

IBM Enterprise Metal C for z/OS, V3.1



Programming Guide and Reference

Version 3 Release 1

IBM Enterprise Metal C for z/OS, V3.1



Programming Guide and Reference

Version 3 Release 1

Note

Before using this information and the product it supports, read the information in “Notices” on page 133.

This edition applies to Version 3 Release 1 of IBM Enterprise Metal C for z/OS (5655-MCE) and to all subsequent releases and modifications until otherwise indicated in new editions.

Last updated: June 12, 2018

© **Copyright IBM Corporation 2018.**

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

Contents

Figures	v
--------------------------	----------

Tables	vii
-------------------------	------------

About this document	ix
--------------------------------------	-----------

Who should read this document	ix
Where to find more information	x
z/OS Basic Skills in IBM Knowledge Center	x
How to read syntax diagrams	x
Technical support	xii
How to send your comments to IBM	xiii
If you have a technical problem	xiii

Chapter 1. About IBM Enterprise Metal C for z/OS. 1

Metal C environment	1
Programming with Enterprise Metal C for z/OS	2
Metal C and MVS linkage conventions	2
Compiler-generated HLASM source code	4
Prolog and epilog code	13
Supplying your own HLASM statements	21
Inserting HLASM instructions into the generated source code	21
AMODE-switching support	31
RENT mode support	32
argc argv parsing support	34
AR-mode programming support	35
Metal C function descriptor support	43
Defining an alternative name for function main	43
Building Metal C programs	44
Summary of useful references for the Metal C programmer	52

Chapter 2. Header files. 55

builtins.h — Declare built-in functions	55
ctype.h — Declare character classification functions	55
float.h — Define ANSI constants for floating-point data types	55
inttypes.h — Define macros for sprintf and scanf family	56
limits.h — Declare symbolic names for resource limits	58
math.h — Define macros for floating-point support	59
metal.h — Define Metal C related function prototypes and data	60
stdarg.h — Define macros for accessing variable-length argument lists in functions	60
stddef.h — Define ptrdiff_t, size_t, and ssize_t data types	60
stdio.h — Define I/O related functions	60
stdint.h — Define integer types and related limits and macros	61
stdlib.h — Define standard library functions	62
string.h — Declare string manipulation functions.	63

Chapter 3. C functions available to Metal C programs 65

Characteristics of Metal C runtime library functions	65
System and static object libraries	65
User-replaceable heap services	67
abs() — Calculate integer absolute value.	69
atoi() — Convert character string to integer.	70
atol() — Convert character string to long	70
atoll() — Convert character string to signed long long	71
calloc() — Reserve and initialize storage.	71
__cinit() - Initialize a Metal C environment	72
__cterm() - Terminate a Metal C environment	75
div() — Calculate quotient and remainder	75
free() — Free a block of storage.	76
isalnum() to isxdigit() — Test integer value.	76
isalpha() — Test for alphabetic character classification	77
isblank() — Test for blank character classification.	78
iscntrl() — Test for control classification	78
isdigit() — Test for decimal-digit classification.	78
isgraph() — Test for graphic classification	78
islower() — Test for lowercase	78
isprint() — Test for printable character classification	78
ispunct() — Test for punctuation classification.	78
isspace() — Test for space character classification.	78
isupper() — Test for uppercase letter classification	78
isxdigit() — Test for hexadecimal digit Classification	78
labs() — Calculate long absolute value	79
ldiv() — Compute quotient and remainder of integral division.	79
llabs() — Calculate absolute value of long long integer	79
lldiv() — Compute quotient and remainder of integral division for long long type	80
malloc() — Reserve storage block	80
__malloc31() — Allocate 31-bit storage	81
memcpy() — Copy bytes in memory	82
memchr() — Search buffer	82
memcmp() — Compare bytes	83
memcpy() — Copy buffer	83
memmove() — Move buffer	84
memset() — Set buffer to value.	84
qsort() — Sort array	85
rand() — Generate random number	85
rand_r() — Pseudo-random number generator.	86
realloc() — Change reserved storage block size	86
snprintf() — Format and write data	87
sprintf() — Format and Write Data	88
srand() — Set Seed for rand() Function	94
sscanf() — Read and Format Data	94
strcat() — Concatenate Strings	99
strchr() — Search for Character	100
strcmp() — Compare Strings	101
strcpy() — Copy String	102

strcspn() — Compare Strings	102
strdup() — Duplicate a String	103
strlen() — Determine String Length	103
strncat() — Concatenate Strings	104
strncmp() — Compare Strings	104
strncpy() — Copy String.	105
strpbrk() — Find Characters in String	106
strrchr() — Find Last Occurrence of Character in String	106
strspn() — Search String.	107
strstr() — Locate Substring	107
strtod — Convert Character String to Double.	108
strtof — Convert Character String to Float.	109
strtok() — Tokenize String	110
strtok_r() — Split String into Tokens.	111
strtol() — Convert Character String to Long	112
strtold — Convert Character String to Long Double	113
strtoll() — Convert String to Signed Long Long	114
strtoul() — Convert String to Unsigned Integer	116
strtoull() — Convert String to Unsigned Long Long	117
tolower(), toupper() — Convert Character Case	119
va_arg(), va_copy(), va_end(), va_start() — Access Function Arguments	119

vsnprintf() — Format and print data to fixed length buffer	120
vsprintf() — Format and Print Data to Buffer.	121
vsscanf() — Format Input of a STDARG Argument List.	122

Appendix A. Function stack requirements 125

Appendix B. Accessibility 129

Accessibility features	129
Consult assistive technologies	129
Keyboard navigation of the user interface	129
Dotted decimal syntax diagrams	129

Notices 133

Programming interface information	134
Standards	134
Trademarks	134

Index 135

Figures

1. Prefix data fixed area fields	6	28. Specifying and using the WTO macro (no reentrancy).	29
2. A sample program to generate prefix data	8	29. Support for reentrancy in a code format string	29
3. Prefix data generated.	8	30. Code that supplies specific DSECT mapping macros	30
4. Function entry point marker in generated assembler code.	9	31. Register specification	31
5. Function property block fixed area fields	9	32. AMODE31 program that calls an AMODE64 program	32
6. Function property block in generated assembler code	12	33. Far pointer sizes under different addressing modes	36
7. Debug data block generated	13	34. Built-in functions for setting far-pointer components	38
8. Specification of your own prolog and epilog code for a function	14	35. Built-in functions for getting far-pointer components	38
9. SCJTSAM(CJTZGBL)	17	36. Library functions for use only in AR-mode functions	38
10. SCJTSAM(MYPROLOG)	19	37. Allocation and deallocation routines	40
11. SCJTSAM(MYEPILOG).	20	38. Copying a C string pointer to a far pointer	42
12. Simple code format string.	22	39. Example of a simple dereference of a far pointer	43
13. Code format string with two instructions	22	40. Metal C application build process	45
14. Code format string with two instructions, formatted for readability	22	41. C source file (mycode.c) that builds a Metal C program	46
15. Substitution of a C variable into an output __asm operand	23	42. C compiler invocation to generate mycode.s	46
16. HLASM source code embedded by the compiler	23	43. Command that invokes HLASM to assemble mycode.s	46
17. Substitution of a C pointer into an __asm operand.	24	44. Command that compiles an HLASM source file containing symbols longer than eight characters	46
18. __asm operand lists.	24	45. Command that binds mycode.o and produces a Metal C program in an MVS data set	47
19. Compiler-generated HLASM code from the __asm statement	25	46. Commands that compile and link programs with different addressing modes	47
20. Unsuccessful attempt to specify registers	25	47. Job step that compiles HLQ.SOURCE.C(MYCODE)	48
21. Register specification with clobbers	26	48. Assembly step of HLQ.SOURCE.ASM(MYCODE).	48
22. Incorrect __asm operand definition for both input and output.	26	49. Job step that binds the generated HLASM object into a program	48
23. Incorrect compiler-generated HLASM source code from the incorrect __asm operand definition for both input and output	26	50. The process of building Metal C programs with IPA	49
24. Successful definition of an __asm operand for both input and output	27	51. JCL that invokes the ASMLANGX utility	52
25. Correct compiler-generated HLASM source code from the correct __asm operand definition for both input and output	27		
26. The + constraint to define an __asm operand for both input and output.	28		
27. Error: Redundant definition of an __asm operand.	28		

Tables

1. Typographical conventions	ix	8. Definitions in float.h	56
2. Syntax examples	xi	9. csysenv argument in __cinit()	72
3. Compiler-generated global SET symbols	15	10. csysenvtkn argument in __cterm()	75
4. User modifiable global SET symbols	17	11. Flag Characters for sprintf() Family	89
5. Language constructs that may have special impact on far pointers	36	12. Precision Argument in sprintf()	91
6. Implicit ALET associations for AR-mode-function variables	37	13. Type Characters and their Meanings	92
7. Summary of useful references for the Metal C programmer	52	14. Conversion Specifiers in sscanf()	97
		15. Stack frame requirements for Metal C runtime functions	125

About this document

This document contains reference information that is intended to help you understand the IBM® Enterprise Metal C for z/OS® compiler and use the header files and functions provided by the compiler to write applications that can be compiled using the Enterprise Metal C for z/OS compiler.

Typographical conventions

The following table explains the typographical conventions used in this document.

Table 1. *Typographical conventions*

Typeface	Indicates	Example
bold	Commands, executable names, compiler options and pragma directives that contain lower-case letters.	The metal invocation command invokes the Enterprise Metal C for z/OS compiler.
<i>italics</i>	Parameters or variables whose actual names or values are to be supplied by the user. Italics are also used to introduce new terms.	Make sure that you update the <i>size</i> parameter if you return more than the <i>size</i> requested.
monospace	Programming keywords and library functions, compiler built-in functions, file and directory names, examples of program code, command strings, or user-defined names.	If one or two cases of a <code>switch</code> statement are typically executed much more frequently than other cases, break out those cases by handling them separately before the <code>switch</code> statement.

Softcopy documents

The Enterprise Metal C for z/OS publications are supplied in PDF format and available for download from the Enterprise Metal C for z/OS Knowledge Center home page (www.ibm.com/support/knowledgecenter/en/SSSHGK_3.1.0/com.ibm.metalc.v3r1.doc/welcome.html).

To read a PDF file, use the Adobe Reader. If you do not have the Adobe Reader, you can download it (subject to Adobe license terms) from the Adobe website (www.adobe.com).

You can also browse the documents on the World Wide Web by visiting the Enterprise Metal C for z/OS Knowledge Center home page (www.ibm.com/support/knowledgecenter/en/SSSHGK_3.1.0/com.ibm.metalc.v3r1.doc/welcome.html).

Who should read this document

This document is intended for application programmers interested in writing Metal C applications using the Enterprise Metal C for z/OS compiler.

Where to find more information

For an overview of the information associated with z/OS, see *z/OS Information Roadmap*.

Additional information on Enterprise Metal C for z/OS is available on the Marketplace page for Enterprise Metal C for z/OS (www.ibm.com/us-en/marketplace/xl-cpp-compiler-zos).

z/OS Basic Skills in IBM Knowledge Center

z/OS Basic Skills in IBM Knowledge Center is a Web-based information resource intended to help users learn the basic concepts of z/OS, the operating system that runs most of the IBM mainframe computers in use today. IBM Knowledge Center is designed to introduce a new generation of Information Technology professionals to basic concepts and help them prepare for a career as a z/OS professional, such as a z/OS system programmer.

Specifically, z/OS Basic Skills is intended to achieve the following objectives:

- Provide basic education and information about z/OS without charge
- Shorten the time it takes for people to become productive on the mainframe
- Make it easier for new people to learn z/OS.

z/OS Basic Skills in IBM Knowledge Center (www.ibm.com/support/knowledgecenter/zosbasics/com.ibm.zos.zbasics/homepage.html) is available to all users (no login required).

How to read syntax diagrams

This section describes how to read syntax diagrams. It defines syntax diagram symbols, items that may be contained within the diagrams (keywords, variables, delimiters, operators, fragment references, operands) and provides syntax examples that contain these items.

Syntax diagrams pictorially display the order and parts (options and arguments) that comprise a command statement. They are read from left to right and from top to bottom, following the main path of the horizontal line.

For users accessing the Information Center using a screen reader, syntax diagrams are provided in dotted decimal format.

Symbols

The following symbols may be displayed in syntax diagrams:

Symbol	Definition
--------	------------

- | | |
|-----|--|
| ▶▶— | Indicates the beginning of the syntax diagram. |
| —▶ | Indicates that the syntax diagram is continued to the next line. |
| ▶— | Indicates that the syntax is continued from the previous line. |
| —▶◀ | Indicates the end of the syntax diagram. |

Syntax items

Syntax diagrams contain many different items. Syntax items include:

- Keywords - a command name or any other literal information.
- Variables - variables are italicized, appear in lowercase, and represent the name of values you can supply.
- Delimiters - delimiters indicate the start or end of keywords, variables, or operators. For example, a left parenthesis is a delimiter.
- Operators - operators include add (+), subtract (-), multiply (*), divide (/), equal (=), and other mathematical operations that may need to be performed.
- Fragment references - a part of a syntax diagram, separated from the diagram to show greater detail.
- Separators - a separator separates keywords, variables or operators. For example, a comma (,) is a separator.

Note: If a syntax diagram shows a character that is not alphanumeric (for example, parentheses, periods, commas, equal signs, a blank space), enter the character as part of the syntax.

Keywords, variables, and operators may be displayed as required, optional, or default. Fragments, separators, and delimiters may be displayed as required or optional.

Item type

Definition

Required

Required items are displayed on the main path of the horizontal line.

Optional

Optional items are displayed below the main path of the horizontal line.

Default

Default items are displayed above the main path of the horizontal line.

Syntax examples

The following table provides syntax examples.

Table 2. Syntax examples

Item	Syntax example
Required item.	►►—KEYWORD—required_item—►◄
Required items appear on the main path of the horizontal line. You must specify these items.	
Required choice.	►►—KEYWORD— required_choice1 required_choice2 —►◄
A required choice (two or more items) appears in a vertical stack on the main path of the horizontal line. You must choose one of the items in the stack.	
Optional item.	►►—KEYWORD— optional_item —►◄
Optional items appear below the main path of the horizontal line.	

Table 2. Syntax examples (continued)

Item	Syntax example
Optional choice. An optional choice (two or more items) appears in a vertical stack below the main path of the horizontal line. You may choose one of the items in the stack.	
Default. Default items appear above the main path of the horizontal line. The remaining items (required or optional) appear on (required) or below (optional) the main path of the horizontal line. The following example displays a default with optional items.	
Variable. Variables appear in lowercase italics. They represent names or values.	
Repeatable item. An arrow returning to the left above the main path of the horizontal line indicates an item that can be repeated. A character within the arrow means you must separate repeated items with that character.	
An arrow returning to the left above a group of repeatable items indicates that one of the items can be selected, or a single item can be repeated.	
Fragment. The fragment symbol indicates that a labelled group is described below the main syntax diagram. Syntax is occasionally broken into fragments if the inclusion of the fragment would overly complicate the main syntax diagram.	fragment:

Technical support

Additional technical support is available from the Enterprise Metal C for z/OS Support page (https://www.ibm.com/support/home/product/A032956W76367A04/IBM_Enterprise_Metal_C_for_z/OS). This page provides a portal with search capabilities to technical support FAQs and other support documents.

For the latest information about Enterprise Metal C for z/OS, visit Marketplace page for Enterprise Metal C for z/OS (www.ibm.com/us-en/marketplace/xl-cpp-compiler-zos).

If you cannot find what you need, you can e-mail:

compinfo@cn.ibm.com

How to send your comments to IBM

We appreciate your input on this documentation. Please provide us with any feedback that you have, including comments on the clarity, accuracy, or completeness of the information.

You can send an email to compinfo@cn.ibm.com and include the following information:

- Your name and address
- Your email address
- Your phone or fax number
- The publication title and order number:
Enterprise Metal C for z/OS Programming Guide and Reference
SC27-9049-00
- The topic and page number or URL of the specific information to which your comment relates
- The text of your comment.

When you send comments to IBM, you grant IBM a nonexclusive right to use or distribute the comments in any way appropriate without incurring any obligation to you.

IBM or any other organizations use the personal information that you supply to contact you only about the issues that you submit.

If you have a technical problem

If you have a technical problem, take one or more of the following actions:

- Visit the IBM Support Portal (support.ibm.com).
- Contact your IBM service representative.
- Call IBM technical support.

Chapter 1. About IBM Enterprise Metal C for z/OS

The IBM Enterprise Metal C for z/OS compiler generates code that does not require access to the Language Environment[®] support at run time. The Enterprise Metal C for z/OS compiler provides C-language extensions that allow you to specify assembly statements that call system services directly. Using these language extensions, you can provide almost any assembly macro, and your own function prologs and epilogs, to be embedded in the generated HLASM source file. When you understand how the METAL-generated code uses MVS linkage conventions to interact with HLASM code, you can use this capability to write freestanding programs.

Because a freestanding program does not depend on any supplied runtime environment, it must obtain the system services that it needs by calling assembler services directly. For information about how METAL-generated code uses MVS linkage conventions, see “Metal C and MVS linkage conventions” on page 2. For information about embedding assembly statements in the METAL-generated HLASM source code, see “Inserting HLASM instructions into the generated source code” on page 21.

You do not always have to provide your own libraries. IBM supplies the Metal C runtime library functions. This subset includes commonly used basic functions such as `malloc()`. For more information, see Chapter 3, “C functions available to Metal C programs,” on page 65.

Note: You can use these supplied functions or the ones that you provide yourself.

Metal C environment

Some of the functions require that an environment be created before they are called. You can create the environment by using a new function, `__cinit()`. This function sets up the appropriate control blocks and returns an environment token to the caller. The caller must then ensure that GPR 12 contains this token when calling Metal C functions that require an environment. When the environment is no longer needed, a new function, `__cterm()`, can be used to perform cleanup, freeing all resources that had been obtained by using the token.

An environment created by `__cinit()` can be used in both AMODE 31 and AMODE 64. In conjunction with this, the Metal C run time maintains both a below-the-bar heap and an above-the-bar heap for each environment. Calls to `__malloc31()` always affect the below-the-bar heap. Calls that are made in AMODE 31 to all other functions that obtain storage will affect the below-the-bar heap; calls that are made in AMODE 64 affect the above-the-bar heap.

The storage key for all storage obtained on behalf of the environment is the psw key of the caller. The caller needs to ensure that the environment is always used with the same or compatible key.

Metal C environments are intended to be used serially by a single dispatchable unit of work. If you need to share environments between multiple dispatchable units, you must make sure that the use of each environment is serialized.

Programming with Enterprise Metal C for z/OS

When you want to build a C program that can run in any z/OS environment, you can use the Enterprise Metal C for z/OS compiler as a high level language (HLL) alternative to writing the program in assembly language.

Enterprise Metal C for z/OS facilitates direct use of operating system services. For example, you can use the C programming language to write installation exits.

The Enterprise Metal C for z/OS compiler:

- Generates code that is independent of Language Environment.

Note: Although the compiler generates default prolog and epilog code that allows the Metal C code to run, you might be required to supply your own prolog and epilog code to satisfy runtime environment requirements.

- Generates code that follows MVS linkage conventions. This facilitates inter-operations between the Metal C code and the existing code base. See “Metal C and MVS linkage conventions.”

Note: Metal C also provides a feature that improves the program's runtime performance. See “NAB linkage extension” on page 3.

- Provides support for accessing the data stored in data spaces. See “AR-mode programming support” on page 35.
- Provides support for embedding your assembly statements into the compiler-generated code. See “Inserting HLASM instructions into the generated source code” on page 21.

If you use the optimization capabilities of Enterprise Metal C for z/OS, you can use C to write highly optimized system-level code.

Metal C and MVS linkage conventions

Because Metal C follows MVS linkage conventions, it enables the compiler-generated code to inter-operate directly with the existing code base to facilitate the following operations:

- Passing parameters. See “Parameter passing.”
- Returning values. See “Return values” on page 3.
- Setting up function save areas. See “Function save areas” on page 3.

For detailed information about MVS linkage conventions, see the topic about linkage conventions in *z/OS MVS Programming: Assembler Services Guide*.

Parameter passing

The pointer to the parameter list is in GPR 1.

The parameter list is divided into slots.

- The size of each slot depends on the addressing mode:
 - For 31-bit mode (AMODE 31), each slot is four bytes in length.
 - For 64-bit mode (AMODE 64), each slot is eight bytes in length.
- Metal C derives the content of each slot from the function prototype, which follows C by-value semantics (that is, the value of the parameter is passed into the slot).

Notes:

1. If a parameter needs to be passed by reference, the function prototype must specify a pointer of the type to be passed.
2. Under AMODE 31 only: The high-order bit is set on the last parameter if both of the following are true:
 - The called function is a variable arguments function.
 - The last parameter passed is a pointer.

Return values

For any addressing mode, integral type values are returned in GPR 15. Under AMODE 31 only, a 64-bit integer value is returned in GPR 15 + GPR 0 (that is, the high-half of the 64-bit value is returned in GPR 15 and the low-half is returned in GPR 0). All other types are returned in a buffer whose address is passed as the first parameter.

Function save areas

GPR 13 contains the pointer to the dynamic storage area (DSA).

The DSA includes:

- 72-byte save area size for an AMODE 31 function.
- Parameter area for calling other functions. The default pointer size for a parameter or return value is based on the amode attribute of the function prototype.
- Temporary storage that is preallocated for the compiler-generated code and the user-defined automatic variables.

The save area is set up at the beginning of the DSA.

If the function calls only primary-mode functions, the save area format depends on the AMODE:

- Under AMODE 31, the save area takes the standard 18-word format.
- Under AMODE 64, the save area takes the 36-word F4SA format and the compiler will generate code to set up the F4SA signature in the second word of the save area.

If the function needs to call an AR-mode function, the save area will take the 54-word F7SA format, regardless of the addressing mode.

The F4SA signature generation can be suppressed by setting the &CCN_SASIG global SET symbol to 0 in your prolog code. For information about the &CCN_SASIG global SET symbol, see Table 4 on page 17 User modifiable global SET symbols.

NAB linkage extension

Metal C code needs to use dynamic storage area (DSA) as stack space. Each time a function is called, its prolog code acquires this space and, when control is returned to the calling function, its epilog code releases the stack space.

Metal C avoids excessive acquisition and release operations by providing a mechanism that allows a called function to rely on pre-allocated stack space. This mechanism is the next available byte (NAB). All Metal C runtime library functions, as well as functions with a default prolog code, use it and expect the NAB address

to be set by the calling function. The code that is generated to call a function includes the setup instructions to place the NAB address in the "Address of next save area" field in the save area. The called function simply goes to the calling function's save area to pick up the NAB address that points to its own stack space. As a result, the called function does not need to explicitly obtain and free its own stack space.

Note: If usage of the NAB linkage extension requires more stack space than has been allocated, there will be unexpected results. The program must establish a DSA that is large enough to ensure the availability of stack space to all downstream programs. Downstream programs include all functions that are defined in the program as well as the library functions listed in Appendix A, "Function stack requirements," on page 125.

The location of the "Address of next save area" field depends on the save area format:

- In the standard 72-byte save area, it is the third word.
- In the F4SA or F7SA save area, it is the 18th doubleword.

Compiler-generated HLASM source code

The Enterprise Metal C for z/OS compiler generates code in the HLASM source code format.

Characteristics of compiler-generated HLASM source code

Any assembly instructions that you provide need to work with the instructions that are generated by the compiler. Before you provide those instructions, you need to be aware of the characteristics of compiler-generated HLASM source code.

You need to be aware that:

- Because the compiler uses relative-branching instructions, it is not necessary to establish code base registers. When the compiler detects user-embedded assembly statements, you can use the IEABRCX DEFINE instruction to assist any branching instructions that might rely on establishment of a code base register. For other instructions (such as LA or EX) that rely on the establishment of a code base register, you might need to add code to establish your own code base register. To disable the effect of IEABRCX, you can add the instruction IEABRCX DISABLE. For more information about the IEABRCX macro, see *z/OS MVS Programming: Authorized Assembler Services Reference EDT-IXG*.
- If the compiler needs to produce literals, GPR 3 will be set up as the base register to address the literals. This addressability is established after the prolog code. The literals are organized by the LTORG instruction placed at the end of the epilog code. With the presence of user-embedded assembly statements, the compiler assumes there will be literals and establishes GPR 3 to address those literals.
- If you want code to be naturally reentrant, you must not use writable static or external variables; such variables are part of the code.
- There is only one CSECT for each compilation unit. The CSECT name is controlled by the CSECT option.
- Due to the flat name space and the case insensitivity required by HLASM, the compiler prepends extra qualifiers to user names to maintain the uniqueness of each name seen by HLASM. This is referred to as *name encoding*. External symbols are not subject to the name-encoding scheme as they need to be referenced by the exact symbol names.
- The external symbols are determined by the compiler LONGNAME option.

- If the NOLONGNAME option is in effect:
 - All external symbols are truncated to eight characters.
 - All external symbols are converted to upper case.
 - All "_" characters are replaced with the "@" character.
- If the LONGNAME option is in effect the compiler emits an ALIAS instruction to make the real C name externally visible. Because the length limit supported by the ALIAS instruction depends on the HLASM release, the C compiler does not enforce any length limit here.

Note: The HLASM GOFF option is necessary to allow the ALIAS instructions to be recognized. See Figure 44 on page 46.

- GPR 13 is established as the base of the stack space.
- GPR 10 and GPR 11 may be used exclusively to address static and constant data. They should not be used in the user-embedded assembly statements.
- The compiler will generate code to preserve FPR 8 through FPR 15 if they are altered by the function.
- For AMODE 31 functions: The compiler will generate code to preserve the high halves of the 64-bit GPRs if they are altered or if there are user-embedded assembly statements.
- The addressing mode is determined by the compiler option. When the compiler option LP64 is in effect, the addressing mode is AMODE 64; otherwise it is AMODE 31.

Structure of a compiler-generated HLASM source program

Each compiler-generated HLASM source program has the following elements:

- File-scope header
- For each function:
 - A function header
 - A function entry point marker
 - A function property block (FPB)
 - A function body
 - A function trailer
- File-scope trailer

File-scope header:

Statements in the file-scope header apply to the entire compilation unit and might have the following statements:

- TITLE statement to specify the product information of the compiler and the source file being compiled.
- ALIAS/EXTRN statement to declare the external symbols that are referenced in the program, if the LONGNAME compiler option is in effect.
- CSECT statement to identify the relocatable control section in the program.
- AMODE statement to specify the addressing mode.
- RMODE statement to specify the residency mode for running the module.
- Assembly statements to declare the HLASM global SET symbols used by the compiler-generated code for communicating information to the user-embedded prolog and epilog code, if the compiler detects user-embedded prolog and epilog code.

- SYSSTATE ARCHLVL statement, which identifies the minimum hardware requirement. SYSSTATE ARCHLVL=3, if and only if **ARCH(7)** or up and **OSREL(ZOSV2R1)** or higher are in effect; otherwise, SYSSTATE ARCHLVL=2.
- IEABRCX DEFINE statement ensures that all branch instructions are changed to relative-branching instructions, in the event that the compiler encounters user-embedded assembly statements.
- Prefix data to embed a compiler signature and to record attributes about the compilation.

Prefix data

Prefix data is generated to supply a signature, the timestamp of the compilation date, the compiler version, and some control flags. It is placed at the beginning of the code that follows an instruction for branching around the prefix data.

Note: Program code should reference ENTRY rather than CSECT to avoid unnecessary branching.

The prefix data consists of a fixed part (36 bytes in size) followed by a contiguous optional part, with the presence of optional fields indicated by flag bits in flag set 4. Optional fields, if present, are stored immediately following the fixed part of the prefix data aligned on halfword boundaries in the order specified by the left to right bits in flag set 4.

Figure 1 shows the prefix data fixed area fields and definitions.

+0	X'00C300C300D50000'			
+4	Signature			
+8				
+12	Compile Date in YYYYMMDD format			
0xC				
+16				
0x10	Compile Time in HHMMSS format		Compiler Version	
+20				
0x14				
+24	Compiler Version (continue)		X'0000' Reserved	
0x18				
+28	Flag Set 1	Flag Set 2	Flag Set 3	Flag Set 4
0x1C				
+32	X'00000000' Reserved			
0x20				

Figure 1. Prefix data fixed area fields

Signature

An 8-byte field that is set to 0x00C300C300D50000. The last byte in the signature is the version number which can change in future releases.

Compile date

An 8-byte field that contains the date of the compile in YYYYMMDD format.

Compile time

A 6-byte field that contains the time of the compile in HHMMSS format.

Compiler version

A 4-byte field that contains the binary value of the compiler version and release.

Flag Set 1

Flag definition

'1.....' Compiled with RENT option.

'0.....' Compiled with NORENT option.

'0.000000'

Reserved.

Flag Set 2

Flag definition

'00000000'

Reserved.

Flag Set 3

Flag definition

'00000000'

Reserved.

Flag Set 4

Flag definition

'1.....' Indicates the presence of a user comment string.

'0.....' Indicates no optional user comment string.

'1.....' Indicates the presence of a service string.

'0.....' Indicates no service string.

'..1.....' Indicates the presence of the offset of the end of the current CSECT.

'..0.....' Indicates no offset of the end of the current CSECT.

'...1....' Indicates the presence of the offset of the debug information block.

'...0....' Indicates no offset of the debug information block.


'....0000'

Reserved.

There are four optional prefix data fields, whose presence is indicated by a flag bit in flag set 4.

Length of User Comment (Bit 0)	User Comment String	Variable Length
User Comment String (continue)		

User Comment String: The user comment string comes from the string specified in both or one of #pragma comment(copyright, "...") and #pragma comment(user, "..."). If you have either or both #pragma, the flag bit is set to one, and the user comment string contains the concatenated strings from multiple #pragma.



Service String: The service string comes from the string specified in the SERVICE compiler option.



Signed offset of the end of current CSECT: A 32-bit field that contains the signed offset from PFD (Prefix Data) to the end of the current CSECT.



Signed offset of the debug data block: A 32-bit field that contains the signed offset from PFD to the debug data block.

Examples

Figure 2 and Figure 3 show how prefix data is generated from a sample program that is compiled with the RENT, SERVICE ("Service String"), and DEBUG (or -g) options.

```
#pragma comment(copyright,"copyright comment")
#pragma comment(user,"user comment")
int main(){
    return 0;
}
```

Figure 2. A sample program to generate prefix data

@@PFD@@	DC	XL8'00C300C300D50000'	Prefix Data Marker	000008
	DC	CL8'20180615'	Compiled Date YYYYMMDD	000008
	DC	CL6'152705'	Compiled Time HHMMSS	000008
	DC	XL4'42033010'	Compiler Version	000008
	DC	XL2'0000'	Reserved	000008
	DC	BL1'00000000'	Flag Set 1	000008
	DC	BL1'00000000'	Flag Set 2	000008
	DC	BL1'00000000'	Flag Set 3	000008
	DC	BL1'11110000'	Flag Set 4	000008
	DC	XL4'00000000'		000008
	DS	0H		000008
	DC	AL2(30)		000008
	DC	C'copyright comment user comment'		000008
	DS	0H		000008
	DC	AL2(14)		000008
	DC	C'Service String'		000008
	DC	A(@@END@@-@@PFD@@)		000008
	DC	A(@@DDB@@-@@PFD@@)		000008

Figure 3. Prefix data generated

Function elements: Function header

The function header might have the following statements or code:

- ALIAS/ENTRY statement to define the entry point by associating its C symbol with the generated HLASM name, if the LONGNAME compiler option is in effect.
- Assembly statements to set the values for the declared HLASM global SET symbols, if the compiler detects user-embedded prolog and epilog code.
- Prolog code, which might be either the default prolog code generated by the compiler or user-embedded prolog code.

Function entry point marker

A function entry point marker is generated immediately before the entry point of each function. The function entry point marker is an 8-byte field containing the signature 0x00C300C300D501nn. Immediately following the marker is a 4-byte signed offset from the start of the entry point marker to the function property block belonging to the current function. Figure 4 shows what a function entry point marker looks like in generated assembler code.

```

ENTRY @CCN@2                                000005
@CCN@2 AMODE 31                             000005
      DC XL8'00C300C300D50100'           Function entry point marker 000005
      DC A(@FPB@1-**+8)                   Signed offset to FPB      000005
      DC XL4'00000000'                     Reserved                000005
@CCN@2 DS 0F

```

Figure 4. Function entry point marker in generated assembler code

Function property block

The function property block (FPB) is made up of a fixed part (20 bytes in size) followed by a contiguous optional part, with the presence of optional fields indicated by flag bits. Optional fields, if present, are stored immediately following the fixed part of the FPB aligned on fullword boundaries in the order specified below.

Figure 5 shows the FPB fixed area fields and the definitions.

+0	X'CCD5' Eyecatcher		Saved GPR Mask	
+4	Signed offset to Prefix Data from the start of FPB			
+8	Flag Set 1	Flag Set 2	Flag Set 3	Flag Set 4
+12 0xC	X'00000000' Reserved			
+16 0x10	X'00000000' Reserved			

Figure 5. Function property block fixed area fields

Eyecatcher

A 16-bit field that is set to 0xCCD5.

Saved GPR Mask

A 16-bit mask, indicating which registers are saved and restored by the associated routine. Bit 0 indicates register 0, followed by bits for registers 1 to 15 in order.

Signed offset to Prefix Data from the start of FPB

The offset of the prefix data belonging to the compilation unit containing the function described by this FPB.

Flag Set 1

Flag definition

- '1.....' Function is AMODE 64.
- '0.....' Function is AMODE 31.
- '1.....' Function is AR mode.
- '0.....' Function is primary mode.
- '..00000.'
- Reserved.
- '.....1' A vararg function.
- '.....0' Not a vararg function.

Flag Set 2

Flag definition

- '1.....' External function.
- '0.....' Internal function.
- '.....00' This function has the standard 72-byte save area.
- '.....01' This function has the F4SA 144-byte save area.
- '.....10' This function has the F7SA 216-byte save area.
- '00000..'
- Reserved.

Flag Set 3

Flag definition

- '1.....' Indicates the floating-point registers (FPR) are saved in the DSA and the FPR mask and offset to the FPR save area are present in the optional part of the FPB.
- '0.....' Indicates the floating-point registers (FPR) are not saved in the DSA.
- '1.....' Indicates the high-half of 64-bit general purpose registers (GPR) are saved in the DSA and the HGPR mask and offset to the HGPR save area are present in the optional part of the FPB.
- '0.....' Indicates the high-half of 64-bit general purpose registers (GPR) are not saved in the DSA.
- '..000000'
- Reserved.

Flag Set 4

Flag definition

- '0000000.'
- Reserved.
- '.....1' Indicates that the length of the function name and the function name field are present in the optional part of the FPB.

'.....0' Indicates that the function name field is not present in the FPB.

Note: When the COMPRESS compiler option is in effect, the function name field is not present in the FPB.

There are several optional FPB fields. The presence of each field is indicated by a flag bit in FPB flag set 3 or FPB flag set 4. When an optional field is less than 4 bytes in length, the entire word is present if any of the fields in that word are present. Unused parts of the word are filled with zeroes. The optional fields are fullword aligned and appear in the order listed below.

FPR Mask (Flag Set 3 Bit 0)	HGPR Mask (Flag Set 3 Bit 1)	Length = 4
-----------------------------	------------------------------	------------

FPR Mask

A 16-bit mask indicating which of floating-point registers (FPR) are saved and restored by this function. Bit 0 indicates FPR0, followed by bits for FPR1 to FPR 15.

HGPR Mask

A 16-bit mask indicating which of 64-bit general purposes registers (GPR) whose high-words are saved and restored by this function. Bit 0 indicates GPR0, followed by bits for GPR1 to GPR 15.

Note: If either bit 0 or bit 1 of flag set 3 is on, the fullword variable representing FPR mask and HGPR mask is present.

FPR Savearea Offset (Flag Set 3 Bit 0)	Length = 4
--	------------

FPR Savearea Offset

A 32-bit field containing the offset from the start of the DSA to the FPR save area. The contents of the FPRs indicated by the FPR Mask are stored contiguously from the beginning of the FPR save area regardless of the register number. For example, if the FPR Mask contains 0x00A0 and the FPR save area offset contains 0x100, FPR8 is stored at R13+0x100 and FPR10 is stored at R13+0x108.

HGPR Savearea Offset (Flag Set 3 Bit 1)	Length = 4
---	------------

HGPR Savearea Offset

A 32-bit field containing the offset from the start of the DSA to the HGPR save area. The HGPR save area is 64-byte in size which holds 16 32-bit high-words. The order of the high-words stored in the save area is GPR14, GPR15, and GPR0 - GPR13. Only the slots which correspond to the bits in the HGPR Mask contain the saved high-word contents.

Length of Name (Flag Set 4 Bit 7)	Name of Function	Variable Length
Name of Function (continue)		

Name of Function

The optional function name fields start with a 2-byte length field which contains the actual length of the function name that follows.

Figure 6 shows what a function property block looks like in generated assembler code.

```

@@FPB@ LOCTR
@@FPB@1 DS 0F          Function Property Block          000000
DC XL2'CCD5'           Eyecatcher                      000000
DC BL2'111110000000011' Saved GPR Mask                000000
DC A(@@PFD@-@@FPB@1)  Signed Offset to Prefix Data    000000
DC BL1'00000000'       Flag Set 1                     000000
DC BL1'10000000'       Flag Set 2                     000000
DC BL1'01000000'       Flag Set 3                     000000
DC BL1'00000001'       Flag Set 4                     000000
DC XL4'00000000'       Reserved                        000000
DC XL4'00000000'       Reserved                        000000
DC XL2'0000'           Saved FPR Mask                  000000
DC BL2'111100000000011' Saved HGPR Mask               000000
DC XL4'00000058'       HGPR Save Area Offset           000000
DC AL2(4)
DC C'main'

```

Figure 6. Function property block in generated assembler code

In this example, the @@FPB@ LOCTR instruction tells the assembler to group all FPBs separate from the code and data generated for the functions.

File-scope trailer:

The file-scope trailer might have the following statements or areas:

- DC statements to define static variables with their initial values.
- DSECT statement to provide a map for the static variables.
- DC statements that define constants.
- ALIAS/ENTRY statement to define all external variables with their initial values.
- END statement to specify compiler product information and the compilation date.
- A debug data block to check whether the debug side file matches the object file.
- A label to mark the end of a CSECT.

Debug data block:

When you specify the -g or DEBUG option, a new debug data block is added for each CSECT in the Metal C generated assembly file. The debug data block can be used to check whether the debug side file matches the object file. When a debug session starts, the debugger checks whether the debug data block for the first CSECT exists. If it exists, the debugger retrieves the information from the debug data block.

The debug data block contains the following information:

Debug data block signature

An 8-byte field that is set to 0x'00C300C300D50200'. The last two bytes of the signature is 0x'0200', in which the first byte 0x'02' is the block signature for the debug data block, while the second byte 0x'00' represents the current version of the debug data block.

Size of the debug data block

A 4-byte field that represents the size of the debug data block.

Reserved bytes

A 4-byte field that is reserved for future usage.

MD5 signature

A 16-byte field that is derived from the time stamp.

Source file name

A string with 2-byte prefix length, containing the source file name.

Debug side file name

A string with 2-byte prefix length, containing the debug side file name.

Example

The following example shows a sample debug data block followed by a CSECT end label.

```

@@DDBE@@ DC XL8'00C300C300D50200'
          DC A(@@DBGEE@@-@@DBGGB@@)
          DC XL4'00000000'
@@DSIG@@ DC X16'194ab396eb68ebd68d476285b476abda'
          DC AL2(23)
          DC C'/home/temp/a.c'
@@DSFN@@ DC AL2(25)
          DC C'/home/temp/a.dbg'
@@DDBE@@ EQU *
@@END@@  EQU *

```

Figure 7. Debug data block generated

Labels for @@DDBE@@ and @@END@@ do not have to be the same.

Prolog and epilog code

The primary functions of prolog code are:

- To save the calling function's general-purpose registers in the calling function's save area.
- To obtain the dynamic storage area for this function.
- To chain this function's save area to the calling function's save area, in accordance with the MVS linkage convention.

The primary functions of epilog code are:

- To relinquish this function's dynamic storage area.
- To restore the calling function's general-purpose registers.
- To return control to the calling function.

Note: AR-mode functions require additional prolog and epilog functions. See "AR-mode programming support" on page 35 for details.

Supplying your own prolog and epilog code

If you need the prolog and epilog code to provide additional functionality, you can use `#pragma` directives to instruct the compiler to use your own HLASM prolog and epilog code. Figure 8 provides an example.

```
#pragma prolog(myfunc,"MYPROLOG")
#pragma epilog(myfunc,"MYEPILOG")
int myfunc() {
    return 0;
}
```

Figure 8. Specification of your own prolog and epilog code for a function

To apply the same prolog and epilog code to all your functions in the C source file, use the `PROLOG` and `EPILOG` compiler options. When you use the `PROLOG` and `EPILOG` compiler options, by default, your prolog and epilog code is applied only to the functions that have external linkage. To apply your prolog and epilog code to all functions defined in the compilation unit, use the **all** suboption.

The string you supplied to the `PROLOG` and `EPILOG` options or the `#pragma` directives must contain valid HLASM statements. The compiler does not validate the content of the string but it does take care of some formatting for you:

- If your string contains only a macro name, as shown in Figure 8, you do not need to supply leading blanks.
- If the length of your HLASM statement exceeds 71 characters, you do not need to break it into multiple lines. The compiler will handle that for you.

Your prolog code needs to ensure that:

- The primary functions of the prolog code have been performed.
- Extra DSA space is acquired, in the event that the NAB is needed for the referenced functions.
- Upon exit of your prolog code:
 - GPR 13 points at the DSA for this function.
 - GPR 1 points at the parameter list supplied by the calling function.

Your epilog code needs to ensure that:

- The primary functions of the epilog code have been performed.
- The content of GPR 15, on entry to your epilog code, is preserved.
- If a 64-bit integer value is returned from an AMODE 31 program, the low half of the return value contained in GPR 0 is preserved.

Your prolog and epilog code does not need to perform the following functions:

- Preserve the calling function's floating-point registers.
- Preserve the calling function's vector registers.
- Preserve the high-halves of 64-bit general purpose registers in AMODE 31 functions.
- Preserve the registers used by the compiler generated code.
- Set up the NAB for the called functions.

User reserved DSA space

User reserved DSA space can be enabled by using the compiler option **DSAUSER(value)**. If **DSAUSER** is specified without a suboption, a user field with the

size of a pointer is reserved on the stack. The user field is a 4-byte field for AMODE 31 and an 8-byte field for AMODE 64. This user field can be utilized by your prolog or epilog code.

If a *value* suboption is specified with DSAUSER, a user field with the size of *value* 32-bit words is allocated. Specifying DSAUSER with a suboption requires ARCH(6). The *value* must be an integer in the range of 0 to 50.

The user field can be located by the HLASM global set symbol &CCN_DSAUSER, which provides the offset to the user field. The compiler allocates the field on the stack only, without initializing it.

The following example shows how &CCN_DSAUSER is set by the compiler:

```
&CCN_DSAUSER SETC '#USER_2-@@AUTO@2'
```

The following example shows how &CCN_DSAUSER can be used in your prolog code:

```
STG 0,&CCN_DSAUSER.(,13)
```

Compiler-generated global SET symbols

When you supply your prolog and epilog code, the compiler generates the assembly instructions that set up global SET symbols for communicating compiler-collected information to your prolog and epilog code. Your prolog and epilog code can use this information to determine the code sequence generated by your macros.

Table 3 describes global SET symbols defined by the compiler.

Table 3. Compiler-generated global SET symbols

Global SET symbol	Type	Description
&CCN_DSASZ	Arithmetic	The size of the dynamic storage area for the function.
&CCN_SASZ	Arithmetic	The size of the function save area: • 72 = standard format • 144 = F4SA format • 216 = F7SA format
&CCN_ARGS	Arithmetic	The number of fixed arguments expected by the function.
&CCN_RLOW	Arithmetic	The starting register number to be used in the STORE MULTIPLE instruction for saving the registers of callers if the compiler were to generate that instruction itself.
&CCN_RHIGH	Arithmetic	The ending register number to be used in the STORE MULTIPLE instruction for saving the registers of callers.
&CCN_LP64	Logical	Set to "1" if the LP64 compiler option is specified.
&CCN_NAB	Logical	Set to "1" when there are called programs that depend on the dynamic storage to be pre-allocated. In this case, the prolog code needs to add a generous amount to the size set in &CCN_DSASZ when the dynamic storage is obtained.
&CCN_ALTGPR(16)	Logical	The array representing the general purpose registers. Subscript 1 represents GPR 0 and subscript 16 represents GPR 15. A subscript is set to "1" whenever the corresponding register is altered by the compiler-generated code.
&CCN_STATIC	Logical	Set to "1" if the function is static.
&CCN_MAIN	Logical	Set to "1" if this is function "main".

Table 3. Compiler-generated global SET symbols (continued)

Global SET symbol	Type	Description
&CCN_RENT	Logical	Set to "1" if the RENT compiler option is specified.
&CCN_PRCN	Character	The symbol representing the function.
&CCN_CSECT	Character	The symbol representing the CSECT in effect.
&CCN_DSAUSER	Character	The assembly time computed offset to the user field on the stack of the function.
&CCN_LITN	Character	The symbol representing the LTORG generated by the compiler.
&CCN_BEGIN	Character	The symbol representing the first executable instruction of the function generated by the compiler.
&CCN_ARCHLVL	Character	The symbol representing the architecture level specified in the ARCH option.
&CCN_ASCM	Character	The ASC mode of the function: • A=AR mode • P=Primary mode For information about AR mode, see "AR-mode programming support" on page 35.
&CCN_NAB_OFFSET	Character	The assembly time computed offset to the NAB pointer on the stack of the function. The following example shows how &CCN_NAB_OFFSET is set by the compiler: <pre>&CCN_NAB_OFFSET SETC '#NAB_2-@@AUT0@2'</pre> The following example shows how &CCN_NAB_OFFSET can be used in the prolog code: <pre>STG 0,&CCN_NAB_OFFSET.(,13)</pre>
&CCN_IASM_MACRO	Character	The name of the in-stream macro that contains all the #pragma insert_asm supplied statements. The setting of &CCN_IASM_MACRO only happens in the presence of inserted assembler statements provided by #pragma insert_asm directives. In the presence of such directives, the compiler generates an in-stream HLASM MACRO like this: <pre>MACRO @@IASM@ *from #pragma insert_asm #1 000004 *from #pragma insert_asm #2 000006 *from #pragma insert_asm #3 000008 MEND</pre> The following example shows how &CCN_IASM_MACRO is set by the compiler: <pre>&CCN_IASM_MACRO SETC '@@IASM@'</pre> The MACRO name in &CCN_IASM_MACRO can be placed anywhere within the prolog/epilog code.
&CCN_PRCN_LONG	Character	The actual function name up to the 1024 character HLASM limit. The setting of &CCN_PRCN_LONG is subject to the HLASM limit of 1024 characters on a SETC instruction. When the function name is longer than 1024 characters, the character value set will be truncated to 1021 characters and appended with '...'.

Table 4 describes the global SET symbols that can be set by your prolog and epilog code to conditionally enable or disable code sequences generated by the compiler.

Table 4. User modifiable global SET symbols

Global SET symbol	Type	Default	Description
&CCN_SASIG	Logical	1	Set to "1" to enable the save area signature generation. Set to "0" to disable the save area signature generation.
&CCN_NAB_STORED	Logical	0	Set to "1" to indicate that NAB pointer storing code was done in the prolog code. The following example shows the code that is generated by the compiler to cause the NAB computing and storing code to be conditionally assembled based on the setting of &CCN_NAB_STORED: AIF (&CCN_NAB_STORED).@@NONAB2 LGHI 0,160 ALGR 0,13 STG 0,#NAB_2-@@AUTO02(,13) .@NONAB2 ANOP
&CCN_IASM_FRONT	Logical	0	Set to "1" to indicate that &CCN_IASM_MACRO was already called. The following example shows the code that is generated by the compiler to cause the &CCN_IASM_MACRO to be conditionally assembled based on the setting of &CCN_IASM_STORE: AIF (&CCN_IASM_FRONT).@@NOIASM1 @@IASM@ .@NOIASM1 ANOP
&CCN_WSA_INIT	Character	'CCNZWSAI' for 31-bit and 'CCNZQWSI' for 64-bit	Function name for WSA initialization routine.
&CCN_WSA_TERM	Character	'CCNZWSAT' for 31-bit and 'CCNZQWST' for 64-bit	Function name for WSA termination routine.
&CCN_APARSE	Logical	1	Set to "1" to trigger CCNZINIT call to parse argc and argv. Set to "0" to disable argc and argv parsing.

SCJTSAM(CJTZGBL) macro

Sample macros for prolog and epilog code are supplied in the SCJTSAM data set. The SCJTSAM(CJTZGBL) macro contains assembler instructions to declare all the Global Set Symbols to be referenced. You need to copy the CJTZGBL macro into your prolog and epilog code. Figure 9 shows the sample CJTZGBL macro.

Figure 9. SCJTSAM(CJTZGBL)

```
*****
*
*  MACRO-NAME   = CJTZGBL
*  DESCRIPTIVE-NAME = METAL C GLOBAL SET SYMBOLS
*
*  USAGE =     COPY CJTZGBL
*
*****
```

GBLA &CCN_DSASZ	DSA size of the function
GBLA &CCN_SASZ	Save area size of this function
GBLA &CCN_ARGS	Number of fixed parameters
GBLA &CCN_RLOW	1st GPR on STM/STMG in prolog
GBLA &CCN_RHIGH	2nd GPR on STM/STMG in prolog
GBLB &CCN_MAIN	True if function is main
GBLB &CCN_LP64	True if compiled with LP64
GBLB &CCN_NAB	True if NAB needed
.* &CCN_NAB is to indicate if there are called functions that depend on	
.* stack space being pre-allocated. When &CCN_NAB is true you'll need	
.* to add a generous amount to the size set in &CCN_DSASZ when you	
.* obtain the stack space.	
GBLB &CCN_ALTGPR(16)	Altered GPRs by the function
GBLB &CCN_SASIG	True to gen savearea signature
GBLC &CCN_PRCN	Entry symbol of the function
GBLC &CCN_CSECT	CSECT name of the file
GBLC &CCN_LITN	Symbol name for LTORG
GBLC &CCN_BEGIN	Symbol name for function body
GBLC &CCN_ARCHLVL	n in ARCH(n) option
GBLC &CCN_ASCM	A=AR mode P=Primary mode
GBLC &CCN_IASM_MACRO	MACRO name for all insert_asm
GBLB &CCN_IASM_FRONT	True if insert_asm at front
GBLC &CCN_NAB_OFFSET	Offset to NAB pointer in DSA
GBLB &CCN_NAB_STORED	True if NAB pointer stored
GBLC &CCN_PRCN_LONG	Full func name up to 1024 chars
GBLB &CCN_STATIC	True if function is static
GBLB &CCN_RENT	True if compiled with RENT
GBLC &CCN_WSA_INIT	WSA initialization function name
GBLC &CCN_WSA_TERM	WSA termination function name
GBLB &CCN_APARSE	True to parse OS PARM
GBLC &CCN_DSAUSER	Offset to user field in DSA

SCJTSAM(MYPROLOG) macro

Sample macros for prolog code are supplied in the SCJTSAM data set. Figure 10 on page 19 shows the sample prolog code.

```

MACRO
&NAME MYPROLOG
COPY CJTZGBL
LARL 15,&CCN_LITN
USING &CCN_LITN,15
GBLA &MY_DSASZ
&MY_DSASZ SETA 0
AIF (&CCN_LP64).LP64_1
STM 14,12,12(13)
AGO .NEXT_1
.LP64_1 ANOP
STMG 14,12,8(13)
.NEXT_1 ANOP
AIF (NOT &CCN_RENT).SKIP_R1
AIF (&CCN_LP64).LP64_11
LR 2,0
AGO .SKIP_R1
.LP64_11 ANOP
LGR 2,0
.SKIP_R1 ANOP
AIF (&CCN_DSASZ LE 0).DROP
&MY_DSASZ SETA &CCN_DSASZ
AIF (&CCN_DSASZ GT 32767).USELIT
AIF (&CCN_LP64).LP64_2
LHI 0,&CCN_DSASZ
AGO .NEXT_2
.LP64_2 ANOP
LGHI 0,&CCN_DSASZ
AGO .NEXT_2
.USELIT ANOP
AIF (&CCN_LP64).LP64_3
L 0,=F'&CCN_DSASZ'
AGO .NEXT_2
.LP64_3 ANOP
LGF 0,=F'&CCN_DSASZ'
.NEXT_2 AIF (NOT &CCN_NAB).GETDSA
&MY_DSASZ SETA &MY_DSASZ+1048576
LA 1,1
SLL 1,20
AIF (&CCN_LP64).LP64_4
AR 0,1
AGO .GETDSA
.LP64_4 ANOP
AGR 0,1
.GETDSA ANOP
STORAGE OBTAIN,LENGTH=(0),BNDRY=PAGE
AIF (&CCN_LP64).LP64_5
LR 15,1
ST 15,8(,13)
L 1,24(,13)
ST 13,4(,15)
LR 13,15
AGO .CHECK_R
.LP64_5 ANOP
LGR 15,1
STG 15,136(,13)
LG 1,32(,13)
STG 13,128(,15)
LGR 13,15
.CHECK_R ANOP
AIF (NOT &CCN_RENT).DROP
AIF (&CCN_LP64).LP64_12
LR 0,2
AGO .DROP
.LP64_12 ANOP
LGR 0,2
.DROP ANOP
DROP 15
MEND

```

SCJTSAM(MYEPILOG) macro

Sample macros for epilog code are supplied in the SCJTSAM data set. Figure 11 shows the sample epilog code.

```
MACRO
&NAME MYEPILOG
COPY CJTZGBL
GBLA &MY_DSASZ
AIF (&MY_DSASZ EQ 0).NODSA
AIF (&CCN_LP64).LP64_1
LR 1,13
AGO .NEXT_1
.LP64_1 ANOP
LGR 1,13
.NEXT_1 ANOP
AIF (&CCN_LP64).LP64_2
L 13,4(,13)
AGO .NEXT_2
.LP64_2 ANOP
LG 13,128(,13)
.NEXT_2 ANOP
AIF (&CCN_LP64).LP64_3
ST 15,16(,13)
AGO .NEXT_3
.LP64_3 ANOP
STG 15,16(,13)
.NEXT_3 ANOP
LARL 15,&CCN_LITN
USING &CCN_LITN,15
STORAGE RELEASE,LENGTH=&MY_DSASZ,ADDR=(1)
DROP 15
AIF (&CCN_LP64).LP64_4
L 15,16(,13)
AGO .NEXT_4
.LP64_4 ANOP
LG 15,16(,13)
.NEXT_4 ANOP
.NODSA ANOP
AIF (&CCN_LP64).LP64_5
L 14,12(,13)
LM 1,12,24(13)
AGO .NEXT_5
.LP64_5 ANOP
LG 14,8(,13)
LMG 1,12,32(13)
.NEXT_5 ANOP
BR 14
MEND
```

Figure 11. SCJTSAM(MYEPILOG)

Compiler-generated default prolog and epilog code

The default prolog and epilog code generated for the main function is very much the same as the code produced by the sample prolog and epilog macros in Figure 10 on page 19 and Figure 11. That is, a STORAGE macro is used to obtain and release a dynamic storage area of 1 MB. For functions other than "main", the prolog code simply picks up its DSA pointer (the NAB pointer) from the "Address of next save area" field in the calling function's save area.

Supplying your own HLASM statements

Before you insert your own HLASM statements into your C source file, be aware of the following information:

- The compiler does not recognize either the syntax or the semantics of the HLASM statements embedded in the C `__asm` statement. You need to ensure that the embedded HLASM statements:
 - Meet the requirements of the assembly step that follows the compilation step.
 - Function correctly when embedded in the compiler-generated HLASM source file.
- In the HLASM syntax, the first field is the label field, followed by the op-code, and the rest of the HLASM instruction. If there is no label field, you must leave a blank space at the beginning of the string. Other than this, you can code the rest of the HLASM instruction as you do in HLASM.
- You do not have to consider HLASM line-width requirements. You can code an instruction in the code format string continuously, in accordance with the limitation of the C source file. The C compiler breaks up a code format string that exceeds 71 characters in the HLASM output, inserting continuation characters as required.

Inserting HLASM instructions into the generated source code

You can use the `__asm` language extension to specify assembly instructions to be embedded within the generated HLASM source code. For example, you can embed assembly statements that invoke assembler macros to obtain system services.

Use the `__asm` statement only to embed a short sequence of assembler instructions into a C function, to perform actions that cannot be done using C statements. If you need to use a long routine, put the assembly statements into a source file, assemble it separately, and then call the routine from the C program.

Note: The compiler supports a collection of hardware built-in functions, such as `__csg`. These hardware built-in functions allow the compiler more freedom in blending embedded assembly statements with the rest of the code. For this reason, a hardware built-in function might be better than an `__asm` statement for embedding the assembly instructions that you need.

In addition to the `__asm` language extension, there are language constructs for the following purposes:

- Reserving a register for a global variable of the pointer type. See “Reserving a register for a global variable” on page 30.
- Invoking a macro in the list form. See “Specifying and using the list form of a macro” on page 28.
- Supplying your own function prologs and epilogs. See “Prolog and epilog code” on page 13.

.

Using the `__asm` statement

Within the `__asm` statement, the *code format string* specifies the assembly statement to be embedded in the compiler-generated HLASM source file. Figure 12 on page 22 provides an example of a simple code format string, enclosed in double quotation marks, in an `__asm` statement.

```
void myfunc() {
    __asm ( " AR 1,2" );
}
```

Figure 12. Simple code format string

Treatment of the code format string

The compiler treats the code format string in an `__asm` statement similarly to the way the `printf` function treats a format string, with the following exception: Instead of printing out the string during program execution, the compiler inserts it after the generated sequence of assembly statements, before the `END` statement.

More than one assembler instruction can be put into the code format string. As shown in Figure 13, each assembler statement must be separated by the new line character `\n` (like the new line character that is used in a `printf` format string).

```
void myfunc() {
    __asm ( " AR 1,2\n AR 1,2" );
}
```

Figure 13. Code format string with two instructions

The example in Figure 13 will embed two " AR 1,2" instructions in the HLASM source code. You can make the statement more readable by breaking the string into two. In C, adjacent string literals are automatically concatenated and treated as one. The sample code in Figure 13 and Figure 14 generate the same output.

```
void myfunc() {
    __asm ( " AR 1,2\n"
           " AR 1,2" );
}
```

Figure 14. Code format string with two instructions, formatted for readability

Notes:

1. The character `\n` is still required to delimit statements.
2. The second statement also begins with a blank space.

C expressions as `__asm` operands

You can use substitution specifiers in a code format string just as you can in a `printf` format string. The substitution specifier tells the compiler to substitute the specified C expression into the corresponding *__asm operand* when it embeds the assembly statement in the HLASM source code. You must ensure that the substitution converts the code format string into a valid assembler instruction.

Note: In this document, operands used in a code format string are referred to as `__asm` operands.

An embedded assembly statement can use any C-language expression that is in scope as an `__asm` operand. The *constraint* tells the compiler what to do with the C expression that follows it.

Substitution of a C variable into an `__asm` operand

Figure 15 on page 23 shows an `__asm` statement that substitutes a C variable into an output `__asm` operand. Figure 16 on page 23 shows those assembly instructions.

```

void myfunc() {
    int x;

    __asm ( " ST 12,%0\n" : "=m"(x) :: "r12" );
}

```

Figure 15. Substitution of a C variable into an output `__asm` operand

Notes:

1. A colon that marks the beginning of the list of output `__asm` operands, it follows the code format string.
2. The output `__asm` operand is `"=m"(x)`. The constraint `"m"` communicates the syntactic requirement to the compiler:
 - The symbol `"="` means the (output) `__asm` operand will be modified.
 - The letter `"m"` means that the output `__asm` operand is a memory operand.
3. The C expression is the variable `x`.
4. The compiler does not know that the embedded assembly instruction is `ST`, nor does it know the HLASM syntactic requirement of the second `ST` operand.
5. The variable `x` is the first `__asm` operand in the example, and therefore corresponds to `%0` in the code format string.

From the `__asm` statement used in Figure 15, the compiler embeds the instructions shown in Figure 16 in the generated HLASM source code.

```

1      2
LA      1,@3x
ST      12,0(1)
3

```

Figure 16. HLASM source code embedded by the compiler

Notes:

1. The `LA` instruction is inserted by the C compiler as a result of processing the `"=m"(x)` `__asm` operand.
2. `@3x` is the HLASM symbol name that the compiler assigned to the local variable `x`. Local C symbol names are mapped to HLASM symbol names so that each local variable has a unique name in the HLASM source file.
3. `0(1)` is substituted into `"%0"`, which specified the first `__asm` operand in the code format string in Figure 15 (`ST 12,%0`).

Substitution of a C pointer into an `__asm` operand

The code format string in Figure 17 on page 24 invokes the `WTO` macro by using the execute form of the macro with a user-defined buffer.

In general, you do not control which registers are used during the operand substitution, as illustrated in Figure 17 on page 24. For an example that allows you to specify registers, see Figure 21 on page 26.

```

int main() {
    struct WTO_PARM {
        unsigned short len;
        unsigned short code;
        char text[80];
    } wto_buff = { 4+11, 0, "hello world" };

    __asm( " WTO MF=(E,(%0)) " : : "r"(&wto_buff));
    return 0;
}

```

Figure 17. Substitution of a C pointer into an `__asm` operand

Notes:

1. The absence of a label necessitates that a blank space begin the code format string.
2. There are no output `__asm` operands. The end of the output `__asm` operands list is marked by a colon, which is then followed by a comma-separated list of input `__asm` operands. The colon starting the list of input `__asm` operands is not necessary if there are no input operands (which is the case in Figure 15 on page 23).
3. The input `__asm` operand consists of two components:
 - A constraint "r" that tells the compiler that the operand will be stored in a GPR.
 - An expression (`&wto_buff`) that states that the operand is the address of the message text in the C structure `wto_buff`.

Definition of multiple `__asm` operands

In Figure 18, the compiler is instructed to store the third defined C variable (z) in a register, and then substitute that register into the third `__asm` operand %2.

```

void myfunc() {
    int x, y, z;
    __asm ( " ST 12,%0\n"
           " ST 12,%1\n"
           " AR 12,%2" : "=m"(x), "=m"(y) : "r"(z) : "r12" );
}

```

Figure 18. `__asm` operand lists

Notes:

1. The code format string instructs the compiler to embed an assembly statement that substitutes the register (with contents of the C variable z) into the third `__asm` operand (%2).
2. The constraint "=m" instructs the compiler to use memory operands for the output variables x and y.
3. The constraint "r" instructs the compiler to use a register for the input variable z.

Figure 19 on page 25 shows the compiler-generated HLLASM code from the `__asm` statement in Figure 18. GPR 4 is assigned to the variable z.

```

L 4,@5z      1
LA 2,@4y     2
LA 1,@3x
ST 12,0(1)   3
ST 12,0(2)
AR 12,4

```

Figure 19. Compiler-generated HLASM code from the `__asm` statement

Notes:

1. The first assembly statement `L 4,@5z` is added by the compiler to get `z` into the form specified by the input `__asm` operand constraint `"r"`.
2. The next two instructions are added by the compiler to get the variables `x` and `y` into the form specified by the output `__asm` operand constraints `"=m"`.
3. The contents of the code format string are appended in the last three instructions.

Register specification

In general you do not have control over which registers are used during operand substitutions. The register assignment might change when you use different options or optimization levels, or when the surrounding C code is changed.

In cases where you specify explicit registers to be used in the embedded instructions, you should code a clobber list, as shown in Figure 21 on page 26. Without the clobber list, the `__asm` statement embeds incorrect assembly statements, as shown in Figure 20.

```

__asm ( " LR 0,%0\n"      /* load &p1 */
        " LR 1,%1\n"      /* load &dcB */
        " SVC 21"
        : : "r"(&p1), "r"(&dcB)); 1

```

Figure 20. Unsuccessful attempt to specify registers

Note: The output and input `__asm` operand lists are positional. If there are no output `__asm` operands, the colons separating the output and input operand list are still needed. Because the compiler has no knowledge of assembly instructions and does not understand the `LR` instruction, it does not know that the registers GPR 0 and GPR 1 are being used in the statement. Any connection between the `__asm` statement and the rest of the C code must be specified via the `__asm` operand lists. The information provided in the lists should prevent the compiler from incorrectly moving the other references surrounding the `__asm` statement. In this example, because the compiler doesn't know that GPR 0 and GPR 1 are being used, it will embed incorrect assembly statements.

To prevent the compiler from incorrectly moving the other references surrounding the `__asm` statement, add a clobber list after a colon that follows the input `__asm` operands, as shown in Figure 21 on page 26.

Note: Do not try to use the `__asm` statement to embed a long piece of assembly code with many operand specifiers and stringent register requirements. There is a limited number of registers available for the compiler to use in the operand specifiers, and in the surrounding code generation. If too many registers are clobbered, there may not be enough registers left for the `__asm` statement. The same applies if there are too many specifiers.

```

__asm ( " LR 0,%0\n"      /* load &p1 */
        " LR 1,%1\n"      /* load &dcB */
        " SVC 21"
        : : "r"(&p1), "r"(&dcB) : "r0","r1");

```

1
2

Figure 21. Register specification with clobbers

Notes:

1. This colon is not needed if there is no clobber list.
2. The clobber list specifies the registers that can be modified by the assembly instructions.

C expressions as read-write __asm operands

If you use the same __asm operand for both input and output, you must take care that you tell the compiler that the input __asm operand refers to the same variable as the corresponding output __asm operand. For example, the code format string in Figure 22 uses one register to store a single __asm operand that is used for both input and output.

Definition of __asm operands for both input and output via an operand list

This topic describes how to use a code format string to define __asm operands that can be used for both input and output.

You can use either input and output operand strings both incorrectly (Figure 22) and correctly (Figure 24 on page 27). The code in Figure 22 is incorrect because the AR statement reads the first operand and then modifies it, but the =r constraint specifies the output aspect only.

```

__asm ( " AR %0,%1" : "=r"(x) : "r"(y) );

```

1

Figure 22. Incorrect __asm operand definition for both input and output

Note: No input operand is specified for variable x. The compiler will not know that input and output are stored in the same variable.

The compiler-generated HLASM source code in Figure 23 is the result of the incorrect definition in Figure 22.

```

L 2,@4y
LA 1,@3x
AR 4,2
ST 4,0(,1)

```

1

Figure 23. Incorrect compiler-generated HLASM source code from the incorrect __asm operand definition for both input and output

Note: GPR 4, which is meant for input as well as output, is not loaded from variable x before the code format string is embedded because the code format string in Figure 22 specified variable x as an output operand only.

If a code format string uses a single __asm operand for both input and output, you must ensure that the embedded assembly statements will perform both of the following tasks:

- Define the variable as an input operand as well as an output operand.

- Define both an input operand and an output operand that refers to the same variable. The variable name is not sufficient for this purpose. See Figure 24.

Figure 24 shows the code format string that will embed the correct assembler statements (as shown in Figure 25).

```
__asm ( " AR %0,%1" : "=r"(x) : "r"(y), "0"(x) );
```

1 2 3 4

Figure 24. Successful definition of an `__asm` operand for both input and output

Notes:

1. `%0` is the first operand in the code format string.
2. This example has one output `__asm` operand, `"=r"(x)`.
3. Within the input `__asm` operand list `"r"(y), "0"(x)`, the `__asm` operands are separated by a comma.
4. An input operand `"0"(x)` is added to the input field. The constraint of this `__asm` operand is the `("0")`, which tells the compiler that:
 - This input `__asm` operand is the same as the output `__asm` operand `%0`. (A numeral zero in the constraint `("0")` refers to `%0`; a numeral one in a constraint would refer to `%1`; and so on.)
 - The register needs to be loaded with variable `x`, as shown in Figure 25, before the code format string is embedded in the HLASM output.

The compiler-generated HLASM source code in Figure 25 is the result of the correct definition in Figure 24.

```
L 2,@4y
L 4,@3x                      1
LA 1,@3x
AR 4,2
ST 4,0(,1)
```

Figure 25. Correct compiler-generated HLASM source code from the correct `__asm` operand definition for both input and output

Note: The compiler inserted `L 4,@3x` at the beginning of the instruction sequence because the code format string in Figure 24 included both the output operand `"=r"(x)` and the input operand `"0"(x)`. Together, these statements tell the compiler that the register for the first operand `%0` will be used for variable `x`, which has a value that can be either an input or an output operand.

Definition of an `__asm` operand for both input and output via the "+" constraint

You can also use the "+" constraint to specify that an `__asm` operand is used for both input and output.

In Figure 26 on page 28, the "+" constraint is used to define that the variable `x` is used both as input and output.

```
__asm ( " AR %0,%1" : "+r"(x) : "r"(y));
```

Note: This example is parsed as though the operand list in Figure 24 on page 27 is given.

Figure 26. The + constraint to define an __asm operand for both input and output

Note that an operand can be matched only once. When you use the "+" constraint to implicitly define matching input and output __asm operands, do not explicitly define a corresponding __asm operand.

Figure 27 shows an erroneous example of an __asm operand that is defined both implicitly and explicitly. The notes identify the unnecessary code.

```
__asm ( " AR %0,%1" : "+r"(x) : "r"(y), "0"(x));
```

1
2
3

Notes:

1. %0 is the first operand in the code format string.
2. This example has one output __asm operand, "+r"(x). The "+" constraint implicitly defines a matching input __asm operand.
3. You do not have to define __asm operand "0"(x) explicitly.

Figure 27. Error: Redundant definition of an __asm operand

Specifying and using the list form of a macro

When you specify and use the list form of a macro, you can code for reentrancy by embedding assembly statements that:

1. Allocate space on the stack (that is, use a local variable). See Figure 29 on page 29.
2. Copy the parameter field values from the list form to this allocated space.
3. Invoke the execute form of the macro that will use the allocated space.

Note: The code format string in Figure 17 on page 24 invokes the WTO macro by using the execute form of the macro with a user-defined buffer. That example does allow for reentrancy.

You should not have direct reference to symbols within your code format string as the addressability is not guaranteed. The proper way to use the macro is shown in Figure 29 on page 29, in which all __asm statements are connected through the C variable operands listmsg1 and buff.

Figure 28 on page 29 provides an example that uses the list form of a macro without considering reentrancy.

```

    1 __asm(" WTO 'hello world',MF=L" : "DS"(listmsg1));
    2
    3
int main() {
    4 __asm( " WTO MF=(E,(%0)) " : : "r"(&listmsg1));
    5
    return 0;
}

```

Notes:

1. The first `__asm` statement invokes the macro `WTO` in the list form (`MF=L`). In order for the list form of the macro to be invoked with the values of the parameter fields defined, the `__asm` statement must be specified in the global scope.
2. The message text "hello world" is provided as a macro parameter.
3. The "DS" constraint indicates that this is a data definition, with the name of the C variable defined as the variable `listmsg1`. Because `listmsg1` is implicitly defined as a structure, it can be referenced in subsequent `__asm` statements, therefore the "DS" constraint must be specified in the output operand list. By default, the compiler allocates 256 bytes of memory for the variable `listmsg1`, which should satisfy most requirements. You can change the memory allocation size (for example, "DS:100"(`listmsg1`) to allocate 100 bytes). You can allocate more than 256 bytes of space.
4. The second `__asm` statement invokes the macro `WTO` in the execute form (`MF=(E,(%0))`). It takes the address of the storage defined in the list form.
5. The address of the variable `listmsg1` is defined as an input operand that is stored in a register.

Figure 28. Specifying and using the `WTO` macro (no reentrancy)

Support of reentrancy requirements

If the execute form of the macro needs to change the fields provided in the list form, the assembly statements embedded by the `__asm` statement in Figure 28 will be incorrect when support for reentrancy is required. The proper way to use the macro is shown in Figure 29.

```

__asm(" WTO 'hello world',MF=L" : "DS"(listmsg1)); 1

int main() {
    __asm(" WTO 'hello world',MF=L" : "DS"(buff)); 2
    buff = listmsg1; 3

    __asm( " WTO MF=(E,(%0)) " : : "r"(&buff));
    return 0;
}

```

Notes:

1. The first `__asm` statement uses the list form of the macro `WTO` to define the variable `listmsg1`.
2. The second `__asm` statement, specified within function scope with a "DS" constraint, will allocate stack space for the variable `buff` but will not actually initialize the parameter values.
3. The size of this variable should match that of the corresponding `__asm` statement in global scope. An assignment copies the actual parameter values from the list form to this buffer.

Figure 29. Support for reentrancy in a code format string

Inserting non-executable HLASM statements into the generated source code

You can use the `#pragma insert_asm` directive to supply your own non-executable HLASM statements to the generated source code. The primary purpose of this directive is that you can use it to include the DSECT mapping macros that are required by your embedded assembly statements. The syntax is `#pragma insert_asm("string")`.

The `#pragma insert_asm` directive causes the compiler to insert *string* at an appropriate place in the generated HLASM code. When you use multiple `#pragma insert_asm` directives, they are placed in the same order as they appear in your C source code.

Note: The `#pragma insert_asm` directive can be used with a `_Pragma` operator. If you use the `_Pragma` operator, you must put a slash ("/") character before the double quotation marks that surround the string literal. For example: `_Pragma ("insert_asm(\"MYSTRING\")")`.

Example: Using the #pragma insert_asm directive to map specific DSECT information

Figure 30 uses the `#pragma insert_asm` directive to get the system CVTUSER field to address your specified CVT extension data. Because the CVTPTR and CVTUSER fields are defined in the CVT mapping macro, you can use the `#pragma insert_asm` directive to map specific DSECT information.

```
void myfunc() {
    void * user_cvt;
    __asm(" L 2,CVTPTR\n"
        " L 2,CVTUSER-CVT(2)\n"
        " ST 2,%0"
        : "m"(user_cvt) :: "r2");
}
#pragma insert_asm(" CVT    DSECT=YES,LIST=NO")
```

Figure 30. Code that supplies specific DSECT mapping macros

Reserving a register for a global variable

The *register* storage class specifier is the C-language extension that allows you to specify, for the entire compilation unit, a general purpose register (GPR) for a global variable, as shown in Figure 31 on page 31.

When you use a code format string to specify a GPR for a global variable, be aware that:

- Only GPR 0 through GPR 15 can be specified for storage of a global variable.
- The variable must be declared as a pointer type.
- A declaration with register specifier cannot have an initializer.

```
register int * p __asm("r5");
```

Notes:

1. The variable declaration `int * p` defines the variable as a pointer type.
2. The register "r5" is not initialized.

Figure 31. Register specification

AMODE-switching support

Within a Metal C application, AMODE 31 and AMODE 64 programs can call each other.

To take advantage of the Metal C AMODE-switching support, be aware of the following information:

- The called and calling programs must be in separate source files. Mixing addressing modes within a single C source file is not supported.
- The save area format for the called program is determined by the AMODE and ASC mode of the called program, that is, 72-byte for AMODE 31 programs, F4SA for AMODE 64 programs, F7SA for AR mode programs. The ability for tracing the save areas chain will be interrupted across AMODE switches.
- The parameter list is prepared according to the AMODE of the called program, that is, 4-byte slots for AMODE 31 programs and 8-byte slots for AMODE64 programs.
- It is the user's responsibility to ensure that all storage addresses passed to the AMODE 31 functions are addressable by the AMODE 31 functions. Because the save area and parameter lists are part of the caller's DSA, the caller must have its DSA allocated in the below-the-bar storage.
- The AMODE of the called program can be specified by the new `amode31` and `amode64` type attributes.
- The calling program switches the addressing mode before the call and switches back to its own addressing mode on return from the call.
- The implicit sizes of types `long` and `pointer` in the function prototype are determined by the addressing mode of the called program.
- The `__ptr64` qualifier can be used to specify a 64-bit pointer on an AMODE 31 program; the pointer cannot be dereferenced at the AMODE 31 program.

Example of an AMODE31 program that calls an AMODE64 program

In Figure 32 on page 32, AMODE 31 program "main" in `a31.c` makes calls to AMODE 64 programs `a64a1` and `a64a2` in `a64a.c`. For the commands that compile and link `a31.c` and `a64a.c`, see "Commands that compile and link applications that switch addressing modes" on page 47.

```

a31.c

long a64a1 (long j, int k, short s) __attribute__((amode64));
int a64a2 (long j, int k, short s) __attribute__((amode64));
int main () {
    int a = 40;
    return a64a1(99LL, a, 4) + a64a2(-120LL, -60, -18);
}

a64a.c

long a64a1 (long a, int b, short c) {
    return -(a+b+c);
}

int a64a2 (long a, int b, short c) {
    return -(a+b+c);
}

```

Figure 32. AMODE31 program that calls an AMODE64 program

RENT mode support

The RENT option supports constructed reentrancy for C programs with writable static and external variables. This makes C programs with writable static and external variables to be reentrant so a program can be concurrently used by multiple users. The writable static area (WSA) can be managed by user provided routines. The default of the RENT option is NORENT.

Example

```

metalC -qRENT -S a.c
as -mgoff a.s
export _LD_SYSLIB="//'CJT.SCJT0BJ'"
ld a.o

```

1
2

Notes:

1. The HLASM GOFF option is required to assemble the compiler generated code for RENT.
2. It is necessary to add CJT.SCJT0BJ dataset to the binder SYSLIB for the resolution of CCNZINIT and CCNZTERM.

Linkage convention

General Purpose Register 0 (GPR0) is used to pass the WSA address. The prolog code you supplied needs to preserve the content of GPR0 on exit of the prolog code. Programs compiled with RENT and NORENT can be mixed as long as the NORENT programs do not call RENT programs.

Note: Global variables compiled with RENT and NORENT cannot have the same name.

Assembler code interface

The runtime RENT support is accomplished by additional calls generated for the function "main" between the prolog/epilog code and the function code. The RENT environment initialization and termination routines are called to establish and terminate the dynamically allocated WSA storage with the static initialization data applied. For the AMODE 31 "main" function, CCNZINIT and CCNZTERM are the

names of these routines. While for the AMODE 64 "main" function, CCNZQINI and CCNZQTRM are the function names. For the simplicity of further references, only the names of the 31-bit version are used. The actual WSA storage management is done by user supplied plug-in routines called from CCNZINIT and CCNZTERM. Two user modifiable Global Set Symbols, &CCN_WSA_INIT and &CCN_WSA_TERM, can be used in the user supplied prolog code to set the user supplied WSA initialization and termination function names. The AMODE of the user supplied routines has to be the same as the AMODE of function "main".

Note: CCNZINIT, CCNZTERM, CCNZWSAI and CCNZWSAT require the stack space to be supplied by function "main" prolog code. Both CCNZINIT and CCNZTERM require NAB to be established by function "main". Also, CCNZINIT and CCNZTERM assume stack space to be available for the WSA initialization and termination functions. This arrangement is to ensure that the stack space used by CCNZINIT and CCNZTERM as well as the WSA initialization and termination routines is consistent with the stack space used by function "main". Allocating 1K of extra stack space (in addition to the DSA size suggested by &CCN_DSASZ for "main") by function "main" should be sufficient for AMODE 31. For AMODE 64, the extra stack space is roughly doubled.

The following new Global Set Symbols are introduced for the RENT option.

- &CCN_MAIN
- &CCN_RENT
- &CCN_WSA_INIT
- &CCN_WSA_TERM

For detailed information about these new Global Set Symbols, see "Compiler-generated global SET symbols" on page 15.

You can provide your own WSA initialization and termination routines by setting these Global Set Symbols with the module names of your own routines. For example:

```
GBLC &CCN_WSA_INIT
GBLC &CCN_WSA_TERM
&CCN_WSA_INIT SETC 'MYWSAI'
&CCN_WSA_TERM SETC 'MYWSAT'
```

Your own WSA initialization and termination routines can be object modules, load modules, or program modules, and they need to be supplied to the binder's input.

The compiler generated code for "main" has the following kind of assembly statements in it:

- For AMODE 31:


```
DC V(&CCN_WSA_INIT)
DC V(&CCN_WSA_TERM)
```
- For AMODE 64:


```
DC VD(&CCN_WSA_INIT)
DC VD(&CCN_WSA_TERM)
```

WSA initialization routine

Given the size of the WSA for the whole application and the image of the WSA with initialization data applied, the WSA initialization routine you provided dynamically allocate the WSA storage for the application and copy the WSA image into it. The address of the allocated storage is returned which CCNZINIT saves it

on the function main's stack to be propagated to the rest of the application. You are responsible for ensuring that the allocated storage is addressable to all parts of the application. This particularly means if there are AMODE 31 parts in the application, the WSA storage should not be allocated above the 2G bar if the AMODE 31 parts need to access it. Also, the WSA storage has to be allocated in the primary address space. WSA storage in data spaces is not supported.

The routine you provide is given an address of an area to store whatever extra information you want to keep and pass to the WSA termination routine you provided. The storage area size is the size of a pointer, that is, 4 bytes for AMODE 31, and 8 bytes for AMODE 64.

Function prototype:

```
typedef void * (init_func_t) (void * wsa_image_addr,
                             unsigned long wsa_size, void * *user_info_addr, unsigned int alignment);
```

Input parameters:

- *wsa_image_addr* - the address of the WSA image in the loaded program object
- *wsa_size* - the total size of the application's WSA
- *user_info_addr* - the address of the CCNZINIT provided save area for saving user information
- *alignment* - the minimum required alignment of the allocated WSA storage. For example, "alignment=8" means double-word alignment.

Return value:

The address of the allocated and initialized WSA storage. The default is the IBM supplied routine CCNZWSAI or CCNZQWSI, which allocates storage for both AMODE 31 and AMODE 64 with the following macro:

```
STORAGE OBTAIN, LENGTH=(n), BNDRY=PAGE
```

WSA termination routine

This routine is called from CCNZTERM for the WSA termination and cleanup process. It is passed in the address of the WSA storage allocated by the WSA initialization routine. It is also given the same WSA size that was originally passed to the WSA initialization routine.

Function prototype:

```
typedef void (term_func_t) (void * allocated_wsa_addr,
                             unsigned long wsa_size, void * user_info_addr);
```

Input parameters:

- *allocated_wsa_addr* - the address of the allocated WSA storage
- *wsa_size* - the total size of the application's WSA
- *user_info_addr* - the saved user information

Return value:

The default is the IBM supplied routine CCNZWSAT or CCNZQWST, which frees the storage with the following macro:

```
STORAGE RELEASE, LENGTH=(n), ADDR=(m)
```

argc argv parsing support

If your main() function uses the standard argc and argv arguments, the Metal C initialization routine is called to parse the raw parameter data received from the hosting environment and to convert the parameter to the standard argc and argv

format. If your program is not invoked in or otherwise connected (not dubbed) to the z/OS UNIX System Services (z/OS UNIX) environment, you can use the ARGPARSE or NOARGPARSE options to determine if the EXEC PARM needs to be further parsed into individual arguments; the EXEC PARM must be in this format: a halfword length field followed by a maximum of 100 characters where the length field contains a binary count of the number of bytes in the PARM field.

If your `main()` function uses `argc` and `argv` arguments and you do not want the parsing to be performed, you can set the new Global Set Symbol `&CCN_APARSE` to 0 in your prolog code to conditionally bypass the argument parsing. For detailed information, see Table 4 on page 17.

AR-mode programming support

An AR-mode function can access data stored in data spaces by using the hardware access registers. For more information about AR-mode, see *z/OS MVS Programming: Assembler Services Guide*. A non-AR-mode function is said to be in *primary mode*.

The following sections describe the compiler options, language constructs, and built-in functions that support AR-mode programming.

AR-mode function declaration

You can declare a function to be an AR-mode function with the `armode` attribute. The syntax is:

```
void armode_func() __attribute__((armode));
```

You can also use the `ARMODE` compiler option to declare that all functions in the source program to be AR-mode functions. If you use the `ARMODE` compiler option and you want to single out the functions in the source program to be primary mode functions you can declare the function with the `noarmode` attribute. The syntax is:

```
void nonarmode_func() __attribute__((noarmode));
```

Far pointer declaration, reference, and dereference

The ability to reference data stored in different data spaces is achieved through a C language extension to pointer types called far pointer types. A far pointer type is declared by adding the `__far` qualifier. The syntax is

```
int * __far my_far_pointer;
```

A far pointer can be declared in a function of any mode (AR mode or primary mode). But only an AR-mode function can directly or indirectly dereference a far pointer. In other words, only an AR-mode program can access data stored in data spaces with far pointers.

Note: For an example of a simple dereference of a far pointer, see Figure 39 on page 43.

Regardless of the mode of the function, a far pointer can be manipulated in the following ways:

- It can be passed as a parameter.
- It can be received as a function return value.
- It can be compared with another pointer.
- It can be cast as another pointer type.

- It can be used in pointer arithmetic expressions.

A far pointer consists of ALET and an offset. Although an ALET is always 32 bits in length, the size of a far pointer is twice the size of a regular pointer. The layout of a far pointer in memory depends on the AMODE of the function:

- Under AMODE 31, a far pointer occupies eight bytes.
 - The ALET occupies the first four bytes.
 - The offset occupies the last four bytes.
- Under AMODE 64, a far pointer occupies 16 bytes.
 - The first four bytes are unused.
 - The ALET occupies the second four bytes.
 - The offset occupies the last eight bytes.

This difference in pointer size is illustrated in Figure 33.




Figure 33. Far pointer sizes under different addressing modes

C language constructs and far pointers

Table 5 describes the effects of language constructs that might have special impact on far pointers.

Table 5. Language constructs that may have special impact on far pointers

Language Construct	Effect
Implicit or explicit cast from normal to far pointer	Because the normal pointer is assumed to point to primary address space, the ALET of the far pointer is set to 0.
Explicit cast from far pointer to normal pointer	The offset of the far pointer is extracted and used to form the normal pointer. Unless the ALET of the far pointer was 0, the normal pointer is likely to be invalid.
Operators !=, ==	If either operand is a far pointer, the other operand is implicitly cast to a far pointer before the operands are compared. The comparison is performed on both the ALET and offset components of a far pointer.
Operators <, <=, >, >=	Only the offset of the far pointer is used in the comparison. Unless the ALETs of the far pointers were the same, the result might be meaningless.
Compare to NULL	Because of the implicit cast of NULL to a far pointer, the != and == operators compare both the ALET and the offset to zero. A test of !(p>NULL) is not sufficient to ensure that the ALET is also 0.

Table 5. Language constructs that may have special impact on far pointers (continued)

Language Construct	Effect
Pointer arithmetic	The effects of pointer arithmetic are applied to the offset component of a far pointer only. The ALET component remains unchanged.
Address of Operator, operand of &	The result is a normal pointer, except in the following cases: • If the operand of & is the result of an indirection operator (*), the type of & is the same as the operand of the indirection operator. • If the operand of & is the result of the arrow operator (->, structure member access), the type of & is the same as the left operand of the arrow operator.

Implicit ALET association

In addition to explicitly specifying ALETs that use far pointers to access data in data spaces, the compiler must associate those ALETs with all the memory references contained in the AR-mode function.

In a non-AR-mode function, all variable references are to primary data space (ALET 0). In an AR-mode function, the compiler manages access registers (ARs) so that every memory reference uses an ALET associated with the variable type to reach the appropriate data space. Table 6 lists the ALET associations for different types of variables.

Table 6. Implicit ALET associations for AR-mode-function variables

Variable type	Implied ALET
File-scope variable	ALET 0 (primary data space)
Stack variables (function local variable)	The ALET that is in AR 13 at the time of function entry. This points to the stack frame.
Parameters (function formal parameters)	The ALET that is in AR 1 at the time of function entry. This points to the parameter list.
Data pointed to by regular pointers	ALET 0 (primary data space).
Data pointed to by far pointer	ALET contained in far pointer.

Far pointer construction

The Metal C Runtime Library does not provide functions for allocating or deallocating alternative data spaces. You can use the DSPSERV and ALESERV HLASM macros to allocate space and obtain a valid ALET and offset. For an example, see Figure 37 on page 40. For more information, see *z/OS MVS Programming: Assembler Services Guide*.

Built-in functions that manage far-pointer components

The compiler provides built-in functions for setting and getting the individual components of far pointers. Whenever you use these built-in functions, you must:

- Define the macro `_ML.BUILTIN` to "1".
- Include the header file `builtins.h`.

Figure 34 lists the constructors.

```
void * __far __set_far_ALET_offset(unsigned int alet, void * offset);
void * __far __set_far_ALET(unsigned int alet, void * __far offset);
void * __far __set_far_offset(void * __far alet, void * offset);
```

1
2

Notes:

1. The `__set_far_ALET` function does not modify the far-pointer parameter offset. It simply uses it to provide the offset component of the far pointer being constructed. Its return value is the constructed far pointer.
2. Similarly, the `__set_far_offset` function that uses the far-pointer parameter ALET is not modified; it simply provides the ALET for the far pointer being constructed.

Figure 34. Built-in functions for setting far-pointer components

Figure 35 lists the extractors.

```
unsigned int __get_far_ALET(void * __far p);
void * __get_far_offset(void * __far p);
```

Figure 35. Built-in functions for getting far-pointer components

Library functions that manipulate data stored in data spaces

The Enterprise Metal C for z/OS compiler provides far versions of some of the standard C string and memory library functions. The far versions can be called by either AR-mode or primary-mode functions. If these functions are called by an AR mode function, the compiler will generate inline code for them.

Whenever you use these functions, you must:

- Define the macro `_ML.BUILTIN` to "1".
- Include the header file `builtins.h`.

The semantics of these functions, listed in Figure 36, are identical to the standard version.

```
void * __far __far_memcpy(void * __far s1, const void * __far s2, size_t n);
int __far memcmp(const void * __far s1, const void * __far s2, size_t n);
void * __far __far_memset(void * __far s, int c, size_t n);
void * __far __far_memchr(const void * __far s, int c, size_t n);
char * __far __far_strcpy(char * __far s1, const char * __far s2);
char * __far __far_strncpy(char * __far s1, const char * __far s2, size_t n);
int __far strcmp(const char * __far s1, const char * __far s2);
int __far strncmp(const char * __far s1, const char * __far s2, size_t n);
char * __far __far_strcat(char * __far s1, const char * __far s2);
char * __far __far_strncat(char * __far s1, const char * __far s2, size_t n);
char * __far __far_strchr(const char * __far s, int c);
char * __far __far_strrchr(const char * __far s, int c);
size_t __far strlen(const char * __far s);
```

See Note

Note: For an example that illustrates the use of this function, see Figure 38 on page 42.

Figure 36. Library functions for use only in AR-mode functions

AR-mode function linkage conventions

AR mode functions follow the same linkage conventions as do primary-mode functions, with the following additional requirements:

- Any function that calls an AR-mode function must supply the 54-word F7SA save area for saving the access registers.
- The AR-mode function must preserve the calling function's access registers.
- The AR-mode function is responsible for switching into AR mode on entry and switching back to calling function's ASC mode on exit.

Note: A primary-mode function does not switch the ASC mode when calling an AR-mode function.

- An AR-mode function must switch to primary mode before calling a primary mode function.
- A far pointer is passed and returned as a struct that is based on the layout for the calling function's AMODE.

Default prolog and epilog code for AR-mode functions

If the calling function is in non-AR mode, the DSA and parameter areas are assumed to be located in the primary address space.

For AR-mode functions, the default prolog code generates additional instructions that:

- Save the calling function's access registers in the F7SA save area.
- Save the ASC mode of the calling function in the F7SA save area.
- Switch to AR mode.
- Prime AR 1 and AR 13 with LAE instructions.

For AR-mode functions, the default epilog code generates additional instructions that:

- Restore the calling function's access registers.
- Restore the ASC mode of the calling function.

Data space allocation and deallocation

Figure 37 on page 40 provides examples of routines for allocating and deallocating data space.

```

#define _MI_BUILTN 1
#include builtins.h
#include string.h

/*****
/* Allocation/Deallocation example routines
*****/

int alloc_data_space(void * __far *ret, char dstok[8],
                    long size_blocks, char name[8])
{
    __asm("DSPARMS DSPSERV MF=L\n" : "XL:DS:64"(DSPARMS));
    __asm("ALPARMS ALESERV MF=L\n" : "XL:DS:16"(ALPARMS));
    int res,res2;
    struct _myparms /* To reduce number of operands to __asm */
    {
        unsigned origin; /* +0 */
        unsigned blocks; /* +4 */
        unsigned alet; /* +8 */
        char name[8]; /* +12 */
        char dstok[8]; /* +20 */
    } myparms;

    strncpy(myparms.name,name,8);
    myparms.blocks = size_blocks;

    __asm(
        DSPSERV CREATE,GENNAME=COND,NAME=12(%1),OUTNAME=12(%1),"
        STOKEN=20(%1),ORIGIN=0(%1),BLOCKS=(4(%1)),MF=(E,(%2))\n"
        " ST 15,%0\n"
        : "=m"(res)
        : "a"(&myparms), "a"(&DSPARMS)
        : "r0" , "r1", "r14", "r15");

    if(res==0)
    {
        __asm(
            ALESERV ADD,STOKEN=20(%1),ALET=8(%1),MF=(E,(%2))\n"
            " ST 15,%0\n"
            : "=m"(res2) : "a"(&myparms), "a"(&ALPARMS) : "r0" , "r1", "r14", "r15");

        if(res2!=0)
        {
            __asm(
                DSPSERV DELETE,STOKEN=20(%1),MF=(E,(%2))\n"
                " ST 15,%0\n"
                : "=m"(res2) : "a"(&myparms), "a"(&DSPARMS) : "r0" , "r1", "r14", "r15");
            return -res2;
        }
    }
    else
    {
        return res;
    }

    *ret = __set_far_ALET_offset(myparms.alet,(void *)myparms.origin);
    strncpy(dstok,myparms.dstok,8);
    strncpy(name,myparms.name,8);
    return 0;
}

```

Allocation and deallocation routines (Part 1 of 3)

Figure 37. Allocation and deallocation routines

```

void * __far allocate_far(long size)
{
    void * __far ret;

    ret = NULL;
    if(size > 0)
    {
        int blocks = (size+4095)/4096;
        char name[8];
        char dstok[8];
        strncpy(name,"Z",8); /* provide a prefix */
        alloc_data_space(&ret, dstok, blocks, name);
    }
    return ret;
}

void delete_data_space(void * __far p, char dstok[8])
{
    __asm("DSPARMS DSPSERV MF=L\n" : "XL:DS:64"(DSPARMS));
    __asm("ALPARMS ALESERV MF=L\n" : "XL:DS:16"(ALPARMS));
    int alet;

    if(p!=NULL)
    {
        alet = __get_far_ALET(p);
        __asm(
            ALESERV DELETE,ALET=0(%0),MF=(E,(%1))\n"
            : : "a"(&alet), "a"(&ALPARMS) : "r0" , "r1", "r14", "r15");

        __asm(
            DSPSERV DELETE,STOKEN=0(%0),MF=(E,(%1))\n"
            : : "a"(dstok), "a"(&DSPARMS) : "r0" , "r1", "r14", "r15");
    }
}

int get_data_space_token(void * __far p, char *dstok)
{
    __asm("ALPARMS ALESERV MF=L\n" : "XL:DS:16"(ALPARMS));
    unsigned alet;
    int res;

    if(p!=NULL)
    {
        alet = __get_far_ALET(p);
        __asm(
            ALESERV EXTRACT,ALET=0(%1),STOKEN=0(%2),MF=(E,(%3))\n"
            ST 15,%0\n"
            : "=m"(res) : "a"(&alet), "a"(dstok), "a"(&ALPARMS)
            : "r0" , "r1", "r14", "r15");

        return res;
    }
    return -1;
}

```

Allocation and deallocation routines (Part 2 of 3)

```

void * __far free_far(void * __far p)
{
    int x;
    void * __far ret;

    if(p != NULL)
    {
        char dstok[8];
        x = get_data_space_token(p,dstok);
        if(x!=0)
        {
            delete_data_space(p, dstok);
        }
    }
    return NULL;
}

```

Allocation and deallocation routines (Part 3 of 3)

Copying a string pointer to a far pointer

Figure 38 provides an example of using a built-in function to copy a C string pointer to a far pointer.

```

/*****
/* __far_strcpy example
*****/

char * __far far_strcpy_example() __attribute__((armode));
char * __far far_strcpy_example()
{
    char * __far far_string;
    char * near_string;

    near_string = "Hello World!\n";

    far_string = allocate_far(1024);

    __far_strcpy(far_string,near_string);

    return far_string; /* Assume caller will free allocated data space */
}

```

Figure 38. Copying a C string pointer to a far pointer

Far pointer dereference

The Metal C Runtime Library does not provide functions for allocating or deallocating alternative data spaces. Figure 39 on page 43 provides an example of code that dereferences a far pointer.

```

/*****
/* Simple dereference example
*****/

char get_ith_character(char *__far s, int i) __attribute__((armode));
char get_ith_character(char *__far s, int i)
{
    return s[i];
}

int main()
{
    char c;
    char *__far far_string;

    far_string = far_strcpy_example();

    c = get_ith_character(far_string,1);

    free_far(far_string);

    return c;
}

```

Figure 39. Example of a simple dereference of a far pointer

Metal C function descriptor support

When a library or a set of functions are used by multiple applications, these functions have the same code in all the applications while their associated data is application-specific. In this case, you can use Metal C function descriptors to call these library functions instead of using normal function pointers.

A Metal C function descriptor encapsulates all the information that a function call needs to access both the function and the application-specific data. You can use Metal C function descriptors to point to and call functions with their own set of associated data for the particular program or invocation.

A function pointer that is declared with the `__fdptr` keyword points to a Metal C function descriptor, which is an internal control block that contains two fields:

- The address of the actual function to be called when the function pointer is dereferenced
- The address of the data location or function environment

The length of each field is the same as that of a pointer in the current compilation's unit AMODE.

Defining an alternative name for function main

When building a Metal C program, you might need to define an alternative entry point name for function "main" while maintaining all the characteristics of function "main".

For example, if you want to call your Metal C "main" function as "ANEWMMAIN", you can add the following directive in your source file where "main" is defined:

```

#pragma map(main, "ANEWMMAIN")
void dosomething(char *);
int main(int argc, char *argv[]) {
    int i;

```

```

    for (i=1; i<argc; i++) {
        dosomething(argv[i]);
    }
    return 0;
}

```

The entry point name in the generated prolog/epilog code will be ANEWMAIN. When you link your program, you need to tell the binder that the entry point name is ANEWMAIN. For example:

```
-o a.out a.o -e ANEWMAIN
```

Notes:

1. In your C program, you can have only one "main" function, whether it is called "main" or otherwise. If you use IPA, IPA will terminate with an error message issued when more than one "main" function is detected.
2. The mapped entry point name for function "main" is subject to the effect of the LONGNAME option. If the NOLONGNAME option is in effect, the mapped name will be truncated to maximum of 8 characters, and the name will be in upper case, with "_" in the name converted to "@". For example, "a_newmain" will become "A@NEWMAI".

Building Metal C programs

Because the Enterprise Metal C for z/OS compiler produces the final code in HLASM source code format, the build process needs to include an assembly step to produce the object files. The build process is demonstrated in Figure 40 on page 45. Note that the build process with IPA is more elaborated. For more information, see "Building Metal C programs with IPA" on page 48.




Figure 40. Metal C application build process

In summary, the C source file is sent to the C compiler, which generates the assembler source file. The assembler source file is sent to the HLASM assembler, which generates an object file and a listing.

Examples of building Metal C programs

A set of examples illustrates how to build a Metal C program by using either z/OS UNIX System Services commands or MVS™ JCL procedures. In these examples:

- MYADD is the name of the entry point in the C program.
- The name of the C source file used to generate the HLASM source file is `mycode.c`.
- The name of the HLASM source file is `mycode.s` if it is generated under z/OS UNIX System Services.
- Under MVS, if the C source file is a data set, the compiler uses the C source file name to form the name of the HLASM source file. The high-level qualifier is replaced with the userid under which the compiler is running, and `.ASM` is appended as the low-level qualifier.

C source file

Figure 41 on page 46 shows a C source file `mycode.c` that can be used to generate an HLASM source file. The name of the generated HLASM source file is `mycode.s`

under z/OS UNIX System Services.

```
int myadd(void) {
    int a , b;
    a = 1;
    b = 2;
    __asm(" AR  %0,%1 "
          : "=r"(a)
          : "r"(b), "0"(a)
          );
    return a;
}
```

Figure 41. C source file (*mycode.c*) that builds a Metal C program

Building Metal C programs using z/OS UNIX System Services

There are three steps for building a Metal C program under z/OS UNIX System Services:

1. Use the `metalc` command to generate an HLASM source file.
2. Use the `as` command to generate the object file.
3. Use the `ld` command to generate the program.

Generating an HLASM source file using the `metalc` command

The generated HLASM source file has the same name as the C source file with the suffix derived from the `ssuffix` attribute in the `metal.c` configuration file. The default suffix is `s`, so in the examples in this section, the HLASM source file name is `mycode.s`.

```
metalc mycode.c
```

Figure 42. C compiler invocation to generate *mycode.s*

Generating an object file from the HLASM source using the z/OS UNIX System Services `as` command

The generated object file does not have to be a z/OS UNIX file. The `as` command can write the object file directly to an MVS data set, as shown in Figure 43. The `-o` flag can be used to name the output file, where it can be a UNIX file or an MVS data set.

```
as mycode.s
```

Figure 43. Command that invokes HLASM to assemble *mycode.s*

A successful assemble will produce `mycode.o`.

If the C source file was compiled with the `LONGNAME` compiler option, the generated HLASM source file will contain symbols that are more than eight characters in length. In that case, the HLASM `Goff` option must be specified. Use the `as` utility `-m` flag to specify HLASM options, as shown in Figure 44.

```
as -mgoff mycodelong.s
```

Figure 44. Command that compiles an HLASM source file containing symbols longer than eight characters

A successful assemble will produce `mycodelong.o`.

Creating a program with the z/OS UNIX System Services **ld** command

Use the **ld** command to link the object file produced by the **as** command into a program. The program does not have to be a z/OS UNIX file. The **ld** utility can write the program directly to a specified MVS data set.

Common **ld** command options that control the bind step are:

- **-e** to specify the entry point.
- **-o** to specify the name of the program created by the **ld** utility.
- **-v** to direct the binder listing to stdout.
- **-b** to specify other binder-specific options.

If you compile your C source file with the **LONGNAME** option, you should use **-b case=mixed** and the **-e** option must specify the entry point in its original case, as shown in Figure 45.

```
ld -b case=mixed -e MYADD -o "//LOAD(mycode)" mycode.o
```

Figure 45. Command that binds *mycode.o* and produces a Metal C program in an MVS data set

A successful bind produces **HLQ.LOAD(MYCODE)** with entry point **MYADD**.

Commands that compile and link applications that switch addressing modes

Figure 46 shows the commands that compile and link the programs in Figure 32 on page 32.

```
metalc a31.c
metalc -q64 a64a.c
as -a=a31.lst -mgoff a31.s
as -a=a64a.lst -mgoff a64a.s
ld -o a.out a31.o a64a.o -e MAIN
```

1
2
2
3

Figure 46. Commands that compile and link programs with different addressing modes

Notes:

1. The called program *a64a.c* is an external function in a separate source file.
2. The **-mgoff** command is used to compile an HLASM source file containing symbols longer than eight characters.
3. The **ld** command links the object file produced by the **as** command into a program. The **-e** command specifies the entry point.

Building Metal C programs using JCL

When you build Metal C programs using JCL, you cannot use standard JCL procedures that combine the compilation step with the link step (or link and run steps) because compiling Metal C programs produces HLASM source files that must be assembled by HLASM before they can be linked.

After successful completion of the assembly step, you can use an appropriate binder invocation JCL procedure to produce an program.

Note: Binder invocation JCL procedures are available in the **CEE.SCEEPROC** data set.

Compilation of HLQ.SOURCE.C(MYCODE)

```
//PROC JCLLIB ORDER=(CJT.SCJTPRC)
/*-----
/* Invoke Enterprise Metal C for z/OS compiler
/*-----
//METALCMP EXEC MTCC,
//      INFILE='HLQ.SOURCE.C(MYCODE)',
//      OUTFILE='HLQ.SOURCE.ASM(MYCODE),DISP=SHR',
```

Figure 47. Job step that compiles HLQ.SOURCE.C(MYCODE)

Assembly of HLQ.SOURCE.ASM(MYCODE)

```
/*-----
/* ASSEMBLY STEP:
/*-----
//ASM      EXEC HLASM
//SYSIN    DD DSN=HLQ.SOURCE.ASM(MYCODE),DISP=SHR
//SYSLIN    DD DSN=HLQ.OBJ(MYCODE),DISP=OLD
```

Figure 48. Assembly step of HLQ.SOURCE.ASM(MYCODE)

Bind of HLQ.OBJ(MYCODE) into a Metal C program

```
/*-----
/* BIND STEP:
/*-----
//BIND     EXEC PGM=IEWL,
//          PARM='AMODE=31,MAP,CASE=MIXED'
//SYSLMOD  DD DSN=HLQ.LOAD(MYCODE),DISP=SHR 1
//SYSPRINT DD SYSOUT=*
//OBJECT   DD DSN=HLQ.OBJ,DISP=SHR
//SYSLIN   DD *
//          INCLUDE OBJECT(MYCODE)
//          ENTRY MYADD 2
/*
```

Figure 49. Job step that binds the generated HLASM object into a program

Notes:

1. The program is written to SYSLMOD.
2. The entry point can be specified using the ENTRY binder control statement.

Building Metal C programs with IPA

IPA is an optimization option that enables the compiler to find more optimization opportunities to improve your application performance.




Figure 50. The process of building Metal C programs with IPA

You need to be aware of the following adjustments when invoking IPA.

- The **LONGNAME** option is in effect by default when IPA is specified.
- The IPA compile step only produces IPA object in the output file. Only **IPA(NOOBJECT)** is allowed, which instructs IPA to stop the compile process after the IPA object is produced. It does not produce HLASM source code, so the **GENASM** option cannot be used.
- The output file from IPA link step is one single HLASM source file for the whole program, and the **GENASM** option is required. There could be multiple structures in the HLASM source program, one for each partition. Under z/OS UNIX, the output HLASM source file resides in the directory where the IPA link took place. The default output file name for z/OS UNIX is `a.s`. In BATCH mode, the output HLASM source file goes in the dataset allocated to DD SYSLIN in the IPA link step.
- At the IPA link phase, all external references must be resolved. For Metal C, IPA does not attempt to convert external object modules or load modules into object code for the inclusion in the IPA produced program. You need to provide the same set of library data sets to both IPA link and the binder for symbol resolution.
- If you specify the **PROLOG** and **EPILOG** compiler options to supply your own prolog and epilog macros at compile time, the macros will only be applied to the functions defined in the source file.

- If you have `#pragma insert_asm` in your source file, IPA will assume the strong connection between the string provided by the pragma and the functions in the source file. IPA will not move functions defined in that source file to anywhere else.
- If you use global register variable or the **RESERVE_REGS** option during your compile, IPA link will merge the registers specified in the compile steps and apply the merged set of the originated compilation units to a partition.
- If you use the **DSAUSER** option in any of your compile steps, IPA link applies the option to the entire program.

The following IPA control file directives are not supported:

- **EXPORT**
- **NOEXPORTS**

Example

The following example shows how to compile a Metal C program with IPA.

IPA compile phase:

```
metalc -qipa -c x.c
metalc -qipa -c y.c
```

The above commands produce `x.o` and `y.o`.

Notes:

1. The `-c` option indicates compile.
2. No HLASM output is generated.
3. The objects are IPA objects, which can only be used for IPA link.
4. **LONGNAME** is implicitly turned on.

IPA link phase:

```
metalc -qipa x.o y.o
```

This command produces `a.s`.

Note: The structure of the compiler-generated HLASM source program is similar to that described in “Structure of a compiler-generated HLASM source program” on page 5, except that at IPA link there could be multiple structures in the HLASM source program, one for each partition.

The rest of the build process is similar to building Metal C programs without IPA. You need to add the assembly step to produce the object file from the IPA link generated HLASM source file. You also need to supply the object file produced by the assembler along with all other library data sets to the binder for producing the final executable program.

Assembly phase:

```
as -mgoff a.s
```

This command produces `a.o`.

Note: The HLASM **GOFF** option must be specified because of the **LONGNAME** compiler option requirement with IPA.

Bind/Link phase:

```
ld -b case=mixed -e main a.o
```

This command produces a.out.

Note: Because of the **LONGNAME** compiler option requirement with IPA, you should use the **-b case=mixed** option of the **ld** utility and the **-e** option with the entry point in its original case.

Generation of debugging information

When the **DEBUG** compiler option is specified, debugging information in both ADATA and DWARF format can be generated. The ADATA debug format allows debugging of the generated HLASM source. The DWARF debug format allows debugging of the original C source.

CDAASMC JCL procedure to generate debugging information

The **as** command is a z/OS UNIX System Services utility that invokes the HLASM assembler and can produce debugging information in DWARF format. CDAASMC is the JCL procedure provided with the Enterprise Metal C for z/OS compiler to do the same thing in a batch environment.

Note: If you wish to use the HLASM ASMLANGX debugging utility, you must first assemble your source with the **ADATA** assembler option. The CDAASMC JCL procedure allows you to generate both ADATA and DWARF debugging information.

The cataloged CDAASMC JCL procedure invokes CDAHLASM.

Debugging information for the IDF debugger

The Interactive Debug Facility (IDF) is a symbolic debugging tool for assembly language programs. It uses information from the load module file to determine the locations of a program's control sections and external symbols.

Optionally, IDF can make use of additional information to help disassemble the program. The additional information can be generated by specifying the assembler **TEST** option and the linkage editor **TEST** option.

Note: The Linkage Editor **TEST** option can make the final load module file quite large. If you prefer to suppress them, either omit the linkage editor **TEST** option or specify the **NOTEST** option.

The Linkage Editor **TEST** option increases the size of the load module file, so do not use it for production modules.

ADATA debugging information

The ASMLANGX utility extracts source level information from the ADATA debugging information. The output is an extract file. Although you can create extract files as sequential files, they are typically stored in a PDS.

The recommended format for the extract file is:

```
RECFM(VB) LRECL(1562) BLKSIZE(27998)
```

```
//ASMLANGX EXEC PGM=ASMLANGX,REGION=4096K,
// PARM='member (ASM LOUD ERROR'
//SYSADATA DD DISP=SHR,DSN=h1q..SYSADATA
//ASMLANGX DD DISP=OLD,DSN=h1q..ASMLANGX
```

1
2
3

Notes:

1. The PDS member name of the input and output file is passed as a parameter. For sequential files, this name is ignored.
2. The SYSADATA DD statement specifies the input data set name.
3. The ASMLANGX DD statement specifies the output data set name.

Figure 51. JCL that invokes the ASMLANGX utility

IDF debugger invocation

If you want to use an interactive utility to debug your program, invoke the IDF debugger by performing the following steps:

1. Specify the problem load module and the extract file that contains the debugging information by entering the following commands.

```
ALLOC FI (ASMLANGX) DS('h1q.ASMLANGX') SHR
TSOLIB ACT DS('h1q.LOAD')
```

2. Invoke IDF by entering the following command:

```
ASMIDF MYCODE
```

3. Press F9 to get the **Program Source and Disassembly** view.

Summary of useful references for the Metal C programmer

Table 7 lists topics of interest to the Metal C programmer and, for each topic, lists information found in this document, as well as external references.

Table 7. Summary of useful references for the Metal C programmer

Information	Internal reference	External references
The base linkage conventions that are used by the generated modules.	"Metal C and MVS linkage conventions" on page 2	For detailed information about MVS linkage conventions, see <i>Linkage conventions in z/OS MVS Programming: Assembler Services Guide</i> .
The Metal C Runtime Library.	Chapter 2, "Header files," on page 55	For additional information about the Metal runtime library, see METAL compiler option and Metal C Runtime Library (www.ibm.com/systems/z/os/zos/features/metalc).

Table 7. Summary of useful references for the Metal C programmer (continued)

Information	Internal reference	External references
Using assembler statements within a C program.	• “Inserting HLASM instructions into the generated source code” on page 21 • “Inserting non-executable HLASM statements into the generated source code” on page 30	For detailed information about HLASM programming, see <i>HLASM MVS & VM Programmer’s Guide</i>. For detailed information about inline assembly statements, see Inline assembly statements (IBM extension) in <i>Enterprise Metal C for z/OS Language Reference</i>. For more information about callable system services, see <i>z/OS MVS Programming: Callable Services for High-Level Languages</i>.
Programming with the metal c utility	“Programming with Enterprise Metal C for z/OS” on page 2	
Making access registers available to the Metal C application.	“AR-mode programming support” on page 35	For detailed information about using access registers, see <i>z/OS MVS Programming: Extended Addressability Guide</i> .
Providing prolog and epilog code to customize the environment.	• “Compiler-generated global SET symbols” on page 15 • “Supplying your own prolog and epilog code” on page 14	<i>Not applicable.</i>
Building the application by using JCL procedures.	“Building Metal C programs using JCL” on page 47	<i>Not applicable.</i>
Building the application by using z/OS UNIX System Services.	“Building Metal C programs using z/OS UNIX System Services” on page 46	<i>Not applicable.</i>
Generating the appropriate debugging information.	“Generation of debugging information” on page 51	<i>Not applicable.</i>

Chapter 2. Header files

Header files for the Metal C Runtime Library are located in the z/OS UNIX file system directory: `/usr/include/metal`. The Enterprise Metal C for z/OS compiler automatically searches for the Metal C header files under this directory.

Note: Some Metal C header files such as `stdio.h` have the same names as header files for the Language Environment C/C++ Runtime Library. To avoid including these, or inadvertently including any other headers supported by the LE library and not by Metal C, remove the non-Metal libraries from the search order. Depending on how you specify the system library search path, you need to remove other libraries from the SYSLIB concatenation of the compiler, or specify the **NOSEARCH** compiler option before pointing to `/usr/include/metal/`.

If you are compiling in batch, you can use the **SEARCH** compiler option:

```
SEARCH(/usr/include/metal/)
```

If you are compiling using the **NOSEARCH** compiler option, you have the following options:

- Use the **-I** option of the `metalc` utility.
`-I /usr/include/metal/`
- Use the `cinc` attribute in the `metalc` configuration file.
`cinc = -I /usr/include/metal/`

builtins.h — Declare built-in functions

The `builtins.h` header contains a list of built-in functions supported by the compiler. A built-in function is inline code that is generated in place of an actual function call.

ctype.h — Declare character classification functions

The `ctype.h` header file declares functions used in character classification. The `ctype.h` header file declares the following functions.

<code>isalnum()</code>	<code>isalpha()</code>	<code>isblank()</code>	<code>iscntrl()</code>	<code>isdigit()</code>
<code>isgraph()</code>	<code>islower()</code>	<code>isprint()</code>	<code>ispunct()</code>	<code>isspace()</code>
<code>isupper()</code>	<code>isxdigit()</code>	<code>tolower()</code>	<code>toupper()</code>	

Note: All the functions in the previous table use code page IBM-1047.

float.h — Define ANSI constants for floating-point data types

The `float.h` header file contains definitions of constants listed in ANSI 2.2.4.2.2. The constants describe the characteristics of the internal representations of the three floating-point data types: `float`, `double`, and `long double`. Table 8 on page 56 lists the definitions contained by `float.h`.

Table 8. Definitions in *float.h*

Constant	Description
FLT_RADIX	The radix for a Metal C application. For <code>float(IEEE)</code> , the value is 2.
FLT_MANT_DIG DBL_MANT_DIG LDBL_MANT_DIG	The number of hexadecimal digits stored to represent the significand of a fraction.
FLT_DIG DBL_DIG LDBL_DIG	The number of decimal digits, <i>q</i> , such that any floating-point number with <i>q</i> decimal digits can be rounded into a floating-point number with <i>p</i> radix <code>FLT_RADIX</code> digits, and back again, without any change to the <i>q</i> decimal digits.
FLT_MIN_10_EXP DBL_MIN_10_EXP LDBL_MIN_10_EXP	The minimum negative integer such that 10 raised to that power is in the range of normalized floating-point numbers.
FLT_MAX_EXP DBL_MAX_EXP LDBL_MAX_EXP	The maximum integer such that <code>FLT_RADIX</code> raised to that power minus 1 is a representable finite floating-point number.
FLT_MAX_10_EXP DBL_MAX_10_EXP LDBL_MAX_10_EXP	The maximum integer such that 10 raised to that power is in the range of representable finite floating-point numbers.
FLT_MAX DBL_MAX LDBL_MAX	The maximum representable finite floating-point number.
FLT_EPSILON DBL_EPSILON LDBL_EPSILON	The difference between 1.0 and the least value greater than 1.0 that is representable in the given floating-point type.
FLT_MIN DBL_MIN LDBL_MIN	The minimum normalized positive floating-point number.
DECIMAL_DIG	The minimum number of decimal digits needed to represent all the significant digits for type long double.
FLT_EVAL_METHOD	Describes the evaluation mode for floating point operations. This value is 1, which evaluates • All operations and constants of types <code>float</code> and <code>double</code> to type <code>double</code>. • All operations and constants of long double to type long double.

inttypes.h — Define macros for `sprintf` and `sscanf` family

The following macros are defined in `inttypes.h`. Each expands to a character string literal containing a conversion specifier which can be modified by a length modifier that can be used in the *format* argument of a formatted input/output function when converting the corresponding integer type. These macros have the general form of `PRI` or `SCN`, followed by the conversion specifier, followed by a name corresponding to a similar type name in `inttypes.h`. In these names, the suffix number represents the width of the type. For example, `PRIdFAST32` can be used in a format string to print the value of an integer of type `int_fast32_t`.

Compile requirement: In the following list all macros with the suffix MAX or 64 require long long to be available.

Macros for sprintf family for signed integers.

PRId8	PRId16	PRId32	PRId64
PRIdLEAST8	PRIdLEAST16	PRIdLEAST32	PRIdLEAST64
PRIdFAST8	PRIdFAST16	PRIdFAST32	PRIdFAST64
PRIdMAX			
PRIdPTR			
PRi8	PRi16	PRi32	PRi64
PRiLEAST8	PRiLEAST16	PRiLEAST32	PRiLEAST64
PRiFAST8	PRiFAST16	PRiFAST32	PRiFAST64
PRiMAX			
PRiPTR			

Compile requirement: In the following list all macros with the suffix MAX or 64 require long long to be available.

Macros for sprintf family for unsigned integers.

PRIo8	PRIo16	PRIo32	PRIo64
PRIoLEAST8	PRIoLEAST16	PRIoLEAST32	PRIoLEAST64
PRIoFAST8	PRIoFAST16	PRIoFAST32	PRIoFAST64
PRIoMAX			
PRIoPTR			
PRiU8	PRiU16	PRiU32	PRiU64
PRiULEAST8	PRiULEAST16	PRiULEAST32	PRiULEAST64
PRiUFAST8	PRiUFAST16	PRiUFAST32	PRiUFAST64
PRiUMAX			
PRiUPTR			
PRIx8	PRIx16	PRIx32	PRIx64
PRIxLEAST8	PRIxLEAST16	PRIxLEAST32	PRIxLEAST64
PRIxFAST8	PRIxFAST16	PRIxFAST32	PRIxFAST64
PRIxMAX			
PRIxPTR			
PRIX8	PRIX16	PRIX32	PRIX64
PRIXLEAST8	PRIXLEAST16	PRIXLEAST32	PRIXLEAST64
PRIXFAST8	PRIXFAST16	PRIXFAST32	PRIXFAST64
PRIXMAX			
PRIXPTR			

Compile requirement: In the following list all macros with the suffix MAX or 64 require long long to be available.

Macros for sscanf family for signed integers.

SCNd8	SCNd16	SCNd32	SCNd64
SCNdLEAST8	SCNdLEAST16	SCNdLEAST32	SCNdLEAST64
SCNdFAST8	SCNdFAST16	SCNdFAST32	SCNdFAST64
SCNdMAX			
SCNdPTR			
SCNi8	SCNi16	SCNi32	SCNi64
SCNiLEAST8	SCNiLEAST16	SCNiLEAST32	SCNiLEAST64
SCNiFAST8	SCNiFAST16	SCNiFAST32	SCNiFAST64

SCNiMAX
SCNiPTR

Compile requirement: In the following list all macros with the suffix MAX or 64 require long long to be available.

Macros for sscanf family for unsigned integers.

SCNo8	SCNo16	SCNo32	SCNo64
SCNoLEAST8	SCNoLEAST16	SCNoLEAST32	SCNoLEAST64
SCNoFAST8	SCNoFAST16	SCNoFAST32	SCNoFAST64
SCNoMAX			
SCNoPTR			
SCNu8	SCNu16	SCNu32	SCNu64
SCNuLEAST8	SCNuLEAST16	SCNuLEAST32	SCNuLEAST64
SCNuFAST8	SCNuFAST16	SCNuFAST32	SCNuFAST64
SCNuMAX			
SCNuPTR			
SCNx8	SCNx16	SCNx32	SCNx64
SCNxLEAST8	SCNxLEAST16	SCNxLEAST32	SCNxLEAST64
SCNxFAST8	SCNxFAST16	SCNxFAST32	SCNxFAST64
SCNxMAX			
SCNxPTR			

limits.h — Declare symbolic names for resource limits

The `limits.h` header file contains symbolic names that represent standard values for limits on resources, such as the maximum value for an object of type `char`.

Symbolic name

Resource limit

CHAR_BIT

8

CHAR_MAX

127 (`_CHAR_SIGNED`)

CHAR_MAX

255

CHAR_MIN

(-128) (`_CHAR_SIGNED`)

CHAR_MIN

0

INT_MAX

2147483647

INT_MIN

(-2147483647 - 1)

LLONG_MAX

(9223372036854775807LL)

LLONG_MIN

(-LLONG_MAX-1)

```

LONG_MAX
    2147483647

LONGLONG_MAX
    (9223372036854775807LL)

LONG_MIN
    (-2147483647L - 1)

LONGLONG_MIN
    (-LONGLONG_MAX - 1)

MB_LEN_MAX
    4

SCHAR_MAX
    127

SCHAR_MIN
    (-128)

SHRT_MAX
    32767

SHRT_MIN
    (-32768)

SSIZE_MAX
    2147483647

UCHAR_MAX
    255

UINT_MAX
    4294967295

ULONG_MAX
    4294967295U

ULONGLONG_MAX
    (18446744073709551615ULL)

ULLONG_MAX
    (18446744073709551615ULL)

USHRT_MAX
    65535

```

math.h — Define macros for floating-point support

The `math.h` header file contains macro declarations for use with floating-point support:

No feature test macro is required.

Object-like Macros: The definitions are as follows.

HUGE_VAL
A very large positive number that expands to a double expression.

HUGE_VALF
A very large positive number that expands to a float expression.

HUGE_VALL
A very large positive number that expands to a long double expression.

INFINITY

A constant expression of type float representing positive infinity.

NAN

A constant expression of type float representing a quiet NaN.

metal.h — Define Metal C related function prototypes and data

The `metal.h` header file contains function prototypes and data definitions related to the Metal C runtime library, including the `__cinit()` and `__cterm()` functions.

The `metal.h` header file also includes `__csysenv_s`, which is the structure used to describe the characteristics of a Metal C environment. For more information about the `__csysenv_s` structure, see “`__cinit()` - Initialize a Metal C environment” on page 72.

Note: The `metal.h` header file is automatically included by any Metal C runtime library header file, so it is not necessary to explicitly include it if a header file is being used.

stdarg.h — Define macros for accessing variable-length argument lists in functions

The `stdarg.h` header file defines macros used to access arguments in functions with variable-length argument lists.

`va_arg()` `va_copy()` `va_start()` `va_end()`

The `stdarg.h` header file also defines the structure `va_list`.

The `stdarg.h` header file defines `va_list` as `char *va_list`.

stddef.h — Define `ptrdiff_t`, `size_t`, and `ssize_t` data types

The `stddef.h` header file defines the following types:

`ptrdiff_t`

The signed long type of the result of subtracting two pointers.

`size_t` typedef for the type of the value returned by `sizeof`.

`ssize_t`

similar to `size_t`, but must be a signed type.

The `stddef.h` header defines the macros `NULL` and `offsetof`. `NULL` is a pointer that never points to a data object. The `offsetof` macro expands to the number of bytes between a structure member and the start of the structure. The `offsetof` macro has the form `offsetof(structure_type, member)`.

stdio.h — Define I/O related functions

The `stdio.h` header file declares the following functions:

`snprintf()` `sprintf()` `scanf()` `vsnprintf()` `vsprintf()` `vscanf()`

The `stdio.h` header file also contains definition for the following macro, whose value should not be altered:

`NULL` A pointer which never points to a data object.

stdint.h — Define integer types and related limits and macros

The `stdint.h` header defines integer types, limits of specified width integer types, limits of other integer types, and macros for integer constant expressions.

Note: For the exact width integer types, minimum-width integer types, limits of specified width integer types, and exact width integer constants, *bit sizes* N with the values 8, 16, 32, and 64 are supported.

Requirement: Use of the bit size 64 and greatest-width integer types or macros require LP64 or the long long data type to be available.

Integer types

The following exact width integer types are defined.

- `intN_t`
- `uintN_t`

The following minimum-width integer types are defined.

- `int_leastN_t`
- `uint_leastN_t`

The following fastest minimum-width integer types are defined. These types are the fastest to operate with among all integer types that have at least the specified width.

- `int_fastN_t`
- `uint_fastN_t`

The following greatest-width integer types are defined. These types hold the value of any signed/unsigned integer type.

Note: Requires `long long` to be available.

- `intmax_t`
- `uintmax_t`

The following integer types capable of holding object pointers are defined.

- `intptr_t`
- `uintptr_t`

Object-like macros for limits of integer types:

Macros for limits of exact width integer types.

- `INTN_MAX`
- `INTN_MIN`
- `UINTN_MAX`

Macros for limits of minimum width integer types.

- `INT_LEASTN_MAX`

- INT_LEASTN_MIN
- UINT_LEASTN_MAX

Macros for limits of fastest minimum width integer types.

- INT_FASTN_MAX
- INT_FASTN_MIN
- UINT_FASTN_MAX

Macros for limits of greatest width integer types.

- INTMAX_MAX
- INTMAX_MIN
- UINTMAX_MAX

Macros for limits of pointer integer types.

- INTPTR_MAX
- INTPTR_MIN
- UINTPTR_MAX

Macros for limits of ptrdiff_t.

- PTRDIFF_MAX
- PTRDIFF_MIN

Macro for limit of size_t.

- SIZE_MAX

Function-like macros for integer constants:

Macros for minimum width integer constants.

- INTN_C(*value*)
- UINTN_C(*value*)

Macros for greatest-width integer constants:

- INTMAX_C(*value*)
- UINTMAX_C(*value*)

stdlib.h — Define standard library functions

The `stdlib.h` header file contains declarations for the following functions.

<code>abs()</code> ¹	<code>atoi()</code>	<code>atol()</code>	<code>atoll()</code>	<code>calloc()</code>
<code>div()</code>	<code>free()</code>	<code>labs()</code>	<code>ldiv()</code>	<code>llabs()</code>
<code>lldiv()</code>	<code>malloc()</code>	<code>__malloc31()</code>	<code>qsort()</code>	<code>rand()</code>
<code>rand_r()</code>	<code>realloc()</code>	<code>srand()</code>	<code>strtod()</code>	<code>strtof()</code>
<code>strtol()</code>	<code>strtold()</code>	<code>strtoll()</code>	<code>strtoul()</code>	<code>strtoull()</code>

¹ Built-in function.

Two type definitions are added to `stdlib.h` for the Compare and Swap functions `cs()` and `cds()`. The structures defined are `cs_t` and `cds_t`.

The type `size_t` is declared in the header file. It is used for the type of the value returned by `sizeof`. For more information on the types `size_t`, see “`stddef.h` — Define `ptrdiff_t`, `size_t`, and `ssize_t` data types” on page 60.

The `stdlib.h` declares `div_t`, `ldiv_t`, and `lldiv_t`, which define the structure types that are returned by `div()`, `ldiv()`, and `lldiv()`.

The `stdlib.h` file also contains definitions for the following macros:

NULL The NULL pointer constant (also defined in `stddef.h`).

RAND_MAX

Expands to an integer representing the largest number that the `rand()` or `rand_r()` function can return.

string.h — Declare string manipulation functions

The `string.h` header file declares the string manipulation functions and their built-in versions.

<code>memcpy()</code>	<code>memchr()</code> ¹	<code>memcmp()</code> ¹	<code>memcpy()</code> ¹	<code>memmove()</code>
<code>memset()</code> ¹	<code>strcat()</code> ¹	<code>strchr()</code> ¹	<code>strcmp()</code> ¹	<code>strcpy()</code> ¹
<code>strcspn()</code>	<code>strdup()</code>	<code>strlen()</code> ¹	<code>strncat()</code> ¹	<code>strncmp()</code> ¹
<code>strncpy()</code> ¹	<code>strpbrk()</code>	<code>strrchr()</code> ¹	<code>strspn()</code>	<code>strstr()</code>
<code>strtok()</code>	<code>strtok_r()</code>			

¹ Built-in function.

The `string.h` header file also defines the macro `NULL` and the type `size_t`. For more information see “`stddef.h` — Define `ptrdiff_t`, `size_t`, and `ssize_t` data types” on page 60.

Chapter 3. C functions available to Metal C programs

This topic describes the Metal C runtime library functions.

The linkage conventions used by the Enterprise Metal C for z/OS compiler govern use of the C functions that are available to freestanding programs that are compiled by Enterprise Metal C for z/OS. For more information, see “Metal C and MVS linkage conventions” on page 2.

When you use any of these supplied C functions, be aware of the information provided in “Characteristics of compiler-generated HLASM source code” on page 4.

Characteristics of Metal C runtime library functions

Linkage to each function is through the default linkage provided by the Enterprise Metal C for z/OS compiler. This assumes that GPR 13 points to a stack frame in a contiguous stack, and that the forward pointer in the stack frame contains the address of the next available byte in the stack. The stack frame requirements for each function are documented in Appendix A, “Function stack requirements,” on page 125 so that the caller knows how much space to reserve.

The library functions support AMODE 31 and AMODE 64.

The library functions (with the exception of a few AR mode supporting functions) expect the ASC mode to be Primary on entry. The AR mode support part of Metal C ensures that this is enforced; however, if calling these library functions from within HLASM embedded statements or their own HLASM programs, you need to manage ASC mode to meet this requirement.

The library functions support IEEE floating point numbers.

The library uses code page IBM-1047 and the En_US locale definitions to perform its functions.

System and static object libraries

The Metal C runtime library supports two versions of its library functions: a system library and a static object library. The behavior of the functions within the two versions is the same. What differs is where the functions are located and how the Metal C application interacts with them.

System library

The system library is a version of the Metal C runtime library that exists within the system's link pack area, and is made available during the system IPL process. It is suggested that you use the system library if the Metal C application is run on a level of z/OS that supports the runtime library, and the application runs after the library has been made available. This library has the added advantage of not requiring application module relinks when service is applied to the library.

To use the system library version, you can include the wanted Metal C runtime library headers in the Metal C application source code. The default behavior of the

headers is to generate code within the application that calls this system library. No additional binding is needed in order for these function calls to work.

Note: The compiler does not generate inline code for the following built-in functions when **COMPACT** is specified: `memcmp`, `strcat`, `strchr`, `strncat`, `strncmp`, `strncpy`. In this case, you need to define the `__METAL_SYSVEC` feature test macro.

Static object library

The Static object library is a version of the Metal C runtime library that gets directly bound with a Metal C application load module. The resulting application is self-contained with respect to the library; all library function calls from the application result in the functions that are bound within the load module to be driven.

It is suggested that you use the static object library if the Metal C application is run during system IPL before the system library has been made available.

The static object library functions are provided in two system data sets: `SYS1.SCCR3BND` and `SYS1.SCCR6BND`. `SYS1.SCCR3BND` is used with Metal C applications that have been compiled using ILP32 and run AMODE 31. `SYS1.SCCR6BND` is used with Metal C applications that have been compiled using LP64 and run AMODE 64.

To use the static object library, you must take the following steps:

1. Define the `__METAL_STATIC` feature test macro before including the headers in your Metal C program, and then compile the program. For example:

```
#define __METAL_STATIC
#include <stdio.h>
```

This will cause library function calls in the program to generate external references to the functions contained within the `SCCRnBND` data sets.

2. Bind the compiled object with the corresponding `SCCRnBND` data set. How this is done depends on the environment in which the binding takes place:

Batch When using the binder from a batch job, use the **CALL** option, and use the `SYSLIB DD` to identify the static object library data set that you want to bind with.

z/OS UNIX System Services shell

From the shell, it is suggested that the **ld** shell command be used to bind the application with the library functions.

Note: When service is applied to the static object library, the Metal C application must be relinked to pick up the changes.

General library usage notes

- A Metal C application can use either the system library or the static object library, but not both. The mixing of system library calls and static object library calls within the same application is not supported.
- All static objects that are bound to the application load module must be at compatible service levels.
- Metal C runtime library functions are not supported under Language Environment and must not be used within a Language Environment program, because equivalent functions are already available.

User-replaceable heap services

The Metal C Runtime Library provides the ability to completely replace the underlying heap services at run time. You can use this function if you want the heap services to use a different storage management mechanism, for instance, one that is already in use elsewhere within an application.

A Metal C application replaces the underlying heap services by providing sets of function entry points in the `__csysenv_s` structure that is passed to the `__cinit()` function. To have the function entry point fields available and recognized by the `__cinit()` function, take the following steps:

- Define `__METAL_CSYSENV_VERSION 2` so that the `__csysenv_s` structure contains the heap service function entry point fields.
- In the `__csysenv_s` structure, set field `__cseversion` to `__CSE_VERSION_2`.
- In the `__csysenv_s` structure, provide addresses for heap functions that are to be replaced.
- Call the `__cinit()` function, providing the `__csysenv_s` structure that was initialized.

The Metal C application can provide at environment initialization time 8 bytes of data that can be accessed by the replacement heap services. To reserve the 8 bytes of data, take the following steps:

- Before calling the `__cinit()` function, set the user data of the application in field `__cseheapuserdata` in the `__csysenv_s` structure.
- Use the R12 environment token value as a pointer to the `__csysenvtoken_s` structure. In this structure, field `__csetheapuserdata` contains 8 bytes of data of the application.

During the `__cinit()` call, field `__csetheapuserdata` can only be set from `__cseheapuserdata` if heap services have been replaced; otherwise, field `__csetheapuserdata` will be set to binary zeroes.

Two sets of heap service function entry points are provided, one set for replacing heap services in the AMODE 31 version of the library, and the other set for replacing heap services in the AMODE 64 version of the library.

AMODE 31 heap services

To replace heap services in the AMODE 31 version of the library, consider the following `__csysenv_s` fields:

void * (*__cseamode31malloc) (size_t)

When specified, MCRTL AMODE 31 `malloc()` calls this routine to obtain a piece of below-the-bar heap storage and returns its result to the caller of `malloc()`. `__cseamode31malloc` is treated as having the same function prototype as `malloc()`: `void * malloc (size_t);`

void (* __cseamode31free) (void *)

When specified, MCRTL AMODE 31 `free()` calls this routine to free a piece of heap storage. `__cseamode31free` is treated as having the same function prototype as `free()`: `void free(void *);`

void * (*__cseamode31realloc) (void *, size_t)

When specified, MCRTL AMODE 31 `realloc()` calls this routine to perform a `realloc` for a piece of heap storage and returns its result to the

caller of `realloc()`. `__cseamode31realloc` is treated as having the same function prototype as `realloc()`: `void * realloc (void *, size_t);`

Providing this routine is optional. If `realloc()` is called when a `__cseamode31malloc` routine has been provided but `__cseamode31realloc` has not, `realloc()` will return a zero.

Note: `__cseamode31malloc` and `__cseamode31free` must be provided together. `__cseamode31realloc` is optional, but when it is provided, the application must also include the other AMODE 31 heap services in this set.

AMODE 64 heap services

To replace heap services in the AMODE 64 version of the library, consider the following `__csysenv_s` fields:

void * (* __cseamode64malloc) (size_t)

When specified, MCRTL AMODE 64 `malloc()` calls this routine to obtain a piece of above-the-bar heap storage and returns to the caller of `malloc()`. `__cseamode64malloc` is treated as having the same function prototype as `malloc()`: `void * malloc (size_t);`

void * (* __cseamode64malloc31) (size_t)

When specified, MCRTL AMODE 64 `__malloc31()` calls this routine to obtain a piece of below-the-bar heap storage and returns its result to the caller of `__malloc31()`. `__cseamode64malloc31` is treated as having the same function prototype as `__malloc31()`: `void * __malloc31(size_t);`

void (* __cseamode64free) (void *)

When specified, MCRTL AMODE 64 `free()` calls this routine to free a piece of heap storage. `__cseamode64free` is treated as having the same function prototype as `free()`: `void free(void *)`;

Note that MCRTL AMODE 64 `free()` accepts as input and processes heap storage that is allocated above or below the bar. The user-specified `__cseamode64free` routine must provide the same capability.

void * (* __cseamode64realloc) (void *, size_t)

When specified, MCRTL AMODE 64 `realloc()` calls this routine to perform a `realloc` for a piece of heap storage and returns its result to the caller of `realloc()`. `__cseamode64realloc` is treated as having the same function prototype as `realloc()`: `void * realloc (void *, size_t);`

Providing this routine is optional. If `realloc()` is called when a `__cseamode64malloc` routine has been provided but `__cseamode64realloc` has not, `realloc()` will return a zero.

Note that MCRTL AMODE 64 `realloc()` accepts as input and processes heap storage that is allocated above or below the bar. The user-specified `__cseamode64realloc` routine must provide the same capability.

Note: `__cseamode64malloc`, `__cseamode64malloc31`, and `__cseamode64free` must all be provided together. `__cseamode64realloc` is optional, but when it is provided, the application must also include the other AMODE 64 heap services in this set.

Usage notes

- Each heap service gets control in the AMODE of the calling service. The heap service must return to the calling service in that same AMODE.

- Each heap service is called using standard Metal C linkage conventions, including:
 - GPR 1 containing the address of the function parameter list (using C style parameter passing)
 - GPR 13 containing the address of a stack frame allocated on a contiguous Metal C stack

GPR 12 contains the environment token representing the Metal C environment that is currently in use.

- It is not necessary to provide a replacement for the `calloc()` function. The `calloc()` function calls `malloc()` as part of its processing, so replacing `malloc()` indirectly alters the behavior of `calloc()` as well.
- When user-provided heap services are in use, the Metal C Runtime Library makes no attempt to keep track of any heap storage that has been allocated by the application. The application is entirely responsible for tracking its heap storage, and for freeing it after it calls `__cterm()` to terminate the Metal C environment.
- The heap allocation functions should return NULL when they are unable to obtain storage. The application is responsible for capturing its own diagnostic data when necessary.
- The Metal C Runtime Library expects the following alignment for the storage that is returned by the replacement heap services:
 - Storage returned from the below-the-bar heap (AMODE 64 `__malloc31()`, and AMODE 31 `malloc()`) is doubleword aligned.
 - Storage returned from the above-the-bar heap (AMODE 64 `malloc()`) is quadword aligned.

abs() — Calculate integer absolute value

Format

```
#include <stdlib.h>
```

```
int abs(int n);
```

General description

The `abs()` function returns the absolute value of an argument *n*.

For the integer version of `abs()`, the minimum allowable integer is `INT_MIN+1`. (`INT_MIN` is a macro that is defined in the `limits.h` header file.) For example, with the Metal C compiler, `INT_MIN+1` is -2147483647.

Returned value

The returned value is the absolute value, if the absolute value is possible to represent.

Otherwise the input value is returned.

Related Information

- “limits.h — Declare symbolic names for resource limits” on page 58
- “stdlib.h — Define standard library functions” on page 62
- “labs() — Calculate long absolute value” on page 79

atoi() — Convert character string to integer

Format

```
#include <stdlib.h>
```

```
int atoi(const char *nptr);
```

General description

The `atoi()` function converts the initial portion of the string pointed to by *nptr* to a 'int'. This is equivalent to

```
(int)strtol(nptr, (char **)NULL, 10)
```

Returned value

If successful, `atoi()` returns the converted int value represented in the string.

If unsuccessful, `atoi()` returns an undefined value.

Related Information

- “`stdlib.h` — Define standard library functions” on page 62
- “`atol()` — Convert character string to long”
- “`atoll()` — Convert character string to signed long long” on page 71
- “`strtol()` — Convert Character String to Long” on page 112
- “`strtoll()` — Convert String to Signed Long Long” on page 114
- “`strtoul()` — Convert String to Unsigned Integer” on page 116
- “`strtoull()` — Convert String to Unsigned Long Long” on page 117

atol() — Convert character string to long

Format

```
#include <stdlib.h>
```

```
long int atol(const char *nptr);
```

General description

The `atol()` function converts the initial portion of the string pointed to by *nptr* to a 'long int'. This is equivalent to

```
strtol(nptr, (char **)NULL, 10)
```

Returned value

If successful, `atol()` returns the converted long int value represented in the string.

If unsuccessful, `atol()` returns an undefined value.

Related Information

- “`stdlib.h` — Define standard library functions” on page 62
- “`atoi()` — Convert character string to integer”
- “`atoll()` — Convert character string to signed long long” on page 71
- “`strtol()` — Convert Character String to Long” on page 112
- “`strtoll()` — Convert String to Signed Long Long” on page 114
- “`strtoul()` — Convert String to Unsigned Integer” on page 116
- “`strtoull()` — Convert String to Unsigned Long Long” on page 117

atoll() — Convert character string to signed long long

Format

```
#define _ISOC99_SOURCE
#include <stdlib.h>
long long atoll(const char *nptr);
```

Compile Requirement: Use of this function requires the long long data type.

General description

The function converts the initial portion of the string pointed to by *nptr* to a long long int. This is equivalent to `strtoll(nptr, (char **)NULL, 10)`.

Returned value

If successful, `atoll()` returns the converted signed long long value, represented in the string. If unsuccessful, it returns an undefined value.

Related Information

- “`stdlib.h` — Define standard library functions” on page 62
- “`atoi()` — Convert character string to integer” on page 70
- “`atol()` — Convert character string to long” on page 70
- “`strtol()` — Convert Character String to Long” on page 112
- “`strtoll()` — Convert String to Signed Long Long” on page 114
- “`strtoul()` — Convert String to Unsigned Integer” on page 116
- “`strtoull()` — Convert String to Unsigned Long Long” on page 117

calloc() — Reserve and initialize storage

Format

```
#include <stdlib.h>

void *calloc(size_t num, size_t size);
```

General description

The `calloc()` function reserves storage space for an array of *num* elements, each of length *size* bytes. The `calloc()` function then gives all the bits of each element an initial value of 0.

The `calloc()` function returns a pointer to the reserved space. The storage space to which the returned value points is aligned for storage of any type of object.

Note: Use of this function requires that an environment has been set up through the `__cinit()` function. When the function is called, GPR 12 must contain the environment token created by the `__cinit()` call.

Returned value

If successful, `calloc()` returns the pointer to the area of memory reserved.

If there is not enough space to satisfy the request or if *num* or *size* is 0, `calloc()` returns NULL.

Related Information

- “stdlib.h — Define standard library functions” on page 62
- “free() — Free a block of storage” on page 76
- “malloc() — Reserve storage block” on page 80
- “__malloc31() — Allocate 31-bit storage” on page 81
- “realloc() — Change reserved storage block size” on page 86

__cinit() - Initialize a Metal C environment

Format

```
#include <metal.h>
__csysenv_t __cinit(struct __csysenv_s * csysenv);
```

General description

The `__cinit()` function establishes a Metal C environment based on the characteristics in the input `csysenv` structure. This environment is used when calling Metal C functions that require an environment, such as those related to storage management (`malloc()`, `free()`, and so on). Storage for the environment structures is obtained by using the attributes specified in the input `csysenv` structure.

Use of this function requires the `long long` data type.

The environment token created by `__cinit()` can be used from both AMODE 31 and AMODE 64 programs. Calls to `__malloc31()` always affect the below-the-bar heap. Calls made while in AMODE 31 to all other functions that obtain storage affect the below-the-bar heap; calls made while in AMODE 64 affect the above-the-bar heap.

Table 9. *csysenv* argument in `__cinit()`

Argument	Description
<i>csysenv</i>	A structure describing the characteristics of the environment to be created.

The details on the `csysenv` (`__csysenv_s`) structure is shown as follows:

```
struct __csysenv_s {
    int      __cseversion;      /* Control block version number      */
                                /* Must be set to __CSE_VERSION_1    */

    int      __csubpool;        /* for 31 bit storage                */
    __ptr31(void, __csetcbowner) /* owning TCB for resources           */
                                /* default: TCB mode - caller tcb,    */
                                /* SRB. XMEM - CMRO TCB              */

    int      __csereserved;      /* Reserved field                    */
    char      __csettkowner[16]; /* TCB token of owning TCB for       */
                                /* above the bar storage              */
                                /* default: caller tcbtoken          */
                                /* SRB mode: tcbtoken must be       */
                                /* specified                         */

    unsigned int
        __cseheap31initsize; /* Minimum size, in bytes, to obtain */
                                /* for the initial AMODE 31 heap storage. */
                                /* If 0, defaults to 32768 bytes      */

    unsigned int
        __cseheap31incrsize; /* Minimum size, in bytes, to obtain */
                                /* when expanding the AMODE 31 heap.    */
                                /* If 0, defaults to 32768 bytes      */
};
```

```

#ifdef __LL
    unsigned long long
        __cseheap64initsize; /* Minimum size, in MB, to obtain
                               for the initial AMODE 64 heap storage.
                               If 0, defaults to 1 MB */
    unsigned long long
        __cseheap64incrsz; /* Minimum size, in MB, to obtain
                              When expanding the AMODE 64 heap.
                              If 0, defaults to 1MB */
    unsigned long long
        __cseheap64usertoken; /* usertoken for use with ?iarv64
                               to obtain above the bar storage */
#else
    unsigned int
        __cseheap64initsize_hh;
    unsigned int
        __cseheap64initsize; /* Minimum size, in MB, to obtain
                               for the initial AMODE 64 heap storage.
                               If 0, defaults to 1 MB */
    unsigned int
        __cseheap64incrsz_hh;
    unsigned int
        __cseheap64incrsz; /* Minimum size, in MB, to obtain
                              When expanding the AMODE 64 heap.
                              If 0, defaults to 1MB */
    unsigned int
        __cseheap64usertoken_hh;
    unsigned int
        __cseheap64usertoken; /* usertoken for use with ?iarv64
                               to obtain above the bar storage */
#endif

    unsigned int
        __cseheap64fprot:1; /* AMODE 64 Storage Attributes */
    /* On, AMODE 64 heap storage is to be
       fetch protected
       Off, storage is not fetch
       protected */
    unsigned int
        __cseheap64cntlauth:1; /* On, AMODE 64 heap storage has
                                CONTROL=AUTH attribute
                                Off, storage is CONTROL=UNAUTH */
    int
        __csereserved1[7]; /* Reserved for future use */
};

```

Note: The entire `__csysenv_s` structure must be cleared to binary zeros before initializing specific fields within it.

Returned value

If successful, `__cinit()` returns an environment token that is used on subsequent calls to Metal C functions that require an environment. If unable to create an environment, `__cinit()` returns 0.

Example

```

#include <metal.h>
#include <stdlib.h>
#include <string.h>

#ifdef __LP64
    register void * myenvtkn __asm("r12");
#else
    register void * myetkr12 __asm("r12");
    __csysenv_t myenvtkn;
#endif

```

```

void mymtlfcn(void) {
    struct __csysenv_s mysysenv;

    void * mystg;
    void * my31stg;

    /******
    /* Initialize the csysenv structure.          */
    /******
    memset(&mysysenv, 0x00, sizeof(mysysenv));

    mysysenv.__cseversion = __CSE_VERSION_1;

    mysysenv.__csubpool = 129;

    /******
    /* Set heap initial and increment sizes.      */
    /******
    mysysenv.__cseheap31initsize = 131072;
    mysysenv.__cseheap31incrsz = 8192;
    mysysenv.__cseheap64initsize = 20;
    mysysenv.__cseheap64incrsz = 1;
#ifdef _LP64
    /******
    /* Create a Metal C environment.              */
    /******
    myenvtkn = (void * ) __cinit(&mysysenv);

#else
    /******
    /* Create a Metal C environment.              */
    /******
    myenvtkn = __cinit(&mysysenv);

    /******
    /* Save the high half of R12 and then set R12 to */
    /* the 8 byte environment token.                */
    /******
    __asm(" LG      12,%0\n"
        :
        : "m"(myenvtkn)
        : "r12"
    );
#endif

    /******
    /* Call functions that require an environment. */
    /******
    mystg = malloc(1048576);
    my31stg = __malloc31(100);

    /******
    /* Clean up the environment.                  */
    /******
    __cterm((__csysenv_t) myenvtkn);
}

```

In order to share the environment token to other source files there are 2 options:

- Compile all Metal C files that make up the program by using the **RESERVED_REG("r12")** compiler option. This reserves register 12 so that the environment token will remain untouched by the compiled code.
- Pass myenvtkn by using other methods, and for any source that needs to use the environment token declare the global register variable as in this example and assign the environment token to it.

Output: None.

__cterm() - Terminate a Metal C environment

Format

```
#include <metal.h>
void __cterm(__csysenv_t csysenvtkn);
```

General description

The __cterm() function terminates a Metal C environment, freeing all resources obtained on behalf of the environment.

Table 10. *csysenvtkn* argument in __cterm()

Argument	Description
<i>csysenvtkn</i>	The environment token representing the environment to be terminated.

Returned value

None.

Example

See the example provided for the __cinit() function.

Output: None.

div() — Calculate quotient and remainder

Format

```
#include <stdlib.h>
div_t div(int numerator, int denominator);
```

General description

The div() function calculates the quotient and remainder of the division of *numerator* by *denominator*.

Returned value

The div() function returns a structure of type `div_t`, containing both the quotient `int quot` and the remainder `int rem`. This structure is defined in `stdlib.h`. If the returned value cannot be represented, the behavior of `div()` is undefined. If *denominator* is 0, a divide by 0 exception is raised.

Related Information

- “`stdlib.h` — Define standard library functions” on page 62
- “`ldiv()` — Compute quotient and remainder of integral division” on page 79
- “`lldiv()` — Compute quotient and remainder of integral division for long long type” on page 80

free() — Free a block of storage

Format

```
#include <stdlib.h>
```

```
void free(void *ptr);
```

General description

The `free()` function frees a block of storage pointed to by *ptr*. The *ptr* variable points to a block previously reserved with a call to `calloc()`, `malloc()`, or `realloc()`. The number of bytes freed is the number of bytes specified when you reserved (or reallocated, in the case of `realloc()`), the block of storage. If *ptr* is `NULL`, `free()` simply returns without freeing anything. Since *ptr* is passed by value `free()` will not set *ptr* to `NULL` after freeing the memory to which it points.

Notes:

1. Use of this function requires that an environment has been set up by using the `__cinit()` function. When the function is called, GPR 12 must contain the environment token created by the `__cinit()` call.
2. Attempting to free a block of storage not allocated with `calloc()`, `malloc()`, or `realloc()`, or previously freed storage, can affect the subsequent reserving of storage and lead to an abend.

Returned value

`free()` returns no value.

Related Information

- “`stdlib.h` — Define standard library functions” on page 62
- “`calloc()` — Reserve and initialize storage” on page 71
- “`malloc()` — Reserve storage block” on page 80
- “`__malloc31()` — Allocate 31-bit storage” on page 81
- “`realloc()` — Change reserved storage block size” on page 86

isalnum() to isxdigit() — Test integer value

Format

```
#include <ctype.h>
```

```
int isalnum(int c);  
int isalpha(int c);  
int isblank(int c);  
int iscntrl(int c);  
int isdigit(int c);  
int isgraph(int c);  
int islower(int c);  
int isprint(int c);  
int ispunct(int c);  
int isspace(int c);  
int isupper(int c);  
int isxdigit(int c);
```

General description

The functions listed in the previous section, which are all declared in `ctype.h`, test a given integer value. The valid integer values for *c* are those representable as an *unsigned char* or EOF.

Here are descriptions of each function in this group.

isalnum()

Test for an upper- or lowercase letter, or a decimal digit, as defined by code page IBM-1047.

isalpha()

Test for an alphabetic character, as defined by code page IBM-1047.

isblank()

Test for a blank character, as defined by code page IBM-1047.

iscntrl()

Test for any control character, as defined by code page IBM-1047.

isdigit()

Test for a decimal digit, as defined by code page IBM-1047.

isgraph()

Test for a printable character excluding space, as defined by code page IBM-1047.

islower()

Test for a lowercase character, as defined by code page IBM-1047.

isprint()

Test for a printable character including space, as defined by code page IBM-1047.

ispunct()

Test for any non-alphanumeric printable character, excluding space, as defined by code page IBM-1047.

isspace()

Test for a white space character, as defined by code page IBM-1047.

isupper()

Test for an uppercase character, as defined by code page IBM-1047.

isxdigit()

Test for a hexadecimal digit, as defined by code page IBM-1047.

Returned value

If the integer satisfies the test condition, these functions return nonzero.

If the integer does not satisfy the test condition, these functions return 0.

Related Information

- “`ctype.h` — Declare character classification functions” on page 55
- “`tolower()`, `toupper()` — Convert Character Case” on page 119

isalpha() — Test for alphabetic character classification

The information for this function is included in “`isalnum()` to `isxdigit()` — Test integer value” on page 76.

isblank() — Test for blank character classification

The information for this function is included in “isalnum() to isxdigit() — Test integer value” on page 76.

iscntrl() — Test for control classification

The information for this function is included in “isalnum() to isxdigit() — Test integer value” on page 76.

isdigit() — Test for decimal-digit classification

The information for this function is included in “isalnum() to isxdigit() — Test integer value” on page 76.

isgraph() — Test for graphic classification

The information for this function is included in “isalnum() to isxdigit() — Test integer value” on page 76.

islower() — Test for lowercase

The information for this function is included in “isalnum() to isxdigit() — Test integer value” on page 76.

isprint() — Test for printable character classification

The information for this function is included in “isalnum() to isxdigit() — Test integer value” on page 76.

ispunct() — Test for punctuation classification

The information for this function is included in “isalnum() to isxdigit() — Test integer value” on page 76.

isspace() — Test for space character classification

The information for this function is included in “isalnum() to isxdigit() — Test integer value” on page 76.

isupper() — Test for uppercase letter classification

The information for this function is included in “isalnum() to isxdigit() — Test integer value” on page 76.

isxdigit() — Test for hexadecimal digit Classification

The information for this function is included in “isalnum() to isxdigit() — Test integer value” on page 76.

labs() — Calculate long absolute value

Format

```
#include <stdlib.h>
```

```
long int labs(long int n);
```

General description

The `labs()` function calculates the absolute value of its long integer argument *n*. The result is undefined when the argument is equal to `LONG_MIN`, the smallest available long integer (-2 147 483 648). The value `LONG_MIN` is defined in the `limits.h` header file.

Returned value

The `labs()` function returns the absolute value of the long integer argument *n*.

Related Information

- “limits.h — Declare symbolic names for resource limits” on page 58
- “stdlib.h — Define standard library functions” on page 62
- “abs() — Calculate integer absolute value” on page 69
- “llabs() — Calculate absolute value of long long integer”

ldiv() — Compute quotient and remainder of integral division

Format

```
#include <stdlib.h>
```

```
ldiv_t ldiv(long int numerator, long int denominator);
```

General description

The `ldiv()` function calculates the quotient and remainder of the division of *numerator* by *denominator*.

Returned value

The `ldiv()` function returns a structure of type `ldiv_t`, containing both the quotient `long int quot` and the remainder `long int rem`.

If the value cannot be represented, the returned value is undefined. If *denominator* is 0, a divide by 0 exception is raised.

Related Information

- “stdlib.h — Define standard library functions” on page 62
- “div() — Calculate quotient and remainder” on page 75
- “lldiv() — Compute quotient and remainder of integral division for long long type” on page 80

llabs() — Calculate absolute value of long long integer

Format

```
#include <stdlib.h>
```

```
long long llabs(long long int n);
```

Compile Requirement: Use of this function requires the long long data type.

General description

The `llabs()` function calculates the absolute value of its long long integer argument *n*. The result is undefined when the argument is equal to `LONGLONG_MIN`, the smallest available long long integer (-9 223 372 036 854 775 808). The value `LONGLONG_MIN` is defined in the `limits.h` header file.

Returned value

The `llabs()` function returns the absolute value of the long long integer argument *n*.

Related Information

- “`stdlib.h` — Define standard library functions” on page 62
- “`limits.h` — Declare symbolic names for resource limits” on page 58
- “`abs()` — Calculate integer absolute value” on page 69
- “`labs()` — Calculate long absolute value” on page 79

lldiv() — Compute quotient and remainder of integral division for long long type

Format

```
#include <stdlib.h>
```

```
lldiv_t lldiv (long long number, long long denom);
```

Compile Requirement: Use of this function requires the long long data type.

General description

The `lldiv()` function calculates the quotient and remainder of the division of numerator by denominator.

Returned value

The `lldiv()` function returns a structure of type `lldiv_t`, containing both the quotient long long `quot` and the remainder long long `rem`.

If the value cannot be represented, the returned value is undefined. If *denominator* is 0, a divide by 0 exception is raised.

Related Information

- “`stdlib.h` — Define standard library functions” on page 62
- “`div()` — Calculate quotient and remainder” on page 75
- “`ldiv()` — Compute quotient and remainder of integral division” on page 79

malloc() — Reserve storage block

Format

```
#include <stdlib.h>
```

```
void *malloc(size_t size);
```

General description

The `malloc()` function reserves a block of storage of *size* bytes. Unlike the `calloc()` function, the content of the storage allocated is indeterminate. The storage to which the returned value points is always aligned for storage of any type of object.

Note: Use of this function requires that an environment has been set up through the `__cinit()` function. When the function is called, GPR12 must contain the environment token created by the `__cinit()` call.

Returned value

If successful, `malloc()` returns a pointer to the reserved space.

If not enough storage is available, or if *size* was specified as 0, `malloc()` returns `NULL`.

Related Information

- “`stdlib.h` — Define standard library functions” on page 62
- “`calloc()` — Reserve and initialize storage” on page 71
- “`free()` — Free a block of storage” on page 76
- “`__malloc31()` — Allocate 31-bit storage”
- “`realloc()` — Change reserved storage block size” on page 86

`__malloc31()` — Allocate 31-bit storage

Format

```
#include <stdlib.h>

void *__malloc31(size_t size);
```

General description

The `__malloc31()` function reserves a block of storage of *size* bytes from 31-bit addressable storage. The content of the storage allocated is indeterminate. The storage space to which the returned value points is always suitably aligned for storage of any type of object.

Note: Use of this function requires that an environment has been set up by using the `__cinit()` function. When the function is called, GPR 12 must contain the environment token created by the `__cinit()` call.

Returned value

If successful, `__malloc31()` returns a pointer to the reserved space.

If not enough storage is available, or if *size* was specified as 0, `__malloc31()` returns `NULL`.

Related Information

- “`stdlib.h` — Define standard library functions” on page 62
- “`calloc()` — Reserve and initialize storage” on page 71
- “`free()` — Free a block of storage” on page 76
- “`malloc()` — Reserve storage block” on page 80
- “`realloc()` — Change reserved storage block size” on page 86

memccpy() — Copy bytes in memory

Format

```
#include <string.h>
```

```
void *memccpy(void *__restrict__s1, const void *__restrict__s2, int c, size_t n);
```

General description

The `memccpy()` function copies bytes from memory area `s2` into memory area `s1`, stopping after the first occurrence of byte `c` (converted to an unsigned char) is copied, or after `n` bytes are copied, whichever comes first.

Returned value

If successful, `memccpy()` returns a pointer to the byte after the copy of `c` in `s1`.

If `c` was not found in the first `n` bytes of `s2`, `memccpy()` returns a NULL pointer.

Related Information

- “string.h — Declare string manipulation functions” on page 63
- “memchr() — Search buffer”
- “memcmp() — Compare bytes” on page 83
- “memcpy() — Copy buffer” on page 83
- “memmove() — Move buffer” on page 84
- “memset() — Set buffer to value” on page 84
- “strcpy() — Copy String” on page 102

memchr() — Search buffer

Format

```
#include <string.h>
```

```
void *memchr(const void *buf, int c, size_t count);
```

General description

The `memchr()` built-in function searches the first `count` bytes pointed to by `buf` for the first occurrence of `c` converted to an unsigned character. The search continues until it finds `c` or examines `count` bytes.

Returned value

If successful, `memchr()` returns a pointer to the location of `c` in `buf`.

If `c` is not within the first `count` bytes of `buf`, `memchr()` returns NULL.

Related Information

- “string.h — Declare string manipulation functions” on page 63
- “memccpy() — Copy bytes in memory”
- “memcmp() — Compare bytes” on page 83
- “memcpy() — Copy buffer” on page 83
- “memmove() — Move buffer” on page 84
- “memset() — Set buffer to value” on page 84
- “strchr() — Search for Character” on page 100

memcmp() — Compare bytes

Format

```
#include <string.h>
```

```
int memcmp(const void *buf1, const void *buf2, size_t count);
```

General description

The `memcmp()` built-in function compares the first *count* bytes of *buf1* and *buf2*.

The relation is determined by the sign of the difference between the values of the leftmost first pair of bytes that differ. The values depend on EBCDIC encoding. This function is *not* locale sensitive.

Note: When **COMPACT** is specified, the compiler does not generate inline code for this function. In this case, you need to take either of the following actions:

- Define the `__METAL_SYSVEC` feature test macro to use the system library.
- Define the `__METAL_STATIC` feature test macro to use the static object library and bind in corresponding data set.

Returned value

Indicates the relationship between *buf1* and *buf2* as follows:

Value	Meaning
-------	---------

- | | |
|-----|--|
| < 0 | The contents of the buffer pointed to by <i>buf1</i> less than the contents of the buffer pointed to by <i>buf2</i> |
| = 0 | The contents of the buffer pointed to by <i>buf1</i> identical to the contents of the buffer pointed to by <i>buf2</i> |
| > 0 | The contents of the buffer pointed to by <i>buf1</i> greater than the contents of the buffer pointed to by <i>buf2</i> |

Related Information

- “System and static object libraries” on page 65
- “string.h — Declare string manipulation functions” on page 63
- “memcpy() — Copy bytes in memory” on page 82
- “memchr() — Search buffer” on page 82
- “memcpy() — Copy buffer”
- “memmove() — Move buffer” on page 84
- “memset() — Set buffer to value” on page 84
- “strcmp() — Compare Strings” on page 101

memcpy() — Copy buffer

Format

```
#include <string.h>
```

```
void *memcpy(void * __restrict__ dest, const void * __restrict__ src, size_t count);
```

General description

The `memcpy()` built-in function copies *count* bytes from the object pointed to by *src* to the object pointed to by *dest*. For `memcpy()`, the source characters may be overlaid if copying takes place between objects that overlap. Use the `memmove()`

function to allow copying between objects that overlap.

Returned value

The `memcpy()` function returns the value of *dest*.

Related Information

- “string.h — Declare string manipulation functions” on page 63
- “memcpy() — Copy bytes in memory” on page 82
- “memchr() — Search buffer” on page 82
- “memmove() — Move buffer”
- “memset() — Set buffer to value”
- “strcpy() — Copy String” on page 102

memmove() — Move buffer

Format

```
#include <string.h>
```

```
void *memmove(void *dest, const void *src, size_t count);
```

General description

The `memmove()` function copies *count* bytes from the object pointed to by *src* to the object pointed to by *dest*. The function allows copying between possibly overlapping objects as if the *count* bytes of the object pointed to by *src* must first be copied into a temporary array before being copied to the object pointed to by *dest*.

Returned value

The `memmove()` function returns the value of *dest*.

Related Information

- “string.h — Declare string manipulation functions” on page 63
- “memcpy() — Copy bytes in memory” on page 82
- “memchr() — Search buffer” on page 82
- “memcpy() — Copy buffer” on page 83
- “memset() — Set buffer to value”

memset() — Set buffer to value

Format

```
#include <string.h>
```

```
void *memset(void *dest, int c, size_t count);
```

General description

The `memset()` built-in function sets the first *count* bytes of *dest* to the value *c* converted to an unsigned int.

Returned value

`memset()` returns the value of *dest*.

Related Information

- “string.h — Declare string manipulation functions” on page 63
- “memcpy() — Copy bytes in memory” on page 82
- “memchr() — Search buffer” on page 82
- “memcpy() — Copy buffer” on page 83
- “memmove() — Move buffer” on page 84

qsort() — Sort array**Format**

```
#include <stdlib.h>
```

```
void qsort(void *base, size_t num, size_t width,
           int(*compare)(const void *element1, const void *element2));
```

General description

The `qsort()` function sorts an array of *num* elements, each of *width* bytes in size, where the first element of the array is pointed to by *base*.

The *compare* pointer points to a function, which you supply, that compares two array elements and returns an integer value specifying their relationship. The `qsort()` function calls the comparison function one or more times during the sort, passing pointers to two array elements on each call. The comparison function must compare the elements and return one of the following values:

Value Meaning

<code>< 0</code>	<i>element1</i> less than <i>element2</i>
<code>0</code>	<i>element1</i> equal to <i>element2</i>
<code>> 0</code>	<i>element1</i> greater than <i>element2</i>

The sorted array elements are stored in increasing order, as returned by the comparison function. You can sort in reverse order by reversing the “greater than” and “less than” logic in the comparison function. If two elements are equal, their order in the sorted array is unspecified. The `qsort()` function overwrites the contents of the array with the sorted elements.

Returned value

The `qsort()` function returns no values.

Related information

- “stdlib.h — Define standard library functions” on page 62

rand() — Generate random number**Format**

```
#include <stdlib.h>
```

```
int rand(void);
```

General Description

The `rand()` function generates a pseudo-random integer in the range 0 to `RAND_MAX`. Use the `srand()` function before calling `rand()` to set a seed for the random number generator. If you do not make a call to `srand()`, the default seed is 1.

Note: Use of this function requires that an environment has been set up by using the `__cinit()` function. When the function is called, GPR 12 must contain the environment token created by the `__cinit()` call.

Returned Value

The `rand()` function returns the calculated value.

Related Information

- “`stdlib.h` — Define standard library functions” on page 62
- “`rand_r()` — Pseudo-random number generator”
- “`srand()` — Set Seed for `rand()` Function” on page 94

rand_r() — Pseudo-random number generator

Format

```
#include <stdlib.h>

int rand_r(unsigned int *seed);
```

General Description

The `rand_r()` function generates a sequence of pseudo-random integers in the range 0 to `RAND_MAX`. (The value of the `RAND_MAX` macro will be at least 32767.)

If `rand_r()` is called with the same initial value for the object pointed to by *seed* and that object is not modified between successive returns and calls to `rand_r()`, the same sequence shall be generated.

Returned Value

The `rand_r()` function returns a pseudo-random integer.

Related Information

- “`stdlib.h` — Define standard library functions” on page 62
- “`rand()` — Generate random number” on page 85
- “`srand()` — Set Seed for `rand()` Function” on page 94

realloc() — Change reserved storage block size

Format

```
#include <stdlib.h>

void *realloc(void *ptr, size_t size);
```

General Description

The `realloc()` function changes the size of a previously reserved storage block. The *ptr* argument points to the beginning of the block. The *size* argument gives the new size of the block in bytes. The contents of the block are unchanged up to the shorter of the new and old sizes.

If the *ptr* is `NULL`, `realloc()` reserves a block of storage of *size* bytes. It does not give all bits of each element an initial value of 0.

If *size* is 0 and *ptr* is not `NULL`, the storage pointed to by *ptr* is freed and `NULL` is returned.

If you use `realloc()` with a pointer that does not point to a *ptr* created previously by `malloc()`, `calloc()`, or `realloc()`, or if you pass *ptr* to storage already freed, you get undefined behavior—typically an exception.

If you ask for more storage, the contents of the extension are undefined and are not guaranteed to be 0.

The storage to which the returned value points is aligned for storage of any type of object.

Note: Use of `realloc()` requires that an environment has been set up by using the `__cinit()` function. When the function is called, GPR 12 must contain the environment token created by the `__cinit()` call.

Returned Value

If successful, `realloc()` returns a pointer to the reallocated storage block. The storage location of the block might be moved. Thus, the returned value is not necessarily the same as the *ptr* argument to `realloc()`.

The returned value is `NULL` if *size* is 0. If there is not enough storage to expand the block to the given size, the original block is unchanged and a `NULL` pointer is returned.

Related Information

- “`stdlib.h` — Define standard library functions” on page 62
- “`calloc()` — Reserve and initialize storage” on page 71
- “`free()` — Free a block of storage” on page 76
- “`malloc()` — Reserve storage block” on page 80
- “`__malloc31()` — Allocate 31-bit storage” on page 81

snprintf() — Format and write data

Format

```
#include <stdio.h>
```

```
int  snprintf(char *__restrict__ s, size_t n, const char *__restrict__ format, ...);
```

General Description

The `snprintf()` function formats and writes output to an array (specified by argument *s*). If *n* is zero, nothing is written, and *s* may be a null pointer. Otherwise, output characters beyond the *n*-1st are discarded rather than being

written to the array, and a null character is written at the end of the characters actually written into the array. If copying takes place between objects that overlap, the behavior is undefined.

Note: Use of `snprintf()` requires that an environment has been set up by using the `__cinit()` function. When the function is called, GPR 12 must contain the environment token created by the `__cinit()` call.

Returned Value

The `snprintf()` function returns the number of characters that would have been written had *n* been sufficiently large, not counting the terminating null character, or a negative value if an encoding error occurred. Thus, the null-terminated output has been completely written if and only if the returned value is nonnegative and less than *n*.

Related Information

- “`stdio.h` — Define I/O related functions” on page 60
- “`sprintf()` — Format and Write Data”
- “`scanf()` — Read and Format Data” on page 94

sprintf() — Format and Write Data

Format

```
#include <stdio.h>
```

```
int sprintf(char *__restrict__buffer, const char *__restrict__format-string, ...);
```

General Description


The `sprintf()` function formats and stores a series of characters and values in the array pointed to by *buffer*. Any *argument-list* is converted and put out according to the corresponding format specification in the *format-string*. If the strings pointed to by *buffer* and *format-string* overlap, behavior is undefined.

The *format-string* consists of ordinary characters, escape sequences, and conversion specifications. The ordinary characters are copied in order of their appearance. Conversion specifications, beginning with a percent sign (%) or the sequence (%n\$) where *n* is a decimal integer in the range [1,NL_ARGMAX], determine the output format for any *argument-list* following the *format-string*. The *format-string* can contain multibyte characters beginning and ending in the initial shift state. When the *format-string* includes the use of the optional prefix *ll* to indicate the size expected is a long long datatype then the corresponding value in the argument list should be a long long datatype if correct output is expected.

- If the %n\$ conversion specification is found, the value of the *nth* *argument* after the *format-string* is converted and output according to the conversion specification. Numbered arguments in the argument list can be referenced from *format-string* as many times as required.
- The *format-string* can contain either form of the conversion specification, that is, % or %n\$ but the two forms cannot be mixed within a single *format-string* except that %% can be mixed with the %n\$ form. When numbered conversion specifications are used, specifying the 'nth' argument requires that the first to (n-1)th arguments are specified in the *format-string*.

The *format-string* is read from left to right. When the first format specification is found, the value of the first *argument* after the *format-string* is converted and output

according to the format specification. The second format specification causes the second *argument* after the *format-string* to be converted and output, and so on through the end of the *format-string*. If there are more arguments than there are format specifications, the extra arguments are evaluated and ignored. The results are undefined if there are not enough arguments for all the format specifications. The format specification is illustrated below.



Each field of the format specification is a single character or number signifying a particular format option. The *type* character, which appears after the last optional format field, determines whether the associated argument is interpreted as a character, a string, a number, or pointer. The simplest format specification contains only the percent sign and a *type* character (for example, %s).

The percent sign: If a percent sign (%) is followed by a character that has no meaning as a format field, the character is simply copied to the *buffer*. For example, to print a percent sign character, use %%.

The flag characters: The *flag* characters in Table 11 are used for the justification of output and printing of thousands of grouping characters, signs, blanks, decimal-points, octal prefixes, and hexadecimal prefixes. Note that more than one *flag* can appear in a format specification. This is an optional field.

Table 11. Flag Characters for sprintf() Family

Flag	Meaning	Default
'	The integer portion of the result of a decimal conversion (%i,%d,%u, %f,%g or %G) will be formatted with the thousands' grouping characters.	No grouping.
-	Left-justify the result within the field width.	Right-justify.
+	Prefix the output value with a sign (+ or -) if the output value is of a signed type.	Sign appears only for negative signed values (-).
blank(' ')	Prefix the output value with a blank if the output value is signed and positive. The + flag overrides the <i>blank</i> flag if both appear, and a positive signed value will be output with a sign.	No blank.

Table 11. Flag Characters for sprintf() Family (continued)

Flag	Meaning	Default
#	When used with the o, x, or X formats, the # flag prefixes any nonzero output value with 0, 0x, or 0X, respectively. For o conversion, it increases the precision, if necessary, to force the first digit of the result to be a zero. If the value and precision are both 0, a single 0 is printed. For e, E, f, F, g, and G conversion specifiers, the result always contains a decimal-point, even if no digits follow the decimal-point. Without this flag, a decimal-point appears in the result of these conversions only if a digit follows it. For g and G conversion specifiers, do not remove trailing zeros from the result as they normally are. For other conversion specifiers, the behavior is undefined.	No prefix.
0	When used with the d, i, o, u, x, X, e, E, f, F, g, or G conversion specifiers, leading zeros are used to pad to the field width. If the 0 and - flags both appear, the 0 flag is ignored. For d, i, o, u, x, and X conversion specifiers, if a precision is specified, the 0 flag is ignored. If the 0 and ' flags both appear, the grouping characters are inserted before zero padding. For other conversions, the behavior is undefined.	Space padding.

The code point for the # character varies between the EBCDIC encoded character sets. The Metal C runtime library expects the # character to use the code point for encoded character set IBM-1047.

The # flag should not be used with c, d, i, u, s, or p types.

The Width of the Output: Definition of the *width* specification is as follows.

Width is a nonnegative decimal integer controlling the minimum number of characters printed. If the number of characters in the output value is less than the specified *width*, blanks are added on the left or the right (depending on whether the - flag is specified) until the minimum width is reached.

Width never causes a value to be truncated; if the number of characters in the output value is greater than the specified *width*, or *width* is not given, all characters of the value are output (subject to the *precision* specification).

The *width* specification can be an asterisk (*); if it is, an argument from the argument list supplies the value. The *width* argument must precede the value being formatted in the argument list. This is an optional field.

If *format-string* contains the %n\$ form of conversion specification, *width* can be indicated by the sequence *m\$, where m is a decimal integer in the range [1,NL_ARGMAX] giving the position of an integer argument in the argument list containing the field width.

The Precision of the Output: Definition of the *precision* specification is as follows.

The *precision* specification is a nonnegative decimal integer preceded by a period. It specifies the number of characters to be output, or the number of decimal places. Unlike the *width* specification, the *precision* can cause truncation of the output value.

The *precision* specification can be an asterisk (*); if it is, an argument from the argument list supplies the value. The *precision* argument must precede the value being formatted in the argument list. The *precision* field is optional.

If *format-string* contains the %n\$ form of conversion specification, *precision* can be indicated by the sequence *m\$, where m is a decimal integer in the range [1,NL_ARGMAX] giving the position of an integer argument in the argument list containing the field precision.

The interpretation of the *precision* value and the default when the *precision* is omitted depend upon the *type*, as shown in Table 12.

Table 12. Precision Argument in sprintf()

Type	Meaning	Default
d i o u x X	<i>Precision</i> specifies the minimum number of digits to be output. If the number of digits in the argument is less than <i>precision</i> , the output value is padded on the left with zeros. The value is not truncated when the number of digits exceeds <i>precision</i> .	Default <i>precision</i> is 1. If <i>precision</i> is 0, or if the period (.) appears without a number following it, the <i>precision</i> is set to 0. When <i>precision</i> is 0, conversion of the value zero results in no characters.
c	No effect.	The character is output.
s	<i>Precision</i> specifies the maximum number of characters to be output. Characters in excess of <i>precision</i> are not output.	Characters are output until a NULL character is encountered.
e E f F	<i>Precision</i> specifies the number of digits to be output after the decimal-point. The last digit output is rounded.	Default <i>precision</i> is 6. If <i>precision</i> is 0 or the period appears without a number following it, no decimal-point is output.
g G	<i>Precision</i> specifies the maximum number of significant digits output.	All significant digits are output.

Optional prefix: Used to indicate the size of the argument expected.

- h** A prefix with the integer types d, i, o, u, x, X means the integer is 16 bits long.
- hh** Specifies that a following d, i, o, u, x, or X conversion specifier applies to a signed char or unsigned char argument (the argument will have been promoted according to the integer promotions, but its value shall be converted to signed char or unsigned char before printing); or that a following n conversion specifier applies to a pointer to a signed char argument.

- j** Specifies that a following d, i, o, u, x, or X conversion specifier applies to an intmax_t or uintmax_t argument; or that a following n conversion specifier applies to a pointer to an intmax_t argument.
- l** A prefix with d, i, o, u, x, X, and n types that specifies that the argument is a long int or unsigned long int.
- ll** A prefix with the integer types d, i, o, u, x, X means the integer is 64 bits long.
- L** A prefix with e, E, f, g, or G types that specifies that the argument is long double.
- t** Specifies that a following d, i, o, u, x, or X conversion specifier applies to a ptrdiff_t or the corresponding unsigned type argument; or that a following n conversion specifier applies to a pointer to a ptrdiff_t argument.
- z** Specifies that a following d, i, o, u, x, or X conversion specifier applies to a size_t or the corresponding signed integer type argument; or that a following n conversion specifier applies to a pointer to a signed integer type corresponding to a size_t argument.

Table 13 below shows the meaning of the type characters used in the precision argument.

Table 13. Type Characters and their Meanings

Type	Argument	Output Format
d, i	Integer	Signed decimal integer.
u	Integer	Unsigned decimal integer.
o	Integer	Unsigned octal integer.
x	Integer	Unsigned hexadecimal integer, using abcdef.
X	Integer	Unsigned hexadecimal integer, using ABCDEF.
c	Character	Single character.
s	String	Characters output up to the first NULL character (\0) or until <i>precision</i> is reached.
n	Pointer to integer	Number of characters successfully output so far to the <i>stream</i> or buffer; this value is stored in the integer whose address is given as the argument.
p	Pointer	Pointer to void converted to a sequence of printable characters. See the individual system reference guides for the specific format.

Table 13. Type Characters and their Meanings (continued)

Type	Argument	Output Format
f, F	Double	Signed value having the form [-]dddd.dddd, where dddd is one or more decimal digits. The number of digits before the decimal-point depends on the magnitude of the number. The number of digits after the decimal-point is equal to the requested precision. If the precision is explicitly zero and no # is present, no decimal-point appears. If a decimal-point appears, at least one digit appears before it. Convert a double argument representing an infinity in [+/-]inf: a plus or minus sign with the character sequence inf, followed by a white space character (space, tab, or newline), a NULL character (\0), or EOF. Convert a double argument representing a NaN in one of the styles: • [+/-]nan(n) for a signaling nan. • [+/-]nanq(n) for a quiet nan, where n is an integer and $1 \leq n \leq \text{INT_MAX}-1$. The value of n is determined by the fraction bits of the NaN argument value. For a signaling NaN value, NaN fraction bits are reversed (left to right) to produce bits (right to left) of an even integer value, 2^n. For a quiet NaN value, NaN fraction bits are reversed (left to right) to produce bits (right to left) of an odd integer value, 2^n-1. The F conversion specifier produces INFe, NANS, or NANQ instead of infQ, nans or, nanq respectively.
e, E	Double	Signed value having the form [-]d.dddde[sign]ddd: • d is a single-decimal digit. • dddd is one or more decimal digits. • ddd is 2 or more decimal digits. • sign is + or -. If the <i>precision</i> is zero and no # flag is present, no decimal-point appears. The conversion specifier produces a number with E instead of e to introduce the exponent. A double argument representing an infinity or NaN is converted in the style of an f or F conversion specifier.
g, G	Double	Signed value output in f or e format (or in the F or E format in the case of a G conversion specifier). The e or E format is used only when the exponent of the value is less than -4 or greater than or equal to the <i>precision</i>. Trailing zeros are truncated, and the decimal-point appears only if one or more digits follow it or a # flag is present. A double argument representing an infinity or NaN is converted in the style of an f or F conversion specifier.

Returned Value

If successful, sprintf() returns the number of characters output. The ending NULL character is not counted.

If unsuccessful, `sprintf()` returns a negative value.

Related Information

- “`stdio.h` — Define I/O related functions” on page 60
- “`snprintf()` — Format and write data” on page 87
- “`sscanf()` — Read and Format Data”

srand() — Set Seed for `rand()` Function

Format

```
#include <stdlib.h>

void srand(unsigned int seed);
```

General Description

The `srand()` function uses its argument *seed* as a seed for a new sequence of pseudo-random numbers to be returned by subsequent calls to `rand()`. If `srand()` is not called, the `rand()` seed is set as if `srand(1)` was called at program start. Any other value for *seed* sets the generator to a different starting point. The `rand()` function generates pseudo-random numbers.

Some people find it convenient to use the return value of the `time()` function as the argument to `srand()`, as a way to ensure random sequences of random numbers.

Note: Use of `srand()` requires that an environment has been set up by using the `__cinit()` function. When the function is called, GPR 12 must contain the environment token created by the `__cinit()` call.

Returned Value

`srand()` returns no values.

Related Information

- “`stdlib.h` — Define standard library functions” on page 62
- “`rand()` — Generate random number” on page 85
- “`rand_r()` — Pseudo-random number generator” on page 86

sscanf() — Read and Format Data

Format

```
#include <stdio.h>

int sscanf(const char *__restrict__buffer, const char *__restrict__format-string, ...);
```

General Description

The `sscanf()` function reads data from *buffer* into the locations given by *argument-list*. If the strings pointed to by *buffer* and *format-string* overlap, behavior is undefined.

Each entry in the argument list must be a pointer to a variable of a type that matches the corresponding conversion specification in *format-string*. If the types do not match, the results are undefined.

The *format-string* controls the interpretation of the argument list. The *format-string* can contain multibyte characters beginning and ending in the initial shift state.

The format string pointed to by *format-string* can contain one or more of the following:

- White space characters, as specified by `isspace()`, such as blanks and newline characters. A white space character causes `sscanf()` to read, but not to store, all consecutive white space characters in the input up to the next character that is not white space. One white space character in *format-string* matches any combination of white space characters in the input.
- Characters that are not white space, except for the percent sign character (%). A non-white space character causes `sscanf()` to read, but not to store, a matching non-white space character. If the next character in the input stream does not match, the function ends.
- Conversion specifications which are introduced by the percent sign (%) or the sequence (%n\$) where n is a decimal integer in the range [1,NL_ARGMAX]. A conversion specification causes `sscanf()` to read and convert characters in the input into values of a conversion specifier. The value is assigned to an argument in the argument list.

`sscanf()` reads *format-string* from left to right. Characters outside of conversion specifications are expected to match the sequence of characters in the input stream; the matched characters in the input stream are scanned but not stored. If a character in the input stream conflicts with *format-string*, the function ends, terminating with a “matching” failure. The conflicting character is left in the input stream as if it had not been read.

When the first conversion specification is found, the value of the first *input field* is converted according to the conversion specification and stored in the location specified by the first entry in the argument list. The second conversion specification converts the second input field and stores it in the second entry in the argument list, and so on through the end of *format-string*.

When the %n\$ conversion specification is found, the value of the *input field* is converted according to the conversion specification and stored in the location specified by the nth argument in the argument list. Numbered arguments in the argument list can only be referenced once from *format-string*.


The *format-string* can contain either form of the conversion specification, that is, % or %n\$ but the two forms cannot be mixed within a single *format-string* except that %% or %* can be mixed with the %n\$ form.

An *input field* is defined as:

- All characters until a white space character (space, tab, or newline) is encountered
- All characters until a character is encountered that cannot be converted according to the conversion specification
- All characters until the field *width* is reached.

If there are too many arguments for the conversion specifications, the extra arguments are evaluated but otherwise ignored. The results are undefined if there are not enough arguments for the conversion specifications.

Syntax of Conversion Specification for sscanf()



Each field of the conversion specification is a single character or a number signifying a particular format option. The *conversion specifier*, which appears after the last optional format field, determines whether the input field is interpreted as a character, a string, or a number. The simplest conversion specification contains only the percent sign and a *conversion specifier* (for example, %s).

Each field of the format specification is discussed in detail below.

Other than conversion specifiers, avoid using the percent sign (%), except to specify the percent sign: %%. Currently, the percent sign is treated as the start of a conversion specifier. Any unrecognized specifier is treated as an ordinary sequence of characters. If, in the future, Enterprise Metal C for z/OS permits a new conversion specifier, it could match a section of your format string, be interpreted incorrectly, and result in undefined behavior. See Table 14 on page 97 for a list of conversion specifiers.

An asterisk (*) following the percent sign suppresses assignment of the next input field, which is interpreted as a field of the specified *conversion specifier*. The field is scanned but not stored.

width is a positive decimal integer controlling the maximum number of characters to be read. No more than *width* characters are converted and stored at the corresponding *argument*.

Fewer than *width* characters are read if a white space character (space, tab, or newline), or a character that cannot be converted according to the given format occurs before *width* is reached.

The optional prefix l shows that you use the long version of the following *conversion specifier*, while the prefix h indicates that the short version is to be used. The corresponding *argument* should point to a long or double object (for the l character), a long double object (for the L character), or a short object (with the h character). The l and h modifiers can be used with the d, i, o, x, and u *conversion specifiers*. The l and h modifiers are ignored if specified for any other *conversion specifier*.

Optional prefix: Used to indicate the size of the argument expected.

- h** Specifies that a following d, i, o, u, x, X, or n conversion specifier applies to an argument with type pointer to short or unsigned short.
- hh** Specifies that a following d, i, o, u, x, X or n conversion specifier applies to an argument with type pointer to signed char or unsigned char.

- j** Specifies that a following d, i, o, u, x, X or n conversion specifier applies to an argument with type pointer to intmax_t or uintmax_t.
- l** Specifies that a following e, E, f, F, g, or G conversion specifier applies to an argument with type pointer to double.
- ll** Specifies that a following d, i, o, u, x, X or n conversion specifier applies to an argument with type pointer to long long or unsigned long long.
- L** Specifies that a following e, E, f, g, or G conversion specifier applies to an argument with type pointer to long double.
- t** Specifies that a following d, i, o, u, x, X or n conversion specifier applies to an argument with type pointer to ptrdiff_t or the corresponding unsigned type.
- z** Specifies that a following d, i, o, u, x, X or n conversion specifier applies to an argument with type pointer to size_t or the corresponding signed integer type.

The *type* characters and their meanings are in Table 14.

Table 14. Conversion Specifiers in *sscanf()*

Conversion Specifier	Type of Input Expected	Type of Argument
d	Decimal integer	Pointer to int
o	Octal integer	Pointer to unsigned int
x X	Hexadecimal integer	Pointer to unsigned int
i	Decimal, hexadecimal, or octal integer	Pointer to int
u	Unsigned decimal integer	Pointer to unsigned int
c	Sequence of one or more characters as specified by field width; white space characters that are ordinarily skipped are read when %c is specified. No terminating null is added.	Pointer to char large enough for input field.
s	Like c, a sequence of bytes of type char (signed or unsigned), except that white space characters are not allowed, and a terminating null is always added.	Pointer to character array large enough for input field, plus a terminating NULL character (\0) that is automatically appended.
n	No input read from <i>stream</i> or buffer.	Pointer to int, into which is stored the number of characters successfully read from the <i>stream</i> or buffer up to that point in the call to either fscanf() or to scanf().
p	Pointer to void converted to series of characters. For the specific format of the input, see the individual system reference guides.	Pointer to void.

Table 14. Conversion Specifiers in `sscanf()` (continued)

Conversion Specifier	Type of Input Expected	Type of Argument
[	A non-empty sequence of bytes to be matched against a set of expected bytes (the <i>scanset</i>), which form the conversion specification. White space characters that are ordinarily skipped are read when %[is specified. Consider the following situations: [^<i>bytes</i>]. In this case, the scanset contains all bytes that do not appear between the circumflex and the right square bracket. [<i>abc</i>] or [^<i>abc</i>]. In both these cases the right square bracket is included in the scanset (in the first case: <i>abc</i> and in the second case, <i>not abc</i>) [<i>a-z</i>] In EBCDIC The <i>-</i> is in the scanset, the characters <i>b</i> through <i>y</i> are <i>not</i> in the scanset; in ASCII The <i>-</i> is <i>not</i> in the scanset, the characters <i>b</i> through <i>y</i> are. The code point for the square brackets ([and]) and the caret (^) vary among the EBCDIC encoded character sets. The default C locale expects these characters to use the code points for encoded character set Latin-1 / Open Systems 1047. Conversion proceeds one byte at a time: there is no conversion to wide characters.	Pointer to the initial byte of an array of char, signed char, or unsigned char large enough to accept the sequence and a terminating byte, which will be added automatically.
e	Floating-point value consisting of an optional sign (+ or -), a series of one or more decimal digits possibly containing a decimal-point, and an optional exponent (e or E) followed by a possibly signed integer value.	Pointer to float
E		
f		
F		
g		
G		

The format string passed to `sscanf()` must be encoded as IBM-1047.

To read strings not delimited by space characters, substitute a set of characters in square brackets ([]) for the *s* (string) conversion specifier. The corresponding input field is read up to the first character that does not appear in the bracketed character set. If the first character in the set is a logical not (~), the effect is reversed: the input field is read up to the first character that does appear in the rest of the character set.

To store a string without storing an ending NULL character (\0), use the specification *%ac*, where *a* is a decimal integer. In this instance, the *c* conversion specifier means that the argument is a pointer to a character array. The next *a* characters are read from the input stream into the specified location, and no NULL character is added.

The input for a *%x* conversion specifier is interpreted as a hexadecimal number.

The `sscanf()` function scans each input field character by character. It might stop reading a particular input field either before it reaches a space character, when the specified *width* is reached, or when the next character cannot be converted as

specified. When a conflict occurs between the specification and the input character, the next input field begins at the first unread character. The conflicting character, if there is one, is considered unread and is the first character of the next input field or the first character in subsequent read operations on the input stream.

The sscanf family functions match e, E, f, F, g or, G conversion specifiers to floating-point number substrings in the input stream. The sscanf family functions convert each input substring matched by an e, E, f, F, g, or G conversion specifier to a float, double or long double value depending on the size modifier before the e, E, f, F, g, or G conversion specifier.

Many z/OS Metal C formatted input functions, including the sscanf family of functions, use the IEEE binary floating-point format and recognize special infinity and NaN floating-point number input sequences.

- The special sequence for infinity input is `[+/-]inf` or `[+/-]INF`, where `+` or `-` is optional.
- The special sequence of NaN input is either `[+/-]nan(n)` for a signaling nan or `[+/-]nanq(n)` for a quiet nan, where `n` is an integer and $1 \leq n \leq \text{INT_MAX}-1$. If `(n)` is omitted, `n` is assumed to be 1. The value of `n` determines what IEEE binary floating-point NaN fraction bits are produced by the formatted input functions. For a signaling NaN, these functions produce NaN fraction bits (left to right) by reversing the bits (right to left) of the even integer value 2^n . For a quiet NaN, they produce NaN fraction bits (left to right) by reversing the bits (right to left) of the odd integer value 2^n-1 .

Returned Value

The `sscanf()` function returns the number of input items that were successfully matched and assigned. The returned value does not include conversions that were performed but not assigned (for example, suppressed assignments). The functions return EOF if there is an input failure before any conversion, or if EOF is reached before any conversion. Thus a returned value of 0 means that no fields were assigned: there was a matching failure before any conversion.

Related Information

- “stdio.h — Define I/O related functions” on page 60
- “snprintf() — Format and write data” on page 87
- “sprintf() — Format and Write Data” on page 88
- “vsnprintf() — Format and print data to fixed length buffer” on page 120
- “vsscanf() — Format Input of a STDARG Argument List” on page 122

strcat() — Concatenate Strings

Format

```
#include <string.h>
```

```
char *strcat(char * __restrict_string1, const char * __restrict_string2);
```

General Description

The `strcat()` built-in function concatenates *string2* with *string1* and ends the resulting string with the NULL character. In other words, `strcat()` appends a copy of the string pointed to by *string2*—including the terminating NULL byte—to the end of a string pointed to by *string1*, with its last byte (that is, the terminating NULL byte of *string1*) overwritten by the first byte of the appended string.

Do not use a literal string for a *string1* value, although *string2* might be a literal string.

If the storage of *string1* overlaps the storage of *string2*, the behavior is undefined.

Note: When **COMPACT** is specified, the compiler does not generate inline code for this function. In this case, you need to take either of the following actions:

- Define the `__METAL_SYSVEC` feature test macro to use the system library.
- Define the `__METAL_STATIC` feature test macro to use the static object library and bind in corresponding data set.

Returned Value

The `strcat()` built-in function returns the value of *string1*, the concatenated string.

Related Information

- “System and static object libraries” on page 65
- “string.h — Declare string manipulation functions” on page 63
- “`strchr()` — Search for Character”
- “`strcmp()` — Compare Strings” on page 101
- “`strcpy()` — Copy String” on page 102
- “`strcspn()` — Compare Strings” on page 102
- “`strncat()` — Concatenate Strings” on page 104

strchr() — Search for Character

Format

```
#include <string.h>
```

```
char *strchr(const char *string, int c);
```

General Description

The `strchr()` built-in function finds the first occurrence of *c* converted to `char`, in the string **string*. The character *c* can be the NULL character (`\0`); the ending NULL character of *string* is included in the search.

The `strchr()` function operates on NULL-terminated strings. The string argument to the function *must* contain a NULL character (`\0`) marking the end of the string.

Note: When **COMPACT** is specified, the compiler does not generate inline code for this function. In this case, you need to take either of the following actions:

- Define the `__METAL_SYSVEC` feature test macro to use the system library.
- Define the `__METAL_STATIC` feature test macro to use the static object library and bind in corresponding data set.

Returned Value

If successful, `strchr()` returns a pointer to the first occurrence of *c* (converted to a character) in *string*.

If the character is not found, `strchr()` returns a NULL pointer.

Related Information

- “System and static object libraries” on page 65
- “string.h — Declare string manipulation functions” on page 63
- “memchr() — Search buffer” on page 82
- “strcat() — Concatenate Strings” on page 99
- “strcmp() — Compare Strings”
- “strcpy() — Copy String” on page 102
- “strcspn() — Compare Strings” on page 102
- “strncmp() — Compare Strings” on page 104
- “strpbrk() — Find Characters in String” on page 106
- “strrchr() — Find Last Occurrence of Character in String” on page 106
- “strspn() — Search String” on page 107

strcmp() — Compare Strings

Format

```
#include <string.h>
```

```
int strcmp(const char *string1, const char *string2);
```

General Description

The strcmp() built-in function compares the string pointed to by *string1* to the string pointed to by *string2*. The string arguments to the function must contain a NULL character (`\0`) marking the end of the string.

The relation between the strings is determined by subtracting: *string1*[*i*] - *string2*[*i*], as *i* increases from 0 to *strlen* of the smaller string. The sign of a nonzero return value is determined by the sign of the difference between the values of the first pair of bytes (both interpreted as type unsigned char) that differ in the strings being compared. This function is *not* locale-sensitive.

Returned Value

strcmp() returns a value indicating the relationship between the strings, as listed below.

Value	Meaning
-------	---------

- | | |
|-----|---|
| < 0 | String pointed to by <i>string1</i> less than string pointed to by <i>string2</i> |
| = 0 | String pointed to by <i>string1</i> equivalent to string pointed to by <i>string2</i> |
| > 0 | String pointed to by <i>string1</i> greater than string pointed to by <i>string2</i> |

Related Information

- “string.h — Declare string manipulation functions” on page 63
- “memcmp() — Compare bytes” on page 83
- “strcspn() — Compare Strings” on page 102
- “strncmp() — Compare Strings” on page 104
- “strpbrk() — Find Characters in String” on page 106
- “strrchr() — Find Last Occurrence of Character in String” on page 106
- “strspn() — Search String” on page 107

strcpy() — Copy String

Format

```
#include <string.h>
```

```
char *strcpy(char * __restrict__string1, const char * __restrict__string2);
```

General Description

The `strcpy()` built-in function copies *string2*, including the ending NULL character, to the location specified by *string1*. The *string2* argument to `strcpy()` must contain a NULL character (`\0`) marking the end of the string. You cannot use a literal string for a *string1* value, although *string2* may be a literal string. If the two objects overlap, the behavior is undefined.

Returned Value

The `strcpy()` function returns the value of *string1*.

Related Information

- “string.h — Declare string manipulation functions” on page 63
- “memcpy() — Copy buffer” on page 83
- “strcat() — Concatenate Strings” on page 99
- “strchr() — Search for Character” on page 100
- “strcmp() — Compare Strings” on page 101
- “strcspn() — Compare Strings”
- “strncpy() — Copy String” on page 105
- “strpbrk() — Find Characters in String” on page 106
- “strrchr() — Find Last Occurrence of Character in String” on page 106
- “strspn() — Search String” on page 107

strcspn() — Compare Strings

Format

```
#include <string.h>
```

```
size_t strcspn(const char *string1, const char *string2);
```

General Description

The `strcspn()` function computes the length of the initial portion of the string pointed to by *string1* that contains no characters from the string pointed to by *string2*.

Returned Value

The `strcspn()` function returns the calculated length of the initial portion found.

Related Information

- “string.h — Declare string manipulation functions” on page 63
- “strcat() — Concatenate Strings” on page 99
- “strchr() — Search for Character” on page 100
- “strcmp() — Compare Strings” on page 101
- “strcpy() — Copy String”
- “strncmp() — Compare Strings” on page 104
- “strpbrk() — Find Characters in String” on page 106

- “strchr() — Find Last Occurrence of Character in String” on page 106
- “strspn() — Search String” on page 107

strdup() — Duplicate a String

Format

```
#include <string.h>

char *strdup(const char *string);
```

General Description

The strdup() function creates a duplicate of the string pointed to by *string*.

Note: Use of this function requires that an environment has been set up by using the __cinit() function. When the function is called, GPR 12 must contain the environment token created by the __cinit() call.

Returned Value

If successful, strdup() returns a pointer to a new string which is a duplicate of *string*.

Otherwise, strdup() returns a NULL pointer.

Note: The caller of strdup() should free the storage obtained for the string.

Related Information

- “string.h — Declare string manipulation functions” on page 63
- “free() — Free a block of storage” on page 76
- “malloc() — Reserve storage block” on page 80

strlen() — Determine String Length

Format

```
#include <string.h>

size_t strlen(const char *string);
```

General Description

The strlen() built-in function determines the length of string pointed to by *string*, excluding the terminating NULL character.

Returned Value

The strlen() function returns the length of *string*.

Related Information

- “string.h — Declare string manipulation functions” on page 63
- “strncat() — Concatenate Strings” on page 104
- “strncmp() — Compare Strings” on page 104
- “strncpy() — Copy String” on page 105

strncat() — Concatenate Strings

Format

```
#include <string.h>

char *strncat(char * __restrict__ string1,
const char * __restrict__ string2, size_t count);
```

General Description

The `strncat()` built-in function appends the first *count* characters of *string2* to *string1* and ends the resulting string with a NULL character (`\0`). If *count* is greater than the length of *string2*, `strncat()` appends only the maximum length of *string2* to *string1*. The first character of the appended string overwrites the terminating NULL character of the string pointed to by *string1*.

If copying takes place between overlapping objects, the behavior is undefined.

Note: When **COMPACT** is specified, the compiler does not generate inline code for this function. In this case, you need to take either of the following actions:

- Define the `__METAL_SYSVEC` feature test macro to use the system library.
- Define the `__METAL_STATIC` feature test macro to use the static object library and bind in corresponding data set.

Returned Value

The `strncat()` function returns the value *string1*, the concatenated string.

Related Information

- “System and static object libraries” on page 65
- “string.h — Declare string manipulation functions” on page 63
- “strcat() — Concatenate Strings” on page 99
- “strncmp() — Compare Strings”
- “strncpy() — Copy String” on page 105
- “strpbrk() — Find Characters in String” on page 106
- “strrchr() — Find Last Occurrence of Character in String” on page 106
- “strspn() — Search String” on page 107

strncmp() — Compare Strings

Format

```
#include <string.h>

int strncmp(const char *string1, const char *string2, size_t count);
```

General Description

The `strncmp()` built-in function compares at most the first *count* characters of the string pointed to by *string1* to the string pointed to by *string2*.

The string arguments to the function should contain a NULL character (`\0`) marking the end of the string.

The relation between the strings is determined by the sign of the difference between the values of the leftmost first pair of characters that differ. The values depend on character encoding. This function is *not* locale sensitive.

Note: When **COMPACT** is specified, the compiler does not generate inline code for this function. In this case, you need to take either of the following actions:

- Define the `__METAL_SYSVEC` feature test macro to use the system library.
- Define the `__METAL_STATIC` feature test macro to use the static object library and bind in corresponding data set.

Returned Value

The `strncmp()` function returns a value indicating the relationship between the sub-strings, as follows:

Value Meaning

- < 0 String pointed to by *substring1* less than string pointed to by *substring2*
- = 0 String pointed to by *substring1* equivalent to string pointed to by *substring2*
- > 0 String pointed to by *substring1* greater than string pointed to by *substring2*

Related Information

- “System and static object libraries” on page 65
- “string.h — Declare string manipulation functions” on page 63
- “memcmp() — Compare bytes” on page 83
- “strcmp() — Compare Strings” on page 101
- “strcspn() — Compare Strings” on page 102
- “strncat() — Concatenate Strings” on page 104
- “strncpy() — Copy String”
- “strpbrk() — Find Characters in String” on page 106
- “strrchr() — Find Last Occurrence of Character in String” on page 106
- “strspn() — Search String” on page 107

strncpy() — Copy String

Format

```
#include <string.h>

char *strncpy(char * __restrict_string1,
const char * __restrict_string2, size_t count);
```

General Description

The `strncpy()` built-in function copies at most *count* characters of *string2* to *string1*. If *count* is less than or equal to the length of *string2*, a NULL character (`\0`) is *not* appended to the copied string. If *count* is greater than the length of *string2*, the *string1* result is padded with NULL characters (`\0`) up to length *count*.

If copying takes place between objects that overlap, the behavior is undefined.

Note: When **COMPACT** is specified, the compiler does not generate inline code for this function. In this case, you need to take either of the following actions:

- Define the `__METAL_SYSVEC` feature test macro to use the system library.
- Define the `__METAL_STATIC` feature test macro to use the static object library and bind in corresponding data set.

Returned Value

The `strncpy()` function returns *string1*.

Related Information

- “System and static object libraries” on page 65
- “string.h — Declare string manipulation functions” on page 63
- “memcpy() — Copy buffer” on page 83
- “strcpy() — Copy String” on page 102
- “strncat() — Concatenate Strings” on page 104
- “strncmp() — Compare Strings” on page 104
- “strpbrk() — Find Characters in String”
- “strrchr() — Find Last Occurrence of Character in String”
- “strspn() — Search String” on page 107

strpbrk() — Find Characters in String**Format**

```
#include <string.h>
```

```
char *strpbrk(const char *string1, const char *string2);
```

General Description

The `strpbrk()` function locates the first occurrence in the string pointed to by *string1* of any character from the string pointed to by *string2*.

Returned Value

If successful, `strpbrk()` returns a pointer to the character.

If *string1* and *string2* have no characters in common, `strpbrk()` returns a NULL pointer.

Related Information

- “string.h — Declare string manipulation functions” on page 63
- “strchr() — Search for Character” on page 100
- “strcspn() — Compare Strings” on page 102
- “strncmp() — Compare Strings” on page 104
- “strrchr() — Find Last Occurrence of Character in String”
- “strspn() — Search String” on page 107

strrchr() — Find Last Occurrence of Character in String**Format**

```
#include <string.h>
```

```
char *strrchr(const char *string, int c);
```

General Description

The `strrchr()` function finds the last occurrence of *c* (converted to a char) in *string*. The ending NULL character is considered part of the *string*.

Note: When **COMPACT** is specified, the compiler does not generate inline code for this function. In this case, you need to take either of the following actions:

- Define the `__METAL_SYSVEC` feature test macro to use the system library.
- Define the `__METAL_STATIC` feature test macro to use the static object library and bind in corresponding data set.

Returned Value

If successful, `strchr()` returns a pointer to the last occurrence of *c* in *string*.

If the given character is not found, `strchr()` returns a NULL pointer.

Related Information

- “System and static object libraries” on page 65
- “string.h — Declare string manipulation functions” on page 63
- “memchr() — Search buffer” on page 82
- “strchr() — Search for Character” on page 100
- “strcspn() — Compare Strings” on page 102
- “strncmp() — Compare Strings” on page 104
- “strpbrk() — Find Characters in String” on page 106
- “strspn() — Search String”

strspn() — Search String

Format

```
#include <string.h>
```

```
size_t strspn(const char *string1, const char *string2);
```

General Description

The `strspn()` function calculates the length of the maximum initial portion of the string pointed to by *string1* that consists entirely of the characters contained in the string pointed to by *string2*.

Returned Value

The `strspn()` function returns the length of the substring found.

Related Information

- “string.h — Declare string manipulation functions” on page 63
- “strcat() — Concatenate Strings” on page 99
- “strchr() — Search for Character” on page 100
- “strcmp() — Compare Strings” on page 101
- “strcpy() — Copy String” on page 102
- “strcspn() — Compare Strings” on page 102
- “strpbrk() — Find Characters in String” on page 106
- “strchr() — Find Last Occurrence of Character in String” on page 106

strstr() — Locate Substring

Format

```
#include <string.h>
```

```
char *strstr(const char *string1, const char *string2);
```

General Description

The `strstr()` function finds the first occurrence of the string pointed to by *string2* (excluding the NULL character) in the string pointed to by *string1*.

Returned Value

If successful, `strstr()` returns a pointer to the beginning of the first occurrence of *string2* in *string1*.

If *string2* does not appear in *string1*, `strstr()` returns `NULL`.

If *string2* points to a string with zero length, `strstr()` returns *string1*.

Related Information

- “string.h — Declare string manipulation functions” on page 63
- “strchr() — Search for Character” on page 100
- “strcmp() — Compare Strings” on page 101
- “strcspn() — Compare Strings” on page 102
- “strncmp() — Compare Strings” on page 104
- “strpbrk() — Find Characters in String” on page 106
- “strrchr() — Find Last Occurrence of Character in String” on page 106
- “strspn() — Search String” on page 107

strtod — Convert Character String to Double

Format

```
#include <stdlib.h>
double strtod(const char * __restrict__nptr, char ** __restrict__endptr);
```

General Description

The `strtod()` function converts part of a character string, pointed to by *nptr*, to a double. The parameter *nptr* points to a sequence of characters that can be interpreted as a numerical value of the type double.

The `strtod()` function breaks the string into three parts:

1. An initial, possibly empty, sequence of white-space characters, as specified by `isspace()`.
2. A subject sequence interpreted as a floating-point constant or representing infinity or a NAN.
3. A final string of one or more unrecognized characters, including the terminating null byte of the input string.

The subject string is the longest string that matches the expected form.

The expected form of the subject sequence is an optional plus or minus sign with one of the following parts:

- A non-empty sequence of decimal digits optionally containing a radix character followed by an optional exponent part. A radix character is the character that separates the integer part of a number from the fractional part.
- A `0x` or `0X`, a non-empty sequence of hexadecimal digits optionally containing a radix character, a base 2 decimal exponent part with a `p` or `P` as prefix, a plus or minus sign, and then a sequence of at least one decimal digit, for example, `[-]0xh.hhhhp+/-d`.
- An `INF`, ignoring case.
- One of `NANQ` or `NANQ(n)`, ignoring case.
- One of `NANS` or `NANS(n)`, ignoring case.
- One of `NAN` or `NAN(n)`, ignoring case.

See “`sscanf()` — Read and Format Data” on page 94 for a description of special infinity and NAN sequences recognized by Enterprise Metal C for z/OS.

The pointer to the last string that was successfully converted is stored in the object pointed to by *endptr*, if *endptr* is not a NULL pointer. If the subject string is empty or it does not have the expected form, no conversion is performed. The value of *nptr* is stored in the object pointed to by *endptr*.

Returned Value

If successful, `strtod()` returns the value of the floating-point number in IEEE Binary Floating-Point format.

In an overflow, `strtod()` returns `+/-HUGE_VAL`. In an underflow, it returns 0. If no conversion is performed, `strtod()` returns 0.

Related Information

- “`stdlib.h` — Define standard library functions” on page 62
- “`atoi()` — Convert character string to integer” on page 70
- “`atol()` — Convert character string to long” on page 70
- “`sscanf()` — Read and Format Data” on page 94
- “`strtol()` — Convert Character String to Long” on page 112
- “`strtof` — Convert Character String to Float”
- “`strtold` — Convert Character String to Long Double” on page 113
- “`strtoul()` — Convert String to Unsigned Integer” on page 116
- “`vsscanf()` — Format Input of a STDARG Argument List” on page 122

strtof — Convert Character String to Float

Format

```
#include <stdlib.h>
float strtof(const char * __restrict_nptr, char ** __restrict_endptr);
```

General Description

The `strtof()` function converts part of a character string, pointed to by *nptr*, to a float. The parameter *nptr* points to a sequence of characters that can be interpreted as a numerical value of the type float.

The `strtof()` function breaks the string into three parts:

1. An initial, possibly empty, sequence of white-space characters, as specified by `isspace()`.
2. A subject sequence interpreted as a floating-point constant or representing infinity or a NAN.
3. A final string of one or more unrecognized characters, including the terminating null byte of the input string.

The subject string is the longest string that matches the expected form.

The expected form of the subject sequence is an optional plus or minus sign with one of the following parts:

- A non-empty sequence of decimal digits optionally containing a radix character followed by an optional exponent part. A radix character is the character that separates the integer part of a number from the fractional part.

- A 0x or 0X, a non-empty sequence of hexadecimal digits optionally containing a radix character, a base 2 decimal exponent part with a p or P as prefix, a plus or minus sign, and then a sequence of at least one decimal digit, for example, [-]0xh.hhhhp+/-d.
- An INF, ignoring case.
- One of NANQ or NANQ(n), ignoring case.
- One of NANS or NANS(n), ignoring case.
- One of NAN or NAN(n), ignoring case.

In Enterprise Metal C for z/OS, represent the radix character as a period (.).

See “`sscanf()` — Read and Format Data” on page 94 for a description of special infinity and NAN sequences recognized by Enterprise Metal C for z/OS.

The pointer to the last string that was successfully converted is stored in the object pointed to by *endptr*, if *endptr* is not a NULL pointer. If the subject string is empty or it does not have the expected form, no conversion is performed. The value of *nptr* is stored in the object pointed to by *endptr*.

Returned Value

If successful, `strtouf()` returns the value of the floating-point number in IEEE Binary Floating-Point format.

In an overflow, `strtouf()` returns +/-HUGE_VALF. In an underflow, it returns 0. If no conversion is performed, `strtouf()` returns 0.

Related Information

- “`stdlib.h` — Define standard library functions” on page 62
- “`atoi()` — Convert character string to integer” on page 70
- “`atol()` — Convert character string to long” on page 70
- “`sscanf()` — Read and Format Data” on page 94
- “`strtod` — Convert Character String to Double” on page 108
- “`strtoul()` — Convert Character String to Long” on page 112
- “`strtold` — Convert Character String to Long Double” on page 113
- “`strtoul()` — Convert String to Unsigned Integer” on page 116
- “`vsscanf()` — Format Input of a STDARG Argument List” on page 122

strtok() — Tokenize String

Format

```
#include <string.h>
```

```
char *strtok(char * __restrict __string1, const char * __restrict __string2);
```

General Description

The `strtok()` function breaks a character string, pointed to by *string*, into a sequence of tokens. The tokens are separated from one another by the characters in the string pointed to by *string2*.

The token starts with the first character not in the string pointed to by *string2*. If such a character is not found, there are no tokens in the string. `strtok()` returns a NULL pointer. The token ends with the first character contained in the string pointed to by *string2*. If such a character is not found, the token ends at the

terminating NULL character. Subsequent calls to `strtok()` will return the NULL pointer. If such a character *is* found, then it is overwritten by a NULL character, which terminates the token.

If the next call to `strtok()` specifies a NULL pointer for *string1*, the tokenization resumes at the first character following the found and overwritten character from the previous call. For example:

```
/* Here are two calls */
strtok(string, " ")
strtok(NULL, " ")

/* Here is the string they are processing */
      abc defg hij
first call finds  ↑
                  ↑ second call starts
```

Note: To use the `strtok()` function, set up an environment by using the `__cinit()` function. When the function is called, GPR 12 must contain the environment token created by the `__cinit()` call.

Returned Value

The first time `strtok()` is called, it returns a pointer to the first token in *string1*. In later calls with the same token string, `strtok()` returns a pointer to the next token in the string. A NULL pointer is returned when there are no more tokens. All tokens are NULL-terminated.

Related Information

- “string.h — Declare string manipulation functions” on page 63
- “strcat() — Concatenate Strings” on page 99
- “strchr() — Search for Character” on page 100
- “strcmp() — Compare Strings” on page 101
- “strcpy() — Copy String” on page 102
- “strcspn() — Compare Strings” on page 102
- “strspn() — Search String” on page 107

strtok_r() — Split String into Tokens

Format

```
#define _XOPEN_SOURCE 500
#include <string.h>

char *strtok_r(char *s, const char *sep, char **lasts);
```

General Description

The function `strtok_r()` considers the NULL-terminated string *s* as a sequence of zero or more text tokens separated by spans of one or more characters from the separator string *sep*. The argument *lasts* points to a user-provided pointer which points to stored information necessary for `strtok_r()` to continue scanning the same string.

In the first call to `strtok_r()`, *s* points to a NULL-terminated string, *sep* to a NULL-terminated string of separator characters and the value pointed to by *lasts* is ignored. The function `strtok_r()` returns a pointer to the first character of the first

token, writes a NULL character into *s* immediately following the returned token, and updates the pointer to which *lasts* points.

In subsequent calls, *s* is a NULL pointer and *lasts* will be unchanged from the previous call so that subsequent calls will move through the string *s*, returning successive tokens until no tokens remain. The separator string *sep* may be different from call to call. When no token remains in *s*, a NULL pointer is returned.

Note: To use the `strtok_r()` function, set up an environment by using the `__cinit()` function. When the function is called, GPR 12 must contain the environment token created by the `__cinit()` call.

Returned Value

If successful, `strtok_r()` returns a pointer to the token found.

When no token is found, `strtok_r()` returns a NULL pointer.

Related Information

- “string.h — Declare string manipulation functions” on page 63
- “strcat() — Concatenate Strings” on page 99
- “strchr() — Search for Character” on page 100
- “strcmp() — Compare Strings” on page 101
- “strcpy() — Copy String” on page 102
- “strcspn() — Compare Strings” on page 102
- “strspn() — Search String” on page 107

strtol() — Convert Character String to Long

Format

```
#include <stdlib.h>
```

```
long int strtol(const char * __restrict __nptr,  
char ** __restrict __endptr, int base);
```

General Description

The `strtol()` function converts *nptr*, a character string, to a long int value.

The function decomposes the entire string into three parts:


1. A sequence of white space characters as defined by the IBM-1047 codepage.
2. A sequence of characters interpreted as integer in some base notation. This is the *subject sequence*.
3. A sequence of unrecognized characters.

The base notation is determined by *base*, if *base* is greater than zero. If *base* is zero, the base notation is determined by the format of the sequence of characters that follow an optional plus—or optional minus—sign.

- | | |
|----|---|
| 10 | Sequence starts with nonzero decimal digit. |
| 8 | Sequence starts with 0, followed by a sequence of digits with values from 0 to 7. |
| 16 | Sequence starts with either 0x or 0X, followed by digits, and letters A through F or a through f. |

If the base is greater than zero, the subject sequence contains decimal digits and letters, possibly preceded by either a plus or a minus sign. The letters a (or A) through z (or Z) represent values from 10 through 36, but only those letters whose value is less than the value of the base are allowed.

When you use the `strtol()` function, *nptr* should point to a string with the following form:



The pointer to the converted characters, even if conversion was unsuccessful, is stored in the object pointed to by *endptr*, if *endptr* is not a NULL pointer.

Returned Value

If successful, `strtol()` returns the converted long int value.

If unsuccessful, `strtol()` returns 0 if no conversion could be performed. If the correct value is outside the range of representable values, `strtol()` returns `LONG_MAX` or `LONG_MIN`, according to the sign of the value. If the value of `base` is not supported, `strtol()` returns 0.

Related Information

- “stdlib.h — Define standard library functions” on page 62
- “atoi() — Convert character string to integer” on page 70
- “atol() — Convert character string to long” on page 70
- “atoll() — Convert character string to signed long long” on page 71
- “sscanf() — Read and Format Data” on page 94
- “strtoul() — Convert String to Unsigned Integer” on page 116

strtold — Convert Character String to Long Double

Format

```
#include <stdlib.h>
long double strtold(const char * restrict nptr, char ** restrict endptr);
```

General Description

The `strtold()` function converts part of a character string, pointed to by *nptr*, to long double. The parameter *nptr* points to a sequence of characters that can be interpreted as a numerical value of the type long double.

The `strtold()` function breaks the string into three parts:

1. An initial, possibly empty, sequence of white-space characters, as specified by `isspace()`.
2. A subject sequence interpreted as a floating-point constant or representing infinity or a NAN.

3. A final string of one or more unrecognized characters, including the terminating null byte of the input string.

The function then attempts to convert the subject string into the floating-point number, and returns the result.

The expected form of the subject sequence is an optional plus or minus sign with one of the following parts:

- A non-empty sequence of decimal digits optionally containing a radix character followed by an optional exponent part. A radix character is the character that separates the integer part of a number from the fractional part.
- A 0x or 0X, a non-empty sequence of hexadecimal digits optionally containing a radix character, a base 2 decimal exponent part with a p or P as prefix, a plus or minus sign, and then a sequence of at least one decimal digit, for example, [-]0xh.hhhhp+/-d.
- An INF, ignoring case.
- One of NANQ or NANQ(n), ignoring case.
- One of NANS or NANS(n), ignoring case.
- One of NAN or NAN(n), ignoring case.

See “`sscanf()` — Read and Format Data” on page 94 for a description of special infinity and NAN sequences recognized by Enterprise Metal C for z/OS.

The pointer to the last string that was successfully converted is stored in the object pointed to by *endptr*, if *endptr* is not a NULL pointer. If the subject string is empty or it does not have the expected form, no conversion is performed. The value of *nptr* is stored in the object pointed to by *endptr*.

Returned Value

If successful, `strtold()` returns the value of the floating-point number in IEEE Binary Floating-Point format.

In an overflow, `strtold()` returns +/-HUGE_VAL. In an underflow, it returns 0. If no conversion is performed, `strtold()` returns 0.

Related Information

- “`stdlib.h` — Define standard library functions” on page 62
- “`atoi()` — Convert character string to integer” on page 70
- “`atol()` — Convert character string to long” on page 70
- “`sscanf()` — Read and Format Data” on page 94
- “`strtod` — Convert Character String to Double” on page 108
- “`strtof` — Convert Character String to Float” on page 109
- “`strtol()` — Convert Character String to Long” on page 112
- “`strtoul()` — Convert String to Unsigned Integer” on page 116
- “`vsscanf()` — Format Input of a STDARG Argument List” on page 122

strtoll() — Convert String to Signed Long Long

Format

```
#include <stdlib.h>
```

```
long long strtoll(const char * __restrict nptr,
char ** __restrict endptr, int base);
```

Compile Requirement: Use of this function requires the long long data type.

General Description

The `strtol()` function converts *nptr*, a character string, to a signed long long value.

The function decomposes the entire string into three parts:


1. A sequence of white space characters as defined by the IBM-1047 codepage.
2. A sequence of characters interpreted as an unsigned integer in some base notation. This is the *subject sequence*.
3. A sequence of unrecognized characters.

The base notation is determined by *base*, if *base* is greater than zero. If *base* is zero, the base notation is determined by the format of the sequence of characters that follow an optional plus or optional minus sign.

- 10 Sequence starts with nonzero decimal digit.
- 8 Sequence starts with 0, followed by a sequence of digits with values from 0 to 7.
- 16 Sequence starts with either 0x or 0X, followed by digits, and letters A through F or a through f.

If the base is greater than zero, the subject sequence contains decimal digits and letters, possibly preceded by either a plus or a minus sign. The letters a (or A) through z (or Z) represent values from 10 through 36, but only those letters whose value is less than the value of the base are allowed.

When you are using `strtol()`, *nptr* should point to a string with the following form:



The pointer to the converted characters, even if conversion was unsuccessful, is stored in the object pointed to by *endptr*, if *endptr* is not a NULL pointer.

Returned Value

If successful, `strtol()` returns the converted signed long long value, represented in the string.

If unsuccessful, `strtol()` returns 0 if no conversion could be performed. If the correct value is outside the range of representable values, `strtol()` returns `LLONG_MAX` (`LONGLONG_MAX`) or `LLONG_MIN` (`LONGLONG_MIN`), according to the sign of the value. If the value of *base* is not supported, `strtol()` returns 0.

Related Information

- “stdlib.h — Define standard library functions” on page 62
- “atoi() — Convert character string to integer” on page 70
- “atol() — Convert character string to long” on page 70
- “atoll() — Convert character string to signed long long” on page 71
- “sscanf() — Read and Format Data” on page 94
- “strtoul() — Convert String to Unsigned Integer”

strtoul() — Convert String to Unsigned Integer

Format

```
#include <stdlib.h>
```

```
unsigned long int strtoul(const char * __restrict__ string1,  
char ** __restrict__ string2, int base);
```

General Description

The `strtoul()` function converts *string1*, a character string, to an unsigned long int value.

The function decomposes the entire string into three parts:


1. A sequence of white space characters as defined by the IBM-1047 codepage.
2. A sequence of characters interpreted as an unsigned integer in some base notation. This is the *subject sequence*.
3. A sequence of unrecognized characters.

The base notation is determined by *base*, if *base* is greater than zero. If *base* is zero, the base notation is determined by the format of the sequence of characters that follow an optional plus or optional minus sign.

- | | |
|----|---|
| 10 | Sequence starts with nonzero decimal digit. |
| 8 | Sequence starts with 0, followed by a sequence of digits with values from 0 to 7. |
| 16 | Sequence starts with either 0x or 0X, followed by digits, and letters A through F or a through f. |

If the base is greater than zero, the subject sequence contains decimal digits and letters, possibly preceded by either a plus or a minus sign. The letters a (or A) through z (or Z) represent values from 10 through 36, but only those letters whose value is less than the value of the base are allowed. The function stops reading the string at the first character that it cannot recognize as part of a number. This character can be the first numeric character greater than or equal to the *base*. The `strtoul()` function sets *string2* to point to the end of the resulting output string if a conversion is performed and provided that *string2* is not a NULL pointer.

When you are using the `strtoul()` function, *string1* should point to a string with the following form:



If *base* is in the range of 2-36, it becomes the base of the number. If *base* is 0, the prefix determines the base (8, 16, or 10): the prefix 0 means base 8; the prefix 0x or 0X means base 16; using any other digit without a prefix means decimal.

The pointer to the converted characters, even if conversion was unsuccessful, is stored in the object pointed to by *string2*, if *string2* is not a NULL pointer.

Returned Value

If successful, `strtoul()` returns the converted unsigned long int value, represented in the string.

If unsuccessful, `strtoul()` returns 0 if no conversion could be performed. If the correct value is outside the range of representable values, `strtoul()` returns `ULONG_MAX`. If the value of *base* is not supported, `strtoul()` returns 0.

Related Information

- “`stdlib.h` — Define standard library functions” on page 62
- “`atoi()` — Convert character string to integer” on page 70
- “`atol()` — Convert character string to long” on page 70
- “`atoll()` — Convert character string to signed long long” on page 71
- “`sscanf()` — Read and Format Data” on page 94
- “`strtol()` — Convert Character String to Long” on page 112

strtoull() — Convert String to Unsigned Long Long

Format

```
#include <stdlib.h>
```

```
unsigned long long strtoull(register const char * __restrict__ nptr,
char ** __restrict__ endptr, int base);
```

Compile Requirement: Use of this function requires the long long data type.

General Description

The `strtoull()` function converts *nptr*, a character string, to an unsigned long long value.

The function decomposes the entire string into three parts:


1. A sequence of white space characters as defined by the IBM-1047 codepage.
2. A sequence of characters interpreted as an unsigned integer in some base notation. This is the *subject sequence*.
3. A sequence of unrecognized characters.

The base notation is determined by *base*, if *base* is greater than zero. If *base* is zero, the base notation is determined by the format of the sequence of characters that follow an optional plus or optional minus sign.

- 10 Sequence starts with nonzero decimal digit.
- 8 Sequence starts with 0, followed by a sequence of digits with values from 0 to 7.
- 16 Sequence starts with either 0x or 0X, followed by digits, and letters A through F or a through f.

If the base is greater than zero, the subject sequence contains decimal digits and letters, possibly preceded by either a plus or a minus sign. The letters a (or A) through z (or Z) represent values from 10 through 36, but only those letters whose value is less than the value of the base are allowed. The function stops reading the string at the first character that it cannot recognize as part of a number. This character can be the first numeric character greater than or equal to the *base*. The `strtoull()` function sets *endptr* to point to the end of the resulting output string if a conversion is performed and provided that *endptr* is not a NULL pointer.

When you are using the `strtoull()` function, *nptr* should point to a string with the following form:



If *base* is in the range of 2-36, it becomes the base of the number. If *base* is 0, the prefix determines the base (8, 16 or 10): the prefix 0 means base 8; the prefix 0x or 0X means base 16; using any other digit without a prefix means decimal.

The pointer to the converted characters, even if conversion was unsuccessful, is stored in the object pointed to by *endptr*, if *endptr* is not a NULL pointer.

Returned Value

If successful, `strtoull()` returns the converted unsigned long long value, represented in the string.

If unsuccessful, `strtoull()` returns 0 if no conversion could be performed. If the correct value is outside the range of representable values, `strtoull()` returns `ULLONG_MAX` (`ULONGLONG_MAX`). If the value of *base* is not supported, `strtoull()` returns 0.

Related Information

- “stdlib.h — Define standard library functions” on page 62
- “atoi() — Convert character string to integer” on page 70
- “atol() — Convert character string to long” on page 70
- “atoll() — Convert character string to signed long long” on page 71
- “sscanf() — Read and Format Data” on page 94
- “strtol() — Convert String to Signed Integer” on page 116

tolower(), toupper() — Convert Character Case

Format

```
#include <ctype.h>

int tolower(int c); /* Convert c to lowercase if appropriate */
int toupper(int c); /* Convert c to uppercase if appropriate */
```

General Description

The `tolower()` function converts *c* to a lowercase letter, if possible. Conversely, the `toupper()` function converts *c* to an uppercase letter, if possible.

Returned Value

If successful, `tolower()` and `toupper()` return the corresponding character, as defined in the IBM-1047 code page, if such a character exists.

If unsuccessful, `tolower()` and `toupper()` return the unchanged value *c*.

Related Information

- “`ctype.h` — Declare character classification functions” on page 55
- “`isalnum()` to `isxdigit()` — Test integer value” on page 76

va_arg(), va_copy(), va_end(), va_start() — Access Function Arguments

Format

```
#include <stdarg.h>

var_type va_arg(va_list arg_ptr, var_type);
void va_end(va_list arg_ptr);
void va_start(va_list arg_ptr, variable_name);
```

C99: See the sample code below.

```
#define _ISOC99_SOURCE
#include <stdarg.h>

var_type va_arg(va_list arg_ptr, var_type);
void va_end(va_list arg_ptr);
void va_start(va_list arg_ptr, variable_name);
void va_copy(va_list dest, va_list src);
```

General Description

The `va_arg()`, `va_end()`, and `va_start()` macros access the arguments to a function when it takes a fixed number of required arguments and a variable number of optional arguments. You declare required arguments as ordinary parameters to the function and access the arguments through the parameter names.

The `va_start()` macro initializes the *arg_ptr* pointer for subsequent calls to `va_arg()` and `va_end()`.

The argument *variable_name* is the identifier of the rightmost named parameter in the parameter list (preceding `,` ...). Use the `va_start()` macro before the `va_arg()` macro. Corresponding `va_start()` and `va_end()` macro calls must be in the same function. If *variable_name* is declared as a register, with a function or an array type,

va_arg, va_copy, va_end, va_start

or with a type that is not compatible with the type that results after application of the default argument promotions, then the behavior is undefined.

The `va_arg()` macro retrieves a value of the given *var_type* from the location given by *arg_ptr* and increases *arg_ptr* to point to the next argument in the list. The `va_arg()` macro can retrieve arguments from the list any number of times within the function.

The macros also provide fixed-point decimal support under Enterprise Metal C for z/OS. The `sizeof(xx)` operator is used to determine the size and type casting that is used to generate the values. Therefore, a call, such as, `x = va_arg(ap, _Decimal(5,2));` is valid. The size of a fixed-point decimal number, however, cannot be made a variable. Therefore, a call, such as, `z = va_arg(ap, _Decimal(x,y))` where `x = 5` and `y = 2` is not valid.

The `va_end()` macro is needed by some systems to indicate the end of parameter scanning.

`va_start()` and `va_arg()` do not work with parameter lists of functions whose linkages were changed with the `#pragma linkage` directive.

`stdarg.h` and `varargs.h` are mutually exclusive. Whichever `#include` comes first, determines the form of macro that is visible.

The type definition for the `va_list` type in this implementation is "char *va_list".

The `va_copy()` function creates a copy (*dest*) of a variable of type `va_list` (*src*). The copy appear as if it has gone through a `va_start()` and the exact set of sequences of `va_arg()` as that of *src*.

After `va_copy()` initializes *dest*, the `va_copy()` macro shall not be invoked to reinitialize *dest* without an intervening invocation of the `va_end()` macro for the same *dest*.

Returned Value

The `va_arg()` macro returns the current argument.

The `va_end()`, `va_copy()`, and `va_start()` macros return no values.

Related Information

- “`stdarg.h` — Define macros for accessing variable-length argument lists in functions” on page 60
- “`vsnprintf()` — Format and print data to fixed length buffer”
- “`vsprintf()` — Format and Print Data to Buffer” on page 121

vsnprintf() — Format and print data to fixed length buffer

Format

```
#include <stdarg.h>
#include <stdio.h>
```

```
int vsnprintf(char *__restrict__ s, size_t n,
              const char *__restrict__ format, va_list arg);
```

General Description

The `vsnprintf()` function is equivalent to `sprintf()`, except that instead of being called with a variable number of arguments, it is called with an argument list as defined by `stdarg.h`. For a specification of the *format* string, see “`sprintf()` — Format and Write Data” on page 88.

Initialize the argument list by using the `va_start` macro before each call. These functions do not invoke the `va_end` macro, but instead invoke the `va_arg` macro causing the value of `arg` after the return to be unspecified.

Notes:

1. Use of `vsnprintf()` requires that an environment has been set up by using the `__cinit()` function. When the function is called, GPR 12 must contain the environment token created by the `__cinit()` call.
2. In contrast to some UNIX-based implementations of the C language, the Enterprise Metal C for z/OS implementation of the `vprintf()` family increments the pointer to the variable arguments list. To control whether the pointer is incremented, call the `va_end` macro after each function call.

Returned Value

The `vsnprintf()` function returns the number of characters that would have been written had *n* been sufficiently large, not counting the terminating null character, or a negative value if an encoding error occurred. Thus, the null-terminated output has been completely written if and only if the returned value is non-negative and less than *n*.

Related Information

- “`stdarg.h` — Define macros for accessing variable-length argument lists in functions” on page 60
- “`stdio.h` — Define I/O related functions” on page 60
- “`va_arg()`, `va_copy()`, `va_end()`, `va_start()` — Access Function Arguments” on page 119

vsprintf() — Format and Print Data to Buffer

Format

```
#include <stdarg.h>
#include <stdio.h>

int vsprintf(char * __restrict__target-string,
             const char * __restrict__format, va_list arg_ptr);
```

General Description

The `vsprintf()` function is equivalent to the `sprintf()` function, except that instead of being called with a variable number of arguments, it is called with an argument list as defined in `stdarg.h`. For a specification of the *format* string, see “`sprintf()` — Format and Write Data” on page 88.

Initialize the argument list by using the `va_start` macro before each call. These functions do not invoke the `va_end` macro, but instead invoke the `va_arg` macro causing the value of `arg` after the return to be unspecified.

Notes:

vsprintf

1. Use of `vsprintf()` requires that an environment has been set up by using the `__cinit()` function. When the function is called, GPR 12 must contain the environment token created by the `__cinit()` call.
2. In contrast to some UNIX-based implementations of the C language, the Enterprise Metal C for z/OS implementation of the `vprintf()` family increments the pointer to the variable arguments list. To control whether the pointer to the argument is incremented, call the `va_end` macro after each call to `vsprintf()`.

Returned Value

If successful, `vsprintf()` returns the number of characters written *target-string*.

If unsuccessful, `vsprintf()` returns a negative value.

Related Information

- “`stdarg.h` — Define macros for accessing variable-length argument lists in functions” on page 60
- “`stdio.h` — Define I/O related functions” on page 60
- “`va_arg()`, `va_copy()`, `va_end()`, `va_start()` — Access Function Arguments” on page 119

vsscanf() — Format Input of a STDARG Argument List

Format

```
#define _ISOC99_SOURCE
#include <stdarg.h>
#include <stdio.h>

int vsscanf(const char *__restrict__ s,
            const char *__restrict__ format, va_list arg);
```

General Description

The `vsscanf()` function is equivalent to the `sscanf()` function, except that instead of being called with a variable number of arguments, it is called with an argument list as defined in `stdarg.h`.

Initialize the argument list by using the `va_start` macro before each call. These functions do not invoke the `va_end` macro, but instead invoke the `va_arg` macro causing the value of *arg* after the return to be unspecified.

Notes:

1. Use of `vsscanf()` requires that an environment has been set up by using the `__cinit()` function. When the function is called, GPR 12 must contain the environment token created by the `__cinit()` call.
2. In contrast to some UNIX-based implementations of the C language, the Enterprise Metal C for z/OS implementation of the `vscanf()` family increments the pointer to the variable arguments list. To control whether the pointer is incremented, call the `va_end` macro after each function call.

Returned Value

See “`sscanf()` — Read and Format Data” on page 94.

Related Information

- “stdarg.h — Define macros for accessing variable-length argument lists in functions” on page 60
- “stdio.h — Define I/O related functions” on page 60
- “sscanf() — Read and Format Data” on page 94

Appendix A. Function stack requirements

Table 15 lists the stack frame requirements for each Metal C runtime function. All sizes are in bytes.

Table 15. Stack frame requirements for Metal C runtime functions

Function	AMODE 31 stack size	AMODE 64 stack size
abs	256	512
atoi	256	512
atol	256	512
atoll	1280	1536
calloc	1024	1536
__cinit	512	512
__cterm	1024	1024
div	256	512
free	512	1536
isalnum	256	512
isalpha	256	512
isblank	256	512
iscntrl	256	512
isdigit	256	512
isgraph	256	512
islower	256	512
isprint	256	512
ispunct	256	512
isspace	256	512
isupper	256	512
isxdigit	256	512
labs	256	512
ldiv	256	512
llabs	512	512
lldiv	512	512
malloc	768	1024
__malloc31	768	1024
memcpy	512	512
memchr	512	512
memcmp	512	512
memcpy	512	512
memmove	512	512
memset	256	512
qsort	1280 ¹	1792 ¹

Table 15. Stack frame requirements for Metal C runtime functions (continued)

Function	AMODE 31 stack size	AMODE 64 stack size
rand	256	512
rand_r	256	512
realloc	1024	2048
snprintf	3072	3584
snprintf when using e, E, f, F, g, G conversion specifiers	32000	32768
snprintf when using e, E, f, F, g, G conversion specifiers with the L conversion prefix	48896	49920
sprintf	3072	3584
sprintf when using e, E, f, F, g, G conversion specifiers	32000	32768
sprintf when using e, E, f, F, g, G conversion specifiers with the L conversion prefix	48896	49920
srand	256	512
sscanf	2304	2560
sscanf when using e, E, f, F, g, G conversion specifiers	4864	5632
sscanf when using e, E, f, F, g, G conversion specifiers	5888	6656
sscanf when using e, E, f, F, g, G conversion specifiers with the L conversion prefix	23040	23552
strcat	512	512
strchr	512	512
strcmp	512	512
strcpy	512	512
strcspn	768	768
strdup	1024	1536
strlen	512	512
strncat	512	512
strncmp	512	512
strncpy	512	512
strpbrk	768	768
strrchr	512	512
strspn	768	768
strstr	512	512
strtod	4096	4352
strtof	3072	3328
strtok	768	1024
strtok_r	1024	1536
strtol	1024	1024

Table 15. Stack frame requirements for Metal C runtime functions (continued)

Function	AMODE 31 stack size	AMODE 64 stack size
strtold	21248	21248
strtoll	1024	1024
strtoul	1024	1024
strtoull	768	1024
tolower	256	512
toupper	256	512
vsnprintf	3072	3584
vsnprintf when using e, E, f, F, g, G conversion specifiers	32000	32768
vsnprintf when using e, E, f, F, g, G conversion specifiers with the L conversion prefix	48896	49920
vsprintf	3072	3584
vsprintf when using e, E, f, F, g, G conversion specifiers	32000	32768
vsprintf when using e, E, f, F, g, G conversion specifiers with the L conversion prefix	48896	49920
vsscanf	2304	2560
vsscanf when using e, E, f, F, g, G conversion specifiers	4864	5632
vsscanf when using e, E, f, F, g, G conversion specifiers with the l conversion prefix	5888	6656
vsscanf when using e, E, f, F, g, G conversion specifiers with the L conversion prefix	23040	23552
Note: You must add the stack size of the comparison function you supply to the base value.		

Appendix B. Accessibility

Accessible publications for this product are offered through IBM Knowledge Center (www.ibm.com/support/knowledgecenter/SSLTBW/welcome).

If you experience difficulty with the accessibility of any z/OS information, send a detailed message to the Contact z/OS web page (www.ibm.com/systems/z/os/zos/webqs.html) or use the following mailing address.

IBM Corporation
Attention: MHVRCFS Reader Comments
Department H6MA, Building 707
2455 South Road
Poughkeepsie, NY 12601-5400
United States

Accessibility features

Accessibility features help users who have physical disabilities such as restricted mobility or limited vision use software products successfully. The accessibility features in z/OS can help users do the following tasks:

- Run assistive technology such as screen readers and screen magnifier software.
- Operate specific or equivalent features by using the keyboard.
- Customize display attributes such as color, contrast, and font size.

Consult assistive technologies

Assistive technology products such as screen readers function with the user interfaces found in z/OS. Consult the product information for the specific assistive technology product that is used to access z/OS interfaces.

Keyboard navigation of the user interface

You can access z/OS user interfaces with TSO/E or ISPF. The following information describes how to use TSO/E and ISPF, including the use of keyboard shortcuts and function keys (PF keys). Each guide includes the default settings for the PF keys.

- *z/OS TSO/E Primer*
- *z/OS TSO/E User's Guide*
- *z/OS ISPF User's Guide Vol I*

Dotted decimal syntax diagrams

Syntax diagrams are provided in dotted decimal format for users who access IBM Knowledge Center with a screen reader. In dotted decimal format, each syntax element is written on a separate line. If two or more syntax elements are always present together (or always absent together), they can appear on the same line because they are considered a single compound syntax element.

Each line starts with a dotted decimal number; for example, 3 or 3.1 or 3.1.1. To hear these numbers correctly, make sure that the screen reader is set to read out

punctuation. All the syntax elements that have the same dotted decimal number (for example, all the syntax elements that have the number 3.1) are mutually exclusive alternatives. If you hear the lines 3.1 USERID and 3.1 SYSTEMID, your syntax can include either USERID or SYSTEMID, but not both.

The dotted decimal numbering level denotes the level of nesting. For example, if a syntax element with dotted decimal number 3 is followed by a series of syntax elements with dotted decimal number 3.1, all the syntax elements numbered 3.1 are subordinate to the syntax element numbered 3.

Certain words and symbols are used next to the dotted decimal numbers to add information about the syntax elements. Occasionally, these words and symbols might occur at the beginning of the element itself. For ease of identification, if the word or symbol is a part of the syntax element, it is preceded by the backslash (\) character. The * symbol is placed next to a dotted decimal number to indicate that the syntax element repeats. For example, syntax element *FILE with dotted decimal number 3 is given the format 3 * FILE. Format 3* FILE indicates that syntax element FILE repeats. Format 3* * FILE indicates that syntax element * FILE repeats.

Characters such as commas, which are used to separate a string of syntax elements, are shown in the syntax just before the items they separate. These characters can appear on the same line as each item, or on a separate line with the same dotted decimal number as the relevant items. The line can also show another symbol to provide information about the syntax elements. For example, the lines 5.1*, 5.1 LASTRUN, and 5.1 DELETE mean that if you use more than one of the LASTRUN and DELETE syntax elements, the elements must be separated by a comma. If no separator is given, assume that you use a blank to separate each syntax element.

If a syntax element is preceded by the % symbol, it indicates a reference that is defined elsewhere. The string that follows the % symbol is the name of a syntax fragment rather than a literal. For example, the line 2.1 %OP1 means that you must refer to separate syntax fragment OP1.

The following symbols are used next to the dotted decimal numbers.

? indicates an optional syntax element

The question mark (?) symbol indicates an optional syntax element. A dotted decimal number followed by the question mark symbol (?) indicates that all the syntax elements with a corresponding dotted decimal number, and any subordinate syntax elements, are optional. If there is only one syntax element with a dotted decimal number, the ? symbol is displayed on the same line as the syntax element, (for example 5? NOTIFY). If there is more than one syntax element with a dotted decimal number, the ? symbol is displayed on a line by itself, followed by the syntax elements that are optional. For example, if you hear the lines 5 ?, 5 NOTIFY, and 5 UPDATE, you know that the syntax elements NOTIFY and UPDATE are optional. That is, you can choose one or none of them. The ? symbol is equivalent to a bypass line in a railroad diagram.

! indicates a default syntax element

The exclamation mark (!) symbol indicates a default syntax element. A dotted decimal number followed by the ! symbol and a syntax element indicate that the syntax element is the default option for all syntax elements that share the same dotted decimal number. Only one of the syntax elements that share the dotted decimal number can specify the ! symbol. For example, if you hear the lines 2? FILE, 2.1! (KEEP), and 2.1 (DELETE), you know that (KEEP) is the

default option for the FILE keyword. In the example, if you include the FILE keyword, but do not specify an option, the default option KEEP is applied. A default option also applies to the next higher dotted decimal number. In this example, if the FILE keyword is omitted, the default FILE(KEEP) is used. However, if you hear the lines 2? FILE, 2.1, 2.1.1! (KEEP), and 2.1.1 (DELETE), the default option KEEP applies only to the next higher dotted decimal number, 2.1 (which does not have an associated keyword), and does not apply to 2? FILE. Nothing is used if the keyword FILE is omitted.

*** indicates an optional syntax element that is repeatable**

The asterisk or glyph (*) symbol indicates a syntax element that can be repeated zero or more times. A dotted decimal number followed by the * symbol indicates that this syntax element can be used zero or more times; that is, it is optional and can be repeated. For example, if you hear the line 5.1* data area, you know that you can include one data area, more than one data area, or no data area. If you hear the lines 3* , 3 HOST, 3 STATE, you know that you can include HOST, STATE, both together, or nothing.

Notes:

1. If a dotted decimal number has an asterisk (*) next to it and there is only one item with that dotted decimal number, you can repeat that same item more than once.
2. If a dotted decimal number has an asterisk next to it and several items have that dotted decimal number, you can use more than one item from the list, but you cannot use the items more than once each. In the previous example, you can write HOST STATE, but you cannot write HOST HOST.
3. The * symbol is equivalent to a loopback line in a railroad syntax diagram.

+ indicates a syntax element that must be included

The plus (+) symbol indicates a syntax element that must be included at least once. A dotted decimal number followed by the + symbol indicates that the syntax element must be included one or more times. That is, it must be included at least once and can be repeated. For example, if you hear the line 6.1+ data area, you must include at least one data area. If you hear the lines 2+, 2 HOST, and 2 STATE, you know that you must include HOST, STATE, or both. Similar to the * symbol, the + symbol can repeat a particular item if it is the only item with that dotted decimal number. The + symbol, like the * symbol, is equivalent to a loopback line in a railroad syntax diagram.

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 Corporation
J74/G4
555 Bailey Avenue
San Jose, CA 95141-1099
U.S.A.

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

Intellectual Property Licensing
Legal and Intellectual Property Law
IBM Japan, Ltd.
3-2-12, Roppongi, Minato-ku, Tokyo 106-8711

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

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

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

This information could include technical inaccuracies or typographical errors.

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

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

Programming interface information

This information documents intended programming interfaces that allow the customer to write programs to obtain the services of z/OS Metal C runtime library.

Standards

The following standards are supported in combination with the z/OS Metal C runtime library:

- The C language is consistent with *Programming languages - C (ISO/IEC 9899:1999)*. This standard has officially replaced American National Standard for Information Systems-Programming Language C (X3.159-1989) and is technically equivalent to the ANSI C standard. The compiler supports the changes adopted into the C Standard by ISO/IEC 9899:1990/Amendment 1:1994. For more information on ISO, see International Organization for Standardization (ISO) (www.iso.org).

Trademarks

IBM, the IBM logo, and ibm.com[®] are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names If these and other IBM trademarked terms are marked on their first occurrence in this information with a trademark symbol ([®] or [™]), these symbols indicate U.S. registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of IBM trademarks is available on the Web at Copyright and Trademark information (www.ibm.com/legal/copytrade.shtml).

Adobe, Acrobat, and PostScript are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

IEEE is a trademark of the Institute of Electrical and Electronics Engineers, Inc. in the United States and other countries.

Other company, product, and service names might be trademarks or service marks of others.

Index

Special characters

- `__asm` operand lists
 - defining read-write `__asm` operands 26
- `__asm` operands
 - C expressions as `__asm` operands 22
 - defining
 - read-write `__asm` operands 26
 - multiple
 - defining 24
 - read-write 26
- `__asm` statement
 - inserting your own assembly instructions 21
 - using
 - code format string 21
- `__asm` statements
 - C expressions as `__asm` operands 22
 - code format string 22
 - constraints 22
 - examples
 - read-write `__asm` operands 26
 - specifiers 22
- `__cinit()` library function 72
- `__far` qualifier
 - far pointer 35
- `__malloc31()` library function 81
- `__cterm()` library function 75
- `_MI_BUILTIN` macro
 - data space allocation 39
- `_MI.BUILTIN` macro
 - AR-mode functions 38
 - far-pointer management 37
- `-mgoff` HLASM option
 - and Metal C programs 46
- `#pragma insert_asm`
 - inserting your own assembly statements 30
- `#pragma directive`
 - MYEPILOG 14
 - MYPROLOG 14
- `+ constraint`
 - defining read-write `__asm` operands 27

A

- `abs()` library function 69
- absolute value 69
 - integer argument 69
- access registers
 - AR mode 37
 - management by compiler 37
 - restoring 39
 - saving 39
- accessibility 129
 - contact IBM 129
 - features 129

- ADATA debugging information
 - additional source-level information
 - output file format 51
 - CDAASMC procedure 51
 - CDAHLASM invocation 51
- addressing mode
 - and global SET symbols 15
 - and passing parameters 2
 - attributes
 - `amode31` 31
 - `amode64` 31
 - recognition of 31
 - switching 31
 - commands 47
 - example 31
- ALESERV HLASM macro
 - allocating alternative data spaces 37
- ALET
 - far pointer 35
 - implicit association 37
- ALIAS instructions
 - recognition of 4
- allocating
 - `realloc()` 86
- allocation
 - `_MI_BUILTIN` macro 39
 - of data space 39
- alphabetic character attribute 77
- alternative data spaces
 - accessing 37
 - allocating 37
- AMODE
 - and global SET symbols 15
 - and passing parameters 2
 - function save areas 3
 - return values 3
 - switching
 - commands 47
 - example 31
 - external function calls 31
 - internal function calls 31
- AR mode 35
 - linkage conventions 39
 - programming support 35
 - far-pointer management 37
- AR-mode functions
 - accessing alternative data spaces 37
 - ALET associations 37
 - built-in functions 38
 - C language constructs and far pointers 36
 - data space allocation 39, 42
 - default prolog and epilog code 39
 - far pointers
 - declaration 35
 - dereference 35
 - reference 35
 - memory references 37
- arguments
 - accessing 119
- ARMODE compiler option 35

- armode function attribute 35
- as command
 - building Metal C programs 46
- ASC mode
 - restoring 39
 - switching 39
- ASMLANGX debugging utility
 - debugging information format 51
- ASMLANGX utility
 - additional source-level information
 - ADATA debugging information 51
- assembly job step 48
- assembly language programs
 - debugging 51
 - load module size 51
 - source-level information 51
- assembly statements
 - embedding
 - code format string 22
 - example, simple 21
 - file-scope header 5
 - function entry point markers 9
 - function headers 8
 - function property blocks 9
 - making a C expression available to HLASM 23
 - making a C variable available to HLASM 22
 - making a C variable expression an `__asm` operand 22
 - operands 24
 - inserting
 - executable 21
 - non-executable 30
 - user-supplied 21
- assistive technologies 129
- `atoi()` library function 70
- `atol()` library function 70
- `atoll()` library function 71

B

- batch environment
 - binder invocation procedures 47
 - building Metal C programs 47
 - assembly step 48
 - bind step 48
 - compilation step 48
 - debugging assembly language programs 51
 - debugging information 51
 - extracting source-level information 51
 - f 32
 - IDF debugging information 51
- bind job step 48
- blank character attribute 78
- buffers
 - format and print data 121

- building Metal C programs
 - assembly step
 - symbols longer than eight characters 46
- built-in functions
 - AR-mode functions 38
 - far-pointer management
 - AR-mode programming support 37
- builtins.h header file 55
 - data space allocation 39
 - far versions of library functions 38
 - far-pointer management 37

C

- C expressions
 - used as `__asm` operands 22
- C language constructs
 - far pointers 36
- C memory functions
 - far versions 38
- C string functions
 - far versions 38
- C string pointer
 - copying to far pointer 42
- C symbols
 - name-encoding 4
- calloc() library function 71
- CDAASMC JCL procedure
 - binder invocation 47
 - extracting source-level information 51
 - invoking 48
- CDAHLASM
 - invocation 51
- CEE.SCEEPROC data set
 - binder invocation batch procedures 47
- characters
 - conversions
 - lowercase 119
 - uppercase 119
 - finding in a string 106
- classifying characters 76
- clobber list
 - example 25
- code base registers 4
- code format string
 - description 21
 - in an `__asm` statement 22
 - substitution specifiers 22
 - treatment of 22
- code format strings
 - data space allocation 39
- examples
 - read-write `__asm` operands 26
- command
 - syntax diagrams x
- comparing
 - `strcmp()` 101
 - `strcspn()` 102
 - strings 101, 102, 104
 - `strncmp()` 104
- compilation job step 48

- compiler options
 - AMODE
 - characteristics of
 - compiler-generated assembly source code 4
 - ARMODE|NOARMODE 35
 - EPILOG
 - versus `#pragma` epilog 14
 - LONGNAME
 - entry point definitions 8
 - entry point marker 9
 - external symbols 5
 - function property blocks 9
 - LP64
 - characteristics of
 - compiler-generated assembly source code 4
 - programming with Metal C 2
 - PROLOG
 - versus `#pragma` prolog 14
- concatenating
 - `strcat()` 99
 - strings 99, 104
 - `strncat()` 104
- constants
 - defining 12
- contact
 - z/OS 129
- conversions
 - character
 - to lowercase 119
 - to uppercase 119
 - specifier
 - argument in `scanf()` 97
 - string to unsigned integer 116
- copying
 - `strcpy()` 102
 - strings 102, 105
 - `strncpy()` 105
- CSECT 4
- ctype.h header file 55

D

- data spaces
 - access 35
 - accessing 42
 - allocation 39
 - deallocation 39, 42
 - referencing 42
- debugging
 - assembly language programs 51
 - data formats 51
 - extracting source-level information 51
 - ADATA 51
 - ASMLANGX 51
 - IDF 52
 - in a batch environment 51
 - Interactive Debug Facility (IDF) 51
 - interactive utility 52
 - source-level information 51
- `div_t` structure 75
- `div()` library function 75
- division 75
- DSA
 - acquisition and release 3

- DSA (*continued*)
 - address space 39
 - and global SET symbols 15
 - default
 - AR-mode functions 39
 - function save areas 3
 - location 39
 - obtaining 13
 - obtaining and releasing 20
 - pointer 20
 - preallocation 14
- DSECT statement
 - and file-scope trailers 12
- DSPSERV HLASM macro
 - allocating alternative data spaces 37
- DWARF debugging information
 - CDAASMC procedure 51
 - CDAHLASM invocation 51
- dynamic storage area
 - acquisition and release 3
 - function save areas 3
 - location 39
 - obtaining and releasing 20, 39
 - preallocation 14

E

- entry point markers
 - defining
 - under LONGNAME compiler option 9
- entry points
 - defining
 - under LONGNAME compiler option 8
- epilog code
 - AR-mode functions 13, 39
 - default 20
 - AR-mode functions 39
 - DSA pointer 20
 - NAB pointer 20
 - primary functions 13
 - sample 20
 - supplying your own 14
- EXIT_FAILURE macro 63
- EXIT_SUCCESS macro 63
- external symbols
 - and generated HLASM code 4
- external variables
 - defining 12
 - initializing 12

F

- F4SA save area format
 - and AMODE 3
 - and NAB 3
- F7SA save area format
 - and AMODE 3
 - and NAB 3
- far pointers
 - ALET associations 37
 - C language constructs 36
 - constructing 37
 - copied from C string pointers 42
 - declaration 35

- far pointers (*continued*)
 - dereference 35
 - dereferencing 42
 - passing and returning 39
 - reference 35
 - setting and getting
 - _MLBUILTIN macro 37
 - built-in functions 37
- far_strcpy library function
 - data space allocation 42
- feedback xiii
- file-scope header
 - structure 5
- file-scope trailers
 - structures 12
- float.h header file 55
- fopen() library function 61
- formatted I/O 88
- free() library function 76
- function entry point markers
 - structures 9
- function headers
 - structures 8
- function property blocks
 - defining
 - under LONGNAME compiler
 - option 9
 - structures 9
- function prototypes
 - and AMODE 31
- function save area
 - chaining 13
- function save areas
 - AMODE 3
 - formats 3
 - setup 3
- functions
 - AR-mode
 - prototypes 35
 - arguments 119
 - attributes
 - AR-mode 35
 - prototypes
 - AR-mode 35

G

- global SET symbols
 - and function entry point markers 9
 - and function headers 8
 - function property blocks 9
- global variables
 - register specification 30
 - storage of 30
- GOFF HLASM option
 - and ALIAS instructions 4
 - when to specify 46
- GPRs
 - and global SET symbols 15

H

- header files
 - builtins.h 38
 - data space allocation 39
 - far-pointer management 37

- header files (*continued*)
 - stdint.h header file 61
 - string.h
 - data space allocation 39
 - strings.h
 - data space allocation 39
- heap services
 - user-replaceable 67
- hexadecimal 77
- HLASM
 - as utility
 - invoking 46
 - global SET symbols
 - values 8, 9
 - ld utility
 - invoking 47
- HLASM options
 - GOFF
 - and ALIAS instructions 4
- HLASM options
 - with LONGNAME compiler
 - option 46
- HLASM source program,
 - compiler-generated 4
 - characteristics 4
- HLASM source programs
 - file-scope headers 5
 - file-scope trailers 12
 - debug data block 12
 - function elements 8
 - structures 5

I

- IDF debugger
 - invocation 52
- insert_asm pragma
 - inserting your own assembly
 - statements 30
- integer
 - pseudo-random 85, 86
- Interactive Debug Facility (IDF)
 - generation of information 51
- inttypes.h header file 56
- IPA and HOT options
 - to build Metal C programs 48
- isalnum() library function 76, 77
- isalpha() library function 77
- isblank() library function 77, 78
- isctrl() library function 77
- isdigit() library function 77
- isgraph() library function 77
- islower() library function 77
- isprint() library function 77
- ispunct() library function 77
- isspace() library function 77
- isupper() library function 77
- isxdigit() library function 77

J

- JCL
 - assembly job step 48
 - bind job step 48
 - compilation job step 48

- JCL procedures
 - CEE.SCEEPROC data set 47
 - to build Metal C programs 47

K

- keyboard
 - navigation 129
 - PF keys 129
 - shortcut keys 129

L

- labs() library function 79
- ld command
 - building Metal C programs
 - bind options 47
- ldiv() library function 79
- length function 103
- library functions
 - far versions 38
- limits.h header file 58
- linkage conventions
 - AR-mode functions
 - ASC mode 39
 - MVS and Metal C 2
- Linkage Editor
 - TEST option and load module
 - size 51
- list form of a macro
 - specifying and using 28
- llabs() library function 79
- lldiv() library function 80
- locating storage 76
- LONGNAME compiler option 4
 - and HLASM options 46
 - and Metal C programs 47
- lowercase
 - tolower() 119

M

- mainframe
 - education x
- malloc() library function 80
- matching failure 99
- math.h header file 59
- MB_CUR_MAX macro 63
- memccpy() library function 82
- memchr() library function 82
- memcmp() library function 83
- memcpy() library function 83
- memmove() library function 84
- memory
 - allocation 86
- memory references
 - AR mode 37
- memset() library function 84
- Metal C
 - feature and benefits 2
- Metal C function descriptors 43
- Metal C programs
 - argc argv parsing 34
 - building 44
 - alternative name for "main" 43
 - assembly step 46

- Metal C programs (*continued*)
 - building (*continued*)
 - compilation step 46
 - metalc utility 46
 - z/OS UNIX System Services 46
 - IPA and HOT enablement 48
 - Example 50
 - JCL procedures to build 47
 - ld command 47
 - reentrant Metal C program 32
 - RENT option 32
- metal.h header file 60
- metalc utility
 - and HLASM source file 46
- MVS linkage conventions
 - and Metal C 2
- MYEPILOG #pragma directive
 - using 14
- MYPROLOG #pragma directive
 - using 14

N

- NAB linkage extension
 - description 3
- name encoding
 - and C symbols 4
- navigation
 - keyboard 129
- next available byte (NAB)
 - pre-allocated stack space 3
- noarmode function attribute 35
- NOTEST assembler option
 - and load module size 51
- NULL macro 60, 61
- NULL pointer 60, 61
- NULL pointer constant 63
- numbers 76

O

- object code control
 - address space control 35
 - ASC mode 35
- offsetof macro 60

P

- parameter passing
 - and AMODE 2
- parameters
 - and global SET symbols 15
- PDF documents ix
- pointers
 - storing 23
- precision argument, fprintf() family 91
- printing
 - sprintf() 88
 - vsprintf() 121
- prolog
 - user-supplied
 - global SET symbols 15
- prolog code
 - AR-mode functions 13, 39
 - default 20
 - AR-mode functions 39

- prolog code (*continued*)
 - DSA pointer 20
 - NAB pointer 20
 - primary functions 13
 - sample 18
- ptrdiff_t type in stddef header file 60

Q

- qsort() library function 85

R

- RAND_MAX macro 63
- rand_r() library function 86
- rand() library function 85
- random
 - number generator 85, 86
 - number initializer 94
 - rand_r() 86
 - rand() 85
 - srand() 94
- read-write operands, defining
 - using the + constraint 27
- reading
 - formatted 94
 - scanning 94
- realloc() library function 86
- reallocation of block size 86
- reentrancy 4
- register storage class specifier
 - register specification 30
- registers
 - access 37
 - clobbering 25
 - controlling use of 25
 - hardware access 35
 - specified as __asm operands 22
 - specifying 30
- remainder 75
- resource limits defined 58
- return values
 - AMODE 3
 - formats 3
 - setup 3

S

- save area formats
 - and AMODE 3
 - and NAB 3
- scanning
 - sscanf() 94
- SCJTSAM data set
 - epilog code sample 20
 - prolog code sample 18
- searching
 - strchr() 100
 - strings 100, 106
 - strings for tokens 110, 111
 - strspn() 107
- seed for random numbers 94
- sending to IBM
 - reader comments xiii
- SET symbols
 - and AMODE 15

- SET symbols (*continued*)
 - and DSA 15
 - and GPRs 15
 - and number of fixed parameters 15
 - and storage instructions 15
 - compiler-defined 15
 - for a user-supplied prolog 15
 - shortcut keys 129
 - size_t structure 60
 - sprintf() library function 87
 - source-level information
 - extracting 51
 - extracting in a batch environment
 - CDAASMC 51
 - for disassembly
 - suppressing 51
 - for IDF 51
 - space (white space)
 - characters
 - testing 77
 - sprintf() library function 88
 - srand() library function 94
 - sscanf() library function 94
 - stack
 - allocating space 28
 - pre-allocated stack space 3
 - standard save area format
 - and AMODE 3
 - and NAB 3
 - static variables
 - defining 12
 - mapping 12
 - stdarg.h header file 60
 - stddef.h header file 60
 - stdint.h header file 61
 - stdio.h header file 60
 - stdlib.h header file 62
 - storage
 - allocation 86
 - storage instructions
 - and global SET symbols 15
 - strcat() library function 99
 - strchr() library function 100
 - strcmp() library function 101
 - strcpy() library function 102
 - strcspn() library function 102
 - strdup() library function 103
 - streams
 - formatted I/O 94
 - string.h header file 63
 - strings
 - comparing 102, 104
 - concatenating 99, 104
 - conversions
 - to unsigned integer 116
 - copying 102, 105
 - ignoring case 101, 102
 - initializing 105
 - length of 103
 - searching 100, 106
 - strspn() 107
 - searching for tokens 110, 111
 - substring
 - locating 107
 - strings.h header file
 - data space allocation 39
 - strlen() library function 103

- strncat() library function 104
- strncmp() library function 104
- strncpy library function
 - data space allocation 39
- strncpy() library function 105
- strpbrk() library function 106
- strrchr() library function 106
- strspn() library function 107
- strstr() library function 107
- strtod() library function 108
- strtof() library function 109
- strtok_r() library function 111
- strtok() library function 110
- strtol() library function 112
- strtold() library function 113
- strtoll() library function 114
- strtoul() library function 116
- strtoull() library function 117
- syntax diagrams
 - how to read x
- syntax of format for sprintf() 89

T

- technical support xii
- TEST assembler option
 - and load module size 51
- testing 76, 77, 78
 - characters
 - white space 77
 - numbers
 - hexadecimal 77
- tokens
 - strtok_r() 111
 - strtok() 110
- tolower() library function 119
- toupper() library function 119
- typographical conventions ix

U

- uppercase
 - toupper() 119
- user interface
 - ISPF 129
 - TSO/E 129
- user-replaceable heap services 67

V

- va_arg() macro 119
- va_end() macro 119
- va_start() macro 119
- variables
 - making a C variable available to
 - HLASM 22
- vsnprintf() library function 120
- vsprintf() library function 121
- vsscanf() library function 122

Z

- z/OS Basic Skills Knowledge Center x
- z/OS UNIX System Services
 - as utility 46

- z/OS UNIX System Services (*continued*)
 - bind options 46
 - ld utility
 - bind options 47



Product Number: 5655-MCE

Printed in USA

SC27-9049-00

