system 62
 terminal 323
what does
 CECI display? 450
 EDF display? 420
where does EDF intercept program? 419
which transaction? 201
width
 options, line 344
 page 345
working
 set 56, 57
 storage 38
 specifically for exit and PLT programs,
 program 395
WORKING STORAGE 432
WPMEDIA1–4 option 340
WRITE
 command 178
 commands, 3270 211
 control character 212
 JOURNALNAME command 63
 JOURNALNUM command 63, 136
 create a journal record 135
 MASSINSERT command, sequential adding of
 records 179
WRITEQ
 TD command 403
 TS command 407
writing
 BMS macros, rules for 235
 CICS programs 3
 message to system console 62
 records 178
 to a 3270 terminal 211

WRKAREA parameter 83

X

XCTL command 52, 65, 66, 377, 381
 COMMAREA option 379
 INPUTMSG option
XCTL commands, COMMAREA in EXEC CICS LINK
 and EXEC CICS 66
XOPTS keyword 6
XPCREQ global user exit 121, 125
XREF option 16
XTCOUT exit 78

Y

Year 2000 testing 417

Sending your comments to IBM

CICS® Transaction Server for VSE/ESA™

Application Programming Guide

SC33-1657-00

If you want to send to IBM any comments you have about this book, please use one of the methods listed below. Feel free to comment on anything you regard as a specific error or omission in the subject matter, and on the clarity, organization or completeness of the book itself.

To request additional publications, or to ask questions or make comments about the functions of IBM products or systems, you should talk to your IBM representative or to your IBM authorized remarketer.

When you send comments to IBM, you grant IBM a nonexclusive right to use or distribute your comments in any way it believes appropriate, without incurring any obligation to you.

You can send your comments to IBM in any of the following ways:

- By mail:
 - IBM UK Laboratories
 - Information Development
 - Mail Point 095
 - Hursley Park
 - Winchester, SO21 2JN
 - England
- By fax:
 - From outside the U.K., after your international access code use 44 1962 870229
 - From within the U.K., use 01962 870229
- Electronically, use the appropriate network ID:
 - IBM Mail Exchange: GBIBM2Q9 at IBMMAIL
 - IBMLink: HURSLEY(IDRCF)
 - Email: idrcf@hursley.ibm.com

Whichever method you use, ensure that you include:

- The publication number and title
- The page number or topic to which your comment applies
- Your name and address/telephone number/fax number/network ID.




Program Number: 5648-054



Printed in the United States of America
on recycled paper containing 10%
recovered post-consumer fiber.

SC33-1657-00



Spine information:



CICS TS for VSE/ESA

Application Programming Guide

Release 1