

# DOMAIN System Call Reference

Order No. 007196  
Revision 01

Apollo Computer Inc.  
330 Billerica Road  
Chelmsford, MA 01824

Copyright © 1986 Apollo Computer Inc.

All rights reserved.

Printed in U.S.A.

First Printing: January 1987

This document was produced using the SCRIBE<sup>®</sup> document preparation system. (SCRIBE is a registered trademark of Unilogic, Ltd.)

APOLLO and DOMAIN are registered trademarks of Apollo Computer Inc.

AEGIS, DGR, DOMAIN/BRIDGE, DOMAIN/DFL-100, DOMAIN/DQC-100, DOMAIN/Dialogue, DOMAIN/IX, DOMAIN/Laser-26, DOMAIN/PCI, DOMAIN/SNA, D3M, DPSS, DSEE, GMR, and GPR are trademarks of Apollo Computer Inc.

Apollo Computer Inc. reserves the right to make changes in specifications and other information contained in this publication without prior notice, and the reader should in all cases consult Apollo Computer Inc. to determine whether any such changes have been made.

THE TERMS AND CONDITIONS GOVERNING THE SALE OF APOLLO COMPUTER INC. HARDWARE PRODUCTS AND THE LICENSING OF APOLLO COMPUTER INC. SOFTWARE CONSIST SOLELY OF THOSE SET FORTH IN THE WRITTEN CONTRACTS BETWEEN APOLLO COMPUTER INC. AND ITS CUSTOMERS. NO REPRESENTATION OR OTHER AFFIRMATION OF FACT CONTAINED IN THIS PUBLICATION, INCLUDING BUT NOT LIMITED TO STATEMENTS REGARDING CAPACITY, RESPONSE-TIME PERFORMANCE, SUITABILITY FOR USE OR PERFORMANCE OF PRODUCTS DESCRIBED HEREIN SHALL BE DEEMED TO BE A WARRANTY BY APOLLO COMPUTER INC. FOR ANY PURPOSE, OR GIVE RISE TO ANY LIABILITY BY APOLLO COMPUTER INC. WHATSOEVER.

IN NO EVENT SHALL APOLLO COMPUTER INC. BE LIABLE FOR ANY INCIDENTAL, INDIRECT, SPECIAL OR CONSEQUENTIAL DAMAGES WHATSOEVER (INCLUDING BUT NOT LIMITED TO LOST PROFITS) ARISING OUT OF OR RELATING TO THIS PUBLICATION OR THE INFORMATION CONTAINED IN IT, EVEN IF APOLLO COMPUTER INC. HAS BEEN ADVISED, KNEW OR SHOULD HAVE KNOWN OF THE POSSIBILITY OF SUCH DAMAGES.

THE SOFTWARE PROGRAMS DESCRIBED IN THIS DOCUMENT ARE CONFIDENTIAL INFORMATION AND PROPRIETARY PRODUCTS OF APOLLO COMPUTER INC. OR ITS LICENSORS.

## Preface

This manual consists of a section that introduces the system calls followed by sections that describe a separate operating system manager (e.g., the process manager, stream manager, and variable formatting package). The sections that describe the managers are in **alphabetical order by manager name** and consist of a description of the data types used by the manager, the syntax of the manager's programming calls, and the error messages generated by the manager. Each section is preceded by its own table of contents.

For easy organization, we have numbered the pages of by system manager. For example, the third page in the ACLM section is page ACLM-3.

You should use this manual with the programming handbooks listed under Related Documents. These programming handbooks give detailed instructions about using the programming calls.

### Audience

This manual is intended for programmers who are writing application programs using DOMAIN system calls. Readers of this manual should be familiar with FORTRAN, Pascal, or C and the operating system as described in the *DOMAIN System User's Guide*. This manual is not intended as a tutorial document, but as a reference for programmers who need to use operating system services.

### Summary of Technical Changes

This manual has been reduced in size from two volumes to one. The GM and GPR programming calls, which were previously documented here, are now documented separately. For GM calls, refer to *DOMAIN 2D Graphics Metafile Resource Call Reference*, order no. 009793; for GPR calls, refer to *DOMAIN Graphics Primitive Routines Call Reference*, order no. 007194.

In support of SR 9.5, the following new managers have been added:

CTM  
FPP  
PRF

In addition, the following calls have been added to existing managers:

MBX\_\$TIMED\_OPEN  
MS\_\$ADDMAP  
MS\_\$CRTEMP  
MS\_\$FW\_PARTIAL  
MS\_\$MK\_PERMANENT  
MS\_\$MK\_TEMPORARY  
MS\_\$NEIGHBORS  
PAD\_\$FORCE\_PROMPT  
PAD\_\$IS\_ICON

## Related Documents

The *Programming With General System Calls* handbook, order no. 005506, documents how to write programs that use standard DOMAIN system calls including the ACLM, CAL, EC2, ERROR, MTS, NAME, PAD, PBUFS, PFM, PGM, PM, PROC1, PROC2, RWS, SIO, STREAM, TIME, TONE, TPAD, and VFMT calls.

The *Programming With System Calls for Interprocess Communication* handbook, order no. 005696, documents how to write programs that use the DOMAIN interprocess facilities including the MBX, MS, IPC, MUTEX, and EC2 calls.

## Documentation Conventions

Unless otherwise noted in the text, this manual uses the following symbolic conventions.

UPPERCASE	Uppercase words or characters in formats and command descriptions represent keywords that you must use literally.
lowercase	Lowercase words or characters in formats and command descriptions represent values that you must supply.
[    ]	Square brackets enclose optional items.
{    }	Braces enclose a list from which you must choose an item.
	A vertical bar separates items in a list of choices.
<   >	Angle brackets enclose the name of a key on the keyboard.
CTRL/Z	The notation CTRL/ followed by the name of a key indicates a control character sequence. Hold down <CTRL> while you type the character.
...	Horizontal ellipsis points indicate that you can repeat the preceding item one or more times.
.	Vertical ellipsis points mean that we have omitted irrelevant parts of a figure or example.

## Problems, Questions, and Suggestions

We appreciate comments from the people who use our system. In order to make it easy for you to communicate with us, we provide the User Change Request (UCR) system for software-related comments, and the Reader's Response form for documentation comments. By using these formal channels, you make it easy for us to respond to your comments.

You can get more information about how to submit a UCR by consulting the *DOMAIN System Command Reference* manual. Refer to the CRUCR (Create User Change Request) Shell command description. You can view the same description on-line by typing:

```
$ HELP CRUCR <RETURN>
```

For documentation comments, a Reader's Response form is located at the back of each manual.

# Contents

INTRODUCTION	INTRO-1
ACLM	ACLM-1
CAL	CAL-1
CTM	CTM-1
EC2	EC2-1
ERROR	ERROR-1
FPP	FPP-1
GMF	GMF-1
IOS	IOS-1
IOS_DIR	IOS_DIR-1
IPC	IPC-1
MBX	MBX-1
MS	MS-1
MTS	MTS-1
MUTEX	MUTEX-1
NAME	NAME-1
PAD	PAD-1
PBUFS	PBUFS-1
PFM	PFM-1
PGM	PGM-1
PM	PM-1
PRF	PRF-1
PROC1	PROC1-1
PROC2	PROC2-1
RWS	RWS-1

SIO

SIO-1

SMD

SMD-1

STREAM

STREAM-1

TIME

TIME-1

STONE

STONE-1

TPAD

TPAD-1

VEC

VEC-1

VFMT

VFMT-1

# Introduction

This introductory section describes the DOMAIN system insert files and the format of the information found in the sections that follow. Each of these sections consists of a description of the data types used by a system manager, the syntax of the manager's programming calls, and the error messages generated by the system manager. We have arranged the sections of this manual **alphabetically**, by system manager name.

## DOMAIN Insert Files

The DOMAIN system provides insert files that define data types, constants, values, and routine declarations. The insert files also define the exact form of each system call or routine. (Even the FORTRAN version does this using comments, although the FORTRAN compiler doesn't check the forms that you use.)

The DOMAIN system routines are divided, by function, into several subsystems. Each subsystem is controlled by a system manager. The routines of each subsystem are prefixed for easy identification. A subsystem prefix consists of a number of identifying characters followed by the special underscore and dollar-sign characters, "\_\$." For example, the routines that perform stream functions are prefixed with `STREAM_$`. These subsystem prefixes are also used to distinguish DOMAIN data types and constants that are used by the subsystem routines.

Insert files are located in the directory `/SYS/INS/`. There is one insert file per subsystem for each programming language. Include the appropriate insert file for your programming language. For example, if you are using error routines in a Pascal program, you include the insert file, `/SYS/INS/ERROR.INS.PAS`. Using the same routines in a FORTRAN program, you include `/SYS/INS/ERROR.INS.FTN`. All insert files are specified using the syntax

```
/SYS/INS/subsystem-prefix.INS.language-abbreviation
```

where the language abbreviation is PAS (Pascal), FTN (FORTRAN), or C (C). The listing on the next page shows all the available insert files.

In addition to including required subsystem insert files in a program, you must always include the BASE insert file for your programming language. You specify BASE insert files using the syntax

```
/SYS/INS/BASE.INS.language-abbreviation
```

These files contain some basic definitions that a number of subsystem routines use.

## Summary of Insert Files

Insert File	Operating System Component
/SYS/INS/BASE.INS.lan	Base definitions -- must always be included
/SYS/INS/ACLM.INS.lan	Access control list manager
/SYS/INS/CAL.INS.lan	Calendar
/SYS/INS/CTM.INS.lan	Color Table Manager
/SYS/INS/ERROR.INS.lan	Error reporting
/SYS/INS/EC2.INS.lan	Eventcount
/SYS/INS/FPP.INS.lan	Floating Point Package
/SYS/INS/GMF.INS.lan	Graphics Map Files
/SYS/INS/GPR.INS.lan	Graphics Primitives
/SYS/INS/IOS.INS.lan	I/O Switch Manager
/SYS/INS/IOS_DIR.INS.lan	I/O Switch Directory
/SYS/INS/IPC.INS.lan	Interprocess communications datagrams
/SYS/INS/KBD.INS.lan	[Useful constants for keyboard keys]
/SYS/INS/MBX.INS.lan	Mailbox manager
/SYS/INS/MS.INS.lan	Mapping server
/SYS/INS/MTS.INS.lan	Magtape/streams interface
/SYS/INS/MUTEX.INS.lan	Mutual exclusion lock manager
/SYS/INS/NAME.INS.lan	Naming server
/SYS/INS/PAD.INS.lan	Display Manager
/SYS/INS/PBUFS.INS.lan	Paste buffer manager
/SYS/INS/PFM.INS.lan	Process fault manager
/SYS/INS/PGM.INS.lan	Program manager
/SYS/INS/PM.INS.lan	User process routines
/SYS/INS/PRF.INS.PAS	Print File Manager
/SYS/INS/PROC1.INS.PAS	Process manager (Pascal only)
/SYS/INS/PROC2.INS.lan	User process manager
/SYS/INS/RWS.INS.lan	Read/write storage manager
/SYS/INS/SIO.INS.lan	Serial I/O
/SYS/INS/SMD.INS.lan	Display driver
/SYS/INS/STREAMS.INS.lan	Stream manager
/SYS/INS/TIME.INS.lan	Time
/SYS/INS/TONE.lan	Speaker
/SYS/INS/TPAD.INS.lan	Touchpad manager
/SYS/INS/VEC.INS.lan	Vector arithmetic
/SYS/INS/VFMT.INS.lan	Variable formatter

The suffix ".lan" varies with the high-level language that you're using; it is either ".FTN", ".PAS", or ".C".

### Organizational Information

This introductory section is followed by sections for each subsystem. The material for each subsystem is organized into the following three parts:

1. Detailed data type information (including illustrations of records for the use of FORTRAN programmers).
2. Full description of each system call.
3. List of possible error messages.



## Data Type Sections

A subsystem's data type section precedes the subsystem's individual call descriptions. Each data type section describes the predefined constants and data types for a subsystem. These descriptions include an atomic data type translation (i.e., TIME\_\$REL\_ABS\_T = 4-byte integer) for use by FORTRAN programmers, as well as a brief description of the type's purpose. Where applicable, any predefined values associated with the type are listed and described. Following is an example of a data type description for the TIME\_\$REL\_ABS\_T type:

TIME\_\$REL\_ABS\_T

A 2-byte integer. Indicator of type of time. One of the following predefined values:

TIME\_\$RELATIVE  
Relative time.

TIME\_\$ABSOLUTE  
Absolute time.

In addition, the record data types are illustrated in detail. Primarily, we have geared these illustrations to FORTRAN programmers who need to construct record-like structures, but we've designed the illustrations to convey as much information as possible for all programmers. Each record type illustration:

- Shows FORTRAN programmers the structure of the record that they must construct using standard FORTRAN data type statements. The illustrations show the size and type of each field.
- Describes the fields that make up the record.
- Lists the byte offsets for each field. These offsets are used to access fields individually.
- Indicates whether any fields of the record are, in turn, predefined records.

The following is the description and illustration of the CAL\_\$TIMEDATE\_REC\_T predefined record:

CAL\_\$TIMEDATE\_REC\_T

Readable time format. The diagram below illustrates the CAL\_\$TIMEDATE\_REC\_T data type:

predefined type	byte: offset	field name
	0:	integer year
	2:	integer month
	4:	integer day
	6:	integer hour
	8:	integer minute
	10:	integer second

Field Description:

year  
Integer representing the year.

month  
Integer representing the month.

day  
Integer representing the day.

hour  
Integer representing the hour (24 hr. format).

minute  
Integer representing the minute.

second  
Integer representing the second.

FORTRAN programmers should note that a Pascal variant record is a record structure that may be interpreted differently depending on usage. In the case of variant records, as many illustrations will appear as are necessary to show the number of interpretations.

## System Call Descriptions

We have listed the system call descriptions alphabetically for quick reference. Each system call description contains:

- An abstract of the call's function.
- The order of call parameters.
- A brief description of each parameter.
- A description of the call's function and use.

These descriptions are standardized to make referencing the material as quick as possible.

Each parameter description begins with a phrase describing the parameter. If the parameter can be declared using a predefined data type, the descriptive phrase is followed by the phrase ",in XXX format" where XXX is the predefined data type. Pascal or C programmers, look for this phrase to determine how to declare a parameter.

FORTRAN programmers, use the second sentence of each parameter description for the same purpose. The second sentence describes the data type in atomic terms that you can use, such as "This is a 2-byte integer." In complex cases, FORTRAN programmers are referenced to the respective subsystem's data type section.

The rest of a parameter description describes the use of the parameter and the values it may hold.

The following is an example of a parameter description:

access

New access mode, in MS\_\$ACC\_MODE\_T format. This is a 2-byte integer.  
Specify only one of the following predefined values:

MS_\$R	Read access.
MS_\$WR	Read and write access.
MS_\$RIW	Read with intent to write.

An object which is locked MS\_\$RIW may not be changed to MS\_\$R.

## Error Sections

Each error section lists the status codes that may be returned by subsystem calls. The following information appears for each error:

- Predefined constant for the status code.
- Text associated with the error.



# ACLM

## Table of Contents

ACLM\_ \$DOWN  
ACLM\_ \$UP

ACLM-2  
ACLM-3



## ACLM

The ACLM (Access Control List Manager) programming calls control the protected subsystem privilege level. This section describes their call syntax. The ACLM calls do not use special data types or produce unique error messages. Refer to the Introduction at the beginning of this manual for a description of the call syntax format.

ACLM\_\$DOWN

ACLM\_\$DOWN

Deasserts a program's subsystem manager rights.

**FORMAT**

ACLM\_\$DOWN

**USAGE**

This call deasserts a program's rights to gain access to an object in a protected subsystem, which were asserted by a previous call to ACLM\_\$UP.



**ACLM\_\$UP**

Asserts a program's subsystem manager rights.

**FORMAT**

ACLM\_\$UP

**USAGE**

This call asserts a program's rights to gain access to an object in a protected subsystem, until a corresponding call to ACLM\_\$DOWN is made.

Access Control List manager (ACLM) calls are used by subsystem manager programs in DOMAIN protected subsystems. A protected subsystem is a feature of the operating system that ensures that access to certain objects is restricted to certain programs which are called the managers of the subsystem that contains those objects.

In fact, even a subsystem manager, which has the right to gain access to the protected objects, must call ACLM\_\$UP to assert its rights before it can actually use a protected object. Calling ACLM\_\$DOWN deasserts a program's rights as a subsystem manager.

We recommend that you activate your rights as a subsystem manager for the minimum amount of time you will need them using ACLM\_\$UP and ACLM\_\$DOWN to bracket high-level statements or functions for which you need these rights. This ensures against inadvertent use of the protected objects.

Calling ACLM\_\$UP increments a counter in your process; calling ACLM\_\$DOWN decrements it. Subsystem manager operations are enabled whenever the counter is nonzero. Having a counter instead of a flag ensures that if one routine enables subsystem manager rights, and calls a routine that enables and disables subsystem manager rights, the calling routine does not inadvertently lose its rights.

Calling ACLM\_\$UP obtains all the subsystem rights to which you are entitled. If a program that is not a subsystem manager calls ACLM\_\$UP, it produces no effect, but does not return an error. Likewise, calling ACLM\_\$DOWN when subsystem manager rights were already deasserted has no effect.

Protected subsystems and the reasons for using them are discussed completely in the *Administering Your DOMAIN System*.



# CAL

## Table of Contents

Data Types	CAL-2
CAL_\$ADD_CLOCK	CAL-6
CAL_\$APPLY_LOCAL_OFFSET	CAL-7
CAL_\$CLOCK_TO_SEC	CAL-8
CAL_\$CMP_CLOCK	CAL-9
CAL_\$DECODE_ASCII_DATE	CAL-10
CAL_\$DECODE_ASCII_TIME	CAL-11
CAL_\$DECODE_ASCII_TZDIF	CAL-12
CAL_\$DECODE_LOCAL_TIME	CAL-14
CAL_\$DECODE_TIME	CAL-15
CAL_\$ENCODE_TIME	CAL-16
CAL_\$FLOAT_CLOCK	CAL-17
CAL_\$GET_INFO	CAL-18
CAL_\$GET_LOCAL_TIME	CAL-19
CAL_\$REMOVE_LOCAL_OFFSET	CAL-20
CAL_\$SEC_TO_CLOCK	CAL-21
CAL_\$SUB_CLOCK	CAL-22
CAL_\$WEEKDAY	CAL-23
CAL_\$WRITE_TIMEZONE	CAL-24
Errors	CAL-25



# CAL

The CAL (Calendar) programming calls provide the calendar maintenance services. This section describes their data types, call syntax, and error codes. Refer to the Introduction at the beginning of this manual for a description of data type diagrams and call syntax format.

CAL DATA TYPES

CONSTANTS

CAL\_\$STRING\_SIZE            80            Size of an ASCII string.

DATA TYPES

CAL\_\$DPVAL\_T                            A double-precision floating point value. A 2-element array of 4-byte integers. (REAL\*8 for FORTRAN programs.)

CAL\_\$STRING\_T                          An array of up to CAL\_\$STRING\_SIZE (80) characters. An ASCII string.

CAL\_\$TIMEDATE\_REC\_T

6 integer, readable time format. The diagram below illustrates the cal\_\$timedate\_rec\_t data type:

predefined type	byte: offset	field name
	0:	integer year
	2:	integer month
	4:	integer day
	6:	integer hour
	8:	integer minute
	10:	integer second

Field Description:

**year**  
Integer representing the year.

**month**  
Integer representing the month.

**day**  
Integer representing the day.

**hour**  
Integer representing the hour (24 hr.).

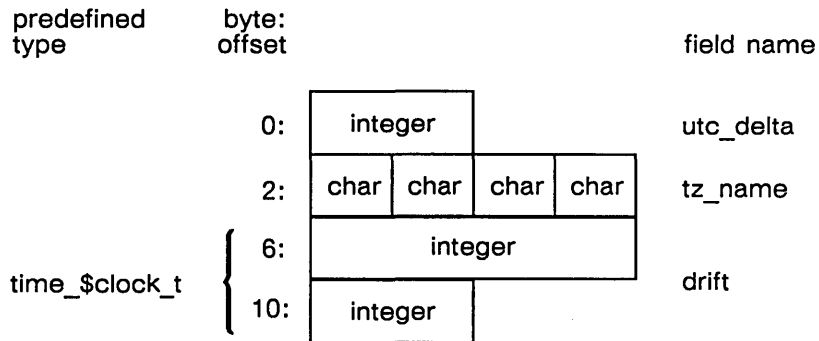
**minute**  
Integer representing the minute.

**second**  
Integer representing the second.

CAL DATA TYPES

CAL\_\$TIMEZONE\_REC\_T

Specifies time difference and timezone name. The diagram below illustrates the cal\_\$timezone\_rec\_t data type:



Field Description:

utc\_delta  
Number of minutes difference from UTC.

tz\_name  
Time zone name.

drift  
Drift adjustment.

CAL\_\$TZ\_NAME\_T

An array of up to 4 characters. Time zone name.

CAL\_\$WEEKDAY\_T

A 2-byte integer. Specifies the day of the week. One of the following predefined values:

CAL\_\$SUN  
Sunday

CAL\_\$MON  
Monday

CAL\_\$TUE  
Tuesday

CAL\_\$WED  
Wednesday

CAL\_\$THU  
Thursday

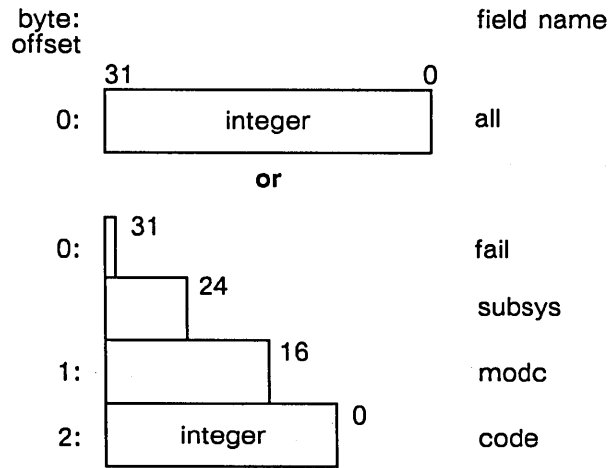
CAL\_\$FRI  
Friday

CAL\_\$SAT  
Saturday



STATUS\_\$T

A status code. The diagram below illustrates the STATUS\_\$T data type:



Field Description:

all

All 32 bits in the status code.

fail

The fail bit. If this bit is set, the error was not within the scope of the module invoked, but occurred within a lower-level module (bit 31).

subsys

The subsystem that encountered the error (bits 24 - 30).

modc

The module that encountered the error (bits 16 - 23).

code

A signed number that identifies the type of error that occurred (bits 0 - 15).

CAL\_\$ADD\_CLOCK

CAL\_\$ADD\_CLOCK

Computes the sum of two times.

### FORMAT

CAL\_\$ADD\_CLOCK (clock1, clock2)

### INPUT/OUTPUT PARAMETERS

#### clock1

Upon input           The Coordinated Universal Time clock value to be added to clock2, in TIME\_\$CLOCK\_T format. This data type is 6 bytes long. See the CAL Data Types section for more information.

Upon output           The sum of clock1 and clock2, in TIME\_\$CLOCK\_T format. This data type is 6 bytes long. See the CAL Data Types section for more information.

### INPUT PARAMETERS

#### clock2

The Coordinated Universal Time clock value to be added to clock1, in TIME\_\$CLOCK\_T format. This data type is 6 bytes long. See the CAL Data Types for more information.

**CAL\_\$APPLY\_LOCAL\_OFFSET**

Computes the local time from a UTC time.

**FORMAT**

CAL\_\$APPLY\_LOCAL\_OFFSET (clock)

**INPUT/OUTPUT PARAMETERS****clock**

Upon input            Coordinated Universal Time clock value to which a local time zone offset will be added, in TIME\_\$CLOCK\_T format. This data type is 6 bytes long. See the CAL Data Types section for more information.

Upon output           Adjusted clock value, representing the local time equivalent of the input parameter, in TIME\_\$CLOCK\_T format. This data type is 6 bytes long. See the CAL Data Types section for more information.

**USAGE**

CAL\_\$APPLY\_LOCAL\_OFFSET adds the local time zone offset to the supplied clock value.

To set the local time zone offset, you may either execute the Shell command TZ (TIME\_ZONE) as described in the *DOMAIN System Command Reference Manual*, or you may use the CAL\_\$WRITE\_TIMEZONE procedure.

CAL\_\$CLOCK\_TO\_SEC

CAL\_\$CLOCK\_TO\_SEC

Converts system clock units to seconds.

### FORMAT

seconds = CAL\_\$CLOCK\_TO\_SEC (clock)

### RETURN VALUE

#### seconds

The computed equivalent of clock, in whole seconds. This is a 4-byte integer.

If the number of seconds calculated from the input value does not represent a whole number, the fractional portion is truncated.

### INPUT PARAMETERS

#### clock

The value to be converted, in TIME\_\$CLOCK\_T format. This data type is 6 bytes long. See the CAL Data Types section for more information.

### USAGE

CAL\_\$CLOCK\_TO\_SEC converts a value in system clock representation (UTC) to an equivalent value in whole seconds.

The system clock value represents a time in units of 4 microseconds.

**CAL\_\$CMP\_CLOCK**

Compares the values of two times.

**FORMAT**

integer = CAL\_\$CMP\_CLOCK (clock1, clock2)

**RETURN VALUE****integer**

The result of the logical compare of clock1 to clock2. This is a 2-byte integer.

Integer  
returned

1	if	clock1 > clock2
0	if	clock1 = clock2
-1	if	clock1 < clock2

**INPUT PARAMETERS****clock1**

The Coordinated Universal Time clock value to be compared to clock2, in TIME\_\$CLOCK\_T format. This data type is 6 bytes long. See the CAL Data Types section for more information.

**clock2**

The Coordinated Universal Time clock value to be compared to clock1, in TIME\_\$CLOCK\_T format. This data type is 6 bytes long. See the CAL Data Types for more information.

CAL\_\$DECODE\_ASCII\_DATE

CAL\_\$DECODE\_ASCII\_DATE

Decodes an ASCII string containing a date specification.

### FORMAT

CAL\_\$DECODE\_ASCII\_DATE (string, stringlength, year, month, day, status)

### INPUT PARAMETERS

#### string

An ASCII character string, of length "stringlength" and in the form "year/month/day". This is an array of up to 80 characters.

#### stringlength

The number of characters in the string. This is a 2-byte integer.

### OUTPUT PARAMETERS

#### year

The year decoded from the string. This is a 2-byte integer.

#### month

The month decoded from the string. This is a 2-byte integer.

#### day

The day decoded from the string. This is a 2-byte integer.

#### status

Completion status, in STATUS\_\$T format. This data type is 4 bytes long. See the CAL Data Types section for more information.

Possible values are:

CAL\_\$BAD\_SYNTAX

The string provided is not in "year/month/day" format.

CAL\_\$EMPTY\_STRING

The string length is zero or the string contains only spaces.

CAL\_\$OUT\_OF\_RANGE

The value for month or day is invalid.

### USAGE

CAL\_\$DECODE\_ASCII\_DATE translates the ASCII date in the supplied string into three integers representing the year, month, and day. The string must contain a year, a month, and a day, separated by slashes.

If a year between 80 and 99 is specified, CAL\_\$DECODE\_ASCII\_DATE adds 1900 to it before returning. If a year between 0 and 79 is specified, 2000 is added.

Leading and trailing spaces are ignored.

**CAL\_\$DECODE\_ASCII\_TIME**

Translates an ASCII string containing a time into integers.

**FORMAT**

CAL\_\$DECODE\_ASCII\_TIME (string, stringlength, hour, minute, second, status)

**INPUT PARAMETERS****string**

An ASCII character string of length "stringlength" in the form "hour:minute" or "hour:minute:second", in 24-hour format. This is an array of up to 80 characters.

**stringlength**

The number of characters in the string. This is a 2 -byte integer.

**OUTPUT PARAMETERS****hour**

The hour decoded from the string. This is a 2-byte integer.

**minute**

The minute decoded from the string. This is a 2-byte integer.

**second**

The second decoded from the string. This is a 2-byte integer.

**status**

Completion status, in STATUS\_\$T format. This data type is 4 bytes long. See the CAL Data Types section for more information. Possible values are:

**CAL\_\$BAD\_SYNTAX**

The string is not in either "hour:minute:second" or "hour:minute" format.

**CAL\_\$EMPTY\_STRING**

String length is zero or the string contains only spaces.

**CAL\_\$OUT\_OF\_RANGE**

The supplied ASCII value for hour, minute, or second is invalid.

**USAGE**

CAL\_\$DECODE\_ASCII\_TIME translates the ASCII string into three integers, representing hours, minutes, and seconds.

If the string specifies only hours and minutes, the value returned for seconds is zero.

Leading and trailing spaces are ignored.

CAL\_\$DECODE\_ASCII\_TZDIF

CAL\_\$DECODE\_ASCII\_TZDIF

Translates an ASCII string specifying a time zone into an offset from UTC.

## FORMAT

CAL\_\$DECODE\_ASCII\_TZDIF (string, stringlength, tz-dif, tz-name, status)

## INPUT PARAMETERS

### string

An ASCII string containing a time zone name or time zone difference. This is an array of up to 80 characters.

A time zone name is one of the following strings: (EST) Eastern Standard Time, (EDT) Eastern Daylight Time, (CST) Central Standard Time, (CDT) Central Daylight Time, (MST) Mountain Standard Time, (MDT) Mountain Daylight Time, (PST) Pacific Standard Time, (PDT) Pacific Daylight Time, (GMT) Greenwich Mean Time, and (UTC) Coordinated Universal Time. These are the eight standard U.S. time zone names, plus those for Greenwich Mean Time and Coordinated Universal Time.

A time zone difference is a value which, when added to Coordinated Universal Time, produces the local time. Specify a time zone difference in the following form:

[ + | - ]hour:minute

The hour must be a number between 0 and 12; the minute must be 0 or 30. The sign is optional. For example, Eastern Daylight Time may be represented as -4:00.

### stringlength

The number of characters in the string. This is a 2-byte integer.

## OUTPUT PARAMETERS

### tz-dif

The difference, in minutes, between the time zone specified in string and UTC. This is a 2-byte integer.

The value of tz-dif is negative for time zones west of the Greenwich Meridian and positive for time zones east of the Greenwich Meridian.

### tz-name

The time zone name, in CAL\_\$TZ\_NAME\_T format. This is an array of up to 4 characters.

If the ASCII string contains a time zone name, this procedure returns that name in tz-name. If the ASCII string contains a time zone difference, this procedure returns spaces in tz-name.



**status**

Completion status, in STATUS\_\$T format. This data type is 4 bytes long. See the CAL Data Types section for more information. Possible values are:

CAL\_\$EMPTY\_STRING

String length is zero or the string contains only spaces.

CAL\_\$UNKNOWN\_TIMEZONE

The string contains a time zone name that is unknown to this procedure.

CAL\_\$BAD\_SYNTAX

The string appears to contain a time zone difference but is syntactically incorrect.

CAL\_\$INVALID\_TZDIF

The string contains a time zone difference, but the number of hours is greater than 12 or the number of minutes is not 0 or 30.

**USAGE**

CAL\_\$DECODE\_ASCII\_TZDIF translates the supplied ASCII string into an offset from UTC, in units of minutes. The ASCII string can contain a time zone name or a time zone difference.

CAL\_\$DECODE\_LOCAL\_TIME

CAL\_\$DECODE\_LOCAL\_TIME

Returns the local time in integer format.

#### FORMAT

CAL\_\$DECODE\_LOCAL\_TIME (decoded\_clock)

#### OUTPUT PARAMETERS

##### decoded\_clock

The local time, in CAL\_\$TIMEDATE\_REC\_T format. This data type is 12 bytes long. See the CAL Data Types section for more information.

#### USAGE

CAL\_\$DECODE\_LOCAL\_TIME returns the local time in "year, month, day, hour, minute, second" format.

**CAL\_\$DECODE\_TIME**

Translates an internal system clock value into a readable date and time.

**FORMAT**

CAL\_\$DECODE\_TIME (clock, decoded\_clock)

**INPUT PARAMETERS****clock**

The value to be translated, in TIME\_\$CLOCK\_T format. This data type is 6 bytes long. See the CAL Data Types section for more information.

**OUTPUT PARAMETERS****decoded\_clock**

A date and time, in CAL\_\$TIMEDATE\_REC\_T format. This data type is 12 bytes long. See the CAL Data Types section for more information.

**USAGE**

CAL\_\$DECODE\_TIME translates a time in TIME\_\$CLOCK\_T form into CAL\_\$TIMEDATE\_REC\_T ("year, month, day, hour, minute, second") form.

This routine translates clock values, such as those returned from the TIME\_\$CLOCK, CAL\_\$GET\_LOCAL\_TIME, CAL\_\$APPLY\_LOCAL\_OFFSET, and CAL\_\$ENCODE TIME routines.

CAL\_\$ENCODE\_TIME

CAL\_\$ENCODE\_TIME

Translates a date and time from integer format into a system clock representation.

#### FORMAT

CAL\_\$ENCODE\_TIME (decoded\_clock, clock)

#### INPUT PARAMETERS

##### decoded\_clock

A date and time, in CAL\_\$TIMEDATE\_REC\_T format. This data type is 12 bytes long. See the CAL Data Types section for more information.

#### OUTPUT PARAMETERS

##### clock

The system clock equivalent value of decoded\_clock, in TIME\_\$CLOCK\_T format. This data type is 6 bytes long. See the CAL Data Types section for more information.

#### USAGE

CAL\_\$ENCODE\_TIME translates the date and time specified by decoded\_clock into the equivalent system representation.

**CAL\_\$FLOAT\_CLOCK**

Converts a system clock representation to the equivalent number of seconds, in double-precision floating-point format.

**FORMAT**

CAL\_\$FLOAT\_CLOCK (clock, float\_seconds)

**INPUT PARAMETERS****clock**

The value to be converted, in TIME\_\$CLOCK\_T format. This data type is 6 bytes long. See the CAL Data Types section for more information.

**OUTPUT PARAMETERS****float\_seconds**

The converted value of clock, in seconds. This is in double-precision floating-point (REAL\*8) format.

**USAGE**

CAL\_\$FLOAT\_CLOCK converts a clock value in UTC format to the equivalent number of seconds expressed as a double-precision floating-point number.

Unlike CAL\_\$CLOCK\_TO\_SEC, CAL\_\$FLOAT\_CLOCK does not truncate the fractional portion of the conversion.

CAL\_\$GET\_INFO

CAL\_\$GET\_INFO

Returns local time zone information.

#### FORMAT

CAL\_\$GET\_INFO (timezone\_info)

#### OUTPUT PARAMETERS

##### timezone\_info

A record containing the name of the local time zone and its offset from UTC, in CAL\_\$TIMEZONE\_REC\_T format. This data type is 12 bytes long. See the CAL Data Types section for more information.

#### USAGE

CAL\_\$GET\_INFO returns the name of the local time zone and the difference between local time and Coordinated Universal Time (UTC).

**CAL\_\$GET\_LOCAL\_TIME**

Returns the current local time in system clock representation.

**FORMAT**

CAL\_\$GET\_LOCAL\_TIME (clock)

**OUTPUT PARAMETERS**

**clock**

The current local time, in TIME\_\$CLOCK\_T format. This data type is 6 bytes long. See the CAL Data Types section for more information.

This is the number of 4-microsecond periods that have elapsed since January 1, 1980, 00:00.

CAL\_\$REMOVE\_LOCAL\_OFFSET

CAL\_\$REMOVE\_LOCAL\_OFFSET

Computes the UTC time from local time.

## FORMAT

CAL\_\$APPLY\_LOCAL\_OFFSET (clock)

## INPUT/OUTPUT PARAMETERS

### clock

Upon input Local time from which the local time offset will be removed, in TIME\_\$CLOCK\_T format. This data type is 6 bytes long. See the CAL Data Types section for more information.

Upon output Adjusted clock value, representing the UTC equivalent of the input parameter, in TIME\_\$CLOCK\_T format. This data type is 6 bytes long. See the CAL Data Types section for more information.

## USAGE

CAL\_\$REMOVE\_LOCAL\_OFFSET subtracts the local time zone offset from the supplied clock value.

To set the local time zone offset, you may either execute the Shell command TZ (TIME\_ZONE) as described in the *DOMAIN System Command Reference*, or you may use the CAL\_\$WRITE\_TIMEZONE procedure.



CAL\_\$SEC\_TO\_CLOCK

Converts seconds to system clock units.

**FORMAT**

CAL\_\$SEC\_TO\_CLOCK (seconds, clock)

**INPUT PARAMETERS**

**seconds**

The value to be converted. This is a 4-byte integer.

**OUTPUT PARAMETERS**

**clock**

The computed equivalent of seconds, in TIME\_\$CLOCK\_T format. This data type is 6 bytes long. See the CAL Data Types section for more information.

**USAGE**

CAL\_\$SEC\_TO\_CLOCK converts a value representing seconds to an equivalent value in 4-microsecond units.

No overflow detection is performed.

CAL\_\$SUB\_CLOCK

CAL\_\$SUB\_CLOCK

Subtracts the values of two times.

#### FORMAT

value = CAL\_\$SUB\_CLOCK (clock1, clock2)

#### RETURN VALUE

##### value

The Boolean result of the subtraction of clock2 from clock1. The returned value is TRUE if the result is  $\geq 0$ .

#### INPUT/OUTPUT PARAMETERS

##### clock1

Upon input        The Coordinated Universal Time clock value from which clock2 is subtracted, in TIME\_\$CLOCK\_T format. This data type is 6 bytes long. See the CAL Data Types section for more information.

Upon output       The difference between clock1 and clock2, in TIME\_\$CLOCK\_T format. This data type is 6 bytes long. See the CAL Data Types section for more information.

#### INPUT PARAMETERS

##### clock2

The Coordinated Universal Time clock value to be subtracted from clock1, in TIME\_\$CLOCK\_T format. This data type is 6 bytes long. See the CAL Data Types section for more information.

**CAL\_ \$WEEKDAY**

Computes the day of the week given a year, month, and day.

**FORMAT**

weekday = CAL\_ \$WEEKDAY (year, month, day)

**RETURN VALUE****weekday**

The computed day of the week, in CAL\_ \$WEEKDAY\_ T format. This is a 2-byte integer. Returns one of the following predefined values:

CAL\_ \$SUN, CAL\_ \$MON, CAL\_ \$TUE, CAL\_ \$WED,  
CAL\_ \$THU, CAL\_ \$FRI, CAL\_ \$SAT.

Their ordinal values are 0 through 6.

**INPUT PARAMETERS****year**

The year for which the weekday is desired. This is a 2-byte integer.

**month**

The month for which the weekday is desired. This is a 2-byte integer.

**day**

The day of the month for which the weekday is desired. This is a 2-byte integer.

**USAGE**

CAL\_ \$WEEKDAY computes the day of the week for any Gregorian date.

CAL\_\$WRITE\_TIMEZONE

CAL\_\$WRITE\_TIMEZONE

Writes local time zone information onto the boot volume.

## FORMAT

CAL\_\$WRITE\_TIMEZONE (timezone\_info, status)

## INPUT PARAMETERS

### timezone\_info

The time zone information to be recorded, in CAL\_\$TIMEZONE\_REC\_T format. This data type is 12 bytes long. See the CAL Data Types section for more information.

The supplied time zone information includes the name of the time zone and its offset from UTC.

## OUTPUT PARAMETERS

### status

Completion status, in STATUS\_\$T format. This data type is 4 bytes long. See the CAL Data Types section for more information.

## USAGE

CAL\_\$WRITE\_TIMEZONE writes the supplied time zone information onto the logical disk volume from which the operating system was started.

This procedure is invalid on a diskless node, and returns a nonzero status.

The time zone information written by this procedure is used by subsequent calls to CAL\_\$DECODE\_LOCAL\_TIME, CAL\_\$GET\_LOCAL\_TIME, CAL\_\$APPLY\_LOCAL\_OFFSET, and CAL\_\$GET\_INFO.

A nonzero status indicates a system problem in reading or writing the volume.

**ERRORS****STATUS\_\$OK**

Successful completion.

**CAL\_\$BAD\_SYNTAX**

Invalid syntax for date or time specification.

**CAL\_\$EMPTY\_STRING**

An empty string was passed to a decode routine.

**CAL\_\$INVALID\_TZDIF**

Invalid time-zone difference.

**CAL\_\$OUT\_OF\_RANGE**

Date or time specification invalid.

**CAL\_\$UNKNOWN\_TIMEZONE**

Timezone specified is unknown.

○

○

○

○

○

# CTM

## Table of Contents

Data Types	CTM-2
CTM_\$ALLOC_PV	CTM-4
CTM_\$FIND_COLOR	CTM-8
CTM_\$INC_USE_COUNT	CTM-9
CTM_\$MARK_READ_ONLY	CTM-10
CTM_\$RELEASE_PV	CTM-11
Errors	CTM-12





# CTM

The CTM (Color Table Manager) programming calls allow different GPR applications running on the same node to share the color map without interfering with each other. This section describes their data types, call syntax, and error codes. Refer to the Introduction at the beginning of this manual for a description of data type diagrams and call syntax format.

**DATA TYPES**

GPR\_\$PIXEL\_VALUE\_T

A 4-byte integer. Defines an index into the color map to identify the color of an individual pixel.

CMT\_\$PIXEL\_VALUE\_VECTOR\_T

An array of GPR\_\$PIXEL\_VALUE\_T that stores multiple pixel values.

CTM\_\$ALLOC\_OPTIONS\_T

A 2-byte integer. Identifies the type of pixel values that can be affected by a CTM call. Specify only one of the following predefined values:

NONE

If you specify 0 for FORTRAN/C or [ ] for Pascal, then any available pixel value is returned.

CTM\_\$CONTIGUOUS

The allocated pixel values must be contiguous. In this case, only the smallest pixel value in the range is returned.

CTM\_\$ZERO\_ONLY

Only pixel values with zero in a given plane will be allocated. If the plane is greater than or equal to zero, then this parameter specifies which plane must have zero bits. If the plane is less than 0, then the CTM package will choose a plane and return it

CTM\_\$ONE\_ONLY


This is similar to CTM\_\$ZERO\_ONLY except that the plane in question must contain only ones.

CTM\_\$BOTH

Pairs of pixel values will be allocated where the two members of each pair differ only in the given plane. If the plane is less than zero, the CTM package selects a plane for you.

STATUS\_\$T

A status code. The diagram below illustrates the STATUS\_\$T data type:



Field Description:

all

All 32 bits in the status code.

fail

The fail bit. If this bit is set, the error was not within the scope of the module invoked, but occurred within a lower-level module (bit 31).

subsys

The subsystem that encountered the error (bits 24 - 30).

modc

The module that encountered the error (bits 16 - 23).

code

A signed number that identifies the type of error that occurred (bits 0 - 15).

CTM\_\$ALLOC\_PV

CTM\_\$ALLOC\_PV

Allocates available pixel values and sets their use counts to one.

## FORMAT

CTM\_\$ALLOC\_PV (count, opt, plane, pixel\_values, status)

## INPUT PARAMETERS

### count

Specifies the number of pixel values (or pixel value pairs). This parameter is a 2-byte integer.

### opt

Specifies the constraints placed on pixel values to be allocated, in CTM\_\$ALLOC\_OPTIONS\_T format. This parameter is a 2-byte integer. Specify only one of the following values:

**NONE** If you specify 0 for FORTRAN /C or [ ] for Pascal, then any available pixel values are returned.

**CTM\_\$CONTIGUOUS**

The allocated pixel values must be contiguous. In this case, only the smallest pixel value in the range is returned.

**CTM\_\$ZERO\_ONLY**

Only pixel values with zero in a given plane will be allocated. If the plane is greater than or equal to zero, then this parameter specifies which plane must have zero bits. If the plane is less than 0, then the CTM package will choose a plane and return it.

**CTM\_\$ONE\_ONLY**

This is similar to CTM\_\$ZERO\_ONLY except that the plane in question must contain only ones.

**CTM\_\$BOTH**

Pairs of pixel values will be allocated where the two members of each pair differ only in the given plane. If the plane is less than zero, the CTM package selects a plane for you. If this option is selected, the count parameter specifies the number of pairs, and the returned pixel\_values parameter will contain only the smallest (i.e., zero bit) number of each pair.

## INPUT/OUTPUT PARAMETERS

### plane

Specifies the bit plane required by the opt parameter. This parameter is a 2-byte integer.

## OUTPUT PARAMETERS

### pixel\_values

The allocated pixel values, in CTM\_\$PIXEL\_VALUE\_VECTOR\_T format. This parameter is an array of 4-byte integers.

**status**

The completion status, in STATUS\_\$T format. This parameter is 4 bytes long. See the CTM Data Types section for more information.

**USAGE**

The CTM\_\$ALLOC\_PV routine searches for pixel values that have use counts of zero and also satisfy the constraints specified by the opt parameter. The allocated values are returned in the pixel\_values parameter, and their use counts are set to one.

The CTM (Color Table Manager) package operates on a data base of pixel values that is shared by all processes using GPR in direct mode or frame mode. The CTM calls allow different applications running on the same node to share the color map without interfering with one another.

Note that borrow mode applications have a separate, private copy of the color map, which is initialized to default values by GPR\_\$INIT.

The concept of a shared color map corresponds to the fact that there is a single physical color map for the display device. For example, if process A calls GPR\_\$SET\_COLOR\_MAP, and process B calls GPR\_\$INQ\_COLOR\_MAP for the same pixel values, process B will see the changes made by process A.

The CTM package keeps track of how many processes can access a pixel value by maintaining a use count for each pixel value. All pixel values except those preallocated by the Display Manager (see below) have an initial use count of zero which means that they are available for allocation.

As each process gains access to a pixel value, the pixel value's use count is incremented. For example, if five processes are sharing a pixel value, its use count is five. When a process releases a pixel value (using CTM\_\$RELEASE\_PV), its use count is decremented. When the use count reaches zero, the pixel value can be reallocated.

The CTM routines are "advisory" in nature, and do not affect the validity of calls to GPR\_\$SET\_COLOR\_MAP. Any pixel value may be set by any program, including the pixel values used by the Display Manager for window colors. The Display Manager preallocates values 0, 1, and 7 through 15. However, the Display Manager will release pixel values 8 through 15 when you use the MONO command (see *DOMAIN System Command Reference*).

The total number of available pixel values depends on the display type and the number of planes on the node as follows:

Display Type	Available Pixel Values
-----	-----
Monochrome	2
4-plane color	16
8-plane color	256

Even monochrome displays have "color maps". Monochrome color maps control whether the display is operating in white-on-black or black-on-white. For example, applications can invert the display by calling GPR\_\$SET\_COLOR\_MAP for pixel values zero and one, specifying the colors white and black, as desired. This feature is available in direct, frame, and borrow mode.

## Using the CTM Routines

Include the insert file `/sys/ins/ctm.ins.pas`, `/sys/ins/ctm.ins.c`, or `/sys/ins/ctm.ins.ftn` in programs that use the CTM package.

The easiest way to use the CTM package is for all processes to use `CTM_$FIND_COLOR`. In this way, the CTM package takes care of all the "bookkeeping" and you can access any color that you need. For example, if process A wants to use light blue, it supplies the pixel value (200,200,255) to `CTM_$FIND_COLOR`. If the color is in the color map, the CTM package increments its use count, making it available to process A. If light blue is not in the color map, CTM creates a light blue entry in the color map and increments its use count. If process B wants to use the same color, it also makes a call to `CTM_$FIND_COLOR`.

If you want to allocate colors yourself, use `CTM_$ALLOC_PV`. You may want to do this if you plan to change the color values and don't want to affect other processes (and don't want other processes to hinder your use of the color map). When you use `CTM_$ALLOC_PV`, there are two ways that other processes can use the values:

1. Use `CTM_$MARK_READ_ONLY` to make the pixel values available to other processes. This call marks the pixel values as read-only and other processes that use `CTM_$FIND_COLOR` can use them (but they cannot change them).
2. This method involves sharing data between processes or between independent components of the same process (for example, passing parameters between processes or using the C functions `fork` or `pipe`). If you obtain pixel values from another process and that process might terminate first (or might release the pixel values first), you can use `CTM_$INC_USE_COUNT` to indicate your intent to continue using those pixel values.

When you are through using a set of pixel values, you can release them using `CTM_$RELEASE_PV`. This routine decrements the use count of selected pixel values. When a pixel value's use count reaches zero, it is available for reallocation. The CTM package automatically decrements the use count of any pixel values that have been allocated by a program when the program terminates.

## Using the CTM\_\$ALLOC\_PV Options

If your application needs to change colors frequently, the easiest method is to use `CTM_$ALLOC_PV` with no options. Use the options when you need to perform more sophisticated tasks such as redrawing a portion of an image in a different color. An example of this is performing an XOR raster operation when rubberbanding a line or dynamically dragging an object across the screen (see `GPR_$SET_RASTER_OP` for more information on raster operations).

When dragging an object across the screen, an application needs to repeatedly redraw and erase the object without disturbing the existing geometry. In an XOR operation, when a dragged object crosses an existing piece of geometry, each intersection point is redrawn in a different color. After the object passes by, the intersection point is redrawn in its original color.

XOR operations work well on monochrome nodes but on a color node it is faster to explicitly set the draw value to one or zero.

For example, assume that your application is drawing in red, green, and blue, and wants to draw intersection points in black. Use the following procedure:

1. Use CTM\_\$ALLOC\_PV with the CTM\_\$BOTH option to allocate three pairs of pixel values. The values in each pair differ only in one plane.
2. Use GPR\_\$SET\_COLOR\_MAP to assign red, green, and blue to the first elements in each pair and assign black to the second element in each pair. For example, the following illustration shows three pixel value pairs that differ only in one plane:

Pixel Value Pairs	Binary Representation	Assigned Color
1	0 0 1	red
3	0 1 1	black
0	0 0 0	blue
2	0 1 0	black
5	1 0 1	green
7	1 1 1	black

3. Use GPR\_\$SET\_PLANE\_MASK to limit operations to the two planes.
4. Use the default raster operation (GPR\_\$SET\_RASTER\_OP operation 3).

Drawing ones in the allocated plane changes the image from its original color to black. Switching back to zeros reverts the image to its original color.

CTM\_\$FIND\_COLOR

CTM\_\$FIND\_COLOR

Searches for a pixel value (within a specified distance) that already contains a desired color. If one does not exist, a new pixel value is allocated and set to the desired color.

## FORMAT

CTM\_\$FIND\_COLOR (color, max\_distance, pixel\_value, status)

## INPUT PARAMETERS

### color

The desired color, in GPR\_\$COLOR\_T format. This parameter is a 4-byte integer that specifies values for red, green, and blue.

### max\_distance

A distance in the color space (see Usage). This parameter is a 2-byte integer.

## OUTPUT PARAMETERS

### pixel\_value

The resulting color value, in GPR\_\$PIXEL\_VALUE\_T format. This parameter is a 4-byte integer.

### status

The completion status, in STATUS\_\$T format. This parameter is 4 bytes long. See the CTM Data Types section for more information.

## USAGE

This routine treats the RGB color space as though it were ordinary 3D geometry. That is, the distance between two color values is the square root of the sum of the squares of the absolute differences in RGB intensity.

CTM\_\$FIND\_COLOR searches for a pixel value that already contains the desired color within the specified distance. First it searches the pixels marked read-only (see CTM\_\$MARK\_READ\_ONLY). If the value is found, the routine increments the pixel's use count. Otherwise, a new pixel value is allocated, set to the specified color, and marked read-only. The routine uses GPR\_\$INQ\_COLOR\_MAP to determine current color settings and GPR\_\$SET\_COLOR\_MAP to establish new colors.

This procedure finds the best match, not the first match. Therefore, if you want the closest existing pixel value, use lastof(integer) for the max\_distance parameter.



**CTM\_\$INC\_USE\_COUNT**

Increments the use count of previously allocated pixel values.

**FORMAT**

CTM\_\$INC\_USE\_COUNT (count, opt, plane, pixel\_values, status)

**INPUT PARAMETERS****count**

Specifies the number of pixel values (or pixel value pairs). This parameter is a 2-byte integer.

**opt**

Specifies the constraints placed on pixel values to be allocated, in CTM\_\$ALLOC\_OPTIONS\_T format. This parameter is a 2-byte integer. Possible values are the following:

CTM\_\$CONTIGUOUS  
 CTM\_\$ZERO\_ONLY  
 CTM\_\$ONE\_ONLY  
 CTM\_\$BOTH

**plane**

Specifies the bit plane required by the opt parameter. This parameter is a 2-byte integer.

**pixel\_values**

The allocated pixel values, in CTM\_\$PIXEL\_VALUE\_VECTOR\_T format. This parameter is an array of 4-byte integers.

**OUTPUT PARAMETERS****status**

The completion status, in STATUS\_\$T format. This parameter is 4 bytes long.

**USAGE**

If you obtain pixel values from another process, and that process might terminate first (or might release the pixel values first), you can use this routine to indicate your intent to continue using those pixel values.

Only use this call if you are sharing data between processes or between independent components of the same process (for example, passing parameters between processes or using the C functions fork or pipe). If you are accessing pixel values using CTM\_\$FIND\_COLOR, you don't need to use this call since CTM\_\$FIND\_COLOR automatically increments the use count for you.

The use count is a 16-bit word that can take on any non-negative value. If the value is zero, it signifies that the pixel value is available to be allocated (see Usage under CTM\_\$ALLOC\_PV for more information).

The count, opt, and plane arguments are used only to determine the actual pixel values specified in the pixel\_values parameter.

CTM\_\$MARK\_READ\_ONLY

CTM\_\$MARK\_READ\_ONLY

Marks pixels as shareable by other processes. Shareable pixel values may be allocated using CTM\_\$FIND\_COLOR.

## FORMAT

CTM\_\$MARK\_READ\_ONLY (count, opt, plane, pixel\_values, status)

## INPUT PARAMETERS

### count

Specifies the number of pixel values (or pixel value pairs). This parameter is a 2-byte integer.

### opt

Specifies the constraints placed on pixel values to be allocated, in CTM\_\$ALLOC\_OPTIONS\_T format. This parameter is a 2-byte integer. Possible values are the following:

CTM\_\$CONTIGUOUS  
CTM\_\$ZERO\_ONLY  
CTM\_\$ONE\_ONLY  
CTM\_\$BOTH

See CTM\_\$ALLOC\_PV for more information.

### plane

Specifies the bit plane required by the opt parameter. This parameter is a 2-byte integer.

### pixel\_values

The allocated pixel values, in CTM\_\$PIXEL\_VALUE\_VECTOR\_T format. This parameter is an array of 4-byte integers.

## OUTPUT PARAMETERS

### status

The completion status, in STATUS\_\$T format. This parameter is 4 bytes long. See the CTM Data Types section for more information.

## USAGE

If you obtain pixel values using CTM\_\$ALLOC\_PV, the CTM package assumes that you will change the colors assigned to these pixel values from time to time. If you do not plan to change the colors, you may make them available to other programs that call CTM\_\$FIND\_COLOR. To make them available, you call CTM\_\$MARK\_READ\_ONLY *after* calling GPR\_\$SET\_COLOR\_MAP.

**CTM\_\$RELEASE\_PV**

Decrements the use count of selected pixel values. A pixel value is made available to other processes when its use count reaches zero.

**FORMAT**

CTM\_\$RELEASE\_PV (count, opt, plane, pixel\_values, status)

**INPUT PARAMETERS****count**

Specifies the number of pixel values (or pixel value pairs). This parameter is a 2-byte integer.

**opt**

Specifies the constraints placed on pixel values to be released, in CTM\_\$ALLOC\_OPTIONS\_T format. This parameter is a 2-byte integer. Possible values are the following:

CTM\_\$CONTIGUOUS  
 CTM\_\$ZERO\_ONLY  
 CTM\_\$ONE\_ONLY  
 CTM\_\$BOTH

See CTM\_\$ALLOC\_PV for more information.

**plane**

Specifies the bit plane required by the opt parameter. This parameter is a 2-byte integer.

**pixel\_values**

The released pixel values, in CTM\_\$PIXEL\_VALUE\_VECTOR\_T format. This parameter is an array of 4-byte integers.

**OUTPUT PARAMETERS****status**

The completion status, in STATUS\_\$T format. This parameter is 4 bytes long. See the CTM Data Types section for more information.

**USAGE**

When you are finished using a set of pixel values, you should return them using this routine. The routine decrements the use count. When the use count reaches zero, the pixel values will be available to other callers. If you don't make this call, any pixel values you have allocated will be automatically released when your program terminates execution.

An error is returned if you attempt to decrement a use count below zero or if you attempt to decrement a use count that you (your process, at or above the current program level) have not previously incremented.

CTM ERRORS

**ERRORS**

STATUS\_ \$OK

Successful completion.

CTM\_ \$NOT\_ ALLOCATED

The pixel values specified by CTM\_ \$RELEASE\_ PV, CTM\_ \$INC\_ USE\_ COUNT, or CTM\_ \$MARK\_ READ\_ ONLY were not allocated. Pixels values must be allocated by CTM\_ \$ALLOC\_ PV or CTM\_ \$FIND\_ COLOR before they can be used by other CTM calls.

CTM\_ \$CAN'T

The CTM package was unable to allocate the requested pixel values because there is no more space available in the color map. Either use GMR\_ \$FIND\_ COLOR to locate an appropriate color that is already allocated, or use CTM\_ \$RELEASE\_ PV to deallocate values.

CTM\_ \$NO\_ SPACE

The CTM package was unable to allocate memory for its tables.

# EC2

## Table of Contents

Data Types	EC2-2
EC2_\$ADVANCE	EC2-4
EC2_\$INIT	EC2-5
EC2_\$READ	EC2-6
EC2_\$WAIT	EC2-7
EC2_\$WAIT_SVC	EC2-9
Errors	EC2-11



## EC2

The EC2 (Level 2 Eventcount) programming calls create and manage eventcounts for synchronizing program events. This section describes their data types, call syntax, and error codes. Refer to the Introduction at the beginning of this manual for a description of data type diagrams and call syntax format.


EC2 DATA TYPES

CONSTANTS

EC2\_\$ALWAYS\_READY\_EC            1            Replaces an EC2\_\$PTR pointer to indicate that the event is always ready.

DATA TYPES

EC2\_\$EVENTCOUNT\_T            An eventcount. The diagram below illustrates the EC2\_\$EVENTCOUNT\_T data type:



Field Description:

value  
The value of the eventcount.

awaiters  
Reserved for internal use by the EC2 manager.

EC2\_\$PTR\_LIST\_T            An array of up to 32 pointers to eventcounts. Each pointer is a 4-byte integer in EC2\_\$PTR\_T format.


EC2\_\$PTR\_T            A 4-byte integer. A pointer to an eventcount.

EC2\_\$VAL\_LIST\_T            An array of trigger values for each of the eventcounts in an eventcount pointer list. Each trigger value is a positive, 4-byte integer.



STATUS\_\$\$T

A status code. The diagram below illustrates the STATUS\_\$\$T data type:



Field Description:

all

All 32 bits in the status code.

fail

The fail bit. If this bit is set, the error was not within the scope of the module invoked, but occurred within a lower-level module (bit 31).

subsys

The subsystem that encountered the error (bits 24 - 30).

modc

The module that encountered the error (bits 16 - 23).

code

A signed number that identifies the type of error that occurred (bits 0 - 15).

EC2\_ \$ADVANCE

EC2\_ \$ADVANCE

Advances the specified user-defined eventcount by one.

## FORMAT

EC2\_ \$ADVANCE (eventcount, status)

## INPUT/OUTPUT PARAMETERS

### eventcount

Eventcount to be advanced, in EC2\_ \$EVENTCOUNT\_ T format. This data type is 6 bytes long. See the EC2 Data Types section for more information.

## OUTPUT PARAMETERS

### status

Completion status, in STATUS\_ \$T format. This data type is 4 bytes long. See the EC2 Data Types section for more information.

## USAGE

EC2\_ \$ADVANCE advances a user-defined eventcount; do not use this call to advance a system-defined eventcount. A user-defined eventcount is one that a process establishes with EC2\_ \$INIT.

In order to advance an eventcount, you must have read and write access to the emory location where the eventcount is located. Typically, the eventcount is in a mapped file, and you should use MS\_ \$MAPL to map this file.

When you map a file containing an eventcount, you should request a shared write lock. Only processes on the same node can concurrently get shared write locks on the same file. See MS\_ \$MAPL for more information about mapping a file.

**EC2\_\$INIT**

Initializes a user-defined eventcount.

**FORMAT**

EC2\_\$INIT (eventcount)

**INPUT PARAMETERS**

None.

**OUTPUT PARAMETERS****eventcount**

Initialized eventcount, in EC2\_\$EVENTCOUNT\_T format. This data type is 6 bytes long. See the EC2 Data Types section for more information.

**USAGE**

Use this call to initialize a user-defined eventcount. Initialize the eventcount within a file that several programs will share. First, map the file for shared write access (by requesting a shared write lock.) Then use EC2\_\$INIT to initialize the eventcount. All programs that use the eventcount must first map the file containing the eventcount. See the Mapped Segment (MS) calls for more information on mapping.

Do not use EC2\_\$INIT to initialize a system-defined eventcount; the system automatically initializes eventcounts associated with system events. To use a system-defined eventcount, use the system call that gets the address of the eventcount you want to wait on. For example, use MBX\_\$GET\_EC to get the address of a mailbox eventcount.

EC2\_\$READ

EC2\_\$READ

Returns the current value of an eventcount.

#### FORMAT

ec-value = EC2\_\$READ (eventcount)

#### RETURN VALUE

**ec-value**

Value of the eventcount. This is a positive, 4-byte integer.

#### INPUT PARAMETERS

**eventcount**

Eventcount, in EC2\_\$EVENTCOUNT\_T format. This data type is 6 bytes long. See the EC2 Data Types section for more information.

#### OUTPUT PARAMETERS

None.

#### USAGE

Use EC2\_\$READ to read the value of an eventcount.

**EC2\_\$WAIT**

Waits until any of a list of eventcounts reaches or exceeds a trigger value.

**FORMAT**

`ec-satisfied = EC2_$WAIT (ec-plist, ec-vlist, ec-count, status)`

**RETURN VALUE****ec-satisfied**

An ordinal number indicating the eventcount that is satisfied. This is a positive, 2-byte integer.

**INPUT PARAMETERS****ec-plist**

Array of pointers to eventcounts. Each pointer is a 4-byte integer in `EC2_$PTR_T` format. The total number of eventcounts in `ec-plist` lists in any one node cannot exceed 32.

**ec-vlist**

Array of trigger values for each of the eventcounts in the `ec-plist`. Each trigger value is a positive, 4-byte integer. When any of the eventcounts from the `ec-plist` reaches its trigger value, the `EC2_$WAIT` call returns.

**ec-count**

Number of eventcounts in the `ec-plist`. This is a 2-byte integer.

**OUTPUT PARAMETERS****status**

Completion status, in `STATUS_$T` format. This data type is 4 bytes long. See the EC2 Data Types section for more information.

**USAGE**

`EC2_$WAIT` waits until one of the eventcounts in the `ec-plist` reaches its trigger value in the `ec-vlist`. When an eventcount reaches its trigger value, `EC2_$WAIT` returns an index value indicating the position (in the `ec-plist`) of the eventcount that is satisfied. The index starts from 1; that is, `ec-satisfied` equals 1 if the first eventcount in the `ec-plist` is satisfied.

Several eventcounts may have been satisfied by the time this call wakes your program. The index number returned refers to only one of the eventcounts. If more than one eventcount is satisfied, `EC2_$WAIT` returns the one with the smallest subscript.

`EC2_$WAIT` only returns when an eventcount advances, regardless of the asynchronous fault handling setting. An asynchronous fault, such as a "quit", is generated outside your program. If an asynchronous fault occurs during an `EC2_$WAIT` call, your program's response depends on the type of error handling that is in effect.

## EC2\_\$WAIT

If asynchronous faults are enabled, a program can respond to an asynchronous fault with a clean-up handler or fault handler. If an asynchronous fault occurs during an EC2\_\$WAIT call, and asynchronous faults are enabled, the program will perform one of the following:

- Execute the clean-up handler, if the program has one.
- Execute the fault handler, if the program has one. If the fault handler returns control to the interrupted code, EC2\_\$WAIT continues waiting.
- If the program has neither a clean-up handler nor a fault handler, the program aborts if an asynchronous fault occurs.

If a program disables asynchronous faults and such a fault occurs during an EC2\_\$WAIT, then the program ignores the fault and continues waiting.

Note that the call EC2\_\$WAIT\_SVC responds differently to asynchronous faults.

**EC2\_\$WAIT\_SVC**

Waits until any of a list of eventcounts reaches or exceeds a trigger value.

**FORMAT**

ec-satisfied = EC2\_\$WAIT\_SVC (ec-plist, ec-vlist, ec-count, status)

**RETURN VALUE****ec-satisfied**

An ordinal number indicating the eventcount that was satisfied. This is a positive, 2-byte integer.

**INPUT PARAMETERS****ec-plist**

Array of pointers to eventcounts. Each pointer in the array is a 4-byte integer in EC2\_\$PTR\_T format. The total number of eventcounts in ec-plist lists in any one node cannot exceed 32.

**ec-vlist**

Array of trigger values for each of the eventcounts in the ec-plist. Each trigger value is a positive, 4-byte integer. When any of the eventcounts from the ec-plist reaches its trigger value, the EC2\_\$WAIT\_SVC call returns.

**ec-count**

Number of eventcounts in the ec-plist. This is a 2-byte integer.

**OUTPUT PARAMETERS****status**

Completion status, in STATUS\_\$T format. This data type is 4 bytes long. See the EC2 Data Types section for more information.

**USAGE**

EC2\_\$WAIT\_SVC waits until one of the eventcounts in the ec-plist reaches its trigger value in the ec-vlist. When an eventcount reaches its trigger value, EC2\_\$WAIT\_SVC returns an index value indicating the position (in the ec-plist) of the eventcount that is satisfied. The index starts from 1; that is, ec-satisfied equals 1 if the first eventcount in the ec-plist is satisfied.

Several eventcounts may have been satisfied by the time this call wakes your program. The index number returned refers to only one of the eventcounts. If more than one eventcount is satisfied, EC2\_\$WAIT\_SVC returns the one with the smallest subscript.

In certain cases, EC2\_\$WAIT\_SVC returns the error EC2\_\$WAIT\_QUIT if an asynchronous fault occurs during the EC2\_\$WAIT\_SVC call. An asynchronous fault, such as a "quit", is generated outside your program. If an asynchronous fault occurs during an EC2\_\$WAIT\_SVC call, your program's response depends on the type of error handling that is in effect.

If asynchronous faults are enabled, a program can respond to an asynchronous fault with a clean-up handler or a fault handler. If an asynchronous fault occurs during an EC2\_\$WAIT\_SVC call, and asynchronous faults are enabled, the program will perform one of the following:

- Execute the clean-up handler, if the program has one.
- Execute the fault handler, if the program has one. If the fault handler returns control to the interrupted code, EC2\_\$WAIT\_SVC returns the error EC2\_\$WAIT\_QUIT.
- If the program has neither a clean-up handler nor a fault handler, the program aborts if an asynchronous fault occurs.

If a program disables asynchronous faults and such a fault occurs during an EC2\_\$WAIT\_SVC, then the program does not handle the fault. However, EC2\_\$WAIT\_SVC returns the error EC2\_\$WAIT\_QUIT.

Note that the call EC2\_\$WAIT responds differently to asynchronous faults.



**ERRORS**

STATUS\_ \$OK

Successful completion.

EC2\_ \$BAD\_ EVENTCOUNT

Bad eventcount.

EC2\_ \$INTERNAL\_ ERROR

Internal error.

EC2\_ \$NO\_ WAIT\_ ENTRIES

Internal table exhausted.

EC2\_ \$WAIT\_ QUIT

Process quit while waiting.



# ERROR

## Table of Contents

### Data Types

ERROR\_\$CODE

ERROR\_\$FAIL

ERROR\_\$FIND\_TEXT

ERROR\_\$GET\_TEXT

ERROR\_\$INIT\_STD\_FORMAT

ERROR\_\$MODULE

ERROR\_\$PRINT

ERROR\_\$PRINT\_FORMAT

ERROR\_\$PRINT\_NAME

ERROR\_\$STD\_FORMAT

ERROR\_\$SUBSYS

ERROR-2

ERROR-3

ERROR-4

ERROR-5

ERROR-6

ERROR-7

ERROR-8

ERROR-9

ERROR-10

ERROR-12

ERROR-13

ERROR-14



## ERROR

The ERROR programming calls provide status printing and interpretation. This section describes their data types and call syntax. The ERROR calls do not produce unique error messages. Refer to the Introduction at the beginning of this manual for a description of data type diagrams and call syntax format.

## ERROR DATA TYPES

### DATA TYPES

ERROR\_ \$INTEGER32

A 2-byte integer. Possible values are integers from -1 through 2147483647.

ERROR\_ \$STRING\_PTR\_T


A 4-byte integer. A pointer to an ERROR\_ \$STRING\_T data type.

ERROR\_ \$STRING\_T

An array of up to 80 characters.

STATUS\_ \$T

A status code. The diagram below illustrates the STATUS\_ \$T data type:



#### Field Description:

all  
All 32 bits in the status code.

fail  
The fail bit. If this bit is set, the error was not within the scope of the module invoked, but occurred within a lower-level module (bit 31).

subsys  
The subsystem that encountered the error (bits 24 - 30).

modc  
The module that encountered the error (bits 16 - 23).

code  
A signed number that identifies the type of error that occurred (bits 0 - 15).

**ERROR\_ \$CODE**

Returns the module-specific code from a status code.

**FORMAT**

code = ERROR\_ \$CODE (status)

**RETURN VALUE****code**

The module-specific code component of the supplied status code. This is a 2-byte integer.

**INPUT PARAMETERS****status**

A status code returned from DOMAIN software, in STATUS\_ \$T format. This data type is 4 bytes long. See the ERROR Data Types section for more information.

**USAGE**

ERROR\_ \$CODE extracts and returns the module-specific code from the supplied status code. The module-specific code is the rightmost 16 bits of a STATUS\_ \$T data type.

This routine is intended for use by FORTRAN programs that need to check for specific status codes. Pascal programs can refer to this component directly.

**ERROR\_\$FAIL**

**ERROR\_\$FAIL**

Returns the state of the fail bit of a status code.

**FORMAT**

`fail = ERROR_$FAIL (status)`

**RETURN VALUE**

**fail**

The value of the fail bit of the status code. This is a Boolean (logical) value.

**INPUT PARAMETERS**

**status**

A status code returned from DOMAIN software, in STATUS\_\$T format. This data type is 4 bytes long. See the ERROR Data Types section for more information.

**USAGE**

ERROR\_\$FAIL extracts and returns the value of the fail bit of the supplied status code. The fail bit is bit number 31 in the STATUS\_\$T data type.

This routine is intended for use by FORTRAN programs that need to check for specific status codes. Pascal programs can refer to this component directly.



**ERROR\_\$FIND\_TEXT**

Finds the text associated with a status code and returns pointers.

**FORMAT**

ERROR\_\$FIND\_TEXT (status, subsys\_p, subsys\_l, module\_p, module\_l, code\_p, code\_l)

**INPUT PARAMETERS****status**

A status code returned from DOMAIN software, in STATUS\_\$T format. This data type is 4 bytes long. See the ERROR Data Types section for more information.

**OUTPUT PARAMETERS****subsys\_p**

The returned pointer to the subsystem name, in ERROR\_\$STRING\_PTR\_T format. This is a 4-byte integer.

**subsys\_l**

The number of characters in the string pointed to by subsys\_p. This is a 2-byte integer.

**module\_p**

The returned pointer to the module name, in ERROR\_\$STRING\_PTR\_T format. This is a 4-byte integer.

**module\_l**

The number of characters in the string pointed to by module\_p. This is a 2-byte integer.

**code\_p**

The returned pointer to the diagnostic text, in ERROR\_\$STRING\_PTR\_T format. This is a 4-byte integer.

**code\_l**

The number of characters in the string pointed to by code\_p. This is a 2-byte integer.

**USAGE**

ERROR\_\$FIND\_TEXT looks up and returns pointers to the text associated with a status code.

Text is associated with three components of the STATUS\_\$T type: subsystem name ("subsys"), module name ("module"), and error text ("code"). If ERROR\_\$FIND\_TEXT cannot find the text associated with a component in the status code, a string length of zero is returned for the component. In this case, the pointer for that component is not useable.

If the subsystem text length is zero, the status is invalid. If the module text length is zero, both the module and code fields are invalid.

FORTRAN programs should use ERROR\_\$GET\_TEXT instead of this routine.

ERROR\_\$GET\_TEXT

ERROR\_\$GET\_TEXT

Finds the text associated with a status code and returns strings.

## FORMAT

ERROR\_\$GET\_TEXT (status, subsys\_t, subsys\_l, module\_t, module\_l, code\_t, code\_l)

## INPUT PARAMETERS

### status

A status code returned from DOMAIN software, in STATUS\_\$T format. This data type is 4 bytes long. See the ERROR Data Types section for more information.

## OUTPUT PARAMETERS

### subsys\_t

The text string containing the subsystem name, in ERROR\_\$STRING\_T format. This is an array of up to 80 characters.

### subsys\_l

The number of characters in the subsystem name. This is a 2-byte integer.

### module\_t

The text string containing the module name, in ERROR\_\$STRING\_T format. This is an array of up to 80 characters.

### module\_l

The number of characters in the module name. This is a 2-byte integer.

### code\_t

The text string containing the diagnostic text, in ERROR\_\$STRING\_T format. This is an array of up to 80 characters.

### code\_l

The number of characters in the diagnostic text. This is a 2-byte integer.

## USAGE

ERROR\_\$GET\_TEXT looks up and returns the text associated with a status code.

Text is associated with three components of the STATUS\_\$T type: subsystem name ("subsys"), module name ("module"), and error text ("code"). If ERROR\_\$GET\_TEXT cannot find the text associated with a component in the status code, a string length of zero is returned for the component.

If the subsystem text length is zero, the status is invalid. If the module text length is zero, both the module and code fields are invalid.

The returned strings are not blank-filled. They contain only the number of characters necessary to represent the names and diagnostic text.

**ERROR\_ \$INIT\_ STD\_ FORMAT**

Establishes the values to be used in subsequent calls to ERROR\_ \$STD\_ FORMAT.

**FORMAT**

ERROR\_ \$INIT\_ STD\_ FORMAT (stream-id, prefix-char, command, length)

**INPUT PARAMETERS****stream\_id**

The stream on which to write the error output, in STREAM\_ \$ID\_ T format. This is a 2-byte integer. This is usually STREAM\_ \$ERROUT (Stream ID = 3).

**prefix-char**

The prefix element of the error format. This is one character. For system messages, this value is usually a question mark (?).

**command**

The command name, in ERROR\_ \$STRING\_ T format. This is an array of up to 80 characters.

**length**

The length of the command name, in bytes. This value can be zero.

**USAGE**

This call establishes constant values for the standard error reporting format. Subsequent calls to ERROR\_ \$STD\_ FORMAT cause error messages to use the values supplied in this call.

Multiple calls may be made to ERROR\_ \$INIT\_ STD\_ FORMAT, but the information is kept on a per-process-level basis. Thus, successive calls to ERROR\_ \$INIT\_ STD\_ FORMAT on the same process level replace previous error format definitions.

Calling ERROR\_ \$INIT\_ STD\_ FORMAT and ERROR\_ \$STD\_ FORMAT is equivalent to calling ERROR\_ \$PRINT\_ FORMAT. For programs that use common subroutines, the former method provides more flexibility. For example, if an application's command level sets the command name with ERROR\_ \$INIT\_ STD\_ FORMAT, it automatically provides the common lower-level modules with the correct command name for their error messages. Also, because ERROR\_ \$STD\_ FORMAT has fewer parameters, it is easier to code using the pair of calls instead of using ERROR\_ \$PRINT\_ FORMAT.

ERROR\_\$MODULE

ERROR\_\$MODULE

Returns the module component from a status code.

#### FORMAT

module = ERROR\_\$MODULE (status)

#### RETURN VALUE

**module**

The module component of the supplied status code. This is a 2-byte integer.

#### INPUT PARAMETERS

**status**

A status code returned from DOMAIN software, in STATUS\_\$T format. This data type is 4 bytes long. See the ERROR Data Types section for more information.

#### USAGE

ERROR\_\$MODULE extracts and returns the module component of the supplied status code. The module is found in bits 23 through 16 of the STATUS\_\$T data type.

This routine is intended for use by FORTRAN programs that need to check for specific status codes. Pascal programs can refer to this component directly.

**ERROR\_\$PRINT**

Prints error text associated with a status code.

**FORMAT**

ERROR\_\$PRINT (status)

**INPUT PARAMETERS****status**

A status code returned from DOMAIN software, in STATUS\_\$T format. This data type is 4 bytes long. See the ERROR Data Types section for more information.

**USAGE**

ERROR\_\$PRINT looks up the text associated with the status code and writes it to the error output stream.

If text is associated with all fields in the status code (subsystem, module, and code), a line is output containing the subsystem and module names.

If the text for any of the three fields is not found, the status code is displayed in hexadecimal, along with the subsystem and module names, if known.

The STCODE command, which can be used to view error messages, uses ERROR\_\$PRINT to output the error text.

ERROR\_\$PRINT\_FORMAT

ERROR\_\$PRINT\_FORMAT

Prints a status code in the given error format.

## FORMAT

ERROR\_\$PRINT\_FORMAT (status, stream-id, prefix-char, command, length,  
control-string, a1, a2, ... a10)

## INPUT PARAMETERS

### status

The status code to be displayed in standard error format, in STATUS\_\$T format. This data type is 4 bytes long. See the ERROR Data Types section for more information.

If the status code is zero, the dash and following error text are omitted from the message.

### stream-id

The stream on which to write the error output, in STREAM\_\$ID\_T format. This is a 2-byte integer. This value is usually STREAM\_\$ERROUT (Stream ID + 3).

### prefix-char

The prefix element of the error format. This is one character. For system error messages, this value is usually a question mark (?).

### command

The command name, in ERROR\_\$STRING\_T format. This is an array of up to 80 characters.

### length

Length of the command name, in bytes. This is a 2-byte integer. If length is zero, the command name portion of the standard error format is omitted.

### control-string

A character string that contains text and control information for encoding the arguments that follow. It is a VFMT control string that must at least contain the two special characters (% , \$). For detailed information on VFMT control strings, see the chapter on "Formatting Variables with VFMT in *Programming With General System Calls*."

### a1, a2, ... a10

One-to-ten substitution arguments that contain data for encoding using the control-string parameter.

If you are encoding ASCII text strings, you must provide two variables for each text string: a character string containing the string, and a 2-byte integer variable containing the length of the string.

**USAGE**

ERROR\_\$PRINT\_FORMAT prints an error in the standard error format.

ERROR\_\$PRINT\_FORMAT takes a variable number of arguments in the a1...a10 parameters. However, all arguments up to and including the control string must be given.

This routine uses the same control string format as the variable formatting routine VFMT\_\$WRITE.

ERROR\_\$PRINT\_NAME

ERROR\_\$PRINT\_NAME

Prints error text associated with a status code, along with a user-supplied name.

## FORMAT

ERROR\_\$PRINT\_NAME (status, name, namelength)

## INPUT PARAMETERS

### status

A status code returned from DOMAIN software, in STATUS\_\$T format. This data type is 4 bytes long. See the ERROR Data Types section for more information.

### name

The name to be printed. This is an array of up to 80 characters.

### namelength

The length of the name to be printed, in bytes. This is a 2-byte integer.

## USAGE

ERROR\_\$PRINT\_NAME looks up the text associated with the status code and writes it to the error output stream, along with the supplied name.

If text is associated with all fields in the status code (subsystem, module, and code), output appears with the supplied name first, followed by a descriptive error message corresponding to the status code, followed the subsystem and module names in parentheses.

If the text for any of the three fields is not found, the status code is displayed in hexadecimal, along with the subsystem and module names, if known. The supplied name is also displayed, in the form shown above.



**ERROR\_ \$STD\_ FORMAT**

Prints the status code in the standard error format using the values specified in the last call to **ERROR\_ \$INIT\_ STD\_ FORMAT**.

**FORMAT**

**ERROR\_ \$STD\_ FORMAT** (status, control-string, a1, a2, ... a10)

**INPUT PARAMETERS****status**

The status to be printed in standard error format, in **STATUS\_ \$T** format. This data type is 4 bytes long. See the **ERROR Data Types** section for more information.

**control-string**

A character string that contains text and control information for encoding the arguments that follow. It is a VFMT control string that must at least contain the two special characters (**%**, **\$**). For more information on VFMT control strings, see Chapter 3 of *Programming With General System Calls*.

**a1, a2, ... a10**

One-to-ten substitution arguments that contain data for encoding using the control-string parameter.

If you are encoding ASCII text strings, you must provide two variables for each text string: a character string containing the string, and a 2-byte integer variable containing the length of the string.

**USAGE**

Programs using **ERROR\_ \$STD\_ FORMAT** must first call **ERROR\_ \$INIT\_ STD\_ FORMAT** to establish constant values for the standard error reporting format.

This routine uses the same control string format as the variable formatting routine **VFMT\_ \$WRITE**.

**ERROR\_ \$SUBSYS**

**ERROR\_ \$SUBSYS**

Returns the subsystem component from a status code.

**FORMAT**

`subsys = ERROR_ $SUBSYS (status)`

**RETURN VALUE**

**subsys**

The subsystem component of the supplied status code. This is a two-byte integer.

**INPUT PARAMETERS**

**status**

A status code returned from DOMAIN software, in STATUS\_ \$T format. This data type is 4 bytes long. See the ERROR Data Types section for more information.

**USAGE**

ERROR\_ \$SUBSYS extracts and returns the subsystem component of the supplied status code. The subsystem is found in bits 30 through 24 of the STATUS\_ \$T data type.

This routine is intended for use by FORTRAN programs that need to check for specific status codes. Pascal programs can refer to this component directly.

# FPP

## Table of Contents

Data Types	FPP-2
FPP_\$CONTROL	FPP-4
FPP_\$RESTORE	FPP-6
FPP_\$SAVE	FPP-7
FPP_\$SAVE_RESTORE	FPP-8
FPP_\$SAVE_RESTORE_SIZE	FPP-9
FPP_\$STATUS	FPP-10
Errors	FPP-12



# FPP

The FPP (Floating Point Package) programming calls perform various operations on floating point state. This section describes their data types, call syntax, and error codes. Refer to the Introduction at the beginning of this manual for a description of data-type diagrams and call syntax format.

FPP DATA TYPES

DATA TYPES

FPP\_OP\_T

A 2-byte integer. Options available with FPP\_\$CONTROL and FPP\_\$STATUS. One of the following predefined values:

READ\_OP

Reads the FP register into the specified variable.

WRITE\_OP

Writes to the FP register from the specified variable.

EXCH\_OP

Exchanges the FP register with the specified register.

FPP\_\$SAVE\_AREA\_PTR


A 4-byte integer. A pointer to the storage area where the floating point state is saved.

FPP\_\$SAVE\_AREA\_T

An array of up to 104 bytes. The storage area for saving a the floating point state.

STATUS\_\$T

A status code. The diagram below illustrates the STATUS\_\$T data type:



Field Description:

all

All 32 bits in the status code.

fail

The fail bit. If this bit is set, the error was not within the scope of the module invoked, but occurred within a lower-level module (bit 31).

subsys

The subsystem that encountered the error (bits 24 - 30).

modc

The module that encountered the error (bits 16 - 23).

code

A signed number that identifies the type of error that occurred (bits 0 - 15).

FPP\_\$CONTROL

FPP\_\$CONTROL

Reads from or writes to the floating point control register.

### FORMAT

FPP\_\$CONTROL (options, control-reg, status)

### INPUT PARAMETERS

#### options

Specifies the type of action to be taken, in FPP\_OP\_T format. Specify one of the following predefined values:

READ\_OP      Read the FP control register into the variable specified for control-reg.

WRITE\_OP     Write the FP control register from the variable specified for control-reg.

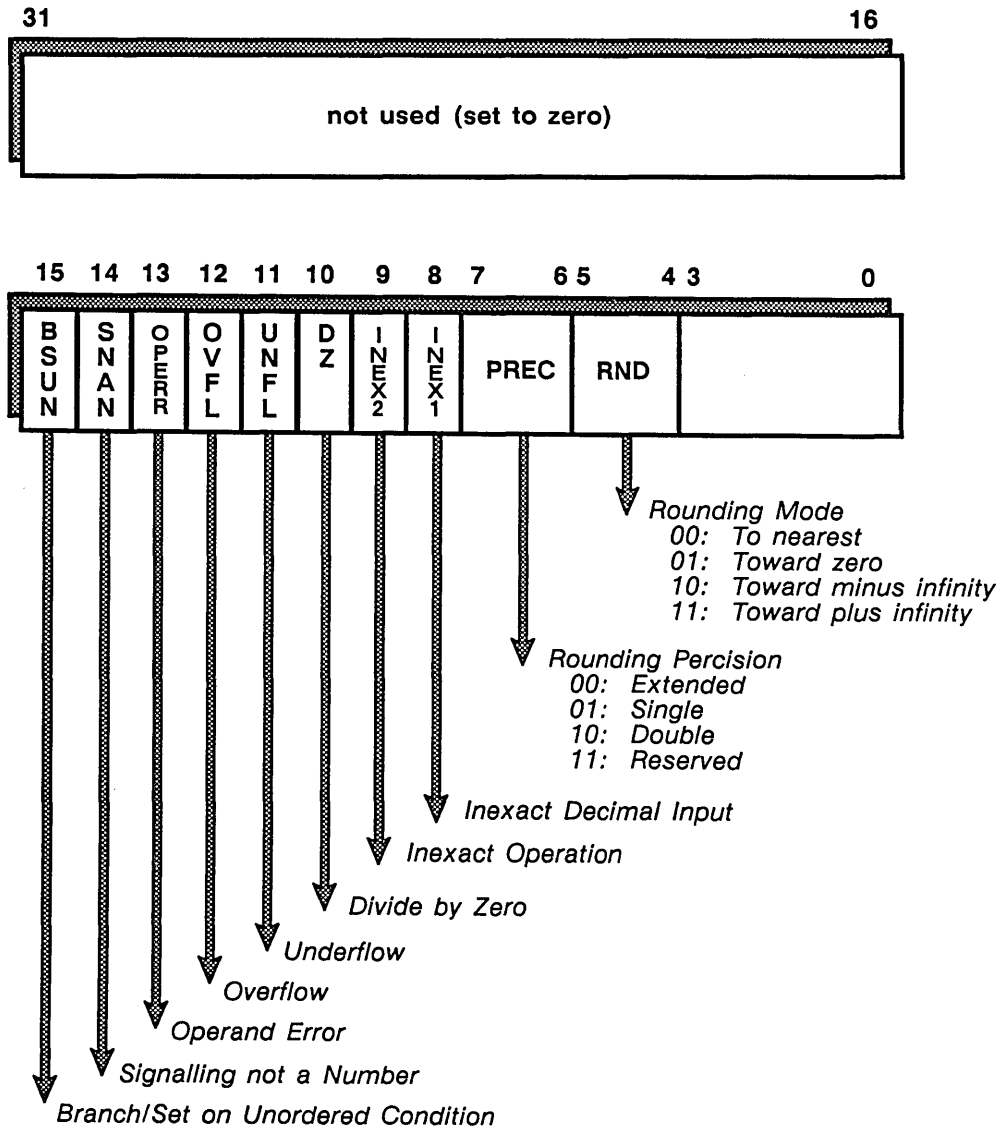
EXCH\_OP      Exchange the FP control register with the variable specified for control-reg.

### OUTPUT PARAMETERS

#### control-reg

A bit mask of 32 bits, of which only the low 16 bits are currently used. They are interpreted as follows:





**status**

Completion status in STATUS\_ \$T format. This data type is 4 bytes long. See the FPP Data Types section for more information.

**USAGE**

FPP\_ \$CONTROL is used to read from or write to the floating point control register to enable/disable various floating point exceptions. This routine is for use only with machines equipped with the MC68881 or FPX units. Use of this routine on other machines will result in a FPP\_ \$UNSUPPORTED\_ FUNCTION error.

FPP\_\$RESTORE

FPP\_\$RESTORE

Restores the floating point state.

#### FORMAT

bytes-restored := FPP\_\$RESTORE(save-area-ptr, status)

#### RETURN VALUE

**bytes-restored**

Number of bytes restored from save area.

#### INPUT PARAMETERS

**save-area-ptr**

A pointer to the storage area from which the floating point state is restored. This is a 4-byte integer.

#### OUTPUT PARAMETERS

**status**

Completion status in STATUS\_\$T format. This data type is 4 bytes long. See the FPP Data Types section for more information.

#### USAGE

This function returns the number of bytes restored from the specified area.

**FPP\_ \$SAVE**

Saves the floating point state.

**FORMAT**

bytes-saved := FPP\_ \$SAVE(save-area-ptr, status)

**RETURN VALUE****bytes-saved**

Number of bytes saved in save area.

**INPUT PARAMETERS****save-area-ptr**

A pointer to the storage area where the floating point state will be saved. This is a 4-byte integer.

**OUTPUT PARAMETERS****status**

Completion status in STATUS\_ \$T format. This data type is 4 bytes long. See the FPP Data Types section for more information.

**USAGE**

This function returns to the caller the number of bytes saved in the specified area. FPP\_ \$SAVE is useful in fault handlers that implement multitasking environments within a single process.

FPP\_\$\$SAVE\_RESTORE

FPP\_\$\$SAVE\_RESTORE

Saves current floating point state, then restores the specified floating point state.

### FORMAT

bytes-moved := FPP\_\$\$SAVE\_RESTORE(save-area-ptr1, save-area-ptr2, status)

### RETURN VALUE

**bytes-moved**

Number of bytes saved in or restored from save area.

### INPUT PARAMETERS

**save-area-ptr1**

A pointer to the storage area where the current floating point state will be saved. This is a 4-byte integer.

### INPUT PARAMETERS

**save-area-ptr2**

A pointer to the storage area from which the specified floating point state is to be restored. This is a 4-byte integer.

### OUTPUT PARAMETERS

**status**

Completion status in STATUS\_\$\$T format. This data type is 4 bytes long. See the FPP Data Types section for more information.

### USAGE

This function returns the number of bytes saved to or restored from the specified area.

**FPP\_ \$SAVE\_ RESTORE\_ SIZE**

Returns the number of bytes required for one state save.

**FORMAT**

save-area-size := FPP\_ \$SAVE\_ RESTORE\_ SIZE

**RETURN VALUE****save-area-size**

Number of bytes required for saving one floating point.

**USAGE**

This function enables the user to determine how many bytes of storage are needed to save the floating point state on the particular machine currently executing the function. (The number of bytes varies according to the the FP machine type.) It is only necessary to specify the returned size amount of space for each state save.

FPP\_\$STATUS

FPP\_\$STATUS

Reads from or writes to the floating point status register.

### FORMAT

FPP\_\$STATUS (options, status-reg, status)

### INPUT PARAMETERS

#### options

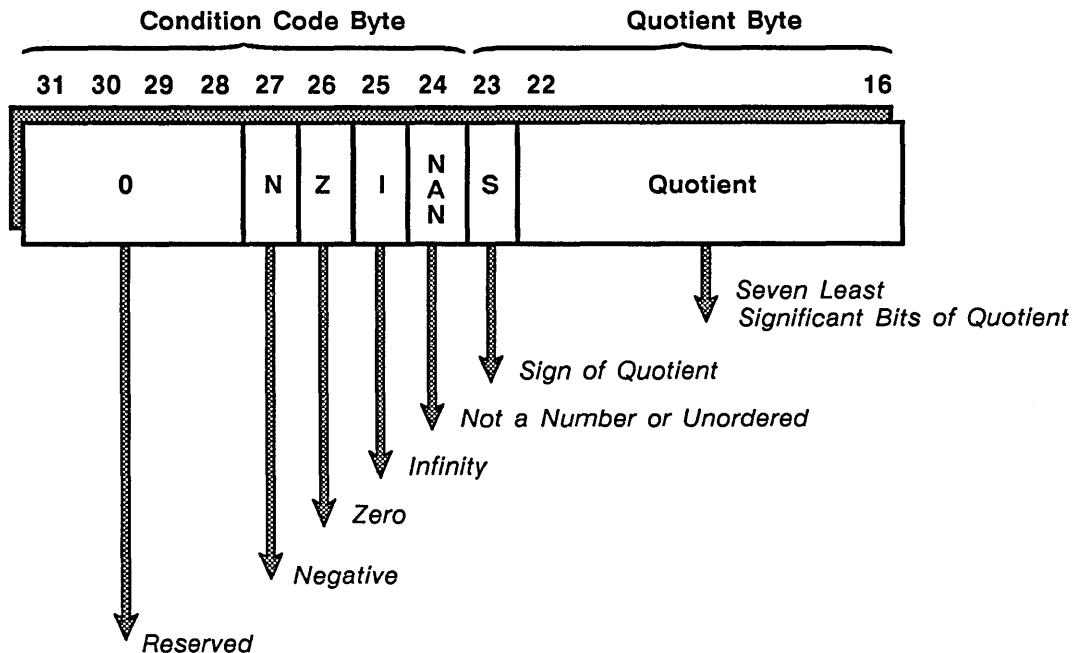
Specifies the type of action to be taken, in FPP\_OP\_T format. Specify one of the following predefined values:

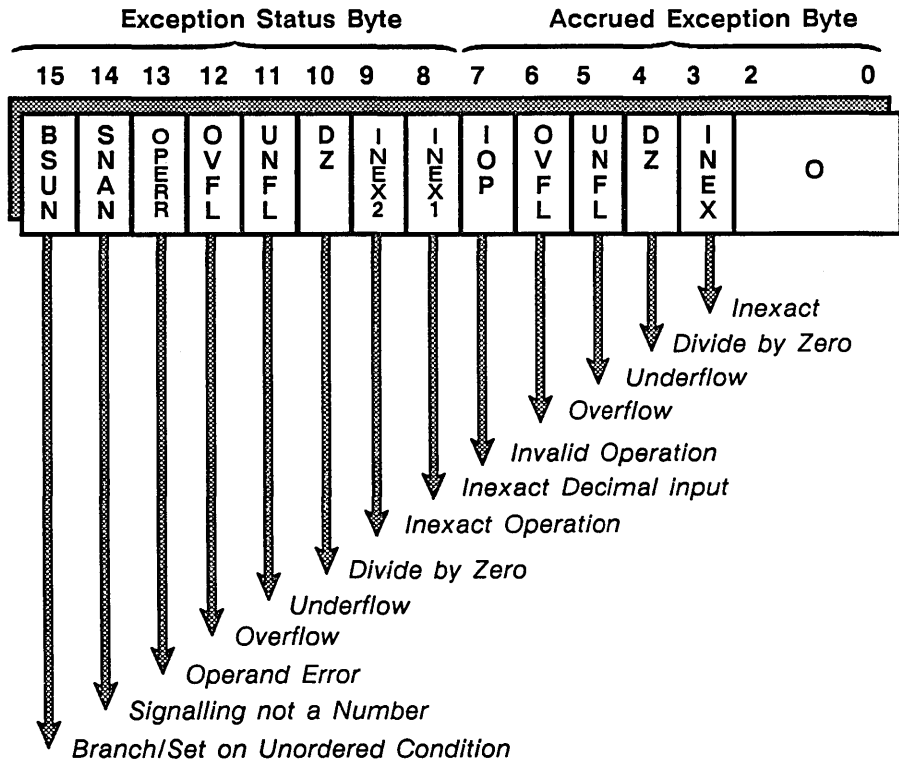
- READ\_OP      Read the FP status register into the variable specified for status-reg.
- WRITE\_OP     Write the FP status register from the variable specified for status-reg.
- EXCH\_OP      Exchange the FP status register with the variable specified for status-reg.

### OUTPUT PARAMETERS

#### status-reg

A bit mask of 32 bits, of which only the low 16 bits are currently used. They are interpreted as follows:





**status**

Completion status in STATUS\_ \$T format. This data type is 4 bytes long. See the FPP Data Types section for more information.

**USAGE**

FPP\_ \$STATUS is used to read from or write to the floating point status register to enable/disable various floating point exceptions. This routine is for use only with machines equipped with the MC68881 or FPX units. Use of this routine on other machines will result in a FPP\_ \$UNSUPPORTED\_ FUNCTION error.

FPP ERRORS

**ERRORS**

FPP\_ \$INVALID\_ OP

Invalid operation attempted.

FPP\_ \$UNSUPPORTED\_ FUNCTION

Operation not supported on current hardware.

STATUS\_ \$OK

Successful completion.



# GMF

## Table of Contents

Data Types	GMF-2
GMF_\$CLOSE	GMF-4
GMF_\$COPY_PLANE	GMF-5
GMF_\$COPY_SUBPLANE	GMF-7
GMF_\$OPEN	GMF-9
GMF_\$RESTORE_PLANE	GMF-11
Errors	GMF-13



## GMF

This section describes the data types, the call syntax, and the error codes for the GMF programming calls. Refer to the Introduction at the beginning of this manual for a description of data type diagrams and call syntax format.

GMF DATA TYPES

DATA TYPES

GMF\_\$OPOS\_T

A 2-byte integer. Specifies the file opening positions. One of the following predefined values:

GMF\_\$OVERWRITE

Provides write access; truncates file to BOF if it already exists.

GMF\_\$APPEND

Provides write access if file exists.

GMF\_\$READ

Provides read access only.

GMF\_\$MEMORY\_T

A 65535-element array of 131070-byte integers. An array of two-byte integers.

GMF\_\$MEMORY\_PTR\_T

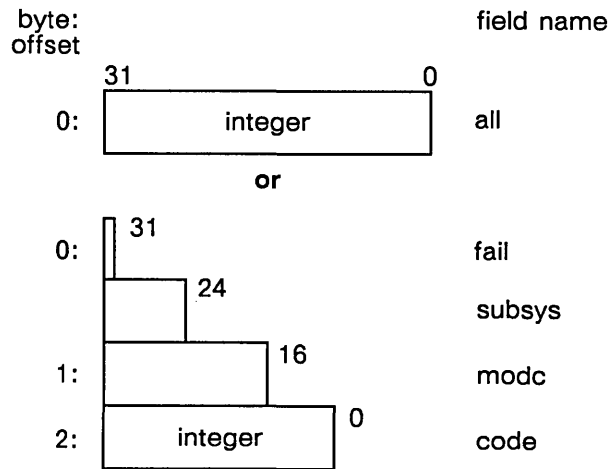
A 4-byte integer. A pointer to an array of type gmf\_\$memory\_t.

STREAM\_\$ID\_T

A 2-byte integer. Open stream identifier.

STATUS\_\$T

A status code. The diagram below illustrates the STATUS\_\$T data type:



Field Description:

all  
All 32 bits in the status code.

fail  
The fail bit. If this bit is set, the error was not within the scope of the module invoked, but occurred within a lower-level module (bit 31).

subsys

The subsystem that encountered the error (bits 24 - 30).

modc

The module that encountered the error (bits 16 - 23).

code

A signed number that identifies the type of error that occurred (bits 0 - 15).

GMF\_\$CLOSE

GMF\_\$CLOSE

Closes a GMF.

## FORMAT

GMF\_\$CLOSE (stream\_id, status)

## INPUT PARAMETERS

### stream\_id

The stream ID of the GMF to be closed, in STREAM\_\$ID\_T format. This is a 2-byte integer. You obtain the stream ID from the call to GMF\_\$OPEN that you used to open the GMF.

## OUTPUT PARAMETERS

### status

Completion status, in STATUS\_\$T format. This data type is 4 bytes long. See the GPR Data Types section for more information.

## USAGE

To open a GMF, use GMF\_\$OPEN.

**GMF\_\$COPY\_PLANE**

Dumps a rectangular area of bits from virtual memory into a GMF.

**FORMAT**

GMF\_\$COPY\_PLANE (stream\_id, black\_or\_white, bpi, bit\_pointer, x\_dim, y\_dim, width, status)

**INPUT PARAMETERS****stream\_id**

The stream ID of the GMF into which the image is to be stored, in STREAM\_\$ID\_T format. This is a 2-byte integer. You obtain the stream ID from the call to GMF\_\$OPEN that you used to open the GMF.

**black\_or\_white**

A Boolean variable. A value of TRUE means "1" bits are black and "0" bits are white. A value of FALSE means "1" bits are white and "0" bits are black. In the GMF, "1" bits are assumed to mean black. Thus if this parameter is false, the bits will be inverted as they are copied.

**bpi**

The number of bits per inch in the GMF. This information is stored in the GMF. It indicates the physical density of the image represented in the GMF. If this parameter is nonzero, a device to which you output the GMF may compress or expand the image to produce a result which is as close as possible to the image's original size. If this parameter is zero, an output device uses one dot to represent each bit from the GMF, regardless of the resulting physical size of the image. This is a 2-byte integer.

**bit\_pointer**

A pointer to the upper left corner of the rectangular area to be stored, in GMF\_\$MEMORY\_PTR\_T format. This is a 4-byte integer. You obtain this value by calling the routine GPR\_\$INQ\_BITMAP\_POINTER.

**x\_dim**

The x dimension of the rectangular area to be stored in the GMF. This is a 2-byte integer.

**y\_dim**

The y dimension of the rectangular area to be stored in the GMF. This is a 2-byte integer.

**width**

The number of 16-bit words per scan line in the source bitmap. The value of this parameter is usually 64. The width must be at least 1/16 of the specified x-dim. For instance, if you are storing an area 400 bits wide in a GMF, the source bitmap must use at least 25 words to represent each scan line (row of dots). This is a 2-byte integer. You obtain this value by calling GPR\_\$INQ\_BITMAP\_POINTER.

GMF\_\$COPY\_PLANE

## OUTPUT PARAMETERS

### **status**

Completion status, in STATUS\_\$T format. This data type is 4 bytes long. See the GPR Data Types section for more information.

## USAGE

To store an image in a GMF, you must have opened the GMF with the GMF\_\$OPEN call.

After storing an image in a GMF, close the GMF with the GMF\_\$CLOSE call.

The GMF\_\$COPY\_PLANE call is a special case of the GMF\_\$COPY\_SUBPLANE call.



**GMF\_\$COPY\_SUBPLANE**

Dumps a rectangular area of bits from virtual memory into a GMF.

**FORMAT**

GMF\_\$COPY\_SUBPLANE (stream\_id, black\_or\_white, bpi, bit\_pointer, x\_dim, y\_dim,  
x\_offset, y\_offset, width, status)

**INPUT PARAMETERS****stream\_id**

The stream ID of the GMF into which the image is to be stored, in STREAM\_\$ID\_T format. This is a 2-byte integer. You obtain the stream ID from the call to GMF\_\$OPEN that you used to open the GMF.

**black\_or\_white**

A Boolean variable. A value of TRUE means "1" bits are black and "0" bits are white. A value of FALSE means "1" bits are white and "0" bits are black. In the GMF, "1" bits are assumed to mean black. Thus if this parameter is false, the bits will be inverted as they are copied.

**bpi**

The number of bits per inch in the GMF. This information is stored in the GMF. It indicates the physical density of the image represented in the GMF. If this parameter is nonzero, a device to which you output the GMF may compress or expand the image to produce a result which is as close as possible to the image's original size. If this parameter is zero, an output device uses one dot to represent each bit from the GMF, regardless of the resulting physical size of the image. This is a 2-byte integer.

**bit\_pointer**

A pointer to a bit which when offset by x\_offset and y\_offset gives the upper left corner of the rectangular area to be stored. This is a 4-byte integer. You obtain this value by calling the routine GPR\_\$INQ\_BITMAP\_POINTER.

**x\_dim**

The x dimension of the rectangular area to be stored in the GMF. This is a 2-byte integer.

**y\_dim**

The y dimension of the rectangular area to be stored in the GMF. This is a 2-byte integer.

**x\_offset**

The x starting position of the rectangular area to be stored in the GMF relative to the bit whose address is given by bit\_pointer. This is a 2-byte integer.

**y\_offset**

The y starting position of the rectangular area to be stored in the GMF relative to the bit whose address is given by bit\_pointer. This is a 2-byte integer.

**width**

The number of 16-bit words per scan line in the source bitmap. The value of this parameter is usually 64. The width must be at least 1/16 of the specified x-dim. For instance, if you are storing an area 400 bits wide in a GMF, the source bitmap must use at least 25 words to represent each scan line (row of dots). This is a 2-byte integer. You obtain this value by calling GPR\_\$INQ\_BITMAP\_POINTER.

GMF\_\$COPY\_SUBPLANE

## OUTPUT PARAMETERS

### **status**

Completion status, in STATUS\_\$T format. This data type is 4 bytes long. See the GPR Data Types section for more information.

## USAGE

To copy a plane into a GMF, you must have opened the GMF with the GMF\_\$OPEN call.

After copying a plane into a GMF, close the GMF with the GMF\_\$CLOSE call.

The GMF\_\$COPY\_SUBPLANE call is a more general form of the GMF\_\$COPY\_PLANE call.

**GMF\_\$OPEN**

Opens or creates a GMF.

**FORMAT**

GMF\_\$OPEN (name, name\_length, start, stream\_id, status)

**INPUT PARAMETERS****name**

Pathname, in NAME\_\$PNAME\_T format.

**name\_length**

The length of the name. This is a 2-byte integer.

**start**

Desired position in the file after open, in GMF\_\$OPOS\_T format. This is a 2-byte integer. If you are opening the GMF to write data to it (to copy a plane or subplane into it), use one of these two constants:

GMF_\$APPEND	sets the initial position to EOF.
GMF_\$OVERWRITE	truncates the object to length 0 and sets the initial position to the beginning.

If you are opening the GMF to read data from it (restoring a plane), use this constant:

GMF_\$READ	sets the initial position to the beginning without truncating the GMF
------------	-----------------------------------------------------------------------

If the specified GMF does not exist and you used GMF\_\$OPEN to create it, it does not matter what value this parameter has.

**OUTPUT PARAMETERS****stream\_id**

The stream ID of the opened GMF, in STREAM\_\$ID\_T format. This is a 2-byte integer. You use this value in subsequent GMF calls that refer to the opened GMF.

**status**

Completion status, in STATUS\_\$T format. This data type is 4 bytes long. See the GPR Data Types section for more information.

GMF\_\$OPEN

## USAGE

If the specified GMF does not exist, the call to GMF\_\$OPEN creates it.

You must call GMF\_\$OPEN before trying to read or write a GMF.

After opening a GMF with GMF\_\$OPEN, you must eventually close it by calling GMF\_\$CLOSE.

**GMF\_\$RESTORE\_PLANE**

Copies an image back to the screen from a GMF.

**FORMAT**

GMF\_\$RESTORE\_PLANE (stream\_id, x\_dim, y\_dim, width, start, bpi, status)

**INPUT PARAMETERS****stream\_id**

The stream ID of the GMF which is to supply the image, in STREAM\_\$ID\_T format. This is a 2-byte integer. You obtain this parameter from the call to GMF\_\$OPEN you used to open the GMF.

**x\_dim**

The x-dimension in bits of the display to which an image is to be restored. This is a 2-byte integer.

**y\_dim**

The y-dimension in bits of the display to which an image is to be restored. This is a 2-byte integer.

**width**

The number of 16-bit words per scanline in the destination bitmap. This is a 2-byte integer.

**start**

The starting address in the destination bitmap. In Pascal this is a UNIV\_PTR. See the GPR Data Types section for more information.

**OUTPUT PARAMETERS****bpi**

Bits per inch as specified in GMF\_\$COPY\_PLANE. This is a 2-byte integer.

**status**

Completion status, in STATUS\_\$T format. This data type is 4 bytes long. See the GPR Data Types section for more information.

GMF\_\$RESTORE\_PLANE

## USAGE

Before calling GMF\_\$RESTORE\_PLANE, you must use GPR\_\$INIT to place the node in borrow-display mode.

The size of the area to be restored is the same as the size of the area you originally copied into the GMF. This information is contained in the GMF.

The area to be restored is determined by the bit-pointer specified in the GMF\_\$RESTORE\_PLANE call and the size data in the GMF. If this area runs off the right side or the bottom of the screen, the GMF manager restores only the portion of the stored image that fits on the screen.

To restore a plane from a GMF, you must have opened the GMF with the GMF\_\$OPEN call.

After restoring a plane from a GMF, you should close the GMF with the GMF\_\$CLOSE call.

**ERRORS**

GMF\_\$BAD\_BPI

Bits/inch parameter is negative.

GMF\_\$BAD\_POS

Opening position parameter is illegal.

GMF\_\$BAD\_WPL

16 bit words/line parameter is too small for x dim.

GMF\_\$BAD\_X\_DIM

X-dimension parameter is not positive.

GMF\_\$BAD\_Y\_DIM

Y-dimension parameter is not positive.

GMF\_\$NOT\_GMF

Opened file not a GMF metafile.

STATUS\_\$OK

Successful completion.





# IOS

## Table of Contents

Data Types	IOS-2
IOS_\$CHANGE_PATH_NAME	IOS-10
IOS_\$CLOSE	IOS-11
IOS_\$CREATE	IOS-12
IOS_\$DELETE	IOS-17
IOS_\$DUP	IOS-18
IOS_\$EQUAL	IOS-20
IOS_\$FORCE_WRITE_FILE	IOS-21
IOS_\$GET	IOS-22
IOS_\$GET_DIR	IOS-24
IOS_\$GET_EC	IOS-25
IOS_\$GET_HANDLE	IOS-27
IOS_\$INQ_BYTE_POS	IOS-28
IOS_\$INQ_CONN_FLAGS	IOS-30
IOS_\$INQ_CUR_REC_LEN	IOS-32
IOS_\$INQ_FILE_ATTR	IOS-33
IOS_\$INQ_FULL_KEY	IOS-34
IOS_\$INQ_MGR_FLAGS	IOS-36
IOS_\$INQ_OBJ_FLAGS	IOS-38
IOS_\$INQ_PATH_NAME	IOS-39
IOS_\$INQ_REC_POS	IOS-41
IOS_\$INQ_REC_REMAINDER	IOS-43
IOS_\$INQ_REC_TYPE	IOS-44
IOS_\$INQ_SHORT_KEY	IOS-45
IOS_\$INQ_TYPE_UID	IOS-47
IOS_\$LOCATE	IOS-49
IOS_\$OPEN	IOS-51
IOS_\$PUT	IOS-53
IOS_\$REPLICATE	IOS-55
IOS_\$SEEK	IOS-57
IOS_\$SEEK_FULL_KEY	IOS-59
IOS_\$SEEK_SHORT_KEY	IOS-60
IOS_\$SEEK_TO_BOF	IOS-61
IOS_\$SEEK_TO_EOF	IOS-62
IOS_\$SET_CONN_FLAG	IOS-63
IOS_\$SET_DIR	IOS-65
IOS_\$SET_LOCATE_BUFFER_SIZE	IOS-66
IOS_\$SET_OBJ_FLAG	IOS-67
IOS_\$SET_REC_TYPE	IOS-69
IOS_\$SWITCH	IOS-70
IOS_\$TRUNCATE	IOS-71
Errors	IOS-72



# IOS

The IOS (I/O Switch) programming calls perform device-independent I/O. This section describes their data types, call syntax, and error codes. Refer to the Introduction at the beginning of this manual for a description of data-type diagrams and call syntax format.

## IOS DATA TYPES

### CONSTANTS

IOS_\$MAX	127	Highest possible number in stream ID.
IOS_\$NO_STREAM	16#7FFF	Placeholder for stream ID.

### VARIABLES

XOID_\$NIL		A variable whose value is the NIL XOID and doesn't change. Used for comparisons and assignments of XOID_\$T variables.
------------	--	------------------------------------------------------------------------------------------------------------------------

### DATA TYPES

IOS_\$ABS_REL_T		A 2-byte integer. Specifies whether seek is relative or absolute. One of the following predefined values:
-----------------	--	-----------------------------------------------------------------------------------------------------------

IOS\_\$RELATIVE  
Seek from the current position.

IOS\_\$ABSOLUTE  
Seek from the beginning of the object (BOF).

IOS_\$CONN_FLAG_T		A 2-byte integer. Attributes associated with a stream connection. One of the following predefined values:
-------------------	--	-----------------------------------------------------------------------------------------------------------

IOS\_\$CF\_TTY  
Connection behaves like a terminal.

IOS\_\$CF\_IPC  
Connection behaves like an interprocess communication (IPC) channel.

IOS\_\$CF\_VT  
Connection behaves like a DOMAIN Display Manager pad.

IOS\_\$CF\_WRITE  
Connection can be written to.

IOS\_\$CF\_APPEND  
Connection's stream marker can be positioned to the end of the object before each put call.

IOS\_\$CF\_UNREGULATED  
Other processes can read and write to the connection.

IOS\_\$CF\_READ\_INTEND\_WRITE  
Connection open for read access, and can later

be open for write access. Other processes can have read access.

IOS\_\$CONN\_FLAG\_SET

A 4-byte integer. A set of connection attributes, in IOS\_\$CONN\_FLAG\_T format, indicating which attributes of the specified connection are set. For a list of options, see IOS\_\$CONN\_FLAG\_T above.

IOS\_\$CREATE\_MODE\_T

A 2-byte integer. Specifies the action to be taken if the name already exists or specifies creation of unnamed objects. One of the following predefined values:

IOS\_\$LOC\_NAME\_ONLY\_MODE

Create a temporary unnamed object, uses pathname to specify location of object, and locates it on the same volume.

IOS\_\$MAKE\_BACKUP\_MODE

Create a backup (.bak) object when closed.

IOS\_\$NO\_PRE\_EXIST\_MODE

Return an error if object already exists.

IOS\_\$PRESERVE\_MODE

Save contents of object, if it exists, opens object, and positions stream marker at the beginning of the object (BOF).

IOS\_\$RECREATE\_MODE

Delete existing object and creates new one of same name.

IOS\_\$TRUNCATE\_MODE

Open object, then truncates the contents.

IOS\_\$DIR\_TYPE\_T

A 2-byte integer. Specifies type of directory. One of the following predefined values:

IOS\_\$WDIR

Current working directory.

IOS\_\$NDIR

Current naming directory

IOS\_\$EC\_KEY\_T

A 2-byte integer. Specifies eventcount key type. One of the following predefined values:

IOS\_\$GET\_EC\_KEY

Key that is advanced with each get call.

IOS\_\$PUT\_EC\_KEY

Key that is advanced with each put call.

## IOS DATA TYPES

IOS\_\$ID\_T

A 2-byte integer, ranging in value from 0 to IOS\_\$MAX. The stream ID.

IOS\_\$MGR\_FLAG\_T

A 2-byte integer. Object attributes associated with an object's manager. One of the following predefined values:

IOS\_\$MF\_CREATE

Manager permits type to create objects.

IOS\_\$MF\_CREATE\_BAK

Manager permits type to create backup (.bak) objects.

IOS\_\$MF\_IMEX

Manager permits type to export streams to new processes.

IOS\_\$MF\_FORK

Manager permits type to pass streams to forked processes.

IOS\_\$MF\_FORCE\_WRITE

Manager permits type to force-write object contents to stable storage (for most object types, this is the disk).

IOS\_\$MF\_WRITE

Manager permits objects to be written to.

IOS\_\$MF\_SEEK\_ABS

Manager permits objects to perform absolute seeks.

IOS\_\$MF\_SEEK\_SHORT

Manager permits objects to seek using short (4-byte) seek keys.

IOS\_\$MF\_SEEK\_FULL

Manager permits objects to seek using full (8-byte) seek keys.

IOS\_\$MF\_SEEK\_BYTE

Manager permits objects to seek to byte positions.

IOS\_\$MF\_SEEK\_REC

Manager permits objects to seek to record positions.

IOS\_\$MF\_SEEK\_BOF

Manager permits objects to seek to the beginning of the object.

IOS\_\$MF\_REC\_TYPE

Manager supports different record type formats.

IOS\_ \$MF\_ TRUNCATE

Manager permits objects to be truncated.

IOS\_ \$MF\_ UNREGULATED

Manager permits objects to have shared (unregulated) concurrency mode.

IOS\_ \$MF\_ SPARSE

Manager permits objects to be as sparse.

IOS\_ \$MF\_ READ\_ INTEND\_ WRITE

Manager permits objects to have read-intend-write access.

IOS\_ \$MGR\_ FLAG\_ SET

A 4-byte integer. A set of object manager attributes, in IOS\_ \$MGR\_ FLAG\_ T format, indicating which attributes of the specified object's manager are set. For a list of options, see IOS\_ \$MGR\_ FLAG\_ T above.

IOS\_ \$NAME\_ TYPE\_ T

A 2-byte integer. Specifies format of pathname. One of the following predefined values:

IOS\_ \$ROOT\_ NAME

Absolute pathname relative to the network root directory (/); for example, //node/sid/file.

IOS\_ \$WDIR\_ NAME

Leaf name if object's name is a name in current working directory; otherwise, specifies absolute pathname.

IOS\_ \$NDIR\_ NAME

Leaf name if object's name is a name in current naming directory; otherwise, specifies absolute pathname.

IOS\_ \$NODE\_ NAME

Name relative to the node's entry directory (/) if object is a name in boot volume; otherwise, specifies absolute pathname; for example, /sid/file.

IOS\_ \$NODE\_ DATA\_ FLAG

Leaf name if object's name is a name in current 'node\_data directory; otherwise, specifies absolute pathname.

IOS\_ \$LEAF\_ NAME

Leaf name regardless of object's name.

IOS DATA TYPES

IOS\_\$RESID\_NAME

Residual name if object is defined using extended naming.

IOS\_\$OBJ\_FLAG\_T

A 2-byte integer. Attributes associated with an object. One of the following predefined values:

IOS\_\$OF\_DELETE\_ON\_CLOSE

Object can be deleted when all its associated connections are closed.

IOS\_\$OF\_SPARSE\_OK

Object can be written as a sparse object.

IOS\_\$OF\_ASCII

Object contains ASCII data.

IOS\_\$OF\_FTNCC

Object uses FORTRAN carriage control characters.

IOS\_\$OF\_COND

Object performs get or put calls conditionally, as if the IOS\_\$COND\_OPT was specified.

IOS\_\$OBJ\_FLAG\_SET

A 4-byte integer. A set of object attributes, in IOS\_\$OBJ\_FLAG\_T format, indicating which attributes of the specified object are set. For a list of options, see IOS\_\$OBJ\_FLAG\_T above.

IOS\_\$OPEN\_OPTIONS\_T

A 2-byte integer. Specifies options for an IOS\_\$OPEN. Any combination of the following predefined values:

IOS\_\$NO\_OPEN\_DELAY\_OPT

Return immediately instead of waiting for open to complete.

IOS\_\$WRITE\_OPT

Permit writing data to a new object.

IOS\_\$UNREGULATED\_OPT

Permit concurrency (unregulated read and write access.) to the object

IOS\_\$POSITION\_TO\_EOF\_OPT

Position stream marker to the end of the object at open.

IOS\_\$INQUIRE\_ONLY\_OPT

Open object for attribute inquiries only.

IOS\_\$READ\_INTEND\_WRITE\_OPT

Object has read-intend-write access, other processes can have read but not write access.



IOS\_\$POS\_OPT\_T

A 2-byte integer. Specify position to return when inquiring about object position. One of the following predefined values:

IOS\_\$CURRENT

Return key for the current stream marker.

IOS\_\$BOF

Return key for beginning of the object (BOF) marker.

IOS\_\$EOF

Return key for end of the object (EOF) marker.

IOS\_\$PUT\_GET\_OPTS\_T

A 2-byte integer. Specifies options for put and get operations. Any combination of the following predefined values:

IOS\_\$COND\_OPT

Read or write data conditionally. If call fails, returns

IOS\_\$xxx\_CONDITIONAL\_FAILED,  
where xxx is either GET or PUT.

IOS\_\$PREVIEW\_OPT

Write data but do not update the stream marker.

IOS\_\$PARTIAL\_RECORD\_OPT

Write a portion of a record but do not terminate it.

IOS\_\$NO\_REC\_BNDRY\_OPT

Ignore record (line) boundaries.

IOS\_\$RTYPE\_T

A 2-byte integer. Specifies the record type format. One of the following predefined values:

IOS\_\$V1

Variable-length record with count fields.

IOS\_\$F2

Fixed-length records with count fields.

IOS\_\$UNDEF

No record structure.

IOS\_\$EXPLICIT\_F2

Fixed-length records that IOS\_\$PUT cannot implicitly change to IOS\_\$V1.

IOS\_\$F1

Fixed-length records without count fields.

IOS DATA TYPES

IOS\_\$SEEK\_KEY\_T

The full seek key. This is an 8-byte integer value.

IOS\_\$SEEK\_TYPE\_T

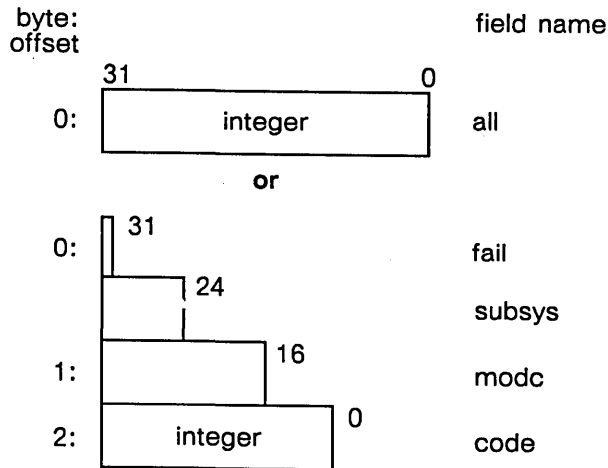
A 2-byte integer. Specifies the type of seek to perform. One of the following predefined values:

IOS\_\$REC\_SEEK  
Record-oriented seek.

IOS\_\$BYTE\_SEEK  
Byte-oriented seek.

STATUS\_\$T

A status code. The diagram below illustrates the STATUS\_\$T data type:



Field Description:

all  
All 32 bits in the status code.

fail  
The fail bit. If this bit is set, the error was not within the scope of the module invoked, but occurred within a lower-level module (bit 31).

subsys  
The subsystem that encountered the error (bits 24 - 30).

modc  
The module that encountered the error (bits 16 - 23).


code  
A signed number that identifies the type of error that occurred (bits 0 - 15).

UID\_ \$T

An object type identifier. This is an 8-byte integer value.

XOID\_ \$T

Unique identifier of an object. Used by type managers only. The diagram below illustrates the XOID\_ \$T data type:



Field Description:

rfu1  
Reserved for future use.

rfu2  
Reserved for future use.

UID  
Unique identifier for an object.

IOS\_\$CHANGE\_PATH\_NAME

IOS\_\$CHANGE\_PATH\_NAME

Changes the pathname of an object.

## FORMAT

IOS\_\$CHANGE\_PATH\_NAME (stream-id, new-pathname, new-namelenh, status)

## INPUT PARAMETERS

### stream-id

Number of the stream on which the object is open, in IOS\_\$ID\_T format. This is a 2-byte integer.

### new-pathname

New name of the object, in NAME\_\$PNAME\_T format. This is an array of up to 256 characters.

### new-namelenh

Length of "new-pathname." This is a 2-byte integer.

## OUTPUT PARAMETERS

### status

Completion status, in STATUS\_\$T format. This data type is 4 bytes long. See the IOS Data Types section for more information.

## USAGE

IOS\_\$CHANGE\_PATH\_NAME changes the pathname of an existing object. The stream ID of the object remains the same.

IOS\_\$CHANGE\_PATH\_NAME permits you to assign a name to a previously unnamed object and, conversely, to remove a name from a previously named object. (To remove a name, specify a null pathname.)

Note that this call can change the delete-on-close object attribute. For example, if you assign a name to an unnamed object, the operation implicitly changes the delete-on-close attribute to FALSE. Likewise, if you specify a null pathname for a previously named object, the operation implicitly changes the delete-on-close attribute to TRUE. Be aware that this behavior can cause unexpected results in cases where you explicitly change the delete-on-close attribute, and then make an unnamed-to-named name change.

**IOS\_\$CLOSE**

Closes a stream.

**FORMAT**

IOS\_\$CLOSE (stream-id, status)

**INPUT PARAMETERS****stream-id**

Number of the stream to be closed, in IOS\_\$ID\_T format. This is a 2-byte integer.

Once IOS\_\$CLOSE closes the stream, the number used for this stream ID becomes available for reuse. If the object is open on more than one stream, IOS\_\$CLOSE closes only "stream-id."

**OUTPUT PARAMETERS****status**

Completion status, in STATUS\_\$T format. This data type is 4 bytes long. See the IOS Data Types section for more information.

**USAGE**

IOS\_\$CLOSE closes the stream so that you can no longer use the stream ID to operate on the object. Closing a stream to an object releases any locks maintained by the stream connection, thus making the object available to other users.

A program can close only the streams that it has opened at the current or lower program levels (that is, streams opened by programs that the calling program has invoked). IOS\_\$CLOSE returns an error status code if you try to close a stream that was opened at a higher program level.

If an object has the delete-on-close attribute (IOS\_\$OF\_DELETE\_ON\_CLOSE), IOS\_\$CLOSE deletes the object. However, the object is not deleted until *all* streams to it are closed. (For details on object attributes, see the IOS\_\$INQ\_\$OBJ\_FLAGS and IOS\_\$SET\_OBJ\_FLAG calls.)

IOS\_\$CREATE

IOS\_\$CREATE

Creates an object and opens a stream to it.

## FORMAT

IOS\_\$CREATE (pathname, namelength, type-uid, create-options,  
open-options, stream-id, status)

## INPUT PARAMETERS

### pathname

Name of the object to be created, in NAME\_\$PNAME\_T format. This is an array of up to 256 characters. To create a temporary object, see the section "Creating an Object in Backup Mode" below.

### namelength

Length of "pathname," in bytes. This is a 2-byte integer. To create a temporary object, see the section "Creating an Object in Backup Mode" below.

### type-uid

UID of the type to be created, in UID\_\$T format. This data type is 8 bytes long. See the IOS Data Types section for more information.

If you specify the predefined UID\_\$NIL, IOS\_\$CREATE creates an object of the default type, which is currently unstructured ASCII (UASC). You can also specify any of the system's predefined type UIDs listed below, or any valid user-created type UID.

DOMAIN currently supports a set of standard object types which include the following types. (Note that objects created by type managers return manager-specific type UIDs.)

Type UID	Object
UASC_\$UID	UASC object
RECORDS_\$UID	Record-oriented object
HDR_UNDEF_\$UID	Nonrecord-oriented object
OBJECT_FILE_\$UID	Object module object (compiler or binder output)
SIO_\$UID	Serial line descriptor object
MT_\$UID	Magnetic tape descriptor object
PAD_\$UID	Saved Display Manager transcript pad
INPUT_PAD_\$UID	Display Manager input pad
MBX_\$UID	Mailbox object

Type UID	Object
DIRECTORY_\$UID	Directory
NULLDEV_\$UID	Null device

**create-options**

Specifies the action to be taken if the object already exists, or specifies the creation of an unnamed object, in IOS\_\$CREATE\_MODE\_T format. This is a 2-byte integer. Specify one of the following predefined values:

IOS_\$NO_PRE_EXIST_MODE	Return the IOS_\$ALREADY_EXISTS error status code if an object with the specified name already exists.
IOS_\$PRESERVE_MODE	Preserve the contents of the object if an object with the specified name already exists. Then open the object and position the stream marker to the beginning of the object (BOF) unless you set the IOS_\$POSITION_TO_EOF open option. Use this mode to append data to an existing object.
IOS_\$RECREATE_MODE	Recreate the object if an object with the specified name already exists. Essentially, this option deletes the existing object and creates a new one. The new object will have the default set of attributes for that object type.
IOS_\$TRUNCATE_MODE	Open the object and delete the contents if an object with the specified name already exists. Use this mode to create an object to preserve the attributes of the specified object.
IOS_\$MAKE_BACKUP_MODE	Create a temporary object with the same type and attributes as the object specified in the pathname if an object with the specified name already exists. Use this mode to create a backup object. (See below for detailed description.)
IOS_\$LOC_NAME_ONLY_MODE	Create a temporary unnamed object. Use the pathname to specify the location of the object. IOS_\$CREATE will locate the temporary object on the same volume as the object specified in the pathname.

IOS\_\$CREATE

**open-options**

Open options, in IOS\_\$OPEN\_OPTIONS\_T format. This is a 2-byte integer. Specify a combination of the following set of predefined values:

- |                             |                                                                                                                                                                                                                                                                                                                          |
|-----------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| IOS_\$NO_OPEN_DELAY_OPT     | Return immediately, instead of waiting for the open call to complete.                                                                                                                                                                                                                                                    |
| IOS_\$WRITE_OPT             | Permit writing data to a new object. If a program tries to write on a stream for which you have not specified this option, it returns an error status. Note that when creating an object, the IOS manager automatically sets this value because it assumes that when you create an object, you will want to write to it. |
| IOS_\$UNREGULATED_OPT       | Permit shared (unregulated) concurrency mode.                                                                                                                                                                                                                                                                            |
| IOS_\$POSITION_TO_EOF_OPT   | Position the stream marker at the end of the object (EOF). Use this to append data to an existing object.                                                                                                                                                                                                                |
| IOS_\$INQUIRE_ONLY_OPT      | Open the object for attribute inquiries only; do not permit reading or writing of data.                                                                                                                                                                                                                                  |
| IOS_\$READ_INTEND_WRITE_OPT | Open the object for read access with the intent to eventually change the object's access to write access. This allows other processes to read the object; but they cannot have write or read-intend-write access.                                                                                                        |

**OUTPUT PARAMETERS**

**stream-id**

Number of the stream on which the object is open, in IOS\_\$ID\_T format. This is a 2-byte integer.

Subsequent IOS calls use this number to identify the stream opened by this call.

**status**

Completion status, in STATUS\_\$T format. This data type is 4 bytes long. See the IOS Data Types section for more information.



## USAGE

If the pathname specifies an object that does not exist, IOS\_\$CREATE creates a new object of the specified type using that pathname and opens a stream to it. If the object already exists, the create mode option specified in the call determines which action IOS\_\$CREATE will perform.

Both IOS\_\$CREATE and IOS\_\$OPEN open a stream to an object. However, IOS\_\$CREATE creates the object if it does not exist, whereas IOS\_\$OPEN returns an error if the object does not exist.

### Inquiring about Object Attributes

When IOS\_\$CREATE creates an object, the object has a default set of attributes (the default attributes depend on the type created). These attributes fall into three categories: manager, object, and connection attributes. To determine which attributes the newly created object has, you can use the following calls:

IOS\_\$INQ\_MGR\_FLAGS

Returns the attributes that the object's type manager defines.

IOS\_\$INQ\_OBJ\_FLAGS

Returns the attributes of the object.

IOS\_\$INQ\_CONN\_FLAGS

Returns the attributes of the stream connection.

To change object or connection attributes, use the IOS\_\$SET\_OBJ\_FLAGS, and IOS\_\$SET\_CONN\_FLAGS calls, respectively. The attributes that you can change depend on the object type. Note that you cannot change manager attributes because the type manager determines them. For details on writing a type manager, see the *Extending the DOMAIN Streams Facility* manual.

### Creating a Temporary Object

IOS\_\$CREATE allows you to create a temporary object two ways. To create a temporary object on your boot volume, specify a null value in "pathname" and a value of 0 in "namelength." To create a temporary object on another volume, specify the pathname of an existing object on that volume with the IOS\_\$LOC\_NAME\_ONLY\_MODE option in "create-options." IOS\_\$CREATE creates a temporary unnamed object on the same volume (node) as the object you specify in "pathname."

### Creating an Object in Backup Mode

You can create a new, *unnamed* temporary object by specifying the create mode option, IOS\_\$MAKE\_BACKUP\_MODE. The call creates the new object with the same type and attributes as the object specified by "pathname" (if it exists), and it is created on the same volume (node).

IOS\_\$CREATE does not open or modify the object specified by "pathname," it merely examines the object to extract its attributes. Even though IOS\_\$CREATE does not modify the "pathname," it conceptually replaces the object, so this operation requires write access to the object.

## IOS\_\$CREATE

When IOS\_\$CLOSE closes the stream created with this call, it changes the object specified by "pathname" to "pathname.bak." It changes the new object (formerly the temporary, unnamed object) to "pathname," and makes the object permanent.

If a ".bak" version of the object already exists, IOS\_\$CLOSE deletes it. (The caller must have either D or P rights to delete the object.) If the ".bak" object is locked at the time IOS\_\$CLOSE is called, the object will be deleted when it is unlocked.

If "pathname" does not exist at the time that IOS\_\$CREATE is called, then IOS\_\$CREATE performs the ordinary functions.

**IOS\_\$DELETE**

Deletes an object and closes the associated stream.

**FORMAT**

IOS\_\$DELETE (stream-id, status)

**INPUT PARAMETERS****stream-id**

Number of a stream on which the object is open, in IOS\_\$ID\_T format. This is a 2-byte integer.

**OUTPUT PARAMETERS****status**

Completion status, in STATUS\_\$T format. This data type is 4 bytes long. See the IOS Data Types section for more information.

**USAGE**

IOS\_\$DELETE deletes the object, then closes the specified stream.

This call actually sets the object attribute IOS\_\$OF\_DELETE\_ON\_CLOSE to TRUE, then closes the stream. So, if the type manager does not allow an object to set the delete-on-close attribute, the delete call fails. In this case, the call closes the stream, but does not delete the object.

If the object is open on more than one stream, IOS\_\$DELETE marks the object for deletion, but the object still exists until *all* streams to that object are closed.

IOS\_ \$DUP

IOS\_ \$DUP

Creates a copy of a specified stream ID.

### FORMAT

return\_stream\_id = IOS\_ \$DUP (stream\_id\_to\_duplicate, copy\_stream\_id, status)

### RETURN VALUE

return\_stream\_id

Number of the new stream created, in IOS\_ \$ID\_ T format. This is a 2-byte integer.

### INPUT PARAMETERS

stream\_id\_to\_duplicate

Number of the stream to duplicate, in IOS\_ \$ID\_ T format. This is a 2-byte integer. This stream number remains a valid connection to the object after IOS\_ \$DUP completes successfully.

copy\_stream\_id

Number of the stream to use as the newly created copy, in IOS\_ \$ID\_ T format. This is a 2-byte integer.

If "copy\_stream\_id" is free, IOS\_ \$DUP returns that value in "return\_stream\_id." If "copy\_stream\_id" is in use, IOS\_ \$DUP begins searching from that number upward (higher numbers) until it finds a free stream number and returns that number in "return\_stream\_id."

If the actual number of "copy\_stream\_id" is insignificant, specify the value 0. This value causes IOS\_ \$DUP to begin searching from the lowest possible stream number and return the first free stream number it finds.

### OUTPUT PARAMETERS

status

Completion status, in STATUS\_ \$T format. This data type is 4 bytes long. See the IOS Data Types section for more information.

### USAGE

Use IOS\_ \$DUP to create a copy of an existing stream ID. The new stream ID refers to the same connection as the existing stream ID. Note that you must close *both* streams with IOS\_ \$CLOSE before the stream connection actually closes.

You can use IOS\_ \$DUP to keep a stream connection open when passing it to a subroutine. Use IOS\_ \$DUP to create a copy of the stream ID before passing it. This way, the subroutine cannot close the connection to the object because *all* copies of the stream connection must be closed before the connection itself closes.

IOS\_ \$DUP is identical to IOS\_ \$REPLICATE except that IOS\_ \$DUP looks for a free stream number in *ascending* order from the specified stream ID, while IOS\_ \$REPLICATE looks in *descending* order. Note that you use IOS\_ \$DUP or IOS\_ \$REPLICATE to *copy* existing stream ID's, both the existing and new stream ID's remain valid connections. However, you use IOS\_ \$SWITCH to *replace* stream IDs; you "switch" the connection from the existing stream ID to the new stream ID.

IOS\_\$EQUAL

IOS\_\$EQUAL

Determines whether two stream IDs refer to the same object.

#### FORMAT

same IOS\_\$EQUAL (stream\_id, stream\_id\_too, status)

#### RETURN VALUE

##### same

Boolean value that indicates whether the specified stream IDs refer to the same object. "Same" is TRUE if the streams refer to the same object, it is FALSE if they do not.

#### INPUT PARAMETERS

##### stream\_id

Number of a stream being compared, in IOS\_\$ID\_T format. This is a 2-byte integer.

##### stream\_id\_too

Number of a stream being compared, in IOS\_\$ID\_T format. This is a 2-byte integer.

#### OUTPUT PARAMETERS

##### status

Completion status, in STATUS\_\$T format. This data type is 4 bytes long. See the IOS Data Types section for more information.

#### USAGE

Use IOS\_\$EQUAL to determine whether two stream IDs refer to the same object. An application program can use this call to avoid using two streams when one is sufficient.

**IOS\_\$FORCE\_WRITE\_FILE**

Forcibly writes an object to permanent storage.

**FORMAT**

IOS\_\$FORCE\_WRITE\_FILE (stream-id, status)

**INPUT PARAMETERS****stream-id**

Number of the stream on which the object is open, in IOS\_\$ID\_T format. This is a 2-byte integer.

**OUTPUT PARAMETERS****status**

Completion status, in STATUS\_\$T format. This data type is 4 bytes long. See the IOS Data Types section for more information.

**USAGE**

IOS\_\$FORCE\_WRITE\_FILE forcibly writes the object to stable storage. Stable storage depends on the object's type, however, in most cases, it is the disk. For example, stable storage for a magnetic tape descriptor is the tape.

Use IOS\_\$FORCE\_WRITE\_FILE before closing the stream to ensure that the object is stored safely in the event of a system crash.

IOS\_\$GET

IOS\_\$GET

Moves data from a stream into a buffer.

## FORMAT

ret-length = IOS\_\$GET (stream-id, get-options, buffer, buffer-size, status)

## RETURN VALUE

### ret-length

Amount of data moved, in bytes. This is a 4-byte integer.

"Ret-length" equals the amount of data read; "ret-length" equals 0 if IOS\_\$GET does not return any data.

If the length of the data read exceeds the amount specified in "data-size," IOS\_\$GET performs the following:

- Reads enough data to fill the requested size
- Sets "ret-length" equal to "data-size"
- Positions the stream marker to the first unread byte
- Returns the IOS\_\$BUFFER\_TOO\_SMALL status code to indicate that this condition has occurred

You can inquire about how many bytes remain to be read in the current record by calling IOS\_\$INQ\_REC\_REMAINDER.

## INPUT PARAMETERS

### stream-id

Number of the stream on which the object is open, in IOS\_\$ID\_T format. This is a 2-byte integer.

### get-options

Options that control how IOS\_\$GET performs the get operation, in IOS\_\$PUT\_GET\_OPTS\_T format. This is a 2-byte integer. Specify a combination of the following set of predefined values:

IOS_\$COND_OPT	Reads data, if available. (For example, data on an SIO line is not always available immediately.) If the data is not available, IOS_\$GET returns the IOS_\$GET_CONDITIONAL_FAILED status code and sets the return value of "ret-length" to 0.
IOS_\$PREVIEW_OPT	Reads data but does not update the stream marker.



**IOS\_\$NO\_REC\_BNDRY\_OPT** Ignores record boundaries while reading data. For example, it ignores NEWLINE characters in a UASC object, which guarantees that the call fills the specified buffer. Some type managers might not support this option.

**IOS\_\$PARTIAL\_RECORD\_OPT** Not meaningful for this call.

**buffer-size**

Maximum number of bytes to be moved to the buffer. This is a 4-byte integer.

**OUTPUT PARAMETERS**

**buffer**

Buffer to store the data. This is a character array.

**status**

Completion status, in STATUS\_\$T format. This data type is 4 bytes long. See the IOS Data Types section for more information.

**USAGE**

You can use either of IOS\_\$LOCATE or IOS\_\$GET to read data from system objects. IOS\_\$GET copies the data into a buffer, while IOS\_\$LOCATE returns the virtual address of the data.

In most cases, use the IOS\_\$LOCATE call to read data because it is faster (IOS\_\$LOCATE does not perform a copy).

You will want to use IOS\_\$GET when you need to read more data than can be obtained in one call, because the pointer remains valid for only one call. For example, use IOS\_\$GET when you need to read and rearrange a number of lines from an object.

IOS\_\$GET\_DIR

IOS\_\$GET\_DIR

Gets the current working or naming directory.

### FORMAT

IOS\_\$GET\_DIR (pathname, namelength, dir\_type, status)

### INPUT PARAMETERS

**dir\_type**

Option specifying which type of directory to get, in IOS\_\$DIR\_TYPE\_T format. Specify one of the predefined values:

IOS\_\$WDIR     Name of the current working directory.

IOS\_\$NDIR     Name of the current naming directory.

### OUTPUT PARAMETERS

**pathname**

Name of the directory to get, in NAME\_\$PNAME\_T format. This is an array of up to 256 characters.

**namelength**

Length of "pathname." This is a 2-byte integer.

**status**

Completion status, in STATUS\_\$T format. This data type is 4 bytes long. See the IOS Data Types section for more information.

### USAGE

Use this call to get the current working or naming directory. It returns the name of the directory in the "pathname" parameter. If you want to change the current working or naming directory, use IOS\_\$SET\_DIR.

**IOS\_\$GET\_EC**

Gets a pointer to an eventcount for a stream event.

**FORMAT**

IOS\_\$GET\_EC (stream-id, stream-key, eventcount-pointer, status)

**INPUT PARAMETERS****stream-id**

Number of stream on which the eventcount is waiting, in IOS\_\$ID\_T format. This is a 2-byte integer.

**stream-key**

The key that specifies which type of eventcount to get a pointer to, in IOS\_\$EC\_KEY\_T format. This is a 2-byte integer. Specify one of the following predefined values:

IOS\_\$GET\_REC\_EC\_KEY      An eventcount that advances when the stream contains data for you to get. This eventcount advances whenever there is anything to get from an open stream.

IOS\_\$PUT\_REC\_EC\_KEY      An eventcount that advances when a previously "full" stream might now be able to accept data. A full stream is a stream that IOS\_\$PUT will block.

**OUTPUT PARAMETERS****eventcount-pointer**

A pointer to the eventcount, in EC2\_\$PTR\_T format. This is a 4-byte integer address that points to an array of eventcounts. See the EC2 Data Types section for more information.

**status**

Completion status, in STATUS\_\$T format. This data type is 4 bytes long. See the IOS Data Types section for more information.

**USAGE**

IOS\_\$GET\_EC is valid for all streams, including those open to objects, pads, mailboxes, and devices. After you use this call to get a stream event, use EC2 calls to read eventcount values and wait for events.

You can wait for two types of events on a stream:

- The IOS-get eventcount indicates that there might be input to get from an open stream.
- The IOS-put eventcount indicates that a previously "full," or blocked, stream might now have enough room to accept the data.

An example of using the IOS-get eventcount is to wait for keyboard input. Whenever the

user types data, the system advances the eventcount associated with the user's input pad. If input pad is in normal (or cooked) mode, the eventcount advances after each carriage return, if the input pad is in raw mode, the eventcount advances after each keystroke. (For details on cooked and raw mode, see the Display Manager chapter in the *Programming with General System Calls* manual.)

An example of using the IOS-put eventcount is to wait on an MBX channel that might get blocked. That is, IOS\_\$PUT blocks streams associated with MBX channels if a server is not ready for the data from the channel. When it's possible to write data without blocking, the system advances the IOS-put eventcount.

For more information on eventcounts, see the *Programming with General System Calls* and the *Programming with System Calls for Interprocess Communication* manuals.

## IOS\_ \$GET\_ HANDLE

Converts a stream ID to a handle pointer.

## FORMAT

handle = IOS\_ \$GET\_ HANDLE (stream-id, type-uid, status)

## RETURN VALUE

**handle**

Pointer to the handle associated with the stream connection, in UNIV\_ PTR format. This is a 4-byte integer.

## INPUT PARAMETERS

**stream-id**

Number of the stream that identifies an open stream, in IOS\_ \$ID\_ T format. This is a 2-byte integer.

**type-uid**

Type UID of the object that the type manager handles, in UID\_ \$T format. Specify the type UID of the manager you are writing. This data type is 8 bytes long. See the IOS Data Types section for more information.

## OUTPUT PARAMETERS

**status**

Completion status, in STATUS\_ \$T format. This data type is 4 bytes long. See the IOS Data Types section for more information.

## USAGE

NOTE: This call is generally of interest to type manager writers only.

Type manager writers use this call to access an object when implementing an operation that is not predefined by the system. When the type manager implements such an operation, it is referred to as a **direct manager call** because the I/O switch does not route the call between the client call and the manager. Without switching, the manager receives a stream ID from a client. To access the object, the manager must then call IOS\_ \$GET\_ HANDLE to obtain the object handle associated with the stream ID.

IOS\_ \$GET\_ HANDLE returns an error if the stream ID is not associated with an object of the type UID specified by "type\_ uid." Specify the type UID of the manager you are writing so that the manager can be sure it has a stream to an object of its type.

See the *Using the Open System Toolkit to Extend the Streams Facility* manual for more information.

IOS\_\$INQ\_BYTE\_POS

IOS\_\$INQ\_BYTE\_POS

Returns the byte position of the stream marker.

## FORMAT

byte-position = IOS\_\$INQ\_BYTE\_POS (stream-id, position-option, status)

## RETURN VALUE

### byte-position

Byte position of the stream marker. This is a 4-byte integer. Note that byte positions are *zero*-based; consequently the byte position of the beginning of an object (BOF) is 0.

## INPUT PARAMETERS

### stream-id

Number of the stream on which the object is open, in IOS\_\$ID\_T format. This is a 2-byte integer.

### position-option

Value specifying the byte position to return, in IOS\_\$POS\_OPT\_T format. This is a 2-byte integer. Specify one of the following predefined values:

- |               |                                                                                                                           |
|---------------|---------------------------------------------------------------------------------------------------------------------------|
| IOS_\$CURRENT | Returns the byte position of the current stream marker.                                                                   |
| IOS_\$EOF     | Returns the byte position of the stream marker at the end of the object (EOF). This is the number of bytes in the object. |
| IOS_\$BOF     | Return the byte position of the stream marker at the beginning of the object (BOF). This value is always 0.               |

## OUTPUT PARAMETERS

### status

Completion status, in STATUS\_\$T format. This data type is 4 bytes long. See the IOS Data Types section for more information.

## USAGE

To obtain the offset of the stream marker, use IOS\_\$INQ\_BYTE\_POS. (Use IOS\_\$INQ\_REC\_POS if your object is record-oriented.)

To get the offset of the stream marker at the beginning or end of the object, specify IOS\_\$BOF or IOS\_\$EOF, in the "position-option" parameter. Specify IOS\_\$CURRENT to get the offset of the stream marker from the beginning of the object. Once you have the returned offset, you can move the stream marker to desired location by calling IOS\_\$SEEK.

This call allows you to perform a nonkeyed seek by specifying an absolute byte position, or by getting an offset from an absolute position, and moving the stream marker to it.

Whether you perform a nonkeyed or keyed seek depends on how the object's data is represented. For example, programs that need to perform "arithmetic" on the data (such as comparing two positions) will use nonkeyed seek operations. Programs that require only the ability to move from one position to another in an object will use keyed seek operations.

IOS\_\$INQ\_CONN\_FLAGS

IOS\_\$INQ\_CONN\_FLAGS

Returns the attributes associated with a connection.

## FORMAT

conn\_flags = IOS\_\$INQ\_CONN\_FLAGS (stream-id, status)

## RETURN VALUE

### conn\_flags

A set (bit mask) indicating which attributes of the specified connection are set, in IOS\_\$CONN\_FLAG\_SET format. This is a 4-byte integer. Any combination of the following set of predefined values, in IOS\_\$CONN\_FLAG\_T format, can be returned. If the set contains the value, the connection has the attribute.

IOS_\$CF_TTY	Connection behaves like a terminal.
IOS_\$CF_IPC	Connection behaves like an interprocess communication (IPC) channel.
IOS_\$CF_VT	Connection behaves like a DOMAIN Display Manager pad.
IOS_\$CF_WRITE	Connection can be written to.
IOS_\$CF_APPEND	Connection's stream marker will be positioned at the end of the object (EOF) before each put call.
IOS_\$CF_UNREGULATED	Connection is open for unregulated (shared) concurrency mode.
IOS_\$CF_READ_INTEND_WRITE	Connection is open for read access, and can be changed to write access. Other connections can have read access, but not write or read-intend-write access.

## INPUT PARAMETERS

### stream-id

Number of the stream on which the object is open, in IOS\_\$ID\_T format. This is a 2-byte integer.

## OUTPUT PARAMETERS

### status

Completion status, in STATUS\_\$T format. This data type is 4 bytes long. See the IOS Data Types section for more information.



**USAGE**

Use this call to determine which *connection* attributes are in effect for the specified stream.

To change object or connection attributes, use the IOS\_\$SET\_OBJ\_FLAGS, and IOS\_\$SET\_CONN\_FLAGS calls respectively. Which attributes you can change depends on the object type.

IOS\_\$INQ\_CUR\_REC\_LEN

IOS\_\$INQ\_CUR\_REC\_LEN

Returns the length of the record at the current stream marker.

#### FORMAT

rec-length = IOS\_\$INQ\_CUR\_REC\_LEN (stream-id, status)

#### RETURN VALUE

##### rec-length

Length of the current record. This is a 4-byte integer.

#### INPUT PARAMETERS

##### stream-id

Number of the stream on which the object is open, in IOS\_\$ID\_T format. This is a 2-byte integer.

#### OUTPUT PARAMETERS

##### status

Completion status, in STATUS\_\$T format. This data type is 4 bytes long. See the IOS Data Types section for more information.

#### USAGE

Use IOS\_\$INQ\_CUR\_REC\_LEN to determine the length of the record at the current stream marker of the specified stream.

The object specified must be record-oriented (for example, RECORDS\_\$UID); otherwise, IOS\_\$INQ\_CUR\_REC\_LEN returns an error.

**IOS\_\$INQ\_FILE\_ATTR**

Returns object usage attributes including date and time created, date and time last used, date and time last modified, number of blocks in the object.

**FORMAT**

IOS\_\$INQ\_FILE\_ATTR (stream\_id, dt-created, dt-modified, dt-used, blocks, status)

**INPUT PARAMETERS****stream-id**

Number of the stream on which the object is open, in IOS\_\$ID\_T format. This is a 2-byte integer.

**OUTPUT PARAMETERS****dt-created**

Date and time the object was created, in TIME\_\$CLOCKH\_T format. This is a 4-byte integer.

**dt-modified**

Date and time the object was last modified, in TIME\_\$CLOCKH\_T format. This is a 4-byte integer.

**dt-used**

Date and time the object was last used, in TIME\_\$CLOCKH\_T format. This is a 4-byte integer.

**blocks**

The number of 1024-byte blocks that the object occupies. This is a 4-byte integer.

**status**

Completion status, in STATUS\_\$T format. This data type is 4 bytes long. See the IOS Data Types section for more information.

**USAGE**

Use IOS\_\$INQ\_FILE\_ATTR to obtain a time stamp for an object and to determine the amount of space that an object occupies.

IOS\_\$INQ\_FULL\_KEY

IOS\_\$INQ\_FULL\_KEY

Returns a full seek key.

## FORMAT

IOS\_\$INQ\_FULL\_KEY (stream-id, position-option, full-key, status)

## INPUT PARAMETERS

### stream-id

Number of the stream on which the object is open, in IOS\_\$ID\_T format. This is a 2-byte integer.

### position-option

Value specifying the position to return a full seek key for, in IOS\_\$POS\_OPT\_T format. This is a 2-byte integer. Specify only one of the following predefined values:

IOS\_\$CURRENT            Return the full seek key of the current marker.

IOS\_\$EOF                Return the full seek key of the end of the object (EOF) marker.

IOS\_\$BOF                Return the full seek key of the beginning of the object (BOF) marker.

## OUTPUT PARAMETERS

### full-key

Full seek key to be used in subsequent seeks, in IOS\_\$SEEK\_KEY\_T format. This data type is 8 bytes long. See the IOS Data Types section for more information.

### status

Completion status, in STATUS\_\$T format. This data type is 4 bytes long. See the IOS Data Types section for more information.

## USAGE

IOS\_\$INQ\_FULL\_KEY returns a seek key based on the position option you specify, the current stream marker position, beginning or end of the object.

Use seek keys to perform random access of data. Typically, you use this call to inquire about a seek key before writing some data, and then store the seek key. To access the data at a later point in time, position the stream marker by calling the IOS\_\$SEEK\_FULL\_KEY call with the stored seek key, and get the data with an IOS get operation (IOS\_\$GET or IOS\_\$LOCATE).

Use seek keys merely as an index -- do not rely on the *contents* of the keys. The contents of seek keys remain private to the IOS manager, which guarantees only that the seek key returns to the position it describes.

Some object types support seek key positioning, but do not support record or byte positioning. Use seek keys for repositioning if your application does not need the "arithmetic" properties of record- or byte-positioning (that is, the ability to compute positions given positions).

The DOMAIN system offers both short (4-byte) and full (8-byte) seek keys. Because short seek keys require half the storage space of full seek keys, you might want to use short seek keys if your application program stores a large number of seek keys. However, short seek keys are limiting in that you can only indicate *record* boundary positions, while full seek keys allow you to indicate *any* position.

IOS\_\$INQ\_MGR\_FLAGS

IOS\_\$INQ\_MGR\_FLAGS

Returns the attribute set of an object's manager.

#### FORMAT

`mgr_flags = IOS_$INQ_MGR_FLAGS (stream-id, status)`

#### RETURN VALUE

##### **mgr\_flags**

A set (bit mask) indicating the attributes of the specified object's manager, in IOS\_\$MGR\_FLAG\_SET format. This is a 4-byte integer. Any combination of the following set of predefined values, in IOS\_\$MGR\_FLAG\_T format, can be returned. If the set contains the value, the manager has the attribute and can perform the following operations:

IOS_\$MF_CREATE	Manager permits type to create objects.
IOS_\$MF_CREATE_BAK	Manager permits type to create backup (.bak) objects.
IOS_\$MF_IMEX	Manager permits type to export streams to new processes.
IOS_\$MF_FORK	Manager permits type to pass streams to forked processes.
IOS_\$MF_FORCE_WRITE	Manager permits type to force-write object contents to stable storage (for most types, this is the disk).
IOS_\$MF_WRITE	Manager permits objects to be written to.
IOS_\$MF_SEEK_ABS	Manager permits objects to perform absolute seeks.
IOS_\$MF_SEEK_SHORT	Manager permits objects to perform seeks using short (4-byte) seek keys.
IOS_\$MF_SEEK_FULL	Manager permits objects to perform seeks using full (8-byte) seek keys.
IOS_\$MF_SEEK_BYTE	Manager permits objects to perform seeks to byte positions.
IOS_\$MF_SEEK_REC	Manager permits objects to perform seeks to record positions.
IOS_\$MF_SEEK_BOF	Manager permits objects to perform seeks to the beginning of the object.
IOS_\$MF_REC_TYPE	Manager supports different record type formats.

IOS_\$MF_TRUNCATE	Manager permits objects to be truncated.
IOS_\$MF_UNREGULATED	Manager permits objects to have unregulated (shared) concurrency mode.
IOS_\$MF_SPARSE	Manager permits objects to be written as sparse objects.
IOS_\$MF_READ_INTEND_WRITE	Manager permits objects to have read-intend-write access.

## INPUT PARAMETERS

### stream-id

Number of the stream on which the object is open, in IOS\_\$ID\_T format. This is a 2-byte integer.

## OUTPUT PARAMETERS

### status

Completion status, in STATUS\_\$T format. This data type is 4 bytes long. See the IOS Data Types section for more information.

## USAGE

Use IOS\_\$INQ\_MGR\_FLAGS to determine what operations an *object's type manager* can perform.

Depending on the nature of the object, a type manager permits some of the operations identified by "mgr-flags." A manager usually will not support operations that are irrelevant for the object type. For example, if you called IOS\_\$INQ\_MGR\_FLAGS specifying a stream open on an SIO line, the set returned would not include any IOS\_\$MF\_SEEK attributes, since serial lines do not support seeking.

Note that even if an object's manager permits an operation, the object itself can prevent the operation because the object's *object* and *connection* attributes must permit the operation as well. For example, a manager's attribute set might contain the attribute that permits writing to a file (IOS\_\$MF\_WRITE), but a specific object's connection attribute set might not include the IOS\_\$CF\_WRITE attribute, which permits writing on the connection. In this case, you *cannot* write to that particular object. However, you could possibly write to another object of the same type if the object's IOS\_\$CF\_WRITE attribute is set for its stream connection.

To change object or connection attributes, use the IOS\_\$SET\_OBJ\_FLAGS and IOS\_\$SET\_CONN\_FLAGS calls, respectively. Which attributes you can change depends on the object type. Note that you cannot change manager attributes because the type manager determines them. For details on writing a type manager, see the *Extending the DOMAIN Streams Facility* manual.

IOS\_\$INQ\_OBJ\_FLAGS

IOS\_\$INQ\_OBJ\_FLAGS

Returns the attribute set associated with an object.

## FORMAT

obj-flags = IOS\_\$INQ\_OBJ\_FLAGS (stream-id, status)

## RETURN VALUE

### obj-flags

A set (bit mask) indicating the attributes of the specified object, in IOS\_\$OBJ\_FLAG\_SET format. This is a 4-byte integer. Any combination of the following set of predefined values, in IOS\_\$OBJ\_FLAG\_T format, can be returned. If the set contains the value, the object has the attribute and can perform the following operations:

IOS_\$OF_DELETE_ON_CLOSE	Object will be deleted when all its associated streams close.
IOS_\$OF_SPARSE_OK	Object can be written as a sparse object.
IOS_\$OF_ASCII	Object contains ASCII data.
IOS_\$OF_FTNCC	Object uses FORTRAN carriage control characters.
IOS_\$OF_COND	Get or put calls to the object will be performed conditionally, as if the IOS_\$COND_OPT was specified on a get or put call.

## INPUT PARAMETERS

### stream-id

Number of the stream on which the object is open, in IOS\_\$ID\_T format. This is a 2-byte integer.

## OUTPUT PARAMETERS

### status

Completion status, in STATUS\_\$T format. This data type is 4 bytes long. See the IOS Data Types section for more information.

## USAGE

Use this call to determine which *object* attributes are in effect for the object on the specified stream.

To change object or connection attributes, use the IOS\_\$SET\_OBJ\_FLAGS, and IOS\_\$SET\_CONN\_FLAGS calls respectively. The attributes that you can change depends on the object type.



## IOS\_\$INQ\_PATH\_NAME

Returns the pathname of the object open on a specified stream.

**FORMAT**

IOS\_\$INQ\_PATH\_NAME (stream-id, name-type, pathname, namelength, status)

**INPUT PARAMETERS****stream-id**

Number of the stream on which the object is open, in IOS\_\$ID\_T format. This is a 2-byte integer.

**name-type**

Format of the returned pathname, in IOS\_\$NAME\_TYPE\_T format. Specify one of the following predefined values:

## IOS\_\$ROOT\_NAME

Return the absolute pathname, relative to the network root directory (/). For example, "//node/sid/file."

## IOS\_\$WDIR\_NAME

Return just the leaf name if the object's pathname is a name in the current working directory. Otherwise, return the absolute pathname.

## IOS\_\$NDIR\_NAME

Return just the leaf name if the object's pathname is a name in the current naming directory. Otherwise, return the absolute pathname.

## IOS\_\$NODE\_NAME

Return a name relative to the node's entry directory (/) if the object's pathname is a name in the boot volume. Otherwise, return the absolute pathname. For example, "/sid/file."

## IOS\_\$NODE\_DATA\_FLAG

Return just the leaf name if the object's pathname is a name in the 'node\_data' directory. Otherwise, return the absolute pathname.

## IOS\_\$LEAF\_NAME

Return just the leaf name regardless of the object's pathname. For example, if the object's pathname is "/a/b/c," it returns "c."

## IOS\_\$RESID\_NAME

Return the residual part of a pathname if the stream is open using extended naming. (Extended naming allows you to add additional text to the end of a pathname.)

**OUTPUT PARAMETERS****pathname**

Name of the object associated with the stream ID, in NAME\_\$PNAME\_T format. This is an array of up to 256 characters.

IOS\_\$INQ\_PATH\_NAME

**namelength**

Length of the pathname. This is a 2-byte integer.

**status**

Completion status, in STATUS\_\$T format. This data type is 4 bytes long. See the IOS Data Types section for more information.

**USAGE**

Use this call to determine the pathname of an object associated with the specified stream ID. Generally, use this call in cases where a program has been passed a stream ID and needs the associated pathname.

**IOS\_\$INQ\_REC\_POS**

Returns the record position of the stream marker.

**FORMAT**

record\_position = IOS\_\$INQ\_REC\_POS (stream-id, position-option, status)

**RETURN VALUE****record-position**

Record position of the stream marker. This is a 4-byte integer. Note that record positions are *zero*-based; consequently, the record position of the beginning of the object is 0.

**INPUT PARAMETERS****stream-id**

Number of the stream on which the object is open, in IOS\_\_\$ID\_\_T format. This is a 2-byte integer.

**position-option**

Value specifying the record position to return, in IOS\_\_\$POS\_\_OPT\_\_T format. This is a 2-byte integer. Specify one of the following predefined values:

IOS__\$CURRENT	Return the record position of the current stream marker.
IOS__\$EOF	Return the record position of the end of the object (EOF) stream marker. This is the number of records in the object.
IOS__\$BOF	Return the record position of the beginning of the object (BOF) stream marker. This value is always 0.

**OUTPUT PARAMETERS****status**

Completion status, in STATUS\_\_\$T format. This data type is 4 bytes long. See the IOS Data Types section for more information.

**USAGE**

To obtain the offset of the stream marker for record-oriented objects, use IOS\_\$INQ\_REC\_POS. (Use IOS\_\_\$INQ\_\_BYTE\_\_POS if your object is *not* record-oriented.)

To get the offset of the stream marker at the beginning or end of the object, specify IOS\_\_\$BOF or IOS\_\_\$EOF, in the "position-option" parameter. Specify IOS\_\_\$CURRENT to get the offset of the stream marker from the beginning of the object. Once you have the returned offset, you can move the stream marker to desired location by calling IOS\_\_\$SEEK.

This call allows you to perform a nonkeyed seek by specifying an absolute byte position, or by getting an offset from an absolute position, and moving the stream marker to it.

Whether you perform a nonkeyed or keyed seek depends on how the object's data is represented. For example, programs that need to perform "arithmetic" on the data (such as comparing two positions) will use nonkeyed seek operations. Programs that require only the ability to move from one position to another in an object will use keyed seek operations.

**IOS\_\$INQ\_REC\_REMAINDER**

Returns the number of bytes remaining in the current record.

**FORMAT**

bytes = IOS\_\$INQ\_REC\_REMAINDER (stream-id, status)

**RETURN VALUE****bytes**

Number of bytes remaining in the current record. This is a 4-byte integer.

**INPUT PARAMETERS****stream-id**

Number of the stream on which the file is open, in IOS\_\$ID\_T format. This is a 2-byte integer.

**OUTPUT PARAMETERS****status**

Completion status, in STATUS\_\$T format. This data type is 4 bytes long. See the IOS Data Types section for more information.

**USAGE**

Use IOS\_\$INQ\_REC\_REMAINDER with the IOS\_\$GET or IOS\_\$LOCATE calls. If IOS\_\$GET or IOS\_\$LOCATE fills the specified buffer, but has not yet finished reading a record, it returns the IOS\_\$BUFFER\_TOO\_SMALL error status code. At this point, use IOS\_\$INQ\_REC\_REMAINDER to determine the number of bytes in the record that remain to be read. If the entire record has been read, the value of "bytes" is undefined.

IOS\_\$INQ\_REC\_TYPE

IOS\_\$INQ\_REC\_TYPE

Returns the record type of an object.

## FORMAT

record-type = IOS\_\$INQ\_REC\_TYPE (stream-id, status)

## RETURN VALUE

### record-type

Type of record format used in the specified object, in IOS\_\$RTYPE\_T format. This is a 2-byte integer. Returns one of the following predefined values:

IOS_\$V1	Variable-length records with count fields.
IOS_\$F1	Fixed-length records without count fields.
IOS_\$F2	Fixed-length records with count fields.
IOS_\$EXPLICIT_F2	Fixed-length records that IOS_\$PUT cannot implicitly change to variable-length records.
IOS_\$UNDEF	No record structure.

## INPUT PARAMETERS

### stream-id

Number of the stream on which the object is open, in IOS\_\$ID\_T format. This is a 2-byte integer.

## OUTPUT PARAMETERS

### status

Completion status, in STATUS\_\$T format. This data type is 4 bytes long. See the IOS Data Types section for more information.

## USAGE

Use IOS\_\$INQ\_REC\_TYPE to determine how records within an object are formatted. You can change the record type of a record-oriented object by calling IOS\_\$SET\_REC\_TYPE.

By default, a record-oriented object has fixed-length records (IOS\_\$F2). They remain fixed-length records until IOS\_\$PUT writes records of different lengths. At this point, IOS\_\$PUT implicitly changes the objects to variable-length type (IOS\_\$V1). In some cases, you might want to explicitly set the record type to IOS\_\$EXPLICIT\_F2 so that an attempt to write a variable-length record results in an error. To do so, use the corresponding call, IOS\_\$SET\_REC\_TYPE.

**IOS\_\$INQ\_SHORT\_KEY**

Returns a short seek key.

**FORMAT**

short-key = IOS\_\$INQ\_SHORT\_SEEK (stream-id, position-option, status)

**RETURN VALUE****short-key**

Short seek key to be used in subsequent seeks. This is a 4-byte integer.

**INPUT PARAMETERS****stream-id**

Number of the stream on which the object is open, in IOS\_\$ID\_T format. This is a 2-byte integer.

**position-option**

Value specifying the position to return, in IOS\_\$POS\_OPT\_T format. This is a 2-byte integer. Specify only one of the following predefined values:

IOS_\$CURRENT	Return the short seek key of the current marker.
IOS_\$EOF	Return the short seek key of the end of the object (EOF) marker.
IOS_\$BOF	Return the short seek key of the beginning of the object (BOF) marker.

**OUTPUT PARAMETERS****status**

Completion status, in STATUS\_\$T format. This data type is 4 bytes long. See the IOS Data Types section for more information.

**USAGE**

IOS\_\$INQ\_SHORT\_KEY returns a seek key based on the position option you specify -- the current stream marker position, beginning or end of the object.

You use seek keys to perform random access of data. Typically, you use this call to inquire about a seek key before writing some data, and then store the seek key. To access the data at a later time, position the stream marker by calling the IOS\_\$SEEK\_SHORT\_KEY call with the stored seek key, and get the data with an IOS get operation (IOS\_\$GET or IOS\_\$LOCATE).

Use seek keys merely as an index -- do not count on the *contents* of the keys. The contents of seek keys remain private to the IOS manager, which guarantees only that the seek key returns to the position it describes.

Some object types support seek key positioning, but not record or byte positioning. Use seek keys for repositioning if your application does not need the "arithmetic" properties of record- or byte-positioning (that is, the ability to compute positions given positions).

The DOMAIN system offers both short (4-byte) and full (8-byte) seek keys. Because short seek keys require half the storage space of full seek keys, you might want to use short seek keys if your application program stores a large number of seek keys. However, short seek keys are limiting in that you can only indicate *record* boundary positions, while full seek keys allow you to indicate *any* position.



**IOS\_\$INQ\_TYPE\_UID**

Returns the type UID of an object.

**FORMAT**

IOS\_\$INQ\_TYPE\_UID (stream-id, type-uid, status)

**INPUT PARAMETERS****stream-id**

Number of the stream on which the object is open, in IOS\_\$ID\_T format. This is a 2-byte integer.

**OUTPUT PARAMETERS****type-uid**

Type UID of the object, in UID\_\$T format. This data type is 8 bytes long. See the IOS Data Types section for more information.

DOMAIN currently supports a set of predefined standard object types which include the following types. (Note that users can also define their own type UIDs by writing a type manager. See the *Using the Open System Toolkit to Extend the Streams Facility* manual for details.)

Type UID	Object
UASC_\$UID	UASC object
RECORDS_\$UID	Record-oriented object
HDR_UNDEF_\$UID	Nonrecord-oriented object
OBJECT_FILE_\$UID	Object module object (compiler or binder output)
SIO_\$UID	Serial line descriptor object
MT_\$UID	Magnetic tape descriptor object
PAD_\$UID	Saved display manager transcript pad
INPUT_PAD_\$UID	Display manager input pad
MBX_\$UID	Mailbox object
DIRECTORY_\$UID	Directory
NULLDEV_\$UID	Null device

**status**

Completion status, in STATUS\_\$T format. This data type is 4 bytes long. See the IOS Data Types section for more information.

IOS\_\$INQ\_TYPE\_UID

## USAGE

Use this call to determine the object's current type UID given its stream ID. You can use the type UID returned by this call as a parameter in the IOS\_\$CREATE call to create another object of the same type.

**IOS\_\$LOCATE**

Reads data from a stream, and returns a pointer to the data.

**FORMAT**

`ret-length = IOS_$LOCATE (stream-id, get-options, data-ptr, data-size,  
status)`

**RETURN VALUE****ret-length**

Amount of data read, in bytes. This is a 4-byte integer.

"Ret-length" equals the amount of data read; "ret-length" equals 0 if IOS\_\$LOCATE does not return any data.

If the length of the data read exceeds the amount specified in "data-size," IOS\_\$LOCATE performs the following:

- Reads enough data to fill the requested size
- Sets "ret-length" equal to "data-size"
- Positions the stream marker to the first unread byte
- Returns the IOS\_\$BUFFER\_TOO\_SMALL status code to indicate that this condition has occurred

You can inquire about how many bytes remain to be read in the current record by calling IOS\_\$INQ\_REC\_REMAINDER.

**INPUT PARAMETERS****stream-id**

Number of the stream on which the object is open, in IOS\_\$ID\_T format. This is a 2-byte integer.

**get-options**

Options that control how IOS\_\$LOCATE performs the get operation, in IOS\_\$PUT\_GET\_OPTS\_T format. This is a 2-byte integer. Specify a combination of the following set of predefined values:

IOS_\$COND_OPT	Reads data, if available. (For example, data on an SIO line is not always available immediately.) If the data is not available, IOS_\$GET returns the IOS_\$GET_CONDITIONAL_FAILED status code and sets the return value of "ret-length" to 0.
IOS_\$PREVIEW_OPT	Reads data but does not update the stream marker.

IOS\_ \$LOCATE

IOS\_ \$NO\_ REC\_ BNDRY\_ OPT      Ignores record boundaries while reading data. For example, it ignores NEWLINE characters in a UASC object, which guarantees that the call fills the specified buffer. Some type managers might not support this call.

IOS\_ \$PARTIAL\_ RECORD\_ OPT      Not meaningful for this call.

**data-size**

Maximum amount of data to be read, in bytes. This is a 4-byte integer.

**OUTPUT PARAMETERS**

**data-ptr**

A pointer to the located data, in UNIV\_ PTR format. This is a 4-byte integer. Note that this pointer remains valid only until the program invokes the next IOS call.

If IOS\_ \$LOCATE is unable to return a pointer to the location of the data, it copies the data into a system buffer and then returns the address of the buffer in "data-ptr." (See the USAGE Section below for more details.)

**status**

Completion status, in STATUS\_ \$T format. This data type is 4 bytes long. See the IOS Data Types section for more information.

**USAGE**

You can use either IOS\_ \$LOCATE or IOS\_ \$GET to read data from system objects. IOS\_ \$LOCATE returns a pointer to the data, while IOS\_ \$GET copies the data into a buffer.

In most cases, use the IOS\_ \$LOCATE call to read data because it is faster (IOS\_ \$LOCATE does not perform a copy).

You will want to use IOS\_ \$GET when you need to read more data than can be obtained in one call, because the pointer remains valid for only one call. For example, when you need to read and rearrange a number of lines from an object.

Normally, IOS\_ \$LOCATE locates data and returns a pointer to the data. However, not all managers support the internal buffering necessary for IOS\_ \$LOCATE to work this way. In these cases, IOS\_ \$LOCATE will not be able to return a pointer to the data.

Instead, IOS\_ \$LOCATE actually creates a buffer and then calls IOS\_ \$GET to perform the get call. In this case, IOS\_ \$LOCATE is no more efficient than IOS\_ \$GET. The size of the buffer that IOS\_ \$LOCATE creates is either the length you specify in "data-size," or 1024 bytes, whichever is the smaller.

Use IOS\_ \$SET\_ LOCATE\_ BUFFER\_ SIZE to specify a buffer larger than 1024 bytes, if necessary. In this case, IOS\_ \$LOCATE is no more efficient than IOS\_ \$GET.

See the IOS\_ \$SET\_ LOCATE\_ BUFFER\_ SIZE call description for more information.

**IOS\_\$OPEN**

Opens a stream to an existing object.

**FORMAT**

stream-id = IOS\_\$OPEN (pathname, namelength, open-options, status)

**RETURN VALUE****stream-id**

Number of the stream on which the object is open, in IOS\_\$ID\_T format. This is a 2-byte integer.

**INPUT PARAMETERS****pathname**

Name of the object to be opened, in NAME\_\$PNAME\_T format. This is an array of up to 256 characters.

**namelength**

Length of the pathname. This is a 2-byte integer.

**open-options**

Options available at open time, in IOS\_\$OPEN\_OPTIONS\_T format. This is a 2-byte integer. Specify a combination of the following set of predefined values:

IOS_\$NO_OPEN_DELAY_OPT	Return immediately, instead of waiting for the open call to complete.
IOS_\$WRITE_OPT	Permit writing data to a new object. If a program tries to write on a stream for which you have not specified this option, it returns an error status.
IOS_\$UNREGULATED_OPT	Permit shared (unregulated) concurrency mode.
IOS_\$POSITION_TO_EOF_OPT	Position the stream marker at the end of the object (EOF). Use this to append data to an existing object.
IOS_\$INQUIRE_ONLY_OPT	Open the object for attribute inquiries only; do not permit reading or writing of data.
IOS_\$READ_INTEND_WRITE_OPT	Open the object for read access with the intent to eventually change the object's access to write access. This allows other processes to read the object; but they cannot have write or read-intend-write access.

IOS\_\$OPEN

## OUTPUT PARAMETERS

### status

Completion status, in STATUS\_\$T format. This data type is 4 bytes long. See the IOS Data Types section for more information.

## USAGE

This routine opens a stream to the named object. It returns the stream ID to be used in subsequent stream activity with the object. An error occurs if the object does not exist. If the object already exists, IOS\_\$OPEN does not change its attributes.

IOS\_\$OPEN does not return information about the object's attributes. To get information about an object, use the calls with the prefix IOS\_\$INQ. To change an object's attributes, use the calls with the prefix IOS\_\$SET.

**IOS\_\$PUT**

Writes data into an object.

**FORMAT**

IOS\_\$PUT (stream-id, put-options, buffer, buffer-size, status)

**INPUT PARAMETERS****stream-id**

Number of the stream on which the object is open, in IOS\_\$ID\_T format. This is a 2-byte integer.

**put-options**

Options that control how IOS\_\$PUT performs the put operation, in IOS\_\$PUT\_GET\_OPTS\_T format. This is a 2-byte integer. Specify any combination of the following set of predefined values:

IOS\_\$COND\_OPT

Write a record only if it can be done without blocking. If the call would block, it returns the IOS\_\$PUT\_CONDITIONAL\_FAILED error status.

IOS\_\$PREVIEW\_OPT

Write data but do not update the stream marker.

IOS\_\$PARTIAL\_RECORD\_OPT

Write a portion of a record but do not terminate it. IOS\_\$PUT terminates the record when you call IOS\_\$PUT without specifying this option. If you do not specify this option, IOS\_\$PUT writes a full record. You can use this option with record-oriented objects only. IOS\_\$PUT ignores this option if you specify it with any other type of objects.

IOS\_\$NO\_REC\_BNDRY\_OPT

Not meaningful for this call.

**buffer**

Buffer to contain the data. This is a character array.

**buffer-size**

Size of the buffer containing the data, in bytes. This is a 4-byte integer.

**OUTPUT PARAMETERS****status**

Completion status, in STATUS\_\$T format. This data type is 4 bytes long. See the IOS Data Types section for more information.

IOS\_\$PUT

## USAGE

IOS\_\$PUT writes data into an object. Use "put-options," which is in the IOS\_\$PUT\_GET\_OPTS\_T format, to write the data to the object in different ways.

If the object is record-oriented, you can write data to it record by record. This is the default action (for record-oriented objects) when you specify the default ([]) value in "put-option."

To write a single record with more than one put operation (for example, to write one field at a time), use the IOS\_\$PARTIAL\_RECORD\_OPT option. If you specify this option, IOS\_\$PUT writes the data, but does not terminate the record. IOS\_\$PUT terminates the record when you call it without specifying this option.

To write to objects which might not always be immediately available (for example, an MBX channel), you perform *conditional* put operations with the IOS\_\$COND\_OPT option.



**IOS\_ \$REPLICATE**

Creates a copy of a specified stream ID.

**FORMAT**

```
return_stream_id = IOS_ $REPLICATE (stream_id_to_replicate, copy_stream_id,
                                     status)
```

**RETURN VALUE****return\_stream\_id**

Number of the new stream created, in IOS\_ \$ID\_ T format. This is a 2-byte integer.

**INPUT PARAMETERS****stream\_id\_to\_replicate**

Number of the stream to replicate, in IOS\_ \$ID\_ T format. This is a 2-byte integer. This stream number remains a valid connection to the object after IOS\_ \$REPLICATE completes successfully.

**copy\_stream\_id**

Number of the stream to use as the copy for "stream\_id\_to\_replicate," in IOS\_ \$ID\_ T format. This is a 2-byte integer.

If "copy\_stream\_id" is free, IOS\_ \$REPLICATE returns that number in "return\_stream\_id." If "copy-stream-id" is in use, IOS\_ \$REPLICATE begins searching from that number downward (lower numbers) until it finds a free stream number, and returns that number in "return\_stream\_id."

If the actual number of the copy stream is insignificant, specify the predefined constant IOS\_ \$MAX. This value causes IOS\_ \$REPLICATE to begin searching at the highest possible stream number and return the first free stream number it finds.

**OUTPUT PARAMETERS****status**

Completion status, in STATUS\_ \$T format. This data type is 4 bytes long. See the IOS Data Types section for more information.

**USAGE**

Use IOS\_ \$REPLICATE to create a copy of an existing stream ID. The new stream ID refers to the same connection as the existing stream ID. Note that you must close *both* streams with IOS\_ \$CLOSE before the stream connection actually closes.

IOS\_ \$REPLICATE is identical to IOS\_ \$DUP except that IOS\_ \$REPLICATE looks for a free stream in *descending* order from the specified stream ID, while IOS\_ \$DUP looks in *ascending* order. Note that you use IOS\_ \$DUP or IOS\_ \$REPLICATE to *copy* existing stream ID's, both the existing and new stream ID's remain valid connections. However, you use IOS\_ \$SWITCH to *replace* stream IDs; you "switch" the connection from the existing stream ID to the new stream ID.

## IOS\_ \$REPLICATE

You can use IOS\_ \$REPLICATE to keep a stream connection open when passing it to a subroutine. Use IOS\_ \$REPLICATE to create a copy of the stream ID before passing it. This way, the subroutine cannot close the connection to the object because *all* copies of the stream connection must be closed before the connection itself gets closed.

IOS\_ \$RELPLICATE is analagous to UNIX DUP.

**IOS\_ \$SEEK**

Performs an absolute or relative seek using byte or record positioning.

**FORMAT**

IOS\_ \$SEEK (stream-id, abs-rel, seek-type, offset, status)

**INPUT PARAMETERS****stream-id**

Number of the stream on which the object is open, in IOS\_ \$ID\_ T format. This is a 2-byte integer.

**abs-rel**

Value specifying the base for the seek operation, in IOS\_ \$ABS\_ REL\_ T format. This is a 2-byte integer. Specify one of the following predefined values:

IOS\_ \$RELATIVE            The seek is relative to the current position.

IOS\_ \$ABSOLUTE            The seek is relative to the beginning of the object (BOF).

**seek-type**

The type of seek to be performed, in IOS\_ \$SEEK\_ TYPE\_ T format. This is a 2-byte integer. Specify one of the following predefined values:

IOS\_ \$REC\_ SEEK            Record-oriented seek.

IOS\_ \$BYTE\_ SEEK          Byte-oriented seek.

**offset**

A signed integer offset value indicating the number of records or bytes from the seek base to position the stream marker. This is a 4-byte integer.

If the integer is a positive number, IOS\_ \$SEEK uses BOF as the seek base and searches forward. If the integer is a negative number, IOS\_ \$SEEK uses EOF as the seek base and searches backward. Whether the offset indicates bytes or records depends on the type of seek you specified in "seek-type."

You can get an offset number to use in an absolute seek with the calls IOS\_ \$INQ\_ BYTE\_ POS and IOS\_ \$INQ\_ REC\_ POS.

Note that both byte and record positions are *zero-based*; consequently, the first byte or record number is 0.

**OUTPUT PARAMETERS****status**

Completion status, in STATUS\_ \$T format. This data type is 4 bytes long. See the IOS Data Types section for more information.

IOS\_\$SEEK

## USAGE

Use IOS\_\$SEEK to seek to an absolute or relative byte or record position within an object.

You can use this call with the IOS\_\$INQ\_BYTE\_POS and IOS\_\$INQ\_REC\_POS calls to perform absolute position seeks.

**IOS\_ \$SEEK\_ FULL\_ KEY**

Performs a seek using a full (8-byte) seek key.

**FORMAT**

IOS\_ \$SEEK\_ FULL\_ KEY (stream-id, full-key, status)

**INPUT PARAMETERS****stream-id**

Number of the stream on which the object is open, in IOS\_ \$ID\_ T format. This is a 2-byte integer.

**full-key**

A full seek key, in IOS\_ \$SEEK\_ KEY\_ T format. This data type is 8 bytes long. See the IOS Data Types section for more information.

**OUTPUT PARAMETERS****status**

Completion status, in STATUS\_ \$T format. This data type is 4 bytes long. See the IOS Data Types section for more information.

**USAGE**

Before performing a full key seek, you must first obtain a full seek key by using the IOS\_ \$INQ\_ FULL\_ KEY call. This call allows you to inquire about a seek key before writing some data, and then store the seek key. To access the data at a later time, position the stream marker by calling the IOS\_ \$SEEK\_ FULL\_ KEY call with the stored seek key, and then get the data with an IOS get call (IOS\_ \$GET or IOS\_ \$LOCATE).

IOS\_\$SEEK\_SHORT\_KEY

IOS\_\$SEEK\_SHORT\_KEY

Performs a seek using a short (4-byte) seek key.

### FORMAT

IOS\_\$SEEK\_SHORT\_KEY (stream-id, short-key, status)

### INPUT PARAMETERS

#### stream-id

Number of the stream on which the object is open, in IOS\_\$ID\_T format. This is a 2-byte integer.

#### short-key

A short seek key. This is a 4-byte integer.

### OUTPUT PARAMETERS

#### status

Completion status, in STATUS\_\$T format. This data type is 4 bytes long. See the IOS Data Types section for more information.

### USAGE

Before performing a short key seek, you must first obtain a short seek key by using the IOS\_\$INQ\_SHORT\_KEY call. This call allows you to inquire about a seek key before writing some data, and then store the seek key. To access the data at a later time, position the stream marker by calling IOS\_\$SEEK\_SHORT\_KEY with the stored seek key, and then get the data with an IOS get call (IOS\_\$GET or IOS\_\$LOCATE).

**IOS\_\$\$SEEK\_TO\_BOF**

Positions the stream marker to the beginning of an object.

**FORMAT**

IOS\_\$\$SEEK\_TO\_BOF (stream-id, status)

**INPUT PARAMETERS****stream-id**

Number of the stream on which the object is open, in IOS\_\$\$ID\_T format. This is a 2-byte integer.

**OUTPUT PARAMETERS****status**

Completion status, in STATUS\_\$\$T format. This data type is 4 bytes long. See the IOS Data Types section for more information.

**USAGE**

Use IOS\_\$\$SEEK\_TO\_BOF to position the stream marker to the beginning of an object (BOF). Use this call when performing a nonkeyed seek on an object.

IOS\_\$SEEK\_TO\_EOF

IOS\_\$SEEK\_TO\_EOF

Positions the stream marker to the end of an object.

### FORMAT

IOS\_\$SEEK\_TO\_EOF (stream-id, status)

### INPUT PARAMETERS

#### stream-id

Number of the stream on which the object is open, in IOS\_\$ID\_T format. This is a 2-byte integer.

### OUTPUT PARAMETERS

#### status

Completion status, in STATUS\_\$T format. This data type is 4 bytes long. See the IOS Data Types section for more information.

### USAGE

Use IOS\_\$SEEK\_TO\_EOF to position the stream marker to the end of an object (EOF). Use this call when performing a nonkeyed seek on an object.



**IOS\_\$SET\_CONN\_FLAG**

Changes the set of connection attributes associated with a stream connection.

**FORMAT**

IOS\_\$SET\_CONN\_FLAG (stream-id, conn-flag, on-off, status)

**INPUT PARAMETERS****stream-id**

Number of the stream on which the object is open, in IOS\_\$ID\_T format. This is a 2-byte integer.

**conn-flag**

Flag indicating which attribute of the specified connection you want to change, in IOS\_\$CONN\_FLAG\_T format. This is a 2-byte integer. Specify one of the following predefined values:

IOS_\$CF_TTY	Connection behaves like a terminal.
IOS_\$CF_IPC	Connection behaves like an interprocess communication (IPC) channel.
IOS_\$CF_VT	Connection behaves like a DOMAIN Display Manager pad.
IOS_\$CF_WRITE	Connection can be written to.
IOS_\$CF_APPEND	Connection's stream marker will be positioned at the end of the object (EOF) before each put call.
IOS_\$CF_UNREGULATED	Connection is open for unregulated (shared) concurrency mode.
IOS_\$CF_READ_INTEND_WRITE	Connection is open for read access, and can be changed to write access. Other connections can have read access, but not write or read-intend-write access.

**on-off**

Boolean value indicating whether the specified attribute should be included in the set (on), or removed from the set (off).

**OUTPUT PARAMETERS****status**

Completion status, in STATUS\_\$T format. This data type is 4 bytes long. See the IOS Data Types section for more information.

## USAGE

Use IOS\_\$SET\_CONN\_FLAG to change the attributes of a connection. Note that objects do not support *all* connection attributes. To determine the connection's current set of attributes, use IOS\_\$INQ\_CONN\_FLAGS before using this call.

To change the set of attributes, you must call IOS\_\$SET\_CONN\_FLAG for *each* connection attribute you want to change. To add an attribute to the set, call IOS\_\$SET\_CONN\_FLAG, specifying the desired attribute, and set the "on-off" parameter to TRUE. To remove an attribute from the set, use this call, specifying the attribute to remove, and set the "on-off" parameter to FALSE.

Before an object can permit the operation indicated by an attribute, the object's *manager* and *connection* attributes must permit the operation as well. For example, a manager's attribute set might contain the attribute that permits writing to an object (IOS\_\$MF\_WRITE), but a specific object's connection attribute set might not include the IOS\_\$CF\_WRITE attribute, which permits writing to the object. In this case, you *cannot* write to that particular object.

**IOS\_\$SET\_DIR**

Changes the current working or naming directory.

**FORMAT**

IOS\_\$SET\_DIR (pathname, namelength, dir\_type, status)

**INPUT PARAMETERS****pathname**

Name of the directory to set, in NAME\_\$PNAME\_T format. This is an array of up to 256 characters.

**namelength**

Length of "pathname." This is a 2-byte integer.

**dir\_type**

Option specifying which type of directory to set, in IOS\_\$DIR\_TYPE\_T format. Specify one of the predefined values:

IOS\_\$WDIR     Name of the current working directory.

IOS\_\$NDIR     Name of the current naming directory.

**OUTPUT PARAMETERS****status**

Completion status, in STATUS\_\$T format. This data type is 4 bytes long. See the IOS Data Types section for more information.

**USAGE**

Use this call to change the current working or naming directory. You can use IOS\_\$GET\_DIR to get the name of the current working or naming directory.

IOS\_\$SET\_LOCATE\_BUFFER\_SIZE

IOS\_\$SET\_LOCATE\_BUFFER\_SIZE

Sets the size of the buffer that IOS\_\$LOCATE allocates.

## FORMAT

IOS\_\$SET\_LOCATE\_BUFFER\_SIZE (stream-id, buffer-size, status)

## INPUT PARAMETERS

### stream-id

Number of the stream on which the object is open, in IOS\_\$ID\_T format. This is a 2-byte integer.

### buffer-size

Size of the buffer you want to allocate. This is a 2-byte integer.

## OUTPUT PARAMETERS

### status

Completion status, in STATUS\_\$T format. This data type is 4 bytes long. See the IOS Data Types section for more information.

## USAGE

Normally, IOS\_\$LOCATE locates data and returns a pointer to the data. However, not all managers support the internal buffering necessary for IOS\_\$LOCATE to work this way. In these cases, IOS\_\$LOCATE will not be able to return a pointer to the data.

Instead, IOS\_\$LOCATE actually creates a buffer and then calls IOS\_\$GET to perform the get call. In this case, IOS\_\$LOCATE is no more efficient than IOS\_\$GET. The size of the buffer that IOS\_\$LOCATE creates is either the length you specify in "data-size," or 1024 bytes, whichever is the smaller.

Use IOS\_\$SET\_LOCATE\_BUFFER\_SIZE to specify a buffer larger than 1024 bytes, if necessary.

For example, if you are using IOS\_\$LOCATE with a data-size parameter of 2000 bytes, and the manager of the object from which you are reading does not support internal buffering, the IOS\_\$LOCATE call, by default, will copy as much of the requested data as it can into a 1024-byte-long buffer and return a pointer to that buffer.

However, if you precede the IOS\_\$LOCATE call with a call to IOS\_\$SET\_LOCATE\_BUFFER\_SIZE, specifying a buffer-size of 2000, the IOS\_\$LOCATE call will use a 2000-byte-long buffer and will be able to copy all the requested data into the buffer. This new buffer size will be valid as long as the stream exists.

**IOS\_\$\$SET\_OBJ\_FLAG**

Changes the set of object attributes associated with an object.

**FORMAT**

IOS\_\$\$SET\_OBJ\_FLAG (stream-id, obj-flag, on-off, status)

**INPUT PARAMETERS****stream-id**

Number of the stream on which the object is open, in IOS\_\$\$ID\_T format. This is a 2-byte integer.

**obj-flag**

Flag indicating which attribute of the specified object you want to change, in IOS\_\$\$OBJ\_FLAG\_T format. This is a 2-byte integer. Specify one of the following predefined values:

IOS_\$\$OF_DELETE_ON_CLOSE	Object will be deleted when all its associated streams close.
IOS_\$\$OF_SPARSE_OK	Object can be written as a sparse object.
IOS_\$\$OF_ASCII	Object contains ASCII data.
IOS_\$\$OF_FTNCC	Object uses FORTRAN carriage control characters.
IOS_\$\$OF_COND	Get or put calls to the object will be performed conditionally, as if the IOS_\$\$COND_OPT was specified on a get or put call.

**on-off**

Boolean value indicating whether the specified attribute should be included in the set (on), or removed from the set (off).

**OUTPUT PARAMETERS****status**

Completion status, in STATUS\_\$\$T format. This data type is 4 bytes long. See the IOS Data Types section for more information.

IOS\_\$SET\_OBJ\_FLAG

## USAGE

Use IOS\_\$SET\_OBJ\_FLAGS to change the attributes of an object. Note that objects do not support *all* object attributes. To determine the object's current attribute set, use the IOS\_\$INQ\_OBJ\_FLAGS call.

To change an object's attribute set, you must call IOS\_\$SET\_OBJ\_FLAG once for each object attribute you want to change. To add an attribute to the set, call IOS\_\$SET\_OBJ\_FLAG, specifying the desired attribute, and set the "on-off" parameter to TRUE. To remove an attribute from the set, use this call, specifying the attribute to remove, and set the "on-off" parameter to FALSE.

Before an object can permit the operation indicated by an attribute, the object's *manager* and *object* attributes must permit the operation as well. For example, a manager's attribute set might contain the attribute that allows the object to perform put and get calls conditionally (IOS\_\$MF\_COND), but a specific object's object attribute set might not include the IOS\_\$OF\_COND attribute. In this case, you *cannot* make conditional put or get calls to that particular object.

**IOS\_\$\$SET\_REC\_TYPE**

Sets the record type format and (optionally) record length of a file.

**FORMAT**

IOS\_\$\$SET\_REC\_TYPE (stream-id, record-type, record-length, status)

**INPUT PARAMETERS****stream-id**

Number of the stream on which the object is open, in IOS\_\$\$ID\_T format. This is a 2-byte integer.

**record-type**

Type of record format to change for the specified object, in IOS\_\$\$RTYPE\_T format. This is a 2-byte integer. Specify one of the following predefined values:

IOS_\$\$V1	Variable-length records with count fields.
IOS_\$\$F1	Fixed-length records without count fields.
IOS_\$\$F2	Fixed-length records with count fields. However, IOS_\$\$PUT can change the IOS_\$\$F2 type to IOS_\$\$V1 implicitly. (See Usage section below.)
IOS_\$\$EXPLICIT_F2	Fixed-length records that IOS_\$\$PUT <i>cannot</i> implicitly change to variable-length records. (IOS_\$\$PUT can change the IOS_\$\$F2 to IOS_\$\$V1 implicitly. See Usage section below.)
IOS_\$\$UNDEF	No record structure.

**record-length**

Length to set for the fixed-length records of the object. This is a 4-byte integer. Specify this value only if the object is empty.

**OUTPUT PARAMETERS****status**

Completion status, in STATUS\_\$\$T format. This data type is 4 bytes long. See the IOS Data Types section for more information.

**USAGE**

By default, a record-oriented object has fixed-length records (IOS\_\$\$F2). They remain fixed-length records until IOS\_\$\$PUT writes records of different lengths. At this point, IOS\_\$\$PUT implicitly changes the objects to variable-length type (IOS\_\$\$V1). In some cases, you might want to explicitly set the record type to IOS\_\$\$EXPLICIT\_F2 so that an attempt to write a variable-length record results in an error. To do so, use this call.

IOS\_\$\$SWITCH

IOS\_\$\$SWITCH

Switches a stream from one stream ID to another stream ID.

## FORMAT

ret-stream-id = IOS\_\$\$SWITCH (stream-id-to-switch, new-stream-id, status)

## RETURN VALUE

### ret-stream-id

Number of the new stream ID that replaces the existing stream ID, in IOS\_\$\$ID\_T format. This is a 2-byte integer.

## INPUT PARAMETERS

### stream-id-to-switch

Number of the stream to switch, in IOS\_\$\$ID\_T format. This is a 2-byte integer.

This stream number becomes *invalid* after the IOS\_\$\$SWITCH call completes successfully.

### new-stream-id

Number of the stream to use as the new stream ID, in IOS\_\$\$ID\_T format. This is a 2-byte integer.

If "new-stream-id" is free, IOS\_\$\$SWITCH returns this value in "ret-stream-id." If "new-stream-id" is in use, IOS\_\$\$SWITCH begins searching from that value downward (lower numbers) until it finds a free stream number and returns that number in "ret-stream-id."

If the actual number of the replacement stream is insignificant, specify the predefined constant IOS\_\$\$MAX. This value causes IOS\_\$\$SWITCH to begin searching at highest possible stream number and return the first free number it finds.

## OUTPUT PARAMETERS

### status

Completion status, in STATUS\_\$\$T format. This data type is 4 bytes long. See the IOS Data Types section for more information.

## USAGE

Use IOS\_\$\$SWITCH to switch one stream ID for another. The new stream ID refers to the *same* connection as the old stream ID, making the old stream ID invalid.

Note that you use IOS\_\$\$SWITCH to *replace* stream IDs; you "switch" the connection from the existing stream ID to the new stream ID. However, you use IOS\_\$\$DUP or IOS\_\$\$REPLICATE to *copy* existing stream IDs, both the existing and new stream IDs remain valid connections.



**IOS\_ \$TRUNCATE**

Deletes the contents of an object following the current stream marker.

**FORMAT**

IOS\_ \$TRUNCATE (stream-id, status)

**INPUT PARAMETERS****stream-id**

Number of a stream on which the object is open, in IOS\_ \$ID\_ T format. This is a 2-byte integer.

**OUTPUT PARAMETERS****status**

Completion status, in STATUS\_ \$T format. This data type is 4 bytes long. See the IOS Data Types section for more information.

**USAGE**

IOS\_ \$TRUNCATE decreases the value of the object's length attribute to match the stream marker. (Writing data to a stream that lengthens the object implicitly increases this attribute's value.) This call sets the stream marker to the end of the object (EOF), effectively deleting any data in the object past the stream marker. If the stream position is already at EOF, IOS\_ \$TRUNCATE has no effect.

Truncating an object does not close the stream.

**ERRORS**

**IOS\_\$ALREADY\_EXISTS**

Object already exists; detected by IOS\_\$CREATE with IOS\_\$NO\_PREXIST option.

**IOS\_\$BAD\_CHAR\_SEEK**

Attempted character seek before start of current (variable-length) record.

**IOS\_\$BAD\_COUNT\_FIELD\_IN\_FILE**

Count field for current record is wrong.

**IOS\_\$BAD\_FILE\_HDR**

Wrong stream file header.

**IOS\_\$BAD\_LOCATION**

Bad location parameter on IOS\_\$CREATE call..

**IOS\_\$BAD\_OPEN\_XP**

OPEN\_XP must reference a stream that is already open in this process.

**IOS\_\$BAD\_SHARED\_CURSOR\_REFCNT**

Reference count on a shared object cursor went below zero.

**IOS\_\$BOF\_ERR**

Attempted seek beyond beginning of object (BOF).

**IOS\_\$BUFFER\_TOO\_BIG**

Buffer size too large on IOS\_\$GET or IOS\_\$LOCATE call.

**IOS\_\$BUFFER\_TOO\_SMALL**

Buffer too small on IOS\_\$GET or IOS\_\$LOCATE call, warning.

**IOS\_\$CANT\_CHANGE\_TYPE**

Cannot change the type as requested, detected by IOS\_\$CREATE.

**IOS\_\$CANT\_DELETE\_OLD\_NAME**

Added new name, but cannot delete old name.

**IOS\_\$CANT\_INITIALIZE**

Cannot initialize an object of this type.

**IOS\_\$CANT\_SET\_ADVISORY\_LOCK**

Advisory lock already set on this object.

**IOS\_\$CONCURRENCY\_VIOLATION**

Requested access violates concurrency constraints, object is in use.

**IOS\_\$DEVICE\_MUST\_BE\_LOCAL**

Cannot open stream to remote device.

**IOS\_\$DIR\_NOT\_FOUND**

Couldn't find directory in pathname on IOS\_\$CREATE.

**IOS\_\$END\_OF\_FILE**

End of file.

**IOS\_\$FILE\_NOT\_EMPTY**

Object not empty.

IOS\_ \$FLAG\_ NOT\_ SUPPORTED  
Flag not supported for this object type.

IOS\_ \$FROM\_ ID\_ NOT\_ OPEN  
Stream ID to switch not open on IOS\_ \$SWITCH.

IOS\_ \$FULL\_ REC\_ UNAVAIL  
IOS\_ \$GET or IOS\_ \$LOCATE requested a full record, but only part of the record was available. The call returns the part that is available along with this warning that there is still more room in the buffer.

IOS\_ \$GET\_ CONDITIONAL\_ FAILED  
Cannot read any data because the stream is empty; detected by IOS\_ \$COND\_ OPT option.

IOS\_ \$ID\_ OOR  
Stream ID is out-of-range or invalid.

IOS\_ \$ILLEGAL\_ NAME\_ REDEFINE  
Attempted name change would require object to be moved, detected by IOS\_ \$CHANGE\_ PATH\_ NAME.

IOS\_ \$ILLEGAL\_ OBJ\_ TYPE  
Cannot open a stream for this type of object.

IOS\_ \$ILLEGAL\_ OPERATION  
Operation illegal on named stream.

IOS\_ \$ILLEGAL\_ PAD\_ CREATE\_ TYPE  
Cannot perform this operation on a pad type.

IOS\_ \$ILLEGAL\_ PARAM\_ COMB  
Illegal parameter combination for this operation.

IOS\_ \$ILLEGAL\_ W\_ VAR\_ LGTH\_ RECS  
Operation illegal with variable-length records.

IOS\_ \$INQ\_ ONLY\_ ERROR  
Can only open this operation for inquiries only.

IOS\_ \$INSUFFICIENT\_ RIGHTS  
Insufficient rights for requested access to object.

IOS\_ \$INSUFF\_ MEMORY  
Not enough address space.

IOS\_ \$INTERNAL\_ FATAL\_ ERR  
Internal fatal error on table re-verify operation.

IOS\_ \$INTERNAL\_ MM\_ ERR  
Internal fatal error in stream memory management (windowing).

IOS\_ \$INVALID\_ DATA  
Cannot write this data to object.

IOS\_ \$NAME\_ NOT\_ FOUND  
Name not found.

IOS\_ \$NAME\_ REQD  
Must specify name on IOS\_ \$OPEN.

IOS ERRORS

- IOS\_ \$NEED\_ MOVE\_ MODE  
IOS\_ \$LOCATE operation refused, try IOS\_ \$GET.
- IOS\_ \$NEVER\_ CLOSED  
System (or process) crash prevented complete close of object.
- IOS\_ \$NO\_ ADVISORY\_ LOCK\_ SET  
No advisory lock to unlock.
- IOS\_ \$NO\_ AVAIL\_ TARGET  
No available target stream to switch to on IOS\_ \$SWITCH.
- IOS\_ \$NO\_ MORE\_ STREAMS  
No more available stream IDs.
- IOS\_ \$NO\_ RIGHTS  
No rights to access object.
- IOS\_ \$NO\_ TABLE\_ SPACE  
Internal error.
- IOS\_ \$NOT\_ A\_ DIRECTORY  
Name specified is not a directory detected by IOS\_ \$GET\_ DIR or \_ \$SET\_ DIR.
- IOS\_ \$NOT\_ AT\_ REC\_ BNDRY  
Cannot perform operation with short key -- must be at a record boundary.
- IOS\_ \$NOT\_ OPEN  
Operation attempted on unopened stream.
- IOS\_ \$OBJ\_ DELETED  
Object has been deleted while open on this stream.
- IOS\_ \$OBJECT\_ NOT\_ FOUND  
Object associated with this name not found even though name exists.
- IOS\_ \$OBJECT\_ READ\_ ONLY  
Cannot open this object for writing.
- IOS\_ \$OUT\_ OF\_ SHARED\_ CURSORS  
Internal error.
- IOS\_ \$PART\_ REC\_ WARN  
Partial record at EOF on IOS\_ \$CLOSE -- warning only.
- IOS\_ \$PERM\_ FILE\_ NEEDS\_ NAME  
Only temporary objects can be unnamed, you must name a permanent object.
- IOS\_ \$PUT\_ BAD\_ REC\_ LEN  
Attempted an IOS\_ \$PUT on a record of the wrong length.
- IOS\_ \$PUT\_ CONDITIONAL\_ FAILED  
Cannot write any data because the stream is full, detected by IOS\_ \$COND\_ OPT option.
- IOS\_ \$READ\_ ONLY\_ ERR  
Attempted to write to read-only stream.
- IOS\_ \$RESOURCE\_ LOCK\_ ERR  
Unable to lock resources required to process request.

IOS\_\$SIO\_NOT\_LOCAL

No stream found in conditional put, or cannot open a remote SIO line.

IOS\_\$SOMETHING\_FAILED

Cannot locate attribute set inquiring about manager, connection or object attributes;  
or cannot change the connection or object attribute requested.

IOS\_\$TARGET\_INUSE

Target ID already in use on IOS\_\$SWITCH, no available stream IDs.

IOS\_\$XP\_BUF\_TOO\_SMALL

Buffer supplied to IOS\_\$EXPORT too small.



# IOS\_DIR

## Table of Contents

Data Types

IOS\_DIR\_\$ISA

IOS\_DIR\_\$OPEN

Errors

IOS\_DIR-2

IOS\_DIR-5

IOS\_DIR-6

IOS\_DIR-8





## IOS\_DIR

The IOS\_DIR (I/O Switch Directory) programming calls provide less expensive streams access to file system objects. This section describes their data types, call syntax, and error codes. Refer to the Introduction at the beginning of this manual for a description of data-type diagrams and call syntax format.

Note that programs use the IOS\_DIR calls to request the services of a type manager's IOS\_DIR trait; a type manager does not usually use the IOS\_DIR calls. Also, when you use the IOS\_DIR calls, you must include the /sys/ins/ios.ins.lan insert file before you include the /sys/ins/ios\_dir.ins.lan insert file.

CONSTANTS

IOS_DIR_\$ENTTYPE_FILE	1	A directory entry for a file.
IOS_DIR_\$ENTTYPE_LINK	3	A directory entry for a link.
IOS_DIR_\$DIR_ENTRY_SIZE	44	Size of a directory entry (ios_dir_\$dir_entry_t).

DATA TYPES

IOS_\$ID_T	A 2-byte integer, ranging in value from 0 to IOS_\$MAX. The stream ID.
IOS_\$OPEN_OPTIONS_T	<p>A 2-byte integer. Specifies options for opening a stream connection. Any combination of the following predefined values:</p> <p>IOS_\$INQUIRE_ONLY_OPT Open the object for attribute inquires only; do not permit reading or writing of data.</p> <p>IOS_\$NO_OPEN_DELAY_OPT Return immediately, instead of waiting for the open call to complete.</p> <p>IOS_\$POSITION_TO_EOF_OPT Position the stream marker to the end of the object (EOF).</p> <p>IOS_\$READ_INTEND_WRITE_OPT Open the object for read access with the intent to eventually change to write access. This allows other processes to read the object, but they cannot have write or read-intend-write access.</p> <p>IOS_\$UNREGULATED_OPT Permit shared (unregulated) concurrency mode.</p> <p>IOS_\$WRITE_OPT Permit writing to the object.</p>
NAME_\$PNAME_T	An array of up to name_\$pnamlen_max (256) characters.
IOS_DIR_\$ENTRY_T	The format of a record in an IOS_DIR object (that is, in a directory). The diagram below illustrates the IOS_DIR_\$ENTRY_T data type:

predefined type	byte: offset	field name
	0: integer	enttype
	2: integer	entlen
name_\$name_t	4: character array	entname
	8: character array	
	12: character array	
	16: character array	
	20: character array	
	24: character array	
	28: character array	
	32: character array	
	36: integer	unused1
	40: integer	unused2

Field Description:

enttype  
The type of directory entry. The value 1 indicates a file; the value 3 indicates a link.

entlen  
Length of the name of the entry.

entname  
Name of the entry.

unused1  
Reserved for future use.


unused2  
Reserved for future use.

IOS\_DIR\_\$MODE\_T

Reserved for internal use.

STATUS\_\$T

A status code. The diagram below illustrates the STATUS\_\$T data type:



Field Description:

all

All 32 bits in the status code.

fail

The fail bit. If this bit is set, the error was not within the scope of the module invoked, but occurred within a lower-level module (bit 31).

subsys

The subsystem that encountered the error (bits 24 - 30).

modc

The module that encountered the error (bits 16 - 23).

code

A signed number that identifies the type of error that occurred (bits 0 - 15).

**IOS\_DIR\_\$ISA**

Determines whether an open stream is associated with an object that is a directory.

**FORMAT**

IOS\_DIR\_\$ISA (stream-id, status)

**INPUT PARAMETERS****stream-id**

Stream ID for an open stream, in ios\_\$id\_t format. This is a 2-byte integer.

**OUTPUT PARAMETERS****status**

Completion status, in STATUS\_\$T format. This data type is 4 bytes long. See the IOS\_DIR Data Types section for more information.

**USAGE**

The ios\_dir\_\$isa call determines whether the IOS\_DIR trait is implemented for an object to which you have an open stream. If this trait is implemented, then the object must be a directory. Therefore, you can use the ios\_dir\_\$isa call to determine whether a particular stream ID is associated with an object that is a directory.

IOS\_DIR\_\$OPEN

IOS\_DIR\_\$OPEN

Opens a stream connection to an object in a currently open directory.

## FORMAT

stream-id = IOS\_DIR\_\$OPEN (dir-id, leaf-name, leaf-namelen, open-options, status)

## RETURN VALUE

### stream-id

Stream ID for the stream to the object, in ios\_\$id\_t format. This is a 2-byte integer.

## INPUT PARAMETERS

### dir-id

Stream ID for the stream to the directory that contains the object, in ios\_\$id\_t format. This is a 2-byte integer.

### leaf-name

Name of the object to open, relative to the directory associated with dir-id parameter, in pname\_\$name\_t format. This is an array of up to name\_\$pnamlen\_max (256) characters. The object must be an element in the specified directory, not in a subdirectory.

### leaf-namelen

Length of the leaf-name. This is a 2-byte integer.

### open-options

A set (bitmask) of options, in ios\_\$open\_options\_t format. This is a 2-byte integer. Specify one or more of the following predefined values:

IOS\_\$INQUIRE\_ONLY\_OPT

Open the object for attribute inquires only; do not permit reading or writing of data.

IOS\_\$NO\_OPEN\_DELAY\_OPT

Return immediately, instead of waiting for the open call to complete.

IOS\_\$POSITION\_TO\_EOF\_OPT

Position the stream marker to the end of the object (EOF).

IOS\_\$READ\_INTEND\_WRITE\_OPT

Open the object for read access with the intent to eventually change to write access. This allows other processes to read the object, but they cannot have write or read-intend-write access.

IOS\_\$UNREGULATED\_OPT

Permit shared (unregulated) concurrency mode.

IOS\_\$WRITE\_OPT

Permit writing to the object.

**status**

Completion status, in STATUS\_\$\$T format. This data type is 4 bytes long. See the IOS\_DIR Data Types section for more information.

**USAGE**

This call lets you open an object by specifying its name relative to a currently open directory. For example, if you have already opened the directory //barn/animals, you can use ios\_dir\_\$open to open an object named horse (in this directory). To do this, specify the stream ID for the directory and the leaf name horse. Note that when you use ios\_dir\_\$open, you must specify the name of a leaf in the specified directory, not in a subdirectory.

IOS\_DIR ERRORS

ERRORS

STATUS\_\$OK

Successful completion.



# IPC

## Table of Contents

Data Types	IPC-2
IPC_\$CLOSE	IPC-4
IPC_\$CREATE	IPC-5
IPC_\$DELETE	IPC-6
IPC_\$GET_EC	IPC-7
IPC_\$OPEN	IPC-8
IPC_\$RCV	IPC-9
IPC_\$RESOLVE	IPC-11
IPC_\$SAR	IPC-12
IPC_\$SEND	IPC-14
IPC_\$WAIT	IPC-15
Errors	IPC-16




## IPC

The IPC (Interprocess Communications) programming calls perform interprocess communications using datagrams addressed to destinations identified with file system names. This section describes their data types, call syntax, and error codes. Refer to the Introduction at the beginning of this manual for a description of data-type diagrams and call syntax format.

IPC DATA TYPES

DATA TYPES

- IPC\_\$DATA\_T                      An array of up to 1024 characters. The data portion of an IPC datagram.
  
- IPC\_\$HDR\_INFO\_T                An array of up to 128 characters. The header portion of an IPC datagram.
  
- IPC\_\$SOCKET\_HANDLE\_T         An array of 20 characters. A handle for an IPC socket.
  
- NAME\_\$PNAME\_T                 An array of up to 256 characters. A DOMAIN pathname.
  
- STATUS\_\$T                      A status code. The diagram below illustrates the STATUS\_\$T data type:



Field Description:

all  
All 32 bits in the status code.

fail  
The fail bit. If this bit is set, the error was not within the scope of the module invoked, but occurred within a lower-level module (bit 31).

subsys  
The subsystem that encountered the error (bits 24 - 30).

modc  
The module that encountered the error (bits 16 - 23).

code

A signed number that identifies the type of error that occurred (bits 0 - 15).

IPC\_\$CLOSE

IPC\_\$CLOSE

Closes an IPC socket.

## FORMAT

IPC\_\$CLOSE (pathname, length, status)

## INPUT PARAMETERS

### pathname

Pathname for the file where the socket handle is stored, in NAME\_\$PNAME\_T format. This is an array of up to 256 characters. Specify a file that was created by a previous IPC\_\$CREATE call.

### length

Length of the pathname. This is a 2-byte integer.

## OUTPUT PARAMETERS

### status

Completion status, in STATUS\_\$T format. This data type is 4 bytes long. See the IPC Data Types section for more information.

## USAGE

IPC\_\$CLOSE closes a socket and removes its handle from the file where the handle is stored. IPC\_\$CLOSE does not, however, delete the socket handle file. To delete this file, use IPC\_\$DELETE.

**IPC\_\$CREATE**

Creates a file where an IPC socket handle can be stored.

**FORMAT**

IPC\_\$CREATE (pathname, length, status)

**INPUT PARAMETERS****pathname**

Pathname for a file where a socket handle can be stored, in NAME\_\$PNAME\_T format. This is an array of up to 256 characters.

**length**

Length of the pathname. This is a 2-byte integer.

**OUTPUT PARAMETERS****status**

Completion status, in STATUS\_\$T format. This data type is 4 bytes long. See the IPC Data Types section for more information.

**USAGE**

IPC\_\$CREATE creates a special type of DOMAIN system object that is used only for socket handles. When you open a socket, the system obtains a handle for the socket and stores this handle in the file that you specify. You can open a socket only if you have previously used IPC\_\$CREATE to create a file for the handle.

IPC\_\$DELETE

IPC\_\$DELETE

Deletes a file that was used to store an IPC socket handle.

## FORMAT

IPC\_\$DELETE (pathname, length, status)

## INPUT PARAMETERS

### pathname

Pathname for the file where the socket handle was stored, in NAME\_\$PNAME\_T format. This is an array of up to 256 characters.

### length

Length of the pathname. This is a 2-byte integer.

## OUTPUT PARAMETERS

### status

Completion status, in STATUS\_\$T format. This data type is 4 bytes long. See the IPC Data Types section for more information.

## USAGE

IPC\_\$DELETE deletes a file that the system used to store a handle for an open socket. You must call IPC\_\$CLOSE to close the socket before you can delete the file containing the socket's handle.



**IPC\_\$GET\_EC**

Gets a pointer to the eventcount associated with an IPC socket.

**FORMAT**

IPC\_\$GET\_EC (handle, ec-ptr, status)

**INPUT PARAMETERS****handle**

Handle for the socket whose eventcount you are getting, in IPC\_\$SOCKET\_HANDLE\_T format. This is an array of 20 characters.

**OUTPUT PARAMETERS****ec-ptr**

Pointer to the eventcount, in EC2\_\$PTR\_T format. This is a 4-byte integer.

**status**

Completion status, in STATUS\_\$T format. This data type is 4 bytes long. See the IPC Data Types section for more information.

**USAGE**

IPC\_\$GET\_EC gets a pointer to the eventcount associated with an IPC socket. You can use this eventcount to wait for incoming datagrams. Use EC2 system calls to read the eventcount value and wait for datagrams.

IPC\_\$GET\_EC is useful when you want to wait for messages arriving in more than one socket. To wait for messages from only one socket, use IPC\_\$WAIT or IPC\_\$SAR.

IPC\_\$OPEN

IPC\_\$OPEN

Opens an available IPC socket, obtains its handle, and places the handle in a file.

## FORMAT

IPC\_\$OPEN (pathname, length, depth, handle, status)

## INPUT PARAMETERS

### pathname

Pathname for the file in which to store the handle, in NAME\_\$PNAME\_T format. This is an array of up to 256 characters. Specify a file that you have created with a previous IPC\_\$CREATE call.

### length

Length of the pathname. This is a 2-byte integer.

### depth

Depth of the socket. The depth defines how many datagrams a socket can hold. Allowable values are one through four.

## OUTPUT PARAMETERS

### handle

Handle for the open socket, in IPC\_\$SOCKET\_HANDLE\_T format. This is an array of 20 characters.

### status

Completion status, in STATUS\_\$T format. This data type is 4 bytes long. See the IPC Data Types section for more information.

## USAGE

IPC\_\$OPEN opens an available IPC socket on your program's local node. In addition, IPC\_\$OPEN places the socket's handle in the file you specify. After opening a socket, you can receive datagrams in it. A program must use your socket's handle to send you a message.

User programs running on a node can open a maximum of eight sockets on that node. Only one program at a time can open any socket.

You must use IPC\_\$CREATE to create a file for the socket handle before you can open a socket.

**IPC\_\$RCV**

Gets a datagram that has been received in an IPC socket. This call copies the datagram to the buffers that you specify.

**FORMAT**

```
IPC_$RCV (handle, hdr-buflen, data-buflen, from-handle,
          hdr-buf, hdr-length, data-buf, data-length, status)
```

**INPUT PARAMETERS****handle**

Handle for the socket that received the datagram, in IPC\_\$SOCKET\_HANDLE\_T format. This is an array of 20 characters.

**hdr-buflen**

Length of the buffer where the datagram header will be copied. This is a 2-byte integer. This value defines the maximum number of header bytes that IPC\_\$RCV will get. An IPC datagram can contain up to 128 header bytes. Specify a length that can accommodate the longest header you expect to receive.

**data-buflen**

Length of the buffer where the data portion of the datagram will be copied. This is a 2-byte integer. This value defines the maximum number of data bytes that IPC\_\$RCV will get. The data portion of an IPC datagram can contain up to 1024 bytes. Specify a length that can accommodate the longest data you expect to receive.

**OUTPUT PARAMETERS****from-handle**

Handle for the socket where the datagram originated, in IPC\_\$SOCKET\_HANDLE\_T format. This is an array of 20 characters. Use this handle to send a reply to the datagram you are currently getting.

**hdr-buf**

Buffer where the datagram header is copied. This buffer can contain up to 128 bytes.

**hdr-length**

Length, in bytes, of the header that is copied. This is a 2-byte integer.

**data-buf**

Buffer where the data portion of the datagram is copied. This buffer can contain up to 1024 bytes.

**data-length**

Length, in bytes, of the data that is copied. This is a 2-byte integer.

**status**

Completion status, in STATUS\_\$T format. This data type is 4 bytes long. See the IPC Data Types section for more information.

## IPC\_\$RCV

IPC\_\$RCV gets datagrams that have been received in a socket and copies them to your buffers. This call returns only the number of header and data bytes that you specify, even if the actual datagram (in the socket) contains more bytes.

IPC\_\$RCV gets datagrams in the order in which they arrive in the socket queue. If the socket queue is full when an incoming datagram arrives, the datagram is lost. You can use IPC\_\$RCV to get datagrams only from a socket that you have previously opened with IPC\_\$OPEN.

Usually, you wait for a datagram to arrive in a socket, and then call IPC\_\$RCV to get the datagram. If you call IPC\_\$RCV when the socket is empty, the call returns immediately with the status IPC\_\$SOCKET\_EMPTY.

**IPC\_\$RESOLVE**

Obtain the handle for an open socket.

**FORMAT**

IPC\_\$RESOLVE (pathname, length, handle, status)

**INPUT PARAMETERS****pathname**

Pathname for the file containing the socket handle, in NAME\_\$PNAME\_T format. This is an array of up to 256 characters. Specify a file that was created by a previous IPC\_\$CREATE call.

**length**

Length of the pathname. This is a 2-byte integer.

**OUTPUT PARAMETERS****handle**

Handle for the socket, in IPC\_\$SOCKET\_HANDLE\_T format. This is an array of 20 characters.

**status**

Completion status, in STATUS\_\$T format. This data type is 4 bytes long. See the IPC Data Types section for more information.

**USAGE**

IPC\_\$RESOLVE returns the handle associated with an open socket. Use this call if you know a socket's pathname, but you need the socket handle to send a datagram.

IPC\_\$RESOLVE returns the error IPC\_\$SOCKET\_NOT\_OPEN if the handle file does not contain the handle for an open socket.

IPC\_\$\$SAR

IPC\_\$\$SAR

Performs a single send/await-reply operation. This call sends a datagram, waits a specified amount of time for a reply, and copies the reply to the buffers you specify.

## FORMAT

IPC\_\$\$SAR (retry-time, to-handle, in-hdr-buf, in-hdr-length, in-data-buf, in-data-length, out-hdr-buflen, out-data-buflen, out-hdr-buf, out-hdr-length, out-data-buf, out-data-length, status)

## INPUT PARAMETERS

### retry-time

Number of quarter-seconds to wait for a reply. This is a 2-byte integer.

### to-handle

Handle for the destination socket, in IPC\_\$\$SOCKET\_HANDLE\_T format. This is an array of 20 characters. The destination socket is where you are sending the datagram.

### in-hdr-buf

Buffer that contains the header for the datagram you are sending. This buffer can contain up to 128 bytes.

### in-hdr-length

Length, in bytes, of the header you are sending. This is a 2-byte integer.

### in-data-buf

Buffer that contains the data portion of the datagram you are sending. This buffer can contain up to 1024 bytes.

### in-data-length

Length, in bytes, of the data you are sending. This is a 2-byte integer.

### out-hdr-buflen

Length of the buffer where the reply datagram header will be copied. This is a 2-byte integer. This value defines the maximum number of header bytes that IPC\_\$\$SAR will get from the reply datagram. The reply can contain up to 128 header bytes. Specify a length that can accommodate the longest header you expect to receive.

### out-data-buflen

Length of the buffer where the data portion of the reply datagram will be copied. This is a 2-byte integer. This value defines the maximum number of data bytes that IPC\_\$\$SAR will get from the reply datagram. The data portion of a reply can contain up to 1024 bytes. Specify a length that can accommodate the longest data you expect to receive.

## OUTPUT PARAMETERS

### out-hdr-buf

Buffer where the header for the reply datagram is copied. This buffer can contain up to 128 bytes.

### out-hdr-length

Length, in bytes, of the header that is copied. This is a 2-byte integer.

**out-data-buf**

Buffer where the data portion of the reply datagram is copied. This buffer can contain up to 1024 bytes.

**out-data-length**

Length, in bytes, of the data that is copied. This is a 2-byte integer.

**status**

Completion status, in STATUS\_\$\$T format. This data type is 4 bytes long. See the IPC Data Types section for more information.

**USAGE**

Use IPC\_\$\$SAR to send a datagram to another process and wait a specified time for a reply. If the reply datagram does not arrive within the specified time, IPC\_\$\$SAR returns the status code IPC\_\$\$TIMEOUT.

IPC\_\$\$SAR returns only the number of header and data bytes that you specify, even if the actual datagram (in the socket) contains more bytes.

When you send a datagram that contains less than 128 bytes of information, you can place all the information in the header buffer. Then specify the data length as zero. It takes less time to send a datagram that contains only a header.

IPC\_\$SEND

IPC\_\$SEND

Sends a datagram to an IPC socket.

## FORMAT

IPC\_\$SEND (to-handle, reply-handle, hdr-buf, hdr-length,  
data-buf, data-length, status)

## INPUT PARAMETERS

### to-handle

Handle for the destination socket, in IPC\_\$SOCKET\_HANDLE\_T format. This is an array of 20 characters. The destination socket is where you are sending the datagram.

### reply-handle

Handle for the reply socket, in IPC\_\$SOCKET\_HANDLE\_T format. This is an array of 20 characters. The reply socket is where you can receive a reply.

### hdr-buf

Buffer that contains the header for the datagram you are sending. This buffer can contain up to 128 bytes.

### hdr-length

Length, in bytes, of the datagram header. This is a 2-byte integer.

### data-buf

Buffer that contains the data portion of the datagram you are sending. This buffer can contain up to 1024 bytes.

### data-length

Length, in bytes, of the data portion of the datagram. This is a 2-byte integer.

## OUTPUT PARAMETERS

### status

Completion status, in STATUS\_\$T format. This data type is 4 bytes long. See the IPC Data Types section for more information.

## USAGE

IPC\_\$SEND sends a datagram to the socket that you specify. To obtain a socket handle from a pathname, use IPC\_\$RESOLVE.

Even if IPC\_\$SEND completes successfully, there is no guarantee that the datagram will be received by the process you are sending it to. The programs using IPC datagrams are responsible for verifying that datagrams are successfully received. Note that you can use IPC\_\$SAR to perform a send/await reply operation with a single call.

When you send a datagram that contains less than 128 bytes of information, you can place all the information in the header buffer. Then specify the data length as zero. It takes less time to send a datagram that contains only a header.



**IPC\_\$WAIT**

Waits for a specified amount of time to receive a datagram in an IPC socket.

**FORMAT**

IPC\_\$WAIT (handle, wait-time, status)

**INPUT PARAMETERS****handle**

Handle for the socket that you are waiting to receive data in, in IPC\_\$SOCKET\_HANDLE\_T format. This is an array of 20 characters.

**wait-time**

Number of quarter-seconds to wait for a reply. This is a 2-byte integer.

**OUTPUT PARAMETERS****status**

Completion status, in STATUS\_\$T format. This data type is 4 bytes long. See the IPC Data Types section for more information.

**USAGE**

IPC\_\$WAIT waits for a specified amount of time to receive a datagram in a socket. If a datagram is received before the time elapses, the call returns with the status STATUS\_\$OK. To get the datagram, use IPC\_\$RCV.

If IPC\_\$WAIT times out before a datagram is received, the call returns with the status IPC\_\$TIMEOUT. If you call IPC\_\$WAIT and there is a datagram already in the socket, the call returns immediately with a success status.

Note that you can use IPC\_\$SAR to perform a send/await reply operation with a single call. Also, if you want to wait for datagrams in more than one socket, use IPC\_\$GET\_EC to get pointers to the appropriate eventcounts. Then use eventcount calls (EC2) to wait for datagrams.

IPC ERRORS

**ERRORS**

IPC\_\$OK

Successful completion.

IPC\_\$NOMORE\_SOCKETS

All the sockets are in use.

IPC\_\$NOT\_IPC\_OBJ

The specified pathname does not belong to an IPC object.

IPC\_\$NOT\_OWNER

You did not open the socket so you cannot close it.

IPC\_\$RANGE\_ERROR

Supplied socket number is outside legal range.

IPC\_\$SOCKET\_ALREADY\_OPEN

Specified socket is already open.

IPC\_\$SOCKET\_EMPTY

There are no datagrams in the socket.

IPC\_\$SOCKET\_NOT\_OPEN

The specified socket is not open.

IPC\_\$TIMEOUT

The call timed out before a datagram was received.

IPC\_\$TOO\_DEEP

Supplied socket depth is too big.

IPC\_\$TOO\_MUCH\_DATA

The data is too long to send.

STATUS\_\$OK

Successful completion.

# MBX

## Table of Contents

Data Types	MBX-2
MBX_\$CLIENT_WINDOW	MBX-7
MBX_\$CLOSE	MBX-8
MBX_\$COND_GET_REC_CHAN	MBX-9
MBX_\$COND_GET_REC_CHAN_SET	MBX-11
MBX_\$CREATE_SERVER	MBX-13
MBX_\$DEALLOCATE	MBX-15
MBX_\$GET_CONDITIONAL	MBX-16
MBX_\$GET_EC	MBX-18
MBX_\$GET_REC	MBX-19
MBX_\$GET_REC_CHAN	MBX-21
MBX_\$GET_REC_CHAN_SET	MBX-23
MBX_\$OPEN	MBX-25
MBX_\$PUT_CHR	MBX-26
MBX_\$PUT_CHR_COND	MBX-27
MBX_\$PUT_REC	MBX-28
MBX_\$PUT_REC_COND	MBX-29
MBX_\$SERVER_WINDOW	MBX-30
MBX_\$TIMED_OPEN	MBX-32
Errors	MBX-33



## MBX

The MBX (Mailbox) programming calls perform interprocess communications using virtual circuits addressed to destinations identified with pathnames. This section describes their data types, call syntax, and error codes. Refer to the Introduction at the beginning of this manual for a description of data-type diagrams and call syntax format.

## MBX DATA TYPES

### CONSTANTS

MBX_\$CHN_MAX	255	Maximum number of channels that can be open to a mailbox.
MBX_\$FIL_MAX	257*32768	Maximum mailbox size.
MBX_\$MIN_CHN_SIZE	64	The minimum size of a channel buffer.
MBX_\$MSG_MAX	1024	A mailbox message that is 1024 bytes long.
MBX_\$MSG_TN	1023	For use when declaring a zero-based array that is MBX_\$MSG_MAX bytes long.
MBX_\$MSG_WMAX	512	A mailbox message that is 512 words long.
MBX_\$MSG_WTN	511	For use when declaring a zero-based array that is MBX_\$MSG_WMAX words long.
MBX_\$REC_DATA_MAX	32760	The maximum length of the data portion of a mailbox message.
MBX_\$REC_MSG_MAX	32766	The maximum length of a server message, including the header and data portions.
MBX_\$SERV_MSG_HDR_LEN	6	Length of the mailbox header for a server message.
MBX_\$SERV_MSG_MAX	1030	A server message that contains 1024 bytes of data plus a 6-byte header.
MBX_\$VERSION	1	Current version of MBX.

### DATA TYPES

EC2_\$PTR_T	A 4-byte integer. A pointer to an eventcount.
MBX_\$CHAN_NUM_T	A channel number. Possible values are integers from 0 through MBX_\$CHN_MAX.

MBX\_\$CHAN\_SET\_T

A set of channel numbers of type MBX\_\$CHAN\_NUM\_T. The following Pascal example specifies channels 1, 4, and 7:

```

VAR
    chan_set : mbx_chan_set_t
    .
    .
    .
    chan_set := [ 1, 4, 7]
    
```

In a FORTRAN program, declare an 8-element array of 4-byte integers to indicate a channel set. Use the array as a mask in which the bits represent mailbox channels.

MBX\_\$EC\_KEY\_T

A 2-byte integer. A mailbox eventcount. One of the following predefined values:

MBX\_\$GETREC\_EC\_KEY  
An eventcount that advances when the mailbox contains messages for you to get.

MBX\_\$PUTREC\_EC\_KEY  
An eventcount that advances when enough room exists in the channel to hold the last message you unsuccessfully tried to put there.

MBX\_\$MTYPE\_T

A 2-byte integer. A message type. One of the following predefined values:

MBX\_\$ACCEPT\_OPEN\_MT  
A response from a server to accept a client's open request.

MBX\_\$CHANNEL\_OPEN\_MT  
A request from a client to open a channel to a mailbox.

MBX\_\$DATA\_MT  
A data transmission.

MBX\_\$DATA\_PARTIAL\_MT  
A partial data transmission.

MBX\_\$EOF\_MT  
An end of transmission notice.

MBX\_\$REJECT\_OPEN\_MT  
A response from a server to reject a client's open request.

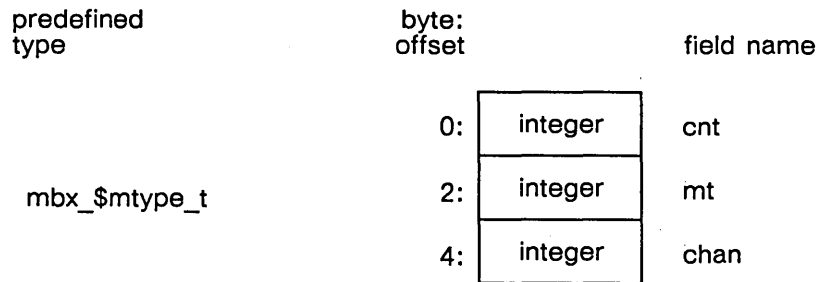
MBX\_\$NAME\_T

An array of up to 256 characters. A mailbox name.

MBX DATA TYPES

MBX\_\$MSG\_HDR\_T

A mailbox message header. The diagram below illustrates the MBX\_\$MSG\_HDR\_T data type:



Field Description:

cnt

The total number of bytes in the message, including the header.

mt

A value representing a message type. This value is one of the predefined values of type MBX\_\$MTYPE\_T.

chan

The channel of the client that sent the message, or that should receive the message.



MBX\_\$SERVER\_MSG\_T


A server message with up to 1024 data bytes. The diagram below illustrates the MBX\_\$SERVER\_MSG\_T data type:

predefined  
type

byte:  
offset

field name

mbx\_\$mtype\_t



Field Description:

cnt

The total number of bytes in the message, including the header.

mt

A value representing a message type. This value is one of the predefined values of type MBX\_\$MTYPE\_T.

chan

The channel of the client that sent the message, or that should receive the message.


data

The data portion of the message. This field can contain up to 1024 bytes.

MBX DATA TYPES

STATUS\_\$\$T

A status code. The diagram below illustrates the STATUS\_\$\$T data type:



Field Description:

**all**  
All 32 bits in the status code.

**fail**  
The fail bit. If this bit is set, the error was not within the scope of the module invoked, but occurred within a lower-level module (bit 31).

**subsys**  
The subsystem that encountered the error (bits 24 - 30).

**modc**  
The module that encountered the error (bits 16 - 23).

**code**  
A signed number that identifies the type of error that occurred (bits 0 - 15).

**MBX\_\$CLIENT\_WINDOW**

Returns the buffer size for the mailbox that a client is using.

**FORMAT**

size = MBX\_\$CLIENT\_WINDOW (handle, status)

**RETURN VALUE****size**

Buffer size for the mailbox. This is a 4-byte integer.

This value defines a window size when a client sends messages to a remote server. That is, the client cannot send messages that are larger than the mailbox's buffer.

**INPUT PARAMETERS****handle**

Identifier for the mailbox, in UNIV\_PTR format. This is a 4-byte integer. Use the handle returned by MBX\_\$CREATE\_SERVER.

**OUTPUT PARAMETERS****status**

Completion status, in STATUS\_\$T format. This data type is 4 bytes long. See the MBX Data Types section for more information.

**USAGE**

When a client sends a message, the message is stored in a channel buffer until the server gets the message. The buffer size defines the maximum number of message bytes that the channel can hold at one time.

A client can use MBX\_\$CLIENT\_WINDOW to get the size of the channel buffer. To get the size, MBX\_\$CLIENT\_WINDOW queries the MBX\_HELPER on the server's node. Note that MBX\_\$CLIENT\_WINDOW returns the actual buffer size, not the number of unused bytes in the buffer.

MBX\_\$CLIENT\_WINDOW only works correctly when the server you are inquiring about is on a node with SR9 or later software. If you call MBX\_\$CLIENT\_WINDOW and the server is on a node with pre-SR9 software, MBX\_\$CLIENT\_WINDOW returns the value 1158. This value is returned, even if the mailbox's actual buffer size is smaller. Therefore, this call does not provide a reliable way to determine the window size when sending messages to a server that is running on a node with pre-SR9 software.

MBX\_\$CLIENT\_WINDOW is for use only by mailbox clients. A server should use MBX\_\$SERVER\_WINDOW.

MBX\_\$CLOSE

MBX\_\$CLOSE

Closes a mailbox or a channel.

#### FORMAT

MBX\_\$CLOSE (handle, status)

#### INPUT PARAMETERS

##### handle

Identifier for the mailbox, in UNIV\_PTR format. This is a 4-byte integer. Use the handle returned by MBX\_\$CREATE\_SERVER or MBX\_\$OPEN.

#### OUTPUT PARAMETERS

##### status

Completion status, in STATUS\_\$T format. This data type is 4 bytes long. See the MBX Data Types section for more information.

#### USAGE

Both clients and servers can use MBX\_\$CLOSE. When called from a client, MBX\_\$CLOSE tells the server that the client is no longer using the channel. When called from a server, MBX\_\$CLOSE closes the mailbox.

After a client calls MBX\_\$CLOSE, the server should call MBX\_\$DEALLOCATE to deallocate the channel and free it for use by other clients. No other client can use the channel until it has been deallocated by the server.

If a server closes a mailbox while there are still active clients, the clients get errors on subsequent attempts to use the mailbox.

## MBX\_\$COND\_GET\_REC\_CHAN

Attempts to get a mailbox message from a specified channel.

**FORMAT**

MBX\_\$COND\_GET\_REC\_CHAN (handle, channel, bufptr, buflen, retptr, retlen, status)

**INPUT PARAMETERS****handle**

Identifier for the mailbox, in UNIV\_PTR format. This is a 4-byte integer. Use the handle you obtained from MBX\_\$CREATE\_SERVER.

**channel**

Channel to read from. This is a 2-byte integer. The mailbox manager assigns a channel number to a client when the client calls MBX\_\$OPEN.

**bufptr**

Pointer to a data buffer where the message can be copied. This is a 4-byte integer.

Your program must allocate a data buffer, although the mailbox manager does not always copy messages to this buffer. Use the output parameter **retptr** to reference the message.

**buflen**

The number of bytes in the data buffer. This is a 4-byte integer. For a server, MBX will never return more than 32766 bytes. For a client, MBX will never return more than 32760 bytes.

**OUTPUT PARAMETERS****retptr**

Pointer to the buffer where the message is copied. This is a 4-byte integer.

**retlen**

Either the number of bytes in the returned message, or the number of message bytes waiting to be returned. This is a 4-byte integer.

MBX\_\$COND\_GET\_REC\_CHAN can get as many bytes as you specify in **buflen**. If the message is less than or equal to **buflen**, then the call gets the entire message and **retlen** specifies the message length. If the message is greater than **buflen**, then the call gets the number of bytes specified in **buflen**. If this occurs, then **retlen** contains a negative value, the absolute value of which is the number of bytes remaining in the message. Get the remaining data with another call.

Note that a server sees the message header each time it gets a piece of the message. The count field contains the total length of the message -- not the length of the returned piece.

**status**

Completion status, in STATUS\_\$T format. This data type is 4 bytes long. See the MBX Data Types section for more information.

MBX\_\$COND\_GET\_REC\_CHAN

## USAGE

MBX\_\$COND\_GET\_REC\_CHAN requests a message from a specified channel. If there is no message, the call returns immediately with the status MBX\_\$CHANNEL\_EMPTY. You can use an eventcount to tell when the status of the mailbox has changed. You get a mailbox eventcount with MBX\_\$GET\_EC.

Only a server can use MBX\_\$COND\_GET\_REC\_CHAN. To perform a conditional get operation from a client, use MBX\_\$GET\_CONDITIONAL.

## MBX\_\$COND\_GET\_REC\_CHAN\_SET

Attempts to get a mailbox record from a set of clients.

**FORMAT**

MBX\_\$COND\_GET\_REC\_CHAN\_SET (handle, chan-set, bufptr, buflen, retptr, retlen, status)

**INPUT PARAMETERS****handle**

Identifier for the mailbox, in UNIV\_PTR format. This is a 4-byte integer. Use the handle returned by MBX\_\$CREATE\_SERVER.

**chan-set**

Set of channels to read from, in MBX\_\$CHAN\_SET\_T format. This is an 8-element array of 4-byte integers. See the MBX Data Types section for more information.

The mailbox manager assigns a channel number to a client when the client calls MBX\_\$OPEN. The channel number can range from 1 through MBX\_\$CHN\_MAX.

**bufptr**

Pointer to a data buffer where the message can be copied. This is a 4-byte integer.

Your program must allocate a data buffer, although the mailbox manager does not always copy messages to this buffer. Use the output parameter **retptr** to reference the message.

**buflen**

Number of bytes in the data buffer. This is a 4-byte integer. For a server, MBX will never return more than 32766 bytes. For a client, MBX will never return more than 32760 bytes.

**OUTPUT PARAMETERS****retptr**

Pointer to the buffer where the message is copied. This is a 4-byte integer.

**retlen**

Either the number of bytes in the returned message, or the number of message bytes waiting to be returned. This is a 4-byte integer.

MBX\_\$COND\_GET\_REC\_CHAN\_SET can get as many bytes as you specify in **buflen**. If the message is less than or equal to **buflen**, then the call gets the entire message and **retlen** specifies the message length. If the message is greater than **buflen**, then the call gets the number of bytes specified in **buflen**. If this occurs, then **retlen** contains a negative value, the absolute value of which is the number of bytes remaining in the message. Get the remaining data with another call.

Note that a server sees the message header each time it gets a piece of the message. The count field contains the total length of the message -- not the length of the returned piece.

**status**

Completion status, in STATUS\_\$T format. This data type is 4 bytes long. See the MBX Data Types section for more information.

MBX\_\$COND\_GET\_REC\_CHAN\_SET

## USAGE

MBX\_\$COND\_GET\_REC\_CHAN\_SET requests a message from a specified set of channels. If there is no message, the call returns immediately with the status MBX\_\$CHANNEL\_EMPTY. You can use an eventcount to tell when the status of the mailbox has changed. You get a mailbox eventcount with MBX\_\$GET\_EC.

Only a mailbox server can use this call. To perform a conditional get operation from a client, use MBX\_\$GET\_CONDITIONAL.



**MBX\_\$CREATE\_SERVER**

Creates and opens a server's mailbox.

**FORMAT**

MBX\_\$CREATE\_SERVER (name, namelen, bufsize, maxchan, handle, status)

**INPUT PARAMETERS****name**

Name of the mailbox, in MBX\_\$NAME\_T format. This is an array of up to 256 characters. Specify the name as a pathname to the mailbox file. If you use the name of a file that already exists, this call deletes the contents of the file. If the file already exists and it is in use, then the call returns an error.

**namelen**

Number of characters in the name. This is a 2-byte integer.

**bufsize**

Number of message bytes that the server and client can each store in a channel. This is a 2-byte integer. You must specify a buffer size of at least MBX\_\$MIN\_CHN\_SIZE (64 bytes). This allocates 128 bytes for each channel -- 64 bytes apiece for the server and the client buffers. The maximum buffer size is 32766. If you specify an odd numbered buffer size, it is rounded down to an even number.

The buffer size should be large enough to store the largest message you plan to send from a server or a client. Note that the maximum message length is MBX\_\$REC\_MSG\_MAX (32766), which includes 32760 data bytes plus a 6-byte header. If you specify a buffer size of less than MBX\_\$REC\_MSG\_MAX, you impose a lower limit on the total length of messages that pass through the mailbox.

Note that if you specify a buffer size that is greater than 1158, and the server is communicating with clients on remote nodes, the length of the transmitted messages may be limited by the MBX\_HELPER on the client node. When a server puts a message into the mailbox and the message is intended for a remote client, the message passes through the system mailbox maintained by the remote node's MBX\_HELPER. By default, this mailbox has a buffer size of 1158 bytes. To allow the remote node's mailbox to handle larger messages, use the -DATASIZE option to specify a larger buffer size when you start the MBX\_HELPER. Specify a value that is at least as large as the largest message the server will send.

**maxchan**

Maximum number of channels that can be simultaneously open to the mailbox. This is a 2-byte integer. You can allow up to MBX\_\$CHN\_MAX (255) channels.

**OUTPUT PARAMETERS****handle**

Identifier for the mailbox, in UNIV\_PTR format. This is a 4-byte integer. Subsequent calls use this handle to send and receive messages.

MBX\_\$CREATE\_SERVER

**status**

Completion status, in STATUS\_\$T format. This data type is 4 bytes long. See the MBX Data Types section for more information.

**USAGE**

A server uses MBX\_\$CREATE\_SERVER to create a mailbox. Once the mailbox is open, clients use MBX\_\$OPEN to open communications channels to the mailbox.

In a secure network, a mailbox gets an access control list (ACL) that is determined by the ACL of the directory in which the mailbox is created. If servers and clients on different nodes use the mailbox, be sure that the server's MBX\_HELPER has read and write access to the mailbox.

**MBX\_ \$DEALLOCATE**

Releases a channel for use by another client.

**FORMAT**

MBX\_ \$DEALLOCATE (handle, channel, status)

**INPUT PARAMETERS****handle**

Identifier for the mailbox, in UNIV\_PTR format. This is a 4-byte integer. Use the handle returned by MBX\_ \$CREATE\_SERVER.

**channel**

Channel to deallocate. This is a 2-byte integer. The mailbox manager assigns a channel number to a client when the client calls MBX\_ \$OPEN.

**OUTPUT PARAMETERS****status**

Completion status, in STATUS\_ \$T format. This data type is 4 bytes long. See the MBX Data Types section for more information.

**USAGE**

Only a server can call MBX\_ \$DEALLOCATE. A client uses MBX\_ \$CLOSE to indicate the end of transmission over a channel. However, the server must deallocate the channel before another client can use it.

A server can deallocate a channel while a client is still using it; this both closes and deallocates the channel. The next time the client tries to use the channel, the client receives the error MBX\_ \$CHANNEL\_NOT\_OPEN.

MBX\_\$GET\_CONDITIONAL

MBX\_\$GET\_CONDITIONAL

Attempts to get a mailbox message.

## FORMAT

MBX\_\$GET\_CONDITIONAL (handle, bufptr, buflen, retptr, retlen, status)

## INPUT PARAMETERS

### handle

Identifier for the mailbox, in UNIV\_PTR format. This is a 4-byte integer. Use the handle returned by MBX\_\$CREATE\_SERVER or MBX\_\$OPEN.

### bufptr

Pointer to a data buffer where the message can be copied. This is a 4-byte integer.

Your program must allocate a data buffer, although the mailbox manager does not always copy messages to this buffer. Use the output parameter retptr to reference the message.

### buflen

Number of bytes in the data buffer. This is a 4-byte integer. For a server, MBX will never return more than 32766 bytes. For a client, MBX will never return more than 32760 bytes.

## OUTPUT PARAMETERS

### retptr

Pointer to the buffer where the message is copied. This is a 4-byte integer.

### retlen

Either the number of bytes in the returned message, or the number of message bytes waiting to be returned. This is a 4-byte integer.

MBX\_\$GET\_CONDITIONAL can get as many bytes as you specify in buflen. If the message is less than or equal to buflen, then the call gets the entire message and retlen specifies the message length. If the message is greater than buflen, then the call gets the number of bytes specified in buflen. If this occurs, then retlen contains a negative value, the absolute value of which is the number of bytes remaining in the message. Get the remaining data with another call.

Note that a server sees the message header each time it gets a piece of the message. The count field contains the total length of the message -- not the length of the returned piece.

### status

Completion status, in STATUS\_\$T format. This data type is 4 bytes long. See the MBX Data Types section for more information.

**USAGE**

MBX\_\$GET\_CONDITIONAL gets a message if one is waiting. Otherwise, the call returns immediately with a completion status of MBX\_\$CHANNEL\_EMPTY. You can use an eventcount to tell when the status of the mailbox has changed. You get a mailbox eventcount with MBX\_\$GET\_EC.

Both servers and clients can use MBX\_\$GET\_CONDITIONAL. When a server calls MBX\_\$GET\_CONDITIONAL, the mailbox manager uses a scheduling algorithm to determine the channels to search for the next message. This algorithm guarantees fair service to each open channel.

MBX\_\$GET\_EC

MBX\_\$GET\_EC

Get pointer to eventcount for mailbox event mbx

## FORMAT

MBX\_\$GET\_EC (mbx-handle, mbx-key, eventcount-pointer, status)

## INPUT PARAMETERS

### mbx-handle

Identifier for the mailbox, in UNIV\_PTR format. This is a 4-byte integer. Use the handle returned by MBX\_\$CREATE\_SERVER or MBX\_\$OPEN.

### mbx-key

Type of eventcount to get a pointer to, in MBX\_\$EC\_KEY\_T format. This is a 2-byte integer. Specify one of these predefined values:

MBX\_\$GETREC\_EC\_KEY

An eventcount that advances when the mailbox may contain messages for you to get. For a server, this eventcount may advance whenever there is anything to get from any open channel.

MBX\_\$PUTREC\_EC\_KEY

An eventcount that advances when there may be enough room in the channel to hold the last message you unsuccessfully tried to put there. A mailbox server sees only one MBX\_\$PUTREC\_EC\_KEY eventcount for the entire mailbox. If puts fail with MBX\_\$NO\_ROOM\_IN\_CHANNEL on two channels of the same mailbox, the event's completion simply says that at least one channel may now take the message. One or both channels may now be capable of taking the respective message.

## OUTPUT PARAMETERS

### eventcount-pointer

A pointer to an eventcount, in EC2\_\$PTR\_T format. This is a 4-byte integer.

### status

Completion status, in STATUS\_\$T format. This data type is 4 bytes long. See the MBX Data Types section for more information.

## USAGE

After you use MBX\_\$GET\_EC to get a mailbox eventcount, use EC2 calls to read eventcount values and wait for events.

**MBX\_\$GET\_REC**

Gets a message from a mailbox.

**FORMAT**

MBX\_\$GET\_REC (handle, bufptr, buflen, retptr, retlen, status)

**INPUT PARAMETERS****handle**

Identifier for the mailbox, in UNIV\_PTR format. This is a 4-byte integer. Use the handle returned by MBX\_\$CREATE\_SERVER or MBX\_\$OPEN.

**bufptr**

Pointer to a data buffer where the message can be copied. This is a 4-byte integer.

Your program must allocate a data buffer, although the mailbox manager does not always copy messages to this buffer. Use the output parameter retptr to reference the message.

**buflen**

Number of bytes in the data buffer. This is a 4-byte integer. For a server, MBX will never return more than 32766 bytes. For a client, MBX will never return more than 32760 bytes.

**OUTPUT PARAMETERS****retptr**

Pointer to the buffer where the message is copied. This is a 4-byte integer.

**retlen**

Either the number of bytes in the returned message, or the number of message bytes waiting to be returned. This is a 4-byte integer.

MBX\_\$GET\_REC can get as many bytes as you specify in buflen. If the message is less than or equal to buflen, then the call gets the entire message and retlen specifies the message length. If the message is greater than buflen, then the call gets the number of bytes specified in buflen. If this occurs, then retlen contains a negative value, the absolute value of which is the number of bytes remaining in the message. Get the remaining data with another call.

Note that a server sees the message header each time it gets a piece of the message. The count field contains the total length of the message -- not the length of the returned piece.

**status**

Completion status, in STATUS\_\$T format. This data type is 4 bytes long. See the MBX Data Types section for more information.

MBX\_\$GET\_REC

## USAGE

MBX\_\$GET\_REC gets a mailbox message. If there is no message in the mailbox, the call waits for one.

Both servers and clients can use MBX\_\$GET\_REC. When a server calls MBX\_\$GET\_REC, the mailbox manager uses a scheduling algorithm to determine the channel to search for the next message. This algorithm guarantees fair service to each open channel.



**MBX\_\$GET\_REC\_CHAN**

Gets a mailbox message from a specified channel.

**FORMAT**

MBX\_\$GET\_REC\_CHAN (handle, channel, bufptr, buflen, retptr, retlen, status)

**INPUT PARAMETERS****handle**

Identifier for the mailbox, in UNIV\_PTR format. This is a 4-byte integer. Use the handle returned by MBX\_\$CREATE\_SERVER.

**channel**

Channel to read from. This is a 2-byte integer. The mailbox manager assigns a channel number to a client when the client calls MBX\_\$OPEN.

**bufptr**

Pointer to a data buffer where the message can be copied. This is a 4-byte integer.

Your program must allocate a data buffer, although the mailbox manager does not always copy messages to this buffer. Use the output parameter retptr to reference the message.

**buflen**

Number of bytes in the data buffer. This is a 4-byte integer. For a server, MBX will never return more than 32766 bytes. For a client, MBX will never return more than 32760 bytes.

**OUTPUT PARAMETERS****retptr**

Pointer to the buffer where the message is copied. This is a 4-byte integer.

**retlen**

Either the number of bytes in the returned message, or the number of message bytes waiting to be returned. This is a 4-byte integer.

MBX\_\$GET\_REC\_CHAN can get as many bytes as you specify in buflen. If the message is less than or equal to buflen, then the call gets the entire message and retlen specifies the message length. If the message is greater than buflen, then the call gets the number of bytes specified in buflen. If this occurs, then retlen contains a negative value, the absolute value of which is the number of bytes remaining in the message. Get the remaining data with another call.

Note that a server sees the message header each time it gets a piece of the message. The count field contains the total length of the message -- not the length of the returned piece.

**status**

Completion status, in STATUS\_\$T format. This data type is 4 bytes long. See the MBX Data Types section for more information.

MBX\_\$GET\_REC\_CHAN

## USAGE

MBX\_\$GET\_REC\_CHAN requests a message from the specified mailbox and channel. If there is no message, the call waits for one. Only a mailbox server can use this call. To perform a get operation from a client, use MBX\_\$GET\_REC.

**MBX\_\$GET\_REC\_CHAN\_SET**

Gets a mailbox message from a specified set of channels.

**FORMAT**

MBX\_\$GET\_REC\_CHAN\_SET (handle, chan-set, bufptr, buflen, retptr, retlen, status)

**INPUT PARAMETERS****handle**

Identifier for the mailbox, in UNIV\_PTR format. This is a 4-byte integer. Use the handle returned by MBX\_\$CREATE\_SERVER.

**chan-set**

Set of channels to read from, in MBX\_\$CHAN\_SET\_T format. This is an 8-element array of 4-byte integers. See the MBX Data Types section for more information.

The mailbox manager assigns a channel number to a client when the client calls MBX\_\$OPEN. The channel number can range from 1 through MBX\_\$CHN\_MAX.

**bufptr**

Pointer to a data buffer where the message can be copied. This is a 4-byte integer.

Your program must allocate a data buffer, although the mailbox manager does not always copy messages to this buffer. Use the output parameter retptr to reference the message.

**buflen**

Number of bytes in the data buffer. This is a 4-byte integer. For a server, MBX will never return more than 32766 bytes. For a client, MBX will never return more than 32760 bytes.

**OUTPUT PARAMETERS****retptr**

Pointer to the buffer where the message is copied. This is a 4-byte integer.

**retlen**

Either the number of bytes in the returned message, or the number of message bytes waiting to be returned. This is a 4-byte integer.

MBX\_\$GET\_REC\_CHAN\_SET can get as many bytes as you specify in buflen. If the message is less than or equal to buflen, then the call gets the entire message and retlen specifies the message length. If the message is greater than buflen, then the call gets the number of bytes specified in buflen. If this occurs, then retlen contains a negative value, the absolute value of which is the number of bytes remaining in the message. Get the remaining data with another call.

Note that a server sees the message header each time it gets a piece of the message. The count field contains the total length of the message -- not the length of the returned piece.

**status**

Completion status, in STATUS\_\$T format. This data type is 4 bytes long. See the MBX Data Types section for more information.

MBX\_\$GET\_REC\_CHAN\_SET

## USAGE

MBX\_\$GET\_REC\_CHAN\_SET requests a message from the specified mailbox and set of channels. If there is no message, the call waits for one. Only a mailbox server can use this call. To perform a get operation from a client, use MBX\_\$GET\_REC.

A group of calls is available for manipulating large sets. The calls are: LIB\_\$INIT\_SET, LIB\_\$ADD\_TO\_SET, LIB\_\$CLR\_FROM\_SET, and LIB\_\$MEMBER\_OF\_SET. The calls are fully described in *Programming with General System Calls*.

**MBX\_\$OPEN**

Opens a client channel to a mailbox.

**FORMAT**

MBX\_\$OPEN (name, namelen, bufptr, buflen, handle, status)

**INPUT PARAMETERS****name**

Name of the mailbox, in MBX\_\$NAME\_T format. This is an array of up to 256 characters. Specify the name as a pathname to the mailbox created by the server.

**namelen**

Number of characters in the name. This is a 2-byte integer.

**bufptr**

Pointer to a buffer containing data to be sent with the open request. This is a 4-byte integer. If you are not sending data, specify a nil pointer.

**buflen**

Number of bytes of data you are sending. This is a 4-byte integer. If you are not sending data, specify a length of 0.

The maximum amount of data you can send with MBX\_\$OPEN is MBX\_\$MSG\_MAX (1024) bytes, even if the mailbox message buffer is larger.

**OUTPUT PARAMETERS****handle**

Identifier for the mailbox, in UNIV\_PTR format. This is a 4-byte integer. Subsequent calls use this handle to send and receive messages.

**status**

Completion status, in STATUS\_\$T format. This data type is 4 bytes long. See the MBX Data Types section for more information.

**USAGE**

MBX\_\$OPEN opens a channel from a client to an existing mailbox. Only a client can use this call.

This call makes the mailbox manager send the server a channel open request. The server must respond by accepting or rejecting the request. After the server responds, MBX\_\$OPEN returns a status code indicating whether the call was successful. The client does not see the acceptance or rejection as a message, but as the completion status of the MBX\_\$OPEN call.

MBX\_\$PUT\_CHR

MBX\_\$PUT\_CHR

Sends a partial message from a client.

## FORMAT

MBX\_\$PUT\_CHR (handle, bufptr, buflen, status)

## INPUT PARAMETERS

### handle

Identifier for the mailbox, in UNIV\_PTR format. This is a 4-byte integer. Use the handle returned by MBX\_\$OPEN.

### bufptr

Pointer to the buffer that contains the message to be sent. This is a 4-byte integer.

### buflen

Length of the message, in bytes. This is a 4-byte integer. For a client, the buffer can contain up to 32760 bytes.

## OUTPUT PARAMETERS

### status

Completion status, in STATUS\_\$T format. This data type is 4 bytes long. See the MBX Data Types section for more information.

## USAGE

MBX\_\$PUT\_CHR is equivalent to MBX\_\$PUT\_REC, except that MBX\_\$PUT\_CHR informs the server that the message is a partial message. If the mailbox is full, this call waits until the mailbox has room for the message.

Only a client can call MBX\_\$PUT\_CHR. A server can send a partial data message by using MBX\_\$PUT\_REC or MBX\_\$PUT\_REC\_COND, and specifying a message type of MBX\_\$DATA\_PARTIAL\_MT. When the client gets such a message, the get call returns a status of MBX\_\$PARTIAL\_RECORD to the client.

**MBX\_\$PUT\_CHR\_COND**

Attempts to send a partial message from a client.

**FORMAT**

MBX\_\$PUT\_CHR\_COND (handle, bufptr, buflen, status)

**INPUT PARAMETERS****handle**

Identifier for the mailbox, in UNIV\_PTR format. This is a 4-byte integer. Use the handle returned by MBX\_\$OPEN.

**bufptr**

Pointer to the buffer that contains the message to be sent. This is a 4-byte integer.

**buflen**

Length of the message, in bytes. This is a 4-byte integer. For a client, the buffer can contain up to 32760 bytes.

**OUTPUT PARAMETERS****status**

Completion status, in STATUS\_\$T format. This data type is 4 bytes long. See the MBX Data Types section for more information.

**USAGE**

MBX\_\$PUT\_CHR\_COND is equivalent to MBX\_\$PUT\_REC\_COND, except that MBX\_\$PUT\_CHR\_COND informs the server that the message is a partial message.

If the client's buffer is full, MBX\_\$PUT\_CHR\_COND returns immediately with a completion status of MBX\_\$NO\_ROOM\_IN\_CHANNEL. You can use an eventcount to tell when the status of the mailbox eventcount has changed. You get a mailbox eventcount with MBX\_\$GET\_EC.

Only a client can call MBX\_\$PUT\_CHR\_COND. A server can send a partial data message by using MBX\_\$PUT\_REC or MBX\_\$PUT\_REC\_COND, and specifying a message type of MBX\_\$DATA\_PARTIAL\_MT. When the client gets such a message, the get call returns a status of MBX\_\$PARTIAL\_RECORD to the client.

MBX\_\$PUT\_REC

MBX\_\$PUT\_REC

Puts a record in the mailbox.

## FORMAT

MBX\_\$PUT\_REC (handle, bufptr, buflen, status)

## INPUT PARAMETERS

### handle

Identifier for the mailbox, in UNIV\_PTR format. This is a 4-byte integer. Use the handle returned by MBX\_\$CREATE\_SERVER or MBX\_\$OPEN.

### bufptr

Pointer to the buffer that contains the message to be sent. This is a 4-byte integer.

### buflen

Length of the message, in bytes. This is a 4-byte integer. For a server, the message can contain up to 32766 bytes. For a client, the buffer can contain up to 32760 bytes.

If a server puts a message that is larger than 1158 bytes, and the client is on a remote node, the client node's MBX\_HELPER must be able to handle the message. To handle the message, the client node's MBX\_HELPER must have a queue data size that is at least as large as the message. Use MBX\_\$SERVER\_WINDOW to determine the client node's queue data size.

## OUTPUT PARAMETERS

### status

Completion status, in STATUS\_\$T format. This data type is 4 bytes long. See the MBX Data Types section for more information.

## USAGE

This call can be used by either servers or clients. Note, however, that servers and clients have different message formats. A server must include the 6-byte message header when sending a message. In contrast, a client sends only data.

If the channel is full, this call waits until there is room for the message.



**MBX\_\$PUT\_REC\_COND**

Attempts to put a message into a mailbox.

**FORMAT**

MBX\_\$PUT\_REC\_COND (handle, bufptr, buflen, status)

**INPUT PARAMETERS****handle**

Identifier for the mailbox, in UNIV\_PTR format. This is a 4-byte integer. Use the handle returned by MBX\_\$CREATE\_SERVER or MBX\_\$OPEN.

**bufptr**

Pointer to the buffer that contains the message to be sent. This is a 4-byte integer.

**buflen**

Length of the message, in bytes. This is a 4-byte integer. For a server, the message can contain up to 32766 bytes. For a client, the buffer can contain up to 32760 bytes.

If a server puts a message that is larger than 1158 bytes, and the client is on a remote node, the client node's MBX\_HELPER must be able to handle the message. To handle the message, the client node's MBX\_HELPER must have a queue data size that is at least as large as the message. Use MBX\_\$SERVER\_WINDOW to determine the client node's queue data size.

**OUTPUT PARAMETERS****status**

Completion status, in STATUS\_\$T format. This data type is 4 bytes long. See the MBX Data Types section for more information.

**USAGE**

MBX\_\$PUT\_REC\_COND can be used by either servers or clients. Note, however, that servers and clients have different message formats. A server must include the 6-byte message header when sending a message. In contrast, a client sends only data.

If the channel is full, MBX\_\$PUT\_REC\_COND returns immediately, with a completion status of MBX\_\$NO\_ROOM\_IN\_CHANNEL. You can use an eventcount to tell when the status of the mailbox eventcount has changed. You get a mailbox eventcount with MBX\_\$GET\_EC.

MBX\_\$SERVER\_WINDOW

MBX\_\$SERVER\_WINDOW

Returns the buffer size for the mailbox maintained by the MBX\_HELPER on a remote client's node.

## FORMAT

size = MBX\_\$SERVER\_WINDOW (handle, channel, status)

## RETURN VALUE

### size

Buffer size for the mailbox maintained by the MBX\_HELPER on the remote client's node. This is a 4-byte integer.

This value defines a window size when a server sends messages to a remote client. That is, the server cannot send messages that are larger than the buffer for the remote MBX\_HELPER's mailbox.

## INPUT PARAMETERS

### handle

Identifier for the mailbox, in UNIV\_PTR format. This is a 4-byte integer. Use the handle returned by MBX\_CREATE\_SERVER.

### channel

Channel belonging to the client whose window size you are inquiring about. This is a 2-byte integer. The mailbox manager assigns a channel number to a client when the client calls MBX\_OPEN.

## OUTPUT PARAMETERS

### status

Completion status, in STATUS\_\$T format. This data type is 4 bytes long. See the MBX Data Types section for more information.

## USAGE

When a server puts a message into a mailbox and the message is intended for a client on a remote node, the message must pass through a system mailbox maintained by the client node's MBX\_HELPER. Thus, the largest message that a server can send depends on the buffer size for the remote client's system mailbox. This buffer size was defined when the client node's MBX\_HELPER was started. (The MBX\_HELPER's -DATASIZE option defines a buffer size for the system mailbox.)

A server can use MBX\_\$SERVER\_WINDOW to determine the buffer size for the remote client's system mailbox. To get this value, MBX\_\$SERVER\_WINDOW queries the MBX\_HELPER on the client's node. Note that MBX\_\$SERVER\_WINDOW returns the actual buffer size, not the number of unused bytes in the buffer.

Note that if a server is communicating with clients on different nodes, the buffer size can differ on each node. Therefore, the server must use MBX\_\$\$SERVER\_WINDOW to obtain the buffer size on each node.

MBX\_\$\$SERVER\_WINDOW correctly returns the buffer size for clients on nodes with SR9 or later software. However, if you call MBX\_\$\$SERVER\_WINDOW and the specified client is on a pre-SR9 node, then the call always returns the value 1158. 1158 is the minimum buffer size for mailboxes maintained by pre-SR9 MBX\_HELPERS.

MBX\_\$\$SERVER\_WINDOW is for use only by mailbox servers. A client should use MBX\_\$\$CLIENT\_WINDOW.

MBX\_\$TIMED\_OPEN

MBX\_\$TIMED\_OPEN

Attempts to open a client channel to a mailbox within a specified time period.

## FORMAT

MBX\_\$TIMED\_OPEN (name, namelen, bufptr, buflen, wait-time, handle, status)

## INPUT PARAMETERS

### name

Name of the mailbox, in MBX\_\$NAME\_T format. This is an array of up to 256 characters. Specify the name as a pathname to the mailbox created by the server.

### namelen

Number of characters in the name. This is a 2-byte integer.

### bufptr

Pointer to a buffer containing data to be sent with the open request. This is a 4-byte integer. If you are not sending data, specify a nil pointer.

### buflen

Number of bytes of data you are sending. This is a 4-byte integer. If you are not sending data, specify a length of 0.

The maximum amount of data you can send with MBX\_\$TIMED\_OPEN is MBX\_\$MSG\_MAX (1024) bytes, even if the mailbox message buffer is larger.

### wait-time

Time to wait for an open reply, in TIME\_\$CLOCKH\_T format. If you specify a value of 0, the call will wait forever for an open reply, just like an MBX\_\$OPEN call.

## OUTPUT PARAMETERS

### handle

Identifier for the mailbox, in UNIV\_PTR format. This is a 4-byte integer. Subsequent calls use this handle to send and receive messages.

### status

Completion status, in STATUS\_\$T format. This data type is 4 bytes long. See the MBX Data Types section for more information.

## USAGE

MBX\_\$TIMED\_OPEN attempts to open a channel from a client to an existing mailbox within a specified time period. If the client does not receive a server response within the wait-time period, the open fails and the client receives the MBX\_\$OPEN\_TIMED\_OUT status message. If the server responds after a client open has timed out, it receives the MBX\_\$CHANNEL\_NOT\_OPEN status message. Only a client can use this call.

No matter what value you specify for wait-time, if no server process is associated with the mailbox, the call will fail and return immediately.

**ERRORS****MBX\_\$BAD\_KEY**

Bad key.

**MBX\_\$BUFFER\_TOO\_SMALL**

A server requested a message using a buffer smaller than 6 bytes. There must be enough room for the message header in all server message requests.

**MBX\_\$CHANNEL\_EMPTY**

There are no messages waiting in the channel. Received in response to an MBX\_\$GET\_COND or MBX\_\$COND\_GET\_CHAN request.

**MBX\_\$CHANNEL\_NOT\_OPEN**

For a server, the channel number given referred to a channel that is not presently open; for a client, the server has deallocated the client's channel.

**MBX\_\$CLIENT\_NO\_RIGHTS**

The client can't access the local MBX\_HELPER's SYSMBX.

**MBX\_\$CLIENT\_SERVER\_DEADLOCK**

A server tried to open a channel to itself; this is illegal.

**MBX\_\$EOF**

The client has sent a message of type MBX\_\$EOF\_MT. Received in response to an MBX\_\$GET\_REC or MBX\_\$GET\_COND request.

**MBX\_\$FILE\_IN\_USE**

An MBX\_\$CREATE\_SERVER or MBX\_\$OPEN request was made giving a mailbox pathname that is the pathname for a file presently in use.

**MBX\_\$HANDLE\_NOT\_VALID**

The handle given does not point to a mailbox.

**MBX\_\$HELPER\_NO\_RIGHTS**

The MBX\_HELPER on the server's node can't access the server's mailbox.

**MBX\_\$ILL\_HANDLE**

The handle given is not a legal handle.

**MBX\_\$MSG\_TOO\_BIG\_FOR\_CHANNEL**

An MBX\_\$PUT\_CONDITIONAL or MBX\_\$PUT\_REC request tried to send a message bigger than the maximum specified when the server created the mailbox.

**MBX\_\$NO\_MORE\_CHANNELS**

An MBX\_\$OPEN was made to a mailbox with no free channels.

**MBX\_\$NO\_MORE\_RESOURCES**

An MBX\_\$CREATE\_SERVER or MBX\_\$OPEN request was made, and the process has insufficient resources left to open the mailbox or the channel.

**MBX\_\$NO\_ROOM\_IN\_CHANNEL**

There is not enough room in the channel for the message. Received in response to an MBX\_\$PUT\_CONDITIONAL request.

**MBX\_\$NO\_SERVERS**

An MBX\_\$OPEN was made to a mailbox without an active server.

MBX ERRORS

MBX\_\$OPEN\_REJECTED

The server rejected an MBX\_\$OPEN request.

MBX\_\$OPEN\_TIMED\_OUT

The server did not respond within the specified period of time.

MBX\_\$PARTIAL\_RECORD

Returned data does not contain a complete record.

MBX\_\$RECORD\_LENGTH\_MISMATCH

A server used one of the MBX\_\$PUT calls, and the value in the buflen parameter did not match the value in the length field of the message header.

MBX\_\$REM\_RCV\_TIMEOUT

A remote operation was attempted, and the network has failed.

MBX\_\$REM\_SEND\_FAILED

A remote operation was attempted, and the network has failed.

MBX\_\$REM\_SERVICE\_UNAVAILABLE

An MBX\_\$OPEN open request was made from a remote node when the MBX\_HELPER program was not running on that node or the server's node.

MBX\_\$REMOTE\_SERVICE\_DENIED

An MBX\_\$OPEN request was made from a remote node, and there are not enough network services free to handle the request.

MBX\_\$SEQUENCED\_SEND\_FAILED

An internal error occurred while sending a message that is larger than 1158 bytes.

MBX\_\$SIZE\_TOO\_LARGE

MBX\_\$CREATE\_SERVER request asked for a mailbox larger than the maximum.

MBX\_\$SIZE\_TOO\_SMALL

An MBX\_\$CREATE\_SERVER request was made with a buffer size smaller than the minimum.

MBX\_\$TOO\_MANY\_CHANNELS

An MBX\_\$CREATE\_SERVER request was made asking for more than the maximum number of channels.

MBX\_\$UNEXPECTED\_CNTL\_MSG

Received by a client when the last message the server sent on that channel had a message type of MBX\_\$ACCEPT\_OPEN\_MT, MBX\_\$REJECT\_OPEN\_MT, or MBX\_\$CHANNEL\_OPEN\_MT when such a message type was inappropriate. (MBX\_\$CHANNEL\_OPEN\_MT should never be used. The other two message types are only used in response to a message of type MBX\_\$CHANNEL\_OPEN\_MT.) Received in response to an MBX\_\$GET\_REC or MBX\_\$GET\_COND request.

MBX\_\$UNKNOWN\_RQST

The client and server are using different versions of the mailbox manager (although the two versions have the same version number), and one of them made a request not recognized by the other manager.

MBX\_\$WRONG\_VERSION\_NUMBER

An MBX\_\$OPEN request was made using a mailbox manager with a different version number than the one used to create the mailbox.

STATUS\_ \$OK  
Successful completion.







# MS

## Table of Contents

Data Types	MS-2
MS_\$ADDMAP	MS-6
MS_\$ADVICE	MS-8
MS_\$ATTRIBUTES	MS-10
MS_\$CRMPL	MS-11
MS_\$CRTEMP	MS-13
MS_\$FW_FILE	MS-15
MS_\$FW_PARTIAL	MS-16
MS_\$MAPL	MS-17
MS_\$MAPL_STREAM	MS-20
MS_\$MK_PERMANENT	MS-24
MS_\$MK_TEMPORARY	MS-25
MS_\$NEIGHBORS	MS-26
MS_\$RELOCK	MS-27
MS_\$REMAP	MS-29
MS_\$TRUNCATE	MS-31
MS_\$UNMAP	MS-32
Errors	MS-33



## MS

The MS (Mapped Segment) programming calls maps to and unmaps from process address space. This section describes their data types, call syntax, and error codes. Refer to the Introduction at the beginning of this manual for a description of data-type diagrams and call syntax format.

MS DATA TYPES

CONSTANTS

MS_\$EXTEND	TRUE	The object can be extended.
MS_\$NO_EXTEND	FALSE	The object cannot be extended.

DATA TYPES

MS\_\$ACC\_MODE\_T A 2-byte integer. Access mode for an object. One of the following predefined values:

MS\_\$R  
Read access.

MS\_\$RX  
Read and execute access.

MS\_\$WR  
Read and write access.

MS\_\$WRX  
Write and execute access.

MS\_\$RIW  
Read with intent to write.

MS\_\$ACCESS\_T A 2-byte integer. Usage patterns for accessing a file. One of the following predefined values:

MS\_\$NORMAL  
Normal use.

MS\_\$RANDOM  
Random access use.

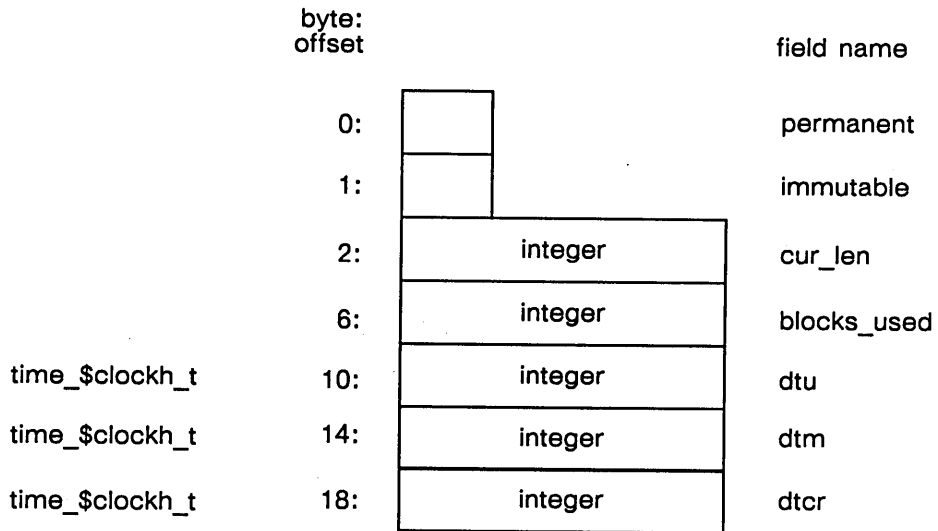
MS\_\$SEQUENTIAL  
Sequential access use.

MS\_\$ADVICE\_OPT\_T Reserved for future use.

MS\_\$ADVICE\_T Four bytes that are reserved for future use.

MS\_\$ATTRIB\_T

An attribute record. The diagram below illustrates the MS\_\$ATTRIB\_T data type:



permanent

A boolean value that indicates whether the object is permanent (TRUE) or temporary (FALSE)

immutable

A boolean value that indicates whether the object can be modified. The value TRUE means that the object is immutable. The value FALSE means that the object is not immutable and can therefore be modified.

cur\_len

Current length, in bytes, of the object.

blocks\_used

The number of blocks used for the object.

dtu

Date-time used, in TIME\_\$CLOCKH\_T format.

dtm

Date-time modified, in TIME\_\$CLOCKH\_T format.

dtr

Date-time created, in TIME\_\$CLOCKH\_T format.

MS DATA TYPES

MS\_\$CONC\_MODE\_T

A 2-byte integer. Concurrency mode for an object. One of the following predefined values:

MS\_\$NR\_XOR\_1W

Allows one writer or any number of readers.

MS\_\$COWRITERS

Allows any number of readers and/or writers.

MS\_\$PERM\_OPT\_T


A 2-byte integer. Available options for MS\_\$MK\_PERMANENT. One of the following predefined values:

MS\_\$MK\_BAK

Makes a backup copy of file.

STATUS\_\$T

A status code. The diagram below illustrates the STATUS\_\$T data type:



Field Description:

all

All 32 bits in the status code.

fail

The fail bit. If this bit is set, the error was not within the scope of the module invoked, but occurred within a lower-level module (bit 31).

subsys

The subsystem that encountered the error (bits 24 - 30).

modc

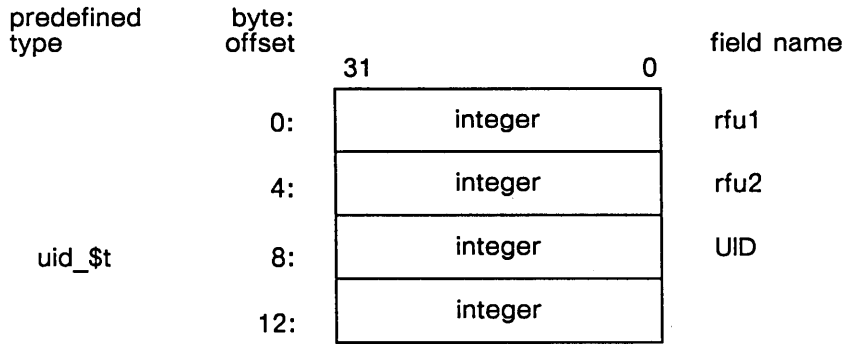
The module that encountered the error (bits 16 - 23).

code

A signed number that identifies the type of error that occurred (bits 0 - 15).

XOID\_\$T

Unique identifier of an object. Used by type managers only. The diagram below illustrates the XOID\_\$T data type:



Field Description:

rfu1  
Reserved for future use.

rfu2  
Reserved for future use.

UID  
Unique identifier for an object.

MS\_\$ADDMAP

MS\_\$ADDMAP

Adds a new map of a different portion of previously mapped object.

### FORMAT

address = MS\_\$ADDMAP (old-address, start, desired-length,  
mapped-length, status)

### RETURN VALUE

#### address

Pointer to the first byte of additional mapped portion, in UNIV\_PTR format. This is a 4-byte integer.

### INPUT PARAMETERS

#### old-address

Pointer to the first byte of the currently mapped portion of the object, in UNIV\_PTR format. This is a 4-byte integer. Use the pointer returned by the most recent call to MS\_\$MAPL, MS\_\$CRMAPL, or MS\_\$REMAP.

#### start

First byte to be mapped. This is a 4-byte integer.

#### desired-length

Number of bytes to map. This is a 4-byte integer.

### OUTPUT PARAMETERS

#### mapped-length

Number of bytes mapped. This is a 4-byte integer.

#### status

Completion status, in STATUS\_\$T format. This data type is 4 bytes long. See the MS Data Types section for more information. Possible values are:

STATUS\_\$OK Completed successfully.

MS\_\$NOT\_MAPPED

No object is mapped at the given virtual address.

MS\_\$BAD\_LENGTH

Desired-length is invalid.



**USAGE**

This call maps a different portion of a previously mapped object. It is unlike MS\_\$REMAP in that it does not move the existing window to the object but adds a new window at a different address from the original mapping. The new window is not contiguous with the old window.

When you add a new map of a file, certain attributes of the mapping (extend, access, concurrency) are left the same as in the original mapping. If you used MS\_\$ADVICE to provide file access advice, the advice in effect for the first part of the currently mapped section is propagated to the newly mapped section. Also, MS\_\$ADDMAP does not change the lock mode of the object.

The locking mode is maintained on the previously mapped object. If you try to unmap the previous window and the object is temporary, you will get an "object not found" error message; if the object is permanent, you will run into concurrency problems because the object will become unlocked. You should unmap the second window before attempting to unmap the previous window.

MS\_\$ADVICE

MS\_\$ADVICE

Provides the operating system with information on how you plan to access an object. This information helps the system optimize performance when managing the object.

## FORMAT

MS\_\$ADVICE (address, length, access, options, record-length, status)

## INPUT PARAMETERS

### address

Pointer to the first byte to provide advice for, UNIV\_PTR format. This is a 4-byte integer. Use the pointer returned by the most recent call to MS\_\$CRMPL, MS\_\$MAPL, or MS\_\$REMAP.

### length

Number of bytes to provide advice for. This is a 4-byte integer.

### access

Method of accessing the object, in MS\_\$ACCESS\_T format. Specify only one of the following predefined values:

MS\_\$NORMAL You do not have a predicted manner for accessing the object. This is the default if a program never uses MS\_\$ADVICE.

MS\_\$RANDOM You access the object randomly.

MS\_\$SEQUENTIAL  
You access the object sequentially.

### options

Reserved for future use, in MS\_\$ADVICE\_T format. This is a 4-byte integer. In Pascal, specify this parameter using the empty set []. In C and FORTRAN, declare a variable and initialize it to 0.

### record-length

Number of bytes in a record in the mapped object. This is a 4-byte integer. If you do not know the record length, or if the object is not record-structured, specify 0.

## OUTPUT PARAMETERS

### status

Completion status, in STATUS\_\$T format. This data type is 4 bytes long. See the MS Data Types section for more information.

## USAGE

MS\_\$ADVICE provides the operating system with information on how you plan to access an object. When you work with a mapped object, the system brings pages into memory as needed. By using MS\_\$ADVICE, you can change the number of pages that the system gets when a page fault occurs. This helps the system provide better performance when managing the object on your behalf.

Although it is not required that you use MS\_\$ADVICE, you should use it whenever you have a predicted type of file access. In addition, you can use MS\_\$ADVICE more than once to change the advice for a mapped object.

If you remap an object with MS\_\$REMAP, the advice in effect for the first part of the currently mapped section is propagated to the newly mapped section.

MS\_\$ATTRIBUTES

MS\_\$ATTRIBUTES

Returns the selected attributes of a mapped object.

#### FORMAT

MS\_\$ATTRIBUTES (address, attrib-buf, attrib-len, attrib-max, status)

#### INPUT PARAMETERS

##### address

Pointer to the first byte of the currently mapped portion of the object, in UNIV\_PTR format. This is a 4-byte integer. Use the pointer returned by the most recent call to MS\_\$MAPL, MS\_\$CRMAPL, or MS\_\$REMAP.

#### OUTPUT PARAMETERS

##### attrib-buf

Buffer in which to receive the attributes, in MS\_\$ATTRIB\_T format. This data type is 22 bytes long. See the MS Data Types section for more information.

##### attrib-len

Length of the attributes returned in the attributes buffer. This is a 2-byte integer.

#### INPUT PARAMETERS

##### attrib-max

Length of the attributes buffer. This is a 2-byte integer. Specify the length of the attributes buffer in the **attrib\_buf** parameter. This value defines the maximum amount of information that MS\_\$ATTRIBUTES can return.

#### OUTPUT PARAMETERS

##### status

Completion status, in STATUS\_\$T format. This data type is 4 bytes long. See the MS Data Types section for more information.

#### USAGE

Use MS\_\$ATTRIBUTES to get information about selected attributes of a mapped object.

**MS\_\$CRMPL**

Creates, maps, and locks an object.

**FORMAT**

**address** = MS\_\$CRMPL (name, name-length, start, desired-length,  
concurrency, status)

**RETURN VALUE****address**

Pointer to the first mapped byte of the object, in UNIV\_PTR format. This is a 4-byte integer.

The first mapped byte is not necessarily the first byte of the object; it is the byte you specify in the start parameter.

**INPUT PARAMETERS****name**

Pathname of the object to be mapped, in NAME\_\$PNAME\_T format. This is an array of up to 256 characters.

**name-length**

Length of the pathname. This is a 2-byte integer.

**start**

First byte to be mapped. This is a 4-byte positive integer. To specify the first byte in an object, provide a start value of 0.

**desired-length**

Number of bytes to map, including the start byte. This is a 4-byte positive integer.

**concurrency**

Concurrency mode for the object, in MS\_\$CONC\_MODE\_T format. This is a 2-byte integer. Specify only one of the following predefined values:

MS\_\$NR\_XOR\_1W

Allows one writer or any number of readers.

MS\_\$COWRITERS

Allows any number of readers and/or writers.

MS\_\$CRMAPL

## OUTPUT PARAMETERS

### status

Completion status, in STATUS\_\$T format. This data type is 4 bytes long. See the MS Data Types section for more information. Possible values are:

STATUS_\$OK	Object created.
MS_\$BAD_ACCESS	Illegal concurrency value.
MS_\$IN_USE	Object is currently locked.
MS_\$NO_SPACE	Insufficient virtual address space to map.
NAME_\$ALREADY_EXISTS	Name given already exists.
Other naming server errors	See the NAME_\$ error codes.

## USAGE

MS\_\$CRMAPL creates a file only if the name you specify does not already exist. The call implicitly uses an MS manager access mode of MS\_\$WR. Thus the object is always mapped for write access. You can get an exclusive write lock (if you specify a concurrency of MS\_\$WR\_XOR\_1W) or you can get a shared write lock (if you specify a concurrency of MS\_\$COWRITERS.) See the description of MS\_\$MAPL for more information on locks.

MS\_\$CRMAPL always uses an extend value of TRUE. Thus you can extend the object to the length you specify in the desired-length parameter.

**MS\_\$CRTEMP**

Creates, maps, and locks an object.

**FORMAT**

address = MS\_\$CRTEMP (volume-name, name-length, start, desired-length, concurrency, status)

**RETURN VALUE****address**

Pointer to the first mapped byte of the temporary object, in UNIV\_PTR format. This is a 4-byte integer.

The first mapped byte is not necessarily the first byte of the object; it is the byte you specify in the start parameter.

**INPUT PARAMETERS****vol-name**

Pathname of the volume where the object is to be mapped, in NAME\_\$PNAME\_T format. This is an array of up to 256 characters.

**name-length**

Length of the pathname. This is a 2-byte integer.

**start**

First byte to be mapped. This is a 4-byte positive integer. To specify the first byte in an object, provide a start value of 0.

**desired-length**

Number of bytes to map, including the start byte. This is a 4-byte positive integer.

**concurrency**

Concurrency mode for the object, in MS\_\$CONC\_MODE\_T format. This is a 2-byte integer. Specify only one of the following predefined values:

MS\_\$NR\_XOR\_1W

Allows one writer or any number of readers.

MS\_\$COWRITERS

Allows any number of readers and/or writers.

## OUTPUT PARAMETERS

### status

Completion status, in STATUS\_ \$T format. This data type is 4 bytes long. See the MS Data Types section for more information. Possible values are:

STATUS_ \$OK	Object created.
MS_ \$BAD_ ACCESS	Illegal concurrency value.
MS_ \$IN_ USE	Object is currently locked.
MS_ \$NO_ SPACE	Insufficient virtual address space to map.
Other naming server errors	See the NAME_ \$ error codes.

## USAGE

MS\_ \$CRTEMP creates a temporary file that will reside on the volume you specify and returns the mapped address. The call implicitly uses an MS manager access mode of MS\_ \$WR. Thus the object is always mapped for write access. You can get an exclusive write lock (if you specify a concurrency of MS\_ \$WR\_ XOR\_ 1W) or you can get a shared write lock (if you specify a concurrency of MS\_ \$COWRITERS.) See the description of MS\_ \$MAPL for more information on locks.

MS\_ \$CRTEMP always uses an extend value of TRUE. Thus you can extend the object to the length you specify in the desired-length parameter.

Temporary files are unlike permanent files in that they do not have pathnames. Also, they disappear when you shut down the system or call MS\_ \$UNMAP. To make a temporary file permanent, you should use MS\_ \$MK\_ PERMANENT.



**MS\_\$FW\_FILE**

Forces the system to write a mapped file onto disk.

**FORMAT**

MS\_\$FW\_FILE (address, status)

**INPUT PARAMETERS****address**

Pointer to the first byte of the currently mapped portion of the object, in UNIV\_PTR format. This is a 4-byte integer. Use the pointer returned by the most recent call to MS\_\$MAPL, MS\_\$CRMPL, or MS\_\$REMAP.

**OUTPUT PARAMETERS****status**

Completion status, in STATUS\_\$T format. This data type is 4 bytes long. See the MS Data Types section for more information.

**USAGE**

When you work with mapped objects, the system uses a predefined set of conditions to determine when to write information (stored in memory) onto the disk. However, if you need to supplement the system's actions, you can use MS\_\$FW\_FILE to force the system to write an object onto disk.

When you use MS\_\$FW\_FILE, the system force writes the entire object, even if the currently mapped portion does not begin at byte 0. However, the system writes only the changed portions of the object onto the disk.

When you force-write a permanent object, the system also force-writes the directory where the object is cataloged.

MS\_\$FW\_PARTIAL

MS\_\$FW\_PARTIAL

Forces the system to write a portion of a mapped file onto disk.

## FORMAT

MS\_\$FW\_FILE (address, deired-length, status)

## INPUT PARAMETERS

### address

Pointer to the first byte of the currently mapped portion of the object, in UNIV\_PTR format. This is a 4-byte integer. Use the pointer returned by the most recent call to MS\_\$MAPL, MS\_\$CRMAPL, or MS\_\$REMAP.

### desired-length

Number of bytes to force write, including the start byte. This is a 4-byte positive integer.

## OUTPUT PARAMETERS

### status

Completion status, in STATUS\_\$T format. This data type is 4 bytes long. See the MS Data Types section for more information.

## USAGE

When you work with mapped objects, the system uses a predefined set of conditions to determine when to write information (stored in memory) onto the disk. However, if you need to supplement the system's actions, you can use MS\_\$FW\_PARTIAL to force the system to write part of an object onto disk. This call is useful if you do not wish to write out the entire database to disk but just a specified portion. If you wish to force write the entire object, use MS\_\$FW\_FILE.

**MS\_\$MAPL**

Maps the specified portion of a file-system object into an available region of the process address space. This call also locks the object.

**FORMAT**

address = MS\_\$MAPL (name, name-length, start, desired-length, concurrency,  
access, extend, length-mapped, status)

**RETURN VALUE****address**

Pointer to the first mapped byte of the object, in UNIV\_PTR format. This is a 4-byte integer.

The first mapped byte is not necessarily the first byte of the object; it is the byte you specify in the start parameter.

**INPUT PARAMETERS****name**

Pathname of the object to be mapped, in NAME\_\$PNAME\_T format. This is an array of up to 256 characters.

**name-length**

Length of the pathname. This is a 2-byte integer.

**start**

First byte to be mapped. This is a 4-byte positive integer. To specify the first byte of an object, provide a start value of 0.

**desired-length**

Number of bytes to map, including the start byte. This is a 4-byte positive integer.

**concurrency**

Concurrency mode for the object, in MS\_\$CONC\_MODE\_T format. This is a 2-byte integer. Specify only one of the following predefined values:

MS\_\$NR\_XOR\_1W

Allows one writer or any number of readers.

MS\_\$COWRITERS

Allows any number of readers and/or writers.

MS\_\$MAPL

**access**

The access mode desired, in MS\_\$ACC\_MODE\_T format. This is a 2-byte integer. Specify only one of the following predefined values:

MS_\$R	Read access.
MS_\$RX	Read and execute access.
MS_\$WR	Read and write access.
MS_\$WRX	Write and execute access.
MS_\$RIW	Read with intent to write.

The access requested must be a subset of the access permitted by the protection for the object.

**extend**

A Boolean value that indicates whether the object can be extended. The value TRUE indicates that the length given in the desired-length parameter should be mapped, even if the object is shorter. Writing beyond the end of the object, but within the space mapped, extends the object. FALSE indicates that the amount mapped should be no greater than the actual length of the file.

**OUTPUT PARAMETERS**

**length-mapped**

Number of bytes actually mapped. This is a 4-byte positive integer.

**status**

Completion status, in STATUS\_\$T format. This data type is 4 bytes long. See the MS Data Types section for more information. Possible values are:

STATUS_\$OK	Object created.
MS_\$BAD_ACCESS	Access type is illegal.
MS_\$IN_USE	Object could not be locked.
NAME_\$NOT_FOUND	No object exists with the given name.

Other naming-server errors  
See the NAME\_\$ error codes.

**USAGE**

Use MS\_\$MAPL to map files that contain data in a user-defined format. For example, font files are a case where MS\_\$MAPL is appropriate. Do not use MS\_\$MAPL to access DOMAIN record structured files; use STREAM\_\$ calls to access these files.

MS\_\$MAPL locks a file, in addition to mapping it. The lock is determined by the concurrency and access modes that you specify. MS\_\$MAPL can obtain the following types of locks:

<i>Lock</i>	<i>Concurrency Mode</i>	<i>Access Mode</i>
Protected Read	MS_\$NR_XOR_1W	MS_\$R or MS_\$RX
Protected RIW	MS_\$NR_XOR_1W	MS_\$RIW
Shared Read	MS_\$COWRITERS	MS_\$R or MS_\$RX or MS_\$RIW
Exclusive Write	MS_\$NR_XOR_1W	MS_\$WR or MS_\$WRX
Shared Write	MS_\$COWRITERS	MS_\$WR or MS

Once you have locked a file, the MS manager allows other processes to map the file only if these processes request a lock that is compatible with your lock. The following lock combinations are allowed and prohibited. (Y means that the combination is allowed; N means that the combination is prohibited.)

<i>Existing Lock</i>	<i>Requested Lock</i>				
	<i>Protected Read</i>	<i>Protected RIW</i>	<i>Shared Read</i>	<i>Exclusive Write</i>	<i>Shared Write</i>
Protected Read	Y	Y	Y	N	N
Protected RIW	Y	N	Y	N	N
Shared Read	Y	Y	Y	N	Y *
Exclusive Write	N	N	N	N	N
Shared Write	N	N	Y *	N	Y *

\* These locks are allowed only if the processes are on the same node.

MS\_\$MAPL\_STREAM

MS\_\$MAPL\_STREAM

Maps the specified filesystem object, given its xoid, into an available region of the process address space. This call also locks the object and protects the mapping on a UNIX EXEC call. For type managers only.

#### FORMAT

address-ptr := MS\_\$MAPL\_STREAM (xoid, start, desired-length, concurrency,  
access, extend, length-mapped, status)

#### RETURN VALUE

##### address-ptr

Pointer to the first mapped byte of the object, in UNIV\_PTR format. This is a 4-byte integer.

The first mapped byte is not necessarily the first byte of the object; it is the byte you specify in the start parameter.

#### INPUT PARAMETERS

##### xoid

Xoid, or unique identifier of an object in XOID\_\$T format. This data type is 16-bytes long. See the MS Data Types section for details.

##### start

First byte to be mapped. This is a 4-byte positive integer. To specify the first byte of an object, provide a start value of 0.

##### desired-length

Number of bytes to map, including the start byte. This is a 4-byte positive integer.

**concurrency**

Concurrency mode for the object, in MS\_\$CONC\_MODE\_T format. This is a 2-byte integer. Specify only one of the following predefined values:

MS\_\$NR\_XOR\_1W

Allows one writer or any number of readers.

MS\_\$COWRITERS

Allows any number of readers and/or writers.

**access**

The access mode desired, in MS\_\$ACC\_MODE\_T format. This is a 2-byte integer. Specify only one of the following predefined values:

MS\_\$R                      Read access.

MS\_\$RX                     Read and execute access.

MS\_\$WR                    Read and write access.

MS\_\$WRX                   Write and execute access.

MS\_\$RIW                   Read with intent to write.

The access requested must be a subset of the access permitted by the protection for the object.

**extend**

A Boolean value that indicates whether the object can be extended. The value TRUE indicates that the length given in the desired-length parameter should be mapped, even if the object is shorter. Writing beyond the end of the object, but within the space mapped, extends the object. FALSE indicates that the amount mapped should be no greater than the actual length of the file.

**OUTPUT PARAMETERS****length-mapped**

Number of bytes actually mapped. This is a 4-byte positive integer.

## MS\_\$MAPL\_STREAM

### status

Completion status, in STATUS\_\$T format. This data type is 4 bytes long. See the MS Data Types section for more information. Possible values are:

STATUS\_\$OK Object created.

MS\_\$BAD\_ACCESS  
Access type is illegal.

MS\_\$IN\_USE Object could not be locked.

### USAGE

Use MS\_\$MAPL\_STREAM to map objects that you access through a type manager. Note that you use MS\_\$MAPL\_STREAM *only* through a type manager. For details, see the *Extending the DOMAIN Streams Facility* manual.

MS\_\$MAPL\_STREAM protects the mapping on a UNIX EXEC call. Conversely, with MS\_\$MAPL, the UNIX EXEC call unmaps any objects on any open streams.

MS\_\$MAPL\_STREAM also locks the object. The lock is determined by the concurrency and access modes that you specify. MS\_\$MAPL\_STREAM can obtain the following types of locks:

<i>Lock</i>	<i>Concurrency Mode</i>	<i>Access Mode</i>
Protected Read	MS_\$NR_XOR_1W	MS_\$R or MS_\$RX
Protected RIW	MS_\$NR_XOR_1W	MS_\$RIW
Shared Read	MS_\$COWRITERS	MS_\$R or MS_\$RX or MS_\$RIW
Exclusive Write	MS_\$NR_XOR_1W	MS_\$WR or MS_\$WRX
Shared Write	MS_\$COWRITERS	MS_\$WR or MS



Once you have locked an object, the MS manager allows other processes to map the object only if these processes request a lock that is compatible with your lock. The following lock combinations are allowed and prohibited. (Y means that the combination is allowed; N means that the combination is prohibited.)

<i>Existing Lock</i>	<i>Requested Lock</i>				
	<i>Protected Read</i>	<i>Protected RIW</i>	<i>Shared Read</i>	<i>Exclusive Write</i>	<i>Shared Write</i>
Protected Read	Y	Y	Y	N	N
Protected RIW	Y	N	Y	N	N
Shared Read	Y	Y	Y	N	Y *
Exclusive Write	N	N	N	N	N
Shared Write	N	N	Y *	N	Y *

\* These locks are allowed only if the processes are on the same node.

MS\_\$MK\_PERMANENT

MS\_\$MK\_PERMANENT

Makes a temporary file permanent and names it.

## FORMAT

MS\_\$MK\_PERMANENT (address, option, name, name-length, status)

## INPUT PARAMETERS

### address

Pointer to the first byte of the currently mapped portion of the temporary object, in UNIV\_PTR format. This is a 4-byte integer. Use the pointer returned by the most recent call to MS\_\$CRTEMP or MS\_\$CRMAPL.

### option

The attribute that is to be set, in MS\_\$PERM\_OPT\_T format. This is a 2-byte integer. Specify only one of the following predefined values:

MS\_\$MK\_BAK Make a backup (".bak") file.

### name

Pathname of the object to be made permanent, in NAME\_\$PNAME\_T format. This is an array of up to 256 characters.

The pathname you specify must be on the same volume you used when creating the temporary file.

### name-length

Length of the pathname. This is a 2-byte integer.

## OUTPUT PARAMETERS

### status

Completion status, in STATUS\_\$T format. This data type is 4 bytes long. See the MS Data Types section for more information.

## USAGE

MS\_\$MK\_PERMANENT takes the address of a temporary object you supply from MS\_\$CRMAPL or MS\_\$CRTEMP and makes the file permanent. The file is also given the name you specify and the name is catalogued.

**MS\_\$MK\_TEMPORARY**

Makes a permanent object temporary.

**FORMAT**

MS\_\$MK\_TEMPORARY (address, status)

**INPUT PARAMETERS****address**

Pointer to the first byte of the currently mapped portion of the permanent object, in UNIV\_PTR format. This is a 4-byte integer. Use the pointer returned by the most recent call to MS\_\$MAPL, MS\_\$CRMAPL, or MS\_\$REMAP.

**OUTPUT PARAMETERS****status**

Completion status, in STATUS\_\$T format. This data type is 4 bytes long. See the MS Data Types section for more information.

**USAGE**

MS\_\$MK\_TEMPORARY makes a permanent object temporary by uncataloguing its name. After this routine is called, the object no longer has a pathname, which means that the object will disappear if you call MS\_\$UNMAP or shut down the system.

MS\_ \$NEIGHBORS

MS\_ \$NEIGHBORS

Indicates whether two temporary objects are neighbors.

### FORMAT

same-volume = MS\_ \$NEIGHBORS (address1, address2, status)

### RETURN VALUE

#### same-volume

A boolean that indicates whether two temporary objects reside on the same volume. TRUE means that they do.

### INPUT PARAMETERS

#### address1

Pointer to the first byte of the first temporary object, in UNIV\_PTR format. This is a 4-byte integer. Use the pointer returned by the most recent call to MS\_ \$CRTEMP or MS\_ \$CRMPL.

#### address2

Pointer to the first byte of the second temporary object, in UNIV\_PTR format. This is a 4-byte integer. Use the pointer returned by the most recent call to MS\_ \$CRTEMP or MS\_ \$CRMPL.

### OUTPUT PARAMETERS

#### status

Completion status, in STATUS\_ \$T format. This data type is 4 bytes long. See the MS Data Types section for more information.

### USAGE

MS\_ \$NEIGHBORS allows you to determine whether two temporary mapped objects reside on the same volume. You supply the addresses of the objects from previous mappings.

**MS\_ \$RELOCK**

Changes the lock on an object.

**FORMAT**

MS\_ \$RELOCK (virtual-address, access, status)

**INPUT PARAMETERS****virtual-address**

Pointer to the first mapped byte of the object whose lock you want to change, in UNIV\_ PTR format. This is a 4-byte integer. Use the pointer returned by an earlier call to MS\_ \$MAPL, MS\_ \$CRMPL, or MS\_ \$REMAP.

**access**

New access mode, in MS\_ \$ACC\_ MODE\_ T format. This is a 2-byte integer. Specify only one of the following predefined values:

MS_ \$R	Read access.
MS_ \$RX	Read and execute access.
MS_ \$WR	Read and write access.
MS_ \$WRX	Read, write, and execute access.
MS_ \$RIW	Read with intent to write.

If you specify an access mode of MS\_ \$RIW when you first lock an object, you cannot relock the object with MS\_ \$R access.

**OUTPUT PARAMETERS****status**

Completion status, in STATUS\_ \$T format. This data type is 4 bytes long. See the MS Data Types section for more information. Possible values are:

STATUS_ \$OK	Completed successfully.
MS_ \$NOT_ MAPPED	No object is mapped at the supplied virtual address.
MS_ \$BAD_ ACCESS	Access mode given is incorrect.
FILE_ \$ILLEGAL_ LOCK_ RQST	Illegal lock request (file server); the access mode given is incorrect.

**USAGE**

MS\_ \$RELOCK changes the lock on an object. With MS\_ \$RELOCK, you specify a new access type. This new access, in combination with the current concurrency mode, forms a new lock. You can relock a file in the following ways:

<i>Current Lock</i>	<i>Changes</i>
Protected read	Change to exclusive write by specifying the access mode MS_\$WR or MS_\$WRX.  Change to protected RIW by specifying the access mode MS_\$RIW.
Protected RIW	Change to exclusive write by specifying the access mode MS_\$WR or MS_\$WRX.  Cannot change to protected read by specifying the access mode MS_\$R.
Shared read	Change to shared write by specifying the access mode MS_\$WR or MS_\$WRX.
Exclusive write	Change to protected read by specifying the access mode MS_\$R.  Change to protected RIW by specifying the access mode MS_\$RIW.
Shared write	Change to shared read by specifying MS_\$R or MS_\$RIW.

See the description of MS\_\$MAPL for a list of the concurrency/access combinations for each lock.

**MS\_\$REMAP**

Maps a different portion of a previously mapped object.

**FORMAT**

address = MS\_\$REMAP (old-address, start, desired-length,  
remapped-length, status)

**RETURN VALUE****address**

Pointer to the first byte of the new mapped section, in UNIV\_PTR format. This is a 4-byte integer.

**INPUT PARAMETERS****old-address**

Pointer to the first byte of the currently mapped portion of the object, in UNIV\_PTR format. This is a 4-byte integer. Use the pointer returned by the most recent call to MS\_\$MAPL, MS\_\$CRMAPL, or MS\_\$REMAP.

**start**

First byte to be mapped. This is a 4-byte integer.

**desired-length**

Number of bytes to remap. This is a 4-byte integer.

**OUTPUT PARAMETERS****remapped-length**

Number of bytes remapped. This is a 4-byte integer.

**status**

Completion status, in STATUS\_\$T format. This data type is 4 bytes long. See the MS Data Types section for more information. Possible values are:

STATUS\_\$OK Completed successfully.

MS\_\$NOT\_MAPPED

No object is mapped at the given virtual address.

MS\_\$BAD\_LENGTH

Desired-length is invalid.

MS\_ \$REMAP

## USAGE

This call maps a different portion of an already mapped object and unmaps the previously mapped portion. This call is useful for moving a sliding window over a big file.

When you remap a file, certain attributes of the mapping (extend, access, concurrency) are left the same as in the original mapping. If you used MS\_ \$ADVICE to provide file access advice, the advice in effect for the first part of the currently mapped section is propagated to the newly mapped section. Also, MS\_ \$REMAP does not change the lock mode of the object.



**MS\_\$TRUNCATE**

Truncates a mapped object to the specified length.

**FORMAT**

MS\_\$TRUNCATE (address, length, status)

**INPUT PARAMETERS****address**

Pointer to the first byte of the currently mapped portion of the object, in UNIV\_PTR format. This is a 4-byte integer. Use the pointer returned by the most recent call to MS\_\$MAPL, MS\_\$CRMAPL, or MS\_\$REMAP.

**length**

Number of bytes to keep in the mapped object, starting at the first byte in the object. This is a 4-byte integer. Everything after this length is truncated.

**OUTPUT PARAMETERS****status**

Completion status, in STATUS\_\$T format. This data type is 4 bytes long. See the MS Data Types section for more information.

**USAGE**

MS\_\$TRUNCATE shortens a mapped file to the length that you specify. In addition, you can use MS\_\$TRUNCATE to define a length for a file, even if you are not throwing away data. For example, when you unmap a file, the system may set the file length to a page-aligned value. (That is, the length will be a multiple of 1024.) However, you can use MS\_\$TRUNCATE to shorten the file to a nonpage-aligned value.

For example, if a file contains only 20 bytes of data, you can use MS\_\$TRUNCATE to set the file length to 20. When you unmap the file, the length will be 20 rather than 1024. Use MS\_\$ATTRIBUTES to determine the current file length and number of blocks used.

MS\_\$UNMAP

MS\_\$UNMAP

Unmaps a previously mapped object.

## FORMAT

MS\_\$UNMAP (address, length, status)

## INPUT PARAMETERS

### address

Pointer to the first byte of the currently mapped portion of the object, in UNIV\_PTR format. This is a 4-byte integer. Use the pointer returned by the most recent call to MS\_\$MAPL, MS\_\$CRMAPL, or MS\_\$REMAP.

### length

Number of mapped bytes. This is a 4-byte integer. Use the length you requested in the most recent call to MS\_\$MAPL, MS\_\$CRMAPL, or MS\_\$REMAP.

## OUTPUT PARAMETERS

### status

Completion status, in STATUS\_\$T format. This data type is 4 bytes long. See the MS Data Types section for more information. Possible values are:

STATUS\_\$OK Completed successfully.

MS\_\$NOT\_MAPPED

Address and length given do not refer to an object mapped with MS\_\$MAPL.

## USAGE

MS\_\$UNMAP unmaps and unlocks an object mapped and locked with MS\_\$MAPL or MS\_\$CRMAPL. You cannot unmap a subset of the object.

If the original object is on a remote node, changes made in the mapped version are written back to the original object when MS\_\$UNMAP is executed. If the original object is on the local node, changes made in the mapped version of the object will be written back to the original object when the space they occupy in memory is needed.

If the object was mapped with the extend parameter equal to TRUE, and your program modified part of the extension space, the original object is extended to include those modifications. Parts of the extension space beyond the last modification are not added to the original object.

An object locked by several calls to MS\_\$MAPL by different processes will remain locked until all the processes have unmapped the object.

**ERRORS****MS\_\$BAD\_ACCESS**

Unsupported access rights requested.

**MS\_\$BAD\_LENGTH**

Bad length.

**MS\_\$IN\_USE**

Object is locked by another process or in an incompatible mode.

**MS\_\$INSUFFICIENT\_RIGHTS**

You have some access rights to the object, but not the ones you requested.

**NAME\_\$NAME\_NOT\_FOUND**

No object exists with the given name.

**MS\_\$NO\_RIGHTS**

You do not have any access rights to the object.

**MS\_\$NO\_SPACE**

No space.

**MS\_\$NOT\_MAPPED**

No object mapped at the virtual address supplied.

**MS\_\$OBJECT\_NOT\_FOUND**

The object does not exist, or it is not accessible over the network.

**MS\_\$WRONG\_LEVEL**


Attempt to release segment mapped by previous level.

**STATUS\_\$OK**

Successful completion.

Other naming server errors.

See the NAME\_\$ error codes.



# MTS

## Table of Contents

Data Types	MTS-2
MTS_\$CLOSE_DESC	MTS-8
MTS_\$COPY_DESC	MTS-9
MTS_\$CREATE_DEFAULT_DESC	MTS-10
MTS_\$GET_ATTR	MTS-11
MTS_\$LABEL	MTS-12
MTS_\$OPEN_DESC	MTS-13
MTS_\$SET_ATTR	MTS-14
Errors	MTS-15



## MTS

The MTS (Streams-Magtape) programming calls perform streams operations to magnetic tape devices. This section describes their data types, call syntax, and error codes. Refer to the Introduction at the beginning of this manual for a description of data-type diagrams and call syntax format.

## MTS DATA TYPES

### CONSTANTS

MTS_\$CURRENT_FILE_SEQ	File sequence number for current file.
MTS_\$END_FILE_SEQ	File sequence number for last tape file.
MTS_\$FIRST_A	MTS_\$UNIT_A First attribute in MTS_\$ATTR_T.
MTS_\$LAST_A	MTS_\$BUFFER_OFFSET_A Last attribute in MTS_\$ATTR_T.

### DATA TYPES

MTS\_\$ATTR\_T                   A 2-byte integer. THE User-modifiable tape file attributes. One of the following predefined values:

MTS\_\$UNIT\_A  
Tape unit number.

MTS\_\$LABELED\_A  
Labeled volume.

MTS\_\$REOPEN\_VOL\_A  
Reopen volume.

MTS\_\$CLOSE\_VOL\_A  
Close file and volume.

MTS\_\$SAVE\_VOL\_POS\_A  
Save position on close.

MTS\_\$VOL\_DEVICE\_A  
Device type.

MTS\_\$VOL\_ID\_A  
Volume ID.

MTS\_\$VOL\_ACCESS\_A  
Volume accessibility.

MTS\_\$OWNER\_ID\_A  
Owner ID.

MTS\_\$FILE\_SEQUENCE\_A  
File sequence number.

MTS\_\$RECORD\_FORMAT\_A  
Record format.

MTS\_\$BLOCK\_LENGTH\_A  
Block length.



MTS\_\$RECORD\_LENGTH\_A  
Maximum record length.

MTS\_\$ASCII\_NL\_A  
ASCII newline head-length

MTS\_\$FILE\_SECTION\_A  
File section number.

MTS\_\$FILE\_ID\_A  
File ID.

MTS\_\$FILE\_SET\_ID\_A  
File set ID.

MTS\_\$GENERATION\_A  
Generation number.

MTS\_\$GENERATION\_VERSION\_A  
Generation version number.

MTS\_\$CREATE\_DATE\_A  
Creation date.

MTS\_\$EXP\_DATE\_A  
Expiration date.

MTS\_\$FILE\_ACCESS\_A  
File accessibility.

MTS\_\$SYSTEM\_CODE\_A  
System code.

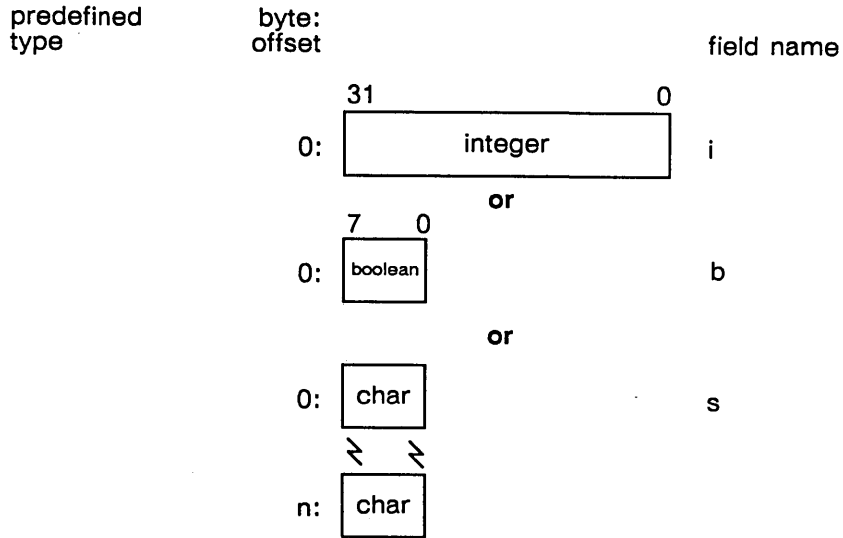
MTS\_\$SYSTEM\_USE\_A  
System use.

MTS\_\$BUFFER\_OFFSET\_A  
Buffer offset.

MTS\_\$ATTR\_VALUE\_T

Attribute values. The diagram below illustrates the  
MTS\_\$ATTR\_VALUE\_T data type:

MTS DATA TYPES



Field Description:

i  
An integer value.

b  
A Boolean value.

s  
A character string.

MTS\_\$DEVICE\_T

A 2-byte integer. Type of device. One of the following predefined values:

MTS\_\$MT  
Magtape device.

MTS\_\$NOT\_REALLY  
Not currently supported.

MTS\_\$CT  
Cartridge tape device.

MTS\_\$HANDLE\_T

A 4-byte integer. A handle to a tape descriptor file.

MTS\_\$RW\_T


A 2-byte integer. Read or write status. One of the following predefined values:

MTS\_\$READ  
Read operation.

MTS\_\$WRITE  
Write operation.

STATUS\_\$T

A status code. The diagram below illustrates the STATUS\_\$T data type:



Field Description:

**all**  
All 32 bits in the status code.

**fail**  
The fail bit. If this bit is set, the error was not within the scope of the module invoked, but occurred within a lower-level module (bit 31).

**subsys**  
The subsystem that encountered the error (bits 24 - 30).

**modc**  
The module that encountered the error (bits 16 - 23).

**code**  
A signed number that identifies the type of error that occurred (bits 0 - 15).

Table MTS-1. Magnetic Tape Volume and File Attributes

<i>Mnemonic</i>	<i>Type</i>	<i>Default</i>	<i>Definition</i>
MTS_\$UNIT_A	int	0	magtape unit number (normally 0)
MTS_\$LABELED_A	t/f	true	true = ANSI labeled volume false = unlabeled volume
MTS_\$REOPEN_VOL_A	t/f	false	true = reopen previously used volume (suppresses rewind) false = do not reopen
MTS_\$CLOSE_VOL_A	t/f	true	true = volume closed when file is closed false = leave volume open
MTS_\$SAVE_VOL_POS_A	t/f	false	true = saves volume position when volume is closed (for reopen) false = rewind volume when closed
MTS_\$VOL_DEVICE_A	int	0	type of device: tfp_\$mt=0 for magtape tfp_\$ct=3 for cartridge
MTS_\$VOL_ID_A	char	-auto	volume identifier (labeled volumes) (Automatically generated.) Six-character string maximum.
MTS_\$VOL_ACCESS_A	char	" "	volume accessibility (labeled volumes only). The default is the space character.
MTS_\$VOL_OWNER_ID_A	char	-auto	volume owner (labeled volumes) Maximum string length is 14.
MTS_\$FILE_SEQUENCE_A	int "cur" "end"	1	file sequence number. Possible values are an integer, "cur" for current file, "end" for new file at end of labeled volume.
MTS_\$RECORD_FORMAT_A	char	D	record format. Possible values: "F" = fixed length "D" = variable length "S" = spanned "U" = undefined
MTS_\$BLOCK_LENGTH_A	int	2048	block length, in bytes
MTS_\$RECORD_LENGTH_A	int	2048	maximum record length, in bytes

Table MTS-1. Magnetic Tape Volume and File Attributes (Continued)

<i>Mnemonic</i>	<i>Type</i>	<i>Default</i>	<i>Definition</i>
MTS_\$ASCII_NL_A	t/f	true	true = ASCII newline handling. Strip newlines on write, supply them on read false = no newline handling
MTS_\$FILE_SECTION_A	int	1	file section number (labeled volumes)
MTS_\$FILE_ID_A	char	" "	file identifier (labeled volumes)
MTS_\$FILE_SET_ID_A	char	" "	file set identifier (labeled volumes)
MTS_\$GENERATION_A	int	1	generation of file (labeled volumes)
MTS_\$GENERATION_VERSION_A	int	1	generation version of file (labeled volumes)
MTS_\$CREATE_DATE_A	date	-auto	creation date of file (labeled volumes)
MTS_\$EXP_DATE_A	date	-auto	expiration date of file (labeled volumes)
MTS_\$FILE_ACCESS_A	char	" "	file accessibility (labeled volumes)
MTS_\$SYSTEM_CODE_A	char	" "	system code (labeled volumes)
MTS_\$SYSTEM_USE_A	char	" "	system use (labeled volumes)
MTS_\$BUFFER_OFFSET_A	int	0	buffer offset (labeled volumes) Must be zero.

MTS\_\$CLOSE\_DESC

MTS\_\$CLOSE\_DESC

Closes a magtape descriptor file.

#### FORMAT

MTS\_\$CLOSE\_DESC (handle, update, status)

#### INPUT PARAMETERS

##### handle

Pointer to the open magtape descriptor file, in MTS\_\$HANDLE\_T format. This is a 4-byte integer. Specify the handle returned by MTS\_\$OPEN\_DESC, MTS\_\$COPY\_DESC, or MTS\_\$CREATE\_DESC.

##### update

Boolean value that determines whether or not the magtape descriptor file is to be modified to reflect the attribute changes specified by calls to MTS\_\$SET\_ATTR. If TRUE, the changes are made.

#### OUTPUT PARAMETERS

##### status

Completion status, in STATUS\_\$T format. This data type is 4 bytes long. See the MTS Data Types section for more information.

#### USAGE

Programs must close magtape descriptor files before calling stream manager routines; an open magtape descriptor file cannot be used by the stream manager.

Closing a magtape descriptor file invalidates its handle.

**MTS\_\$COPY\_DESC**

Copies a magtape descriptor file and opens the destination file.

**FORMAT**

```
handle = MTS_$COPY_DESC (src-pathname, src-namelen, dst-pathname,
                          dst-namelen, status)
```

**RETURN VALUE****handle**

Pointer to the open magtape descriptor file, in MTS\_\$HANDLE\_T format. This is a 4-byte integer.

**INPUT PARAMETERS****src-pathname**

The pathname of the magtape descriptor file to be copied, in NAME\_\$PNAME\_T format. This is an array of up to 256 characters.

**src-namelen**

Length of the source pathname, in bytes. This is a 2-byte integer.

**dst-pathname**

The pathname to which the file is to be copied, in NAME\_\$PNAME\_T format. This is an array of up to 256 characters.

The destination file must not exist before this function is called. To replace a destination file, call the routine NAME\_\$DELETE\_FILE before calling MTS\_\$COPY\_DESC.

**dst-namelen**

Length of the destination pathname, in bytes. This is a 2-byte integer.

**OUTPUT PARAMETERS****status**

Completion status, in STATUS\_\$T format. This data type is 4 bytes long. See the MTS Data Types section for more information.

**USAGE**

This routine copies the specified magtape descriptor file, opens the destination file and returns a pointer to it.

This routine does not affect the source file.

MTS\_\$CREATE\_DEFAULT\_DESC

MTS\_\$CREATE\_DEFAULT\_DESC

Creates a magtape descriptor file with the default volume and file attributes.

#### FORMAT

handle = MTS\_\$CREATE\_DEFAULT\_DESC (pathname, namelen, status)

#### RETURN VALUE

##### handle

A pointer to the open magtape descriptor file, in MTS\_\$HANDLE\_T format. This is a 4-byte integer.

#### INPUT PARAMETERS

##### pathname

The pathname of the descriptor file to be created, in NAME\_\$PNAME\_T format. This is an array of up to 256 characters.

##### namelen

Length of the name, in bytes. This is a 2-byte integer.

#### OUTPUT PARAMETERS

##### status

Completion status, in STATUS\_\$T format. This data type is 4 bytes long. See the MTS Data Types section for more information.

#### USAGE

This routine opens a magtape descriptor file with the default volume and file attribute values and returns a pointer to it. The file must not exist before this routine is called. See the Table in the MTS Data Types section for a list of volume and file attributes and their defaults.



**MTS\_\$GET\_ATTR**

Retrieves a given attribute from a magtape descriptor file.

**FORMAT**

MTS\_\$GET\_ATTR (handle, attribute, value, status)

**INPUT PARAMETERS****handle**

A pointer to the open magtape descriptor file, in MTS\_\$HANDLE\_T format. This is a 4-byte integer. Specify a handle returned by MTS\_\$OPEN\_DESC, MTS\_\$COPY\_DESC, or MTS\_\$CREATE\_DESC.

**attribute**

The attribute to be retrieved, in MTS\_\$ATTR\_T format. This is a 2-byte integer. Specify only one of the following predefined values:

mts_\$unit_a	mts_\$labeled_a	mts_\$reopen_vol_a
mts_\$close_vol_a	mts_\$save_vol_pos_a	mts_\$vol_device_a
mts_\$vol_id_a	mts_\$vol_access_a	mts_\$owner_id_a
mts_\$file_sequence_a	mts_\$record_format_a	mts_\$block_length_a
mts_\$record_length_a	mts_\$ascii_nl_a	mts_\$file_resv1_a
mts_\$file_section_a	mts_\$file_id_a	mts_\$file_set_id_a
mts_\$generation_a	mts_\$generation_version_a	mts_\$create_date_a
mts_\$exp_date_a	mts_\$file_access_a	mts_\$system_code_a
mts_\$system_use_a	mts_\$buffer_offset_a	

See the Table in the MTS Data Types section for a description of volume and file attributes and their defaults.

**OUTPUT PARAMETERS****value**

The current value of the specified attribute, in MTS\_\$ATTR\_VALUE\_T format. Possible values are a 4-byte integer, a Boolean value, or a string, depending upon the attribute requested. See the Table in the MTS Data Types section for a list of volume and file attributes and their corresponding values.

**status**

Completion status, in STATUS\_\$T format. This data type is 4 bytes long. See the MTS Data Types section for more information.

**USAGE**

Programs must call this routine once for each attribute they wish to get.

You can change the attributes within a magtape descriptor file using the MTS\_\$SET\_ATTR system call.

MTS\_\$LABEL

MTS\_\$LABEL

Labels the magtape volume described by the given magtape descriptor file.

#### FORMAT

MTS\_\$LABEL (pathname, namelen, status)

#### INPUT PARAMETERS

##### pathname

The pathname of the magtape descriptor file, in NAME\_\$PNAME\_T format. This is an array of up to 256 characters.

The descriptor file must describe a labeled volume.

##### namelen

Length of the name, in bytes. This is a 2-byte integer.

#### OUTPUT PARAMETERS

##### status

Completion status, in STATUS\_\$T format. This data type is 4 bytes long. See the MTS Data Types section for more information.

#### USAGE

MTS\_\$LABEL causes the volume described by the descriptor file to be labeled according to ANSI x3.27-1978.

The tape volume must not be open (by previous calls to the stream manager).

**MTS\_\$OPEN\_DESC**

Opens the specified magtape descriptor file and returns a pointer to it.

**FORMAT**

handle = MTS\_\$OPEN\_DESC (pathname, namelen, read-write, status)

**RETURN VALUE****handle**

A pointer to the open magtape descriptor file, in MTS\_\$HANDLE\_T format. This is a 4-byte integer.

**INPUT PARAMETERS****pathname**

The pathname of the magtape descriptor file, in NAME\_\$PNAME\_T format. This is an array of up to 256 characters.

**namelen**

Length of the name, in bytes. This is a 2-byte integer.

**read-write**

Read or write status, in MTS\_\$RW\_T format. This is a 2-byte integer. Specify only one of the following predefined values:

MTS\_\$READ Open for reading.

MTS\_\$WRITE Open for writing.

**OUTPUT PARAMETERS****status**

Completion status, in STATUS\_\$T format. This data type is 4 bytes long. See the MTS Data Types section for more information.

**USAGE**

MTS\_\$OPEN\_DESC opens the specified magtape descriptor file for reading or writing and returns a pointer to it.

MTS\_\$SET\_ATTR

MTS\_\$SET\_ATTR

Sets an attribute within the specified magtape descriptor file.

## FORMAT

MTS\_\$SET\_ATTR (handle, attribute, value, status)

## INPUT PARAMETERS

### handle

A pointer to an open magtape descriptor file, in MTS\_\$HANDLE\_T format. This is a 4-byte integer. Specify a handle returned by MTS\_\$OPEN\_DESC, MTS\_\$COPY\_DESC, or MTS\_\$CREATE\_DESC.

### attribute

The volume or file attribute to be set, in MTS\_\$ATTR\_T format. This is a 2-byte integer. Specify only one of the following predefined values:

mts_\$unit_a	mts_\$labeled_a	mts_\$reopen_vol_a
mts_\$close_vol_a	mts_\$save_vol_pos_a	mts_\$vol_device_a
mts_\$vol_id_a	mts_\$vol_access_a	mts_\$owner_id_a
mts_\$file_sequence_a	mts_\$record_format_a	mts_\$block_length_a
mts_\$record_length_a	mts_\$ascii_nl_a	mts_\$file_resv1_a
mts_\$file_section_a	mts_\$file_id_a	mts_\$file_set_id_a
mts_\$generation_a	mts_\$generation_version_a	mts_\$create_date_a
mts_\$exp_date_a	mts_\$file_access_a	mts_\$system_code_a
mts_\$system_use_a	mts_\$buffer_offset_a	

See the Table in the MTS Data Types section for a description of volume and file attributes and their defaults.

### value

The value to assign to the attribute, in MTS\_\$ATTR\_VALUE\_T format. Possible values are a 4-byte integer, a Boolean value, or a string, depending upon the attribute to be changed. See the Table in the MTS Data Types section for a list of volume and file attributes and their corresponding values.

## OUTPUT PARAMETERS

### status

Completion status, in STATUS\_\$T format. This data type is 4 bytes long. See the MTS Data Types section for more information.

## USAGE

Programs must call this routine once for each attribute to be set.

You can change the attributes within a magtape descriptor file using the Shell command EDMTDESC. See the *DOMAIN System Command Reference* for details.

**ERRORS**

STATUS\_ \$OK  
Successful completion.

MTS\_ \$BAD\_ BLOCK\_ LENGTH  
Bad block length.

MTS\_ \$BAD\_ BUFFER\_ OFFSET  
Bad buffer offset.

MTS\_ \$BAD\_ DATA  
Descriptor contains bad data.

MTS\_ \$BAD\_ FILE\_ SECTION  
Bad file section number.

MTS\_ \$BAD\_ FILE\_ SEQUENCE  
Bad file sequence number.

MTS\_ \$BAD\_ GENERATION  
Bad generation number.

MTS\_ \$BAD\_ GENERATION\_ VERSION  
Bad generation version number.

MTS\_ \$BAD\_ RECORD\_ FORMAT  
Bad record format attribute.

MTS\_ \$BAD\_ RECORD\_ LENGTH  
Bad record length.

MTS\_ \$BAD\_ UNIT  
Bad tape unit number.

MTS\_ \$INVALID\_ ATTR  
Invalid attribute to GET\_ ATTR/SET\_ ATTR.

MTS\_ \$INVALID\_ DATE  
Invalid date text string.

MTS\_ \$NOT\_ LABELED  
Attempt to label unlabeled volume.

MTS\_ \$READ\_ ONLY  
SET\_ ATTR on read-only file.

MTS\_ \$VOL\_ IN\_ USE  
Volume in use.

MTS\_ \$WRONG\_ TYPE  
Object is not type MT\_ \$UID.



# MUTEX

## Table of Contents

Data Types

MUTEX\_ \$INIT

MUTEX\_ \$LOCK

MUTEX\_ \$UNLOCK

MUTEX-2

MUTEX-4

MUTEX-5

MUTEX-6





# MUTEX

The MUTEX (Mutual Exclusion) programming calls provide applications with mutual exclusion resource-sharing and synchronization. This section describes their data types and call syntax. The MUTEX calls do not produce unique error messages. Refer to the Introduction at the beginning of this manual for a description of data-type diagrams and call syntax format.

## MUTEX DATA TYPES

### CONSTANTS

MUTEX\_\$WAIT\_FOREVER


integer32(-1)

A value that tells MUTEX\_\$LOCK to wait forever without timing out.

### DATA TYPES

MUTEX\_\$LOCK\_REC\_T

A mutual exclusion lock record. The diagram below illustrates the MUTEX\_\$LOCK\_REC\_T data type:



#### Field Description:

lock\_byte

A Boolean value that indicates whether any programs currently hold a MUTEX lock.

lock\_ec


An eventcount for programs waiting for the MUTEX lock. The lock\_ec field is in EC2\_\$EVENTCOUNT format and has two subfields:

lock\_ec.value      The value of the eventcount.

lock\_ec.awaiters      Used internally by the EC2 manager.

TIME\_\$CLOCK\_T


Internal representation of time. The diagram below illustrates the TIME\_\$CLOCK\_T data type:



Field Description:

high  
High 32 bits of the clock.

low  
Low 16 bits of the clock.



Field Description:

high16  
High 16 bits of the clock.

low32  
Low 32 bits of the clock.

MUTEX\_\$INIT

MUTEX\_\$INIT

Initializes a mutual exclusion lock record.

#### FORMAT

MUTEX\_\$INIT (lock-record)

#### OUTPUT PARAMETERS

##### lock-record

Lock record, in MUTEX\_\$LOCK\_REC\_T format. This data type is 8 bytes long. See the MUTEX Data Types section for more information.

#### USAGE

Use this call to initialize a mutual exclusion (MUTEX) lock record. This lock record allows a program to obtain a MUTEX lock on a file. A MUTEX lock allows a program to get exclusive access to a shared resource.

Initialize a MUTEX lock record within a file. First, map the file with a concurrency mode of MS\_\$COWRITERS and an access type of MS\_\$WR. Then use MUTEX\_\$INIT to initialize the MUTEX lock record.

**MUTEX\_\$LOCK**

Obtains a mutual exclusion lock on a file.

**FORMAT**

lock-status = MUTEX\_\$LOCK (lock-record, wait-time)

**RETURN VALUE****lock-status**

A Boolean value that indicates whether you obtained the lock. TRUE means that you got the lock; FALSE means that the call timed out before obtaining the lock.

**INPUT/OUTPUT PARAMETERS****lock-record**

Lock record, in MUTEX\_\$LOCK\_REC\_T format. This data type is 8 bytes long. See the MUTEX Data Types section for more information.

**INPUT PARAMETERS****wait-time**

The amount of time to wait for the lock, in TIME\_\$CLOCK\_T format. This data type is 6 bytes long. See the MUTEX Data Types section for more information.

If MUTEX\_\$LOCK cannot obtain the lock within the time you specify, the call will time out and return control to your program. Specify the waiting time as a relative time. Use the CAL routines to convert time values to TIME\_\$CLOCK\_T format.

If you specify the waiting time using the constant MUTEX\_\$WAIT\_FOREVER, the MUTEX\_\$LOCK call wait indefinitely to obtain the lock.

**USAGE**

Use MUTEX\_\$LOCK to obtain a mutual exclusion (MUTEX) lock on a file. A MUTEX lock lets you have exclusive access to a shared resource.

MUTEX\_\$LOCK uses the information in a lock record to determine whether you can obtain the lock. (Use MUTEX\_\$INIT to initialize a lock record.) If another program already has the lock, MUTEX\_\$LOCK waits for the amount of time you specify. When MUTEX\_\$LOCK returns, it indicates whether you obtained the lock.

Before calling MUTEX\_\$LOCK, you must map the file containing the lock record. Map the file with a concurrency mode of MS\_\$COWRITERS and an access mode of MS\_\$WR. All programs that map the same MUTEX lock record must be on the same node.

Note that a MUTEX lock is a convention that cooperating programs use to control access to a resource. If a program does not use MUTEX\_\$LOCK and accesses the resource directly, you cannot guarantee mutual exclusion.

MUTEX\_\$UNLOCK

MUTEX\_\$UNLOCK

Terminates a program's mutual exclusion lock on a file.

#### FORMAT

MUTEX\_\$UNLOCK (lock-record)

#### INPUT/OUTPUT PARAMETERS

##### lock-record

Lock record, in MUTEX\_\$LOCK\_REC\_T format. This data type is 8 bytes long. See the MUTEX Data Types section for more information.

#### USAGE

MUTEX\_\$UNLOCK terminates a program's mutual exclusion lock on a file. A waiting program can then obtain the lock.

# NAME

## Table of Contents

Data Types	NAME-2
NAME_\$ADD_LINK	NAME-4
NAME_\$CNAME	NAME-5
NAME_\$CREATE_DIRECTORY	NAME-6
NAME_\$CREATE_FILE	NAME-7
NAME_\$DELETE_DIRECTORY	NAME-8
NAME_\$DELETE_FILE	NAME-9
NAME_\$DROP_LINK	NAME-10
NAME_\$EXTRACT_DATA	NAME-11
NAME_\$GET_NDIR	NAME-12
NAME_\$GET_PATH	NAME-13
NAME_\$GET_WDIR	NAME-14
NAME_\$READ_DIR	NAME-15
NAME_\$READ_LINK	NAME-17
NAME_\$SET_NDIR	NAME-18
NAME_\$SET_WDIR	NAME-19
Errors	NAME-20





## NAME

The NAME (Naming Server) programming calls manipulate text-string object names and translate them to their system (UID) names. This section describes the data types, call syntax, and error codes of the NAME calls. Refer to the Introduction at the beginning of this manual for a description of data-type diagrams and call syntax format.


NAME DATA TYPES

CONSTANTS

NAME_\$COMPLEN_MAX	32	Maximum length of an entry name.
NAME_\$FILE	1	The file type value for the enttype field of the DIR_ENTRY_T record.
NAME_\$LINK	3	The link type value for the enttype field of the DIR_ENTRY_T record.
NAME_\$PNAMLEN_MAX	256	Maximum length of a pathname.

DATA TYPES

NAME\_\$DIR\_ENTRY\_T                      The directory entry returned by NAME\_\$READ\_DIR. The diagram below illustrates the NAME\_\$DIR\_ENTRY\_T data type:



Field Description:

enttype  
Type of the directory entry. Either NAME\_\$FILE or NAME\_\$LINK.

entlen  
Length of the directory entry name.

entname  
Name of the directory entry.

unusedn  
Reserved for future use by Apollo.

NAME\_\$DIR\_LIST\_T

A 1300-element array of NAME\_\$DIR\_ENTRY\_T record structures. The diagram below illustrates a single element:

NAME\_\$NAME\_T

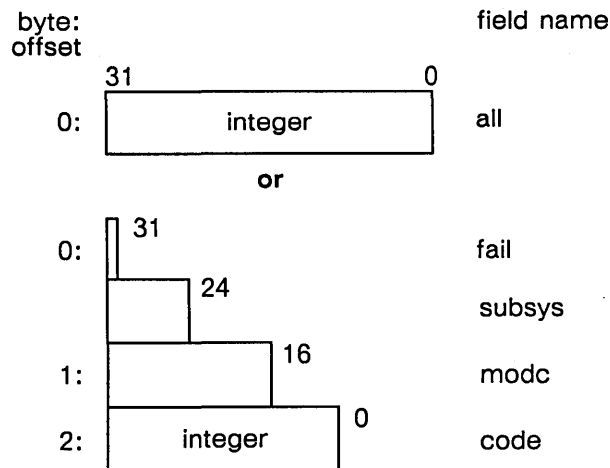
An array of up to NAME\_\$COMPLEN\_MAX (32) characters.

NAME\_\$PNAME\_T

An array of up to NAME\_\$PNAMLEN\_MAX (256) characters.

STATUS\_\$T

A status code. The diagram below illustrates the STATUS\_\$T data type:



Field Description:

all  
All 32 bits in the status code.

fail  
The fail bit. If this bit is set, the error was not within the scope of the module invoked, but occurred within a lower-level module (bit 31).

subsys  
The subsystem that encountered the error (bits 24 - 30).

modc  
The module that encountered the error (bits 16 - 23).

code  
A signed number that identifies the type of error that occurred (bits 0 - 15).

NAME\_\$ADD\_LINK

NAME\_\$ADD\_LINK

Creates a link.

## FORMAT

NAME\_\$ADD\_LINK (linkname, name-length, link-text, text-length, status)

## INPUT PARAMETERS

### linkname

Name of the link, in NAME\_\$PNAME\_T format. This is an array of up to 256 characters.

Specify either an absolute or relative pathname. If a relative pathname is specified, the rest of the pathname defaults to the current working directory. If a pathname is specified beginning with a slash (/), the object is placed in the entry directory of the local node.

### name-length

Length of the linkname, in bytes. This is a 2-byte integer.

### link-text

Pathname to which the link refers, in NAME\_\$PNAME\_T format. This is an array of up to 256 characters.

The link text replaces the linkname when the linkname is used as part of a pathname. For example, suppose a link named YEATS had a link text //MAN/IN/MASK. Using the object name YEATS is exactly equivalent to using the pathname //MAN/IN/MASK directly.

The link text must be a valid filename or pathname. It does not, however, have to refer to an existing object.

### text-length

Length of the link-text pathname, in bytes. This is a 2-byte integer.

## OUTPUT PARAMETERS

### status

Completion status, in STATUS\_\$T format. This data type is 4 bytes long. See the NAME Data Types section for more information.

## USAGE

A link is an object within a directory that points to another object. That is, the link is associated with a pathname that refers to another object. The associated pathname is referred to as the link text. When the link is referenced, the naming server acts as if the link text were given in place of the link name.

To delete a link, you must use the naming server call NAME\_\$DROP\_LINK, or the Shell command DELETE\_LINK (DLL).

This system call corresponds to the CRL Shell command.

**NAME\_ \$CNAME**

Changes the last element of a pathname.

**FORMAT**

NAME\_ \$CNAME (old-pathname, old-length, new-leaf, leaf-length, status)

**INPUT PARAMETERS****old-pathname**

The current pathname, in NAME\_ \$PNAME\_ T format. This is an array of up to 32 characters.

**old-length**

The length of the current pathname, in bytes. This is a 2-byte integer.

**new-leaf**

The name that replaces the right-most element of the current pathname, in NAME\_ \$NAME\_ T format. This is an array of up to 256 characters.

**leaf-length**

The length of the new-leaf name, in bytes. This is a 2-byte integer.

**OUTPUT PARAMETERS****status**

Completion status, in STATUS\_ \$T format. This data type is 4 bytes long. See the NAME Data Types section for more information.

**USAGE**

NAME\_ \$CNAME changes the right-most element of the old-pathname to the string specified by the new-leaf argument.

NAME\_\$CREATE\_DIRECTORY

NAME\_\$CREATE\_DIRECTORY

Creates a directory.

#### FORMAT

NAME\_\$CREATE\_DIRECTORY (directory-name, name-length, status)

#### INPUT PARAMETERS

##### directory-name

Name of the directory, in NAME\_\$PNAME\_T format. This is an array of up to 256 characters.

Specify either an absolute or relative pathname. If a relative pathname is specified, the rest of the pathname defaults to the current working directory. If a pathname is specified beginning with a slash (/), the object is placed in the entry directory of the local node.

##### name-length

Length of directory name, in bytes. This is a 2-byte integer.

#### OUTPUT PARAMETERS

##### status

Completion status, in STATUS\_\$T format. This data type is 4 bytes long. See the NAME Data Types section for more information.

#### USAGE

NAME\_\$CREATE\_DIRECTORY creates a directory using the specified pathname and name length.

This system call corresponds to the CRD Shell command.

**NAME\_\$CREATE\_FILE**

Creates a permanent file.

**FORMAT**

NAME\_\$CREATE\_FILE (filename, name-length, status)

**INPUT PARAMETERS****filename**

Name of the file, in NAME\_\$PNAME\_T format. This is an array of up to 256 characters.

**name-length**

Length of the filename, in bytes. This is a 2-byte integer.

**OUTPUT PARAMETERS****status**

Completion status, in STATUS\_\$T format. This data type is 4 bytes long. See the NAME Data Types section for more information.

**USAGE**

The filename given is treated in the same way as any pathname given to the naming server. For example, a filename beginning with a slash (/) is placed in the entry directory of the local node.

NAME\_\$DELETE\_DIRECTORY

NAME\_\$DELETE\_DIRECTORY

Deletes a directory.

## FORMAT

NAME\_\$DELETE\_DIRECTORY (directory-name, name-length, status)

## INPUT PARAMETERS

### directory-name

Name of the directory, in NAME\_\$PNAME\_T format. This is an array of up to 256 characters.

Specify either an absolute or relative pathname. If a relative pathname is specified, the rest of the pathname defaults to the current working directory. If a pathname is specified beginning with a slash (/), the object is placed in the entry directory of the local node.

### name-length

Length of the name, in bytes. This is a 2-byte integer.

## OUTPUT PARAMETERS

### status

Completion status, in STATUS\_\$T format. This data type is 4 bytes long. See the NAME Data Types section for more information.

## USAGE

NAME\_\$DELETE\_DIRECTORY deletes the specified directory. The directory must be empty for a deletion to succeed.



**NAME\_\$DELETE\_FILE**

Deletes a file.

**FORMAT**

NAME\_\$DELETE\_FILE (filename, name-length, status)

**INPUT PARAMETERS****filename**

Name of the file, in NAME\_\$PNAME\_T format. This is an array of up to 256 characters.

Specify either an absolute or relative pathname. If a relative pathname is specified, the rest of the pathname defaults to the current working directory. If a pathname is specified beginning with a slash (/), the object is placed in the entry directory of the local node.

**name-length**

Length of the filename, in bytes. this is a 2-byte integer.

**OUTPUT PARAMETERS****status**

Completion status, in STATUS\_\$T format. This data type is 4 bytes long. See the NAME Data Types section for more information.

**USAGE**

NAME\_\$DELETE\_FILE deletes the specified file.

This system call corresponds to the DLF Shell command.

NAME\_ \$DROP\_ LINK

NAME\_ \$DROP\_ LINK

Deletes a link.

## FORMAT

NAME\_ \$DROP\_ LINK (linkname, name-length, status)

## INPUT PARAMETERS

### linkname

Name of the link, in NAME\_ \$PNAME\_ T format. This is an array of up to 256 characters.

Specify either an absolute or relative pathname. If a relative pathname is specified, the rest of the pathname defaults to the current working directory. If a pathname is specified beginning with a slash (/), the object is placed in the entry directory of the local node.

### name-length

Length of the link name, in bytes. This is a 2-byte integer.

## OUTPUT PARAMETERS

### status

Completion status, in STATUS\_ \$T format. This data type is 4 bytes long. See the NAME Data Types section for more information.

## USAGE

NAME\_ \$DROP\_ LINK deletes the specified link to an associated object.

This system call corresponds to the DLL Shell command.

**NAME\_\$EXTRACT\_DATA**

Extracts data from a directory entry read by NAME\_\$READ\_DIR. (Intended primarily for use in FORTRAN programs.)

**FORMAT**

NAME\_\$EXTRACT\_DATA (dir-entry, entry-type, entry-length, entry-name)

**INPUT PARAMETERS****dir-entry**

The directory entry for which you wish to extract data, in NAME\_\$ENTRY\_T format. This data type is 44 bytes long. See the NAME Data Types section for more information.

In FORTRAN programs, NAME\_\$READ\_DIR returns the directory entries in a (22,n) INTEGER\*2 array, where n is the maximum number of directory entries your program is prepared to accept. Each column in this array corresponds to an entry in the specified directory and contains information about that entry.

Specify the first element of the column that corresponds to the entry for which you wish to extract data.

**OUTPUT PARAMETERS****entry-type**

Object type of the entry. This is a 2-byte integer with one of the following predefined values:

1 - NAME\_\$FILE

the object is a file. A "file" can be either a streams file or a directory.

2 - NAME\_\$LINK

the object is a link.

**entry-length**

Length of the object's name, in bytes. This is a 2-byte integer.

**entry-name**

The entry name, in NAME\_\$NAME\_T format. This is an array of up to 32 characters.

**USAGE**

This call extracts the description of a single directory entry from the directory entry array (the dir-list parameter) returned by NAME\_\$READ\_DIR. It is intended primarily for use in FORTRAN programs.

In FORTRAN programs, NAME\_\$READ\_DIR returns the directory entries in a (22,n) INTEGER\*2 array, where n is the maximum number of directory entries your program is prepared to accept. Each column in this array corresponds to an entry in the specified directory and contains information about that entry.

The dir-entry parameter for NAME\_\$EXTRACT\_DATA should be one of the array columns. To reference a single column, give the first element of that column.

NAME\_\$GET\_NDIR

NAME\_\$GET\_NDIR

Returns the full pathname of the naming directory.

#### FORMAT

NAME\_\$GET\_NDIR (name, name-length, status)

#### OUTPUT PARAMETERS

##### **name**

Pathname of the naming directory, in NAME\_\$PNAME\_T format. This is an array of up to 256 characters.

##### **name-length**

Length of the name, in bytes. This is a 2-byte integer.

##### **status**

Completion status, in STATUS\_\$T format. This data type is 4 bytes long. See the NAME Data Types section for more information.

#### USAGE

The naming directory is set through the NAME\_\$SET\_NDIR call or the "ND directory-name" Shell command. This system call corresponds to the ND Shell command.

**NAME\_\$GET\_PATH**

Converts a partial pathname into a full pathname.

**FORMAT**

NAME\_\$GET\_PATH (in-name, in-len, out-name, out-len, status)

**INPUT PARAMETERS****in-name**

The relative pathname of an object, in NAME\_\$PNAME\_T format. This is an array of up to 256 characters.

**in-len**

Length of the relative pathname, in bytes. This is a 2-byte integer.

**OUTPUT PARAMETERS****out-name**

The full (absolute) pathname of the object, in NAME\_\$PNAME\_T format. This is an array of up to 256 characters.

**out-len**

Length of the relative pathname, in bytes. This is a 2-byte integer.

**status**

Completion status, in STATUS\_\$T format. This data type is 4 bytes long. See the NAME Data Types section for more information.

**USAGE**

NAME\_\$GET\_PATH converts a partial pathname into a full pathname. For example, if you have been using file FOO, you can call NAME\_\$GET\_PATH to find out that the full pathname of FOO is //FLYNN/PHL/FOO.

NAME\_\$GET\_WDIR

NAME\_\$GET\_WDIR

Returns the full pathname of the working directory.

#### FORMAT

NAME\_\$GET\_WDIR (name, name-length, status)

#### OUTPUT PARAMETERS

##### **name**

Pathname of the working directory, in NAME\_\$PNAME\_T format. This is an array of up to 256 characters.

##### **name-length**

Length of the name, in bytes. This is a 2-byte integer.

##### **status**

Completion status, in STATUS\_\$T format. This data type is 4 bytes long. See the NAME Data Types section for more information.

#### USAGE

The working directory is set through the NAME\_\$SET\_WDIR call or the "WD directory-name" Shell command.

**NAME\_\$READ\_DIR**

Reads a directory.

**FORMAT**

NAME\_\$READ\_DIR (dir-name, name-length, dir-list, index, max-count,  
read-count, status)

**INPUT PARAMETERS****dir-name**

Name of the directory, in NAME\_\$PNAME\_T format. This is an array of up to 256 characters.

Specify either an absolute or relative pathname. If a relative pathname is specified, the rest of the pathname defaults to the current working directory. If a pathname is specified beginning with a slash (/), the entry directory of the local node is searched for the directory.

Specifying a null character (") defaults to the current working directory.

**name-length**

Length of the name, in bytes. This is a 2-byte integer.

If you specify a null character for the directory name, specify zero as the length.

**OUTPUT PARAMETERS****dir-list**

A list of directory entries, in NAME\_\$DIR\_LIST\_T format. This is an array of NAME\_\$DIR\_ENTRY\_T data types. See the NAME Data Types section for more information.

The number of NAME\_\$DIR\_ENTRY\_T data types in the array must equal or exceed max-count.

**INPUT/OUTPUT PARAMETERS****index**

Key indicating the directory entry at which NAME\_\$READ\_DIR begins reading. This is a 4-byte integer.

On input            This number indicates the entry at which reading begins.

On output           This number is adjusted by NAME\_\$READ\_DIR to a number suitable for a subsequent call to NAME\_\$READ\_DIR.

To read from the start of the directory, initialize the index to 1 on your first call to NAME\_\$READ\_DIR.

Because NAME\_\$READ\_DIR adjusts the index parameter to a suitable value for a subsequent call, you should not change the value yourself.

NAME\_\$READ\_DIR

## INPUT PARAMETERS

### max-count

Maximum number of directory entries to read. This is a 2-byte integer.

## OUTPUT PARAMETERS

### read-count

Number of directory entries actually read. This is a 2-byte integer.

If NAME\_\$READ\_DIR reaches the end of the directory before finding the requested number of entries, read-count will contain the actual number of entries read, which will be smaller than the requested max-count. If there are no entries left in the directory prior to being called, NAME\_\$READ\_DIR returns a read-count of 0.

### status

Completion status, in STATUS\_\$T format. This data type is 4 bytes long. See the NAME Data Types section for more information.

## USAGE

NAME\_\$READ\_DIR reads a directory and stores entry names, the length of each entry name, and the type of each entry. Pascal and C programs can access this information directly through the directory entry record structure. FORTRAN programs use the NAME\_\$EXTRACT\_DATA system call to access this information.

The index argument permits a program to make several calls to NAME\_\$READ\_DIR to ensure reading all entries. However, to get an accurate snapshot of a directory, make only one call to NAME\_\$READ\_DIR, using a sufficiently large max-count, because the contents of a directory can change between calls to NAME\_\$READ\_DIR.



**NAME\_\$READ\_LINK**

Returns the link text associated with a link name.

**FORMAT**

NAME\_\$READ\_LINK (linkname, name-length, link-text, text-length, status)

**INPUT PARAMETERS****linkname**

Name of the link, in NAME\_\$PNAME\_T format. This is an array of up to 256 characters.

Specify either an absolute or relative pathname. If a relative pathname is specified, the rest of the pathname defaults to the current working directory. If a pathname is specified beginning with a slash (/), the entry directory of the local node is searched for the link.

**name-length**

Length of the linkname, in bytes. This is a 2-byte integer.

**OUTPUT PARAMETERS****link-text**

Text associated with the linkname, in NAME\_\$PNAME\_T format. This is an array of up to 256 characters.

**text-length**

Length of the text, in bytes. This is a 2-byte integer.

**status**

Completion status, in STATUS\_\$T format. This data type is 4 bytes long. See the NAME Data Types section for more information.

**USAGE**

When you use a linkname, the naming server replaces the link name with the associated link text. NAME\_\$READ\_LINK returns the text associated with a specified link name.

NAME\_\$SET\_NDIR

NAME\_\$SET\_NDIR

Sets the naming directory.

## FORMAT

NAME\_\$SET\_NDIR (name, name-length, status)

## INPUT PARAMETERS

### name

Pathname of the desired naming directory, in NAME\_\$PNAME\_T format. This is an array of up to 256 characters.

Specify either an absolute or relative pathname. If a relative pathname is specified, the rest of the pathname defaults to the current working directory. A directory name beginning with a period (.) indicates a directory within the working directory. You may also specify a period by itself, which sets the naming directory equal to the working directory.

### name-length

Length of the pathname, in bytes. This is a 2-byte integer.

## OUTPUT PARAMETERS

### status

Completion status, in STATUS\_\$T format. This data type is 4 bytes long. See the NAME Data Types section for more information.

## USAGE

NAME\_\$SET\_NDIR sets the naming directory to the specified directory. See the *DOMAIN System Command Reference* for a description of naming directories.

This system call corresponds to the "ND directory-name" Shell command.

**NAME\_ \$SET\_ WDIR**

Sets the working directory.

**FORMAT**

NAME\_ \$SET\_ WDIR (name, name-length, status)

**INPUT PARAMETERS****name**

Pathname of the desired working directory, in NAME\_ \$PNAME\_ T format. This is an array of up to 256 characters.

Specify either an absolute or relative pathname. If a relative pathname is specified, the rest of the pathname defaults to the current working directory. A directory name beginning with a period (.) indicates a directory within the working directory.

**name-length**

Length of the name, in bytes. This is a 2-byte integer.

**OUTPUT PARAMETERS****status**

Completion status, in STATUS\_ \$T format. This data type is 4 bytes long. See the NAME Data Types section for more information.

**USAGE**

NAME\_ \$SET\_ WDIR sets the working directory to the specified directory. See the DOMAIN System Command Reference Manual for a description of working directories.

This system call corresponds to the "WD directory-name" shell command.

NAME ERRORS

**ERRORS**

STATUS\_ \$OK

Successful completion.

NAME\_ \$ALREADY\_EXISTS

Name already exists.

NAME\_ \$BAD\_DIRECTORY

Bad directory.

NAME\_ \$BAD\_LEAF

Invalid leaf.

NAME\_ \$BAD\_LINK

Invalid link.

NAME\_ \$BAD\_PATHNAME

Invalid pathname.

NAME\_ \$DIRECTORY\_FULL

Directory is full.

NAME\_ \$DIRECTORY\_NOT\_EMPTY

Directory is not empty.

NAME\_ \$FILE\_NOT\_DIRECTORY

Branch is not a directory.

NAME\_ \$ILL\_LINK\_OP

Invalid link operation.

NAME\_ \$INSUFFICIENT\_RIGHTS

Insufficient rights.

NAME\_ \$IS\_SYSBOOT

Unable to delete system bootstrap (sysboot).

NAME\_ \$NO\_RIGHTS

No rights.

NAME\_ \$NODE\_UNAVAILABLE

Node is unavailable.

NAME\_ \$NOT\_FILE

Name is not a file.

NAME\_ \$NOT\_FOUND

Name not found.

NAME\_ \$NOT\_LINK

Name is not a link.

# PAD

## Table of Contents

Data Types	PAD-2
PAD_\$CLEAR_FRAME	PAD-9
PAD_\$CLOSE_FRAME	PAD-10
PAD_\$COOKED	PAD-11
PAD_\$CPR_ENABLE	PAD-12
PAD_\$CREATE	PAD-14
PAD_\$CREATE_FRAME	PAD-17
PAD_\$CREATE_ICON	PAD-18
PAD_\$CREATE_WINDOW	PAD-20
PAD_\$DEF_PFK	PAD-22
PAD_\$DELETE_FRAME	PAD-24
PAD_\$DM_CMD	PAD-25
PAD_\$EDIT_WAIT	PAD-26
PAD_\$FORCE_PROMPT	PAD-27
PAD_\$ICON_WAIT	PAD-28
PAD_\$INQ_DISP_TYPE	PAD-29
PAD_\$INQ_FONT	PAD-31
PAD_\$INQ_FULL_WINDOW	PAD-32
PAD_\$INQ_ICON	PAD-33
PAD_\$INQ_ICON_FONT	PAD-34
PAD_\$INQ_KBD	PAD-35
PAD_\$INQ_POSITION	PAD-36
PAD_\$INQ_VIEW	PAD-37
PAD_\$INQ_WINDOWS	PAD-38
PAD_\$IS_ICON	PAD-39
PAD_\$LOAD_FONT	PAD-40
PAD_\$LOCATE	PAD-41
PAD_\$MAKE_ICON	PAD-42
PAD_\$MAKE_INVISIBLE	PAD-43
PAD_\$MOVE	PAD-44
PAD_\$POP_PUSH_WINDOW	PAD-45
PAD_\$RAW	PAD-46
PAD_\$SELECT_WINDOW	PAD-47
PAD_\$SET_AUTO_CLOSE	PAD-48
PAD_\$SET_BORDER	PAD-49
PAD_\$SET_FULL_WINDOW	PAD-50
PAD_\$SET_ICON_FONT	PAD-51
PAD_\$SET_ICON_POS	PAD-52
PAD_\$SET_SCALE	PAD-53
PAD_\$SET_TABS	PAD-54
PAD_\$SET_VIEW	PAD-55
PAD_\$USE_FONT	PAD-56
Errors	PAD-57



# PAD

The PAD (Display Manager) programming calls allow you to create pads, windows, and window panes, and to manipulate text. This section describes their data types, call syntax, and error codes. Refer to the Introduction at the beginning of this manual for a description of data-type diagrams and call syntax format.

## PAD DATA TYPES

### CONSTANTS

MNEMONIC	Value	Explanation
PAD_\$BS	8	Moves the cursor one character position to the left if there is any room in the window.
PAD_\$CPR_ALL	2	Cursor position report: Reports on each raw keystroke.
PAD_\$CPR_CHANGE	1	Cursor position report: Reports only the changed position since the last output call or position report.
PAD_\$CPR_DRAW	4	Cursor position report: Reports on all touchpad data.
PAD_\$CPR_FLAG	16#FF	Cursor position report: Indicates that the next 5 bytes is a report.
PAD_\$CPR_NONE	0	Cursor position report: Does not report any cursor positions.
PAD_\$CPR_PICK	3	Cursor position report: Reports after cursor is settled when it has been moved by the touchpad.
PAD_\$CR	13	Returns cursor to the left edge of the pad at the same line it was on.
PAD_\$ESCAPE	27	For control characters: Tells Display Manager not to interpret the next character as a control character. This precedes ANSI escape sequences.
PAD_\$FF	12	Makes output start at the top of the window or window pane.
PAD_\$LEFT_WINDOW	16#FD	Cursor position report: Indicates that the cursor accompanying the report is outside the window.
PAD_\$MAX_TABSTOPS	100	Defines the maximum number of tabstops allowed to be set.
PAD_\$NEWLINE	10	Marks end of an input or output line, makes next text start on a new line.
PAD_\$NO_KEY	16#FE	Cursor position report: Indicates that no keystroke accompanies the report.
PAD_\$TAB	9	Moves cursor to next tab stop.

### DATA TYPES

PAD\_\$COORDINATE\_T                      2-byte integer for x and y bitmap coordinates.



PAD\_\$CRE\_OPT\_T

A 2-byte integer. Options of a pane. Any combination of the following predefined values:

PAD\_\$ABS\_SIZE

Size parameter is absolute, rather than relative to the size of the existing pad.

PAD\_\$INIT\_RAW

Input pad is initially raw, rather than normal (cooked) processing mode.

PAD\_\$DISPLAY\_TYPE\_T

A 2-byte integer. Type of display associated with the specified stream id. This is a 2-byte integer. One of the following predefined values:

PAD\_\$BW\_15P

Black and white portrait display.

PAD\_\$BW\_19L

Black and white landscape display.

PAD\_\$COLOR\_DISPLAY

Color display (1024 x 1024 pixels).

PAD\_\$800\_COLOR

Color display (1024 x 800 pixels).

PAD\_\$NONE

No display.

PAD\_\$KEY\_DEF\_T

An array of up to 256 characters. Display Manager command to be defined on a program- function key using PAD\_\$DEF\_PFK.

PAD\_\$KEY\_NAME\_T

An array of up to 4 characters. Name of the program-function key to be defined using PAD\_\$DEF\_PFK.

PAD DATA TYPES

PAD\_\$POSITION\_T

X and y coordinates of a point on the display. The diagram below illustrates the PAD\_\$POSITION\_T data type:

predefined  
type

byte:  
offset

field name

0:	integer	x_coord
2:	integer	y_coord

Field Description:

y\_coord

The y coordinate of the point on the display.

x\_coord

The x coordinate of the point on the display.

PAD\_\$REL\_ABS\_T

A 2-byte integer. Indicates whether cursor movement is relative to the last location, or absolute. X and y are scaled. One of the following predefined values:

PAD\_\$ABSOLUTE

X and y are absolute values. Within a frame, movement is relative to the top left corner of the frame. Outside a frame, x is relative to the left end of the current line, and y is undefined.

PAD\_\$RELATIVE

Cursor movement is relative to the last location. X and y denote positive or negative offsets to the current cursor position.

PAD\_ \$SIDE\_ T

A 2-byte integer. Side of a transcript pad that a new pane occupies. One of the following predefined values:

PAD\_ \$BOTTOM

Bottom of transcript pad.

PAD\_ \$LEFT

Left side of transcript pad.

PAD\_ \$RIGHT

Right side of transcript pad.

PAD\_ \$TOP

Top of transcript pad.

PAD\_ \$STRING\_ T

An array of up to 256 characters. String argument to some functions.

PAD DATA TYPES

PAD\_\$TABSTOP\_BUF\_T

A 100-element array of 2-byte integers. Columns for tab stop settings. Each element contains a column number at which a tab stop will be set. Column numbers are scaled.

PAD\_\$TYPE\_T

A 2-byte integer. A type of pad. One of the following predefined values:

PAD\_\$EDIT  
An edit pad.


PAD\_\$INPUT  
An input pad.

PAD\_\$TRANSCRIPT  
A transcript pad.

PAD\_\$READ\_EDIT  
A read/edit pad.

PAD\_\$WINDOW\_DESC\_T

Position of window on display screen. The diagram below illustrates the PAD\_\$WINDOW\_DESC\_T data type:



Field Description:

top  
The y coordinate of the top left corner of the window, in raster units.

left  
The x coordinate of the top left corner of the window, in raster units.

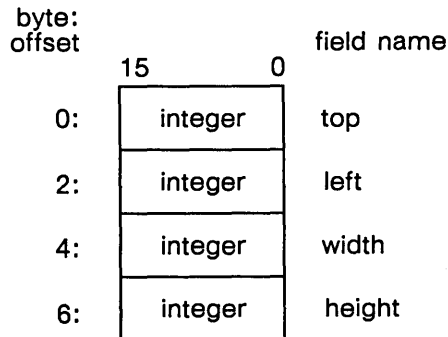
width  
The width of the window, divided by the current x scale factor.

height  
The height of the window, divided by the current y scale factor.

PAD\_\$WINDOW\_LIST\_T

A 10-element array of  
 PAD\_\$WINDOW\_DESC\_T record structures.  
 The diagram below illustrates a single element:

predefined  
 type



top

The y coordinate of the top left corner of the window, in raster units.

left

The x coordinate of the top left corner of the window, in raster units.

width

The width of the window, divided by the current x scale factor.

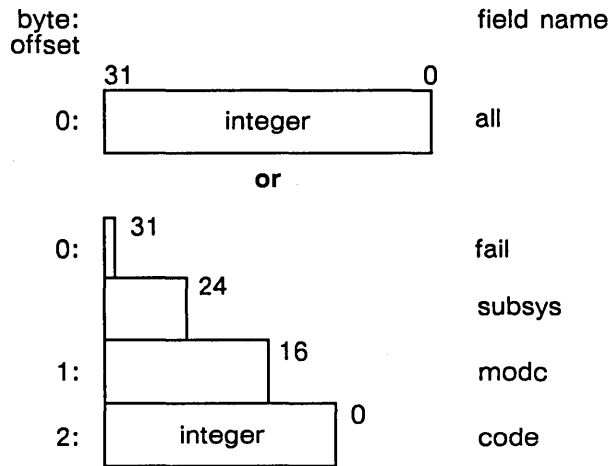
height

The height of the window, divided by the current y scale factor.

PAD DATA TYPES

STATUS\_\$\$T

A status code. The diagram below illustrates the STATUS\_\$\$T data type:



Field Description:

all

All 32 bits are in the status code.

fail

The fail bit. If this bit is set, the error was not within the scope of the module invoked, but occurred within a lower-level module (bit 31).

subsys

The subsystem that encountered the error (bits 24 - 30).

modc

The module that encountered the error (bits 16 - 23).

code

A signed number that identifies the type of error that occurred (bits 0 - 15).

**PAD\_\$CLEAR\_FRAME**

Clears the current frame, leaving it active.

**FORMAT**

PAD\_\$CLEAR\_FRAME (stream-id, seek-key, status)

**INPUT PARAMETERS****stream-id**

Number of the stream on which the pad is open, in STREAM\_\$ID\_T format. This is a 2-byte integer.

**seek-key**

Unique value identifying the record where clearing begins, in STREAM\_\$ID\_T format. This is a three element array of 4-byte integers. See the STREAM Data Types section for more information.

**OUTPUT PARAMETERS****status**

Completion status, in STATUS\_\$T format. This data type is 4-bytes long. See the PAD Data Types section for more information.

**USAGE**

Use this call to clear information from the frame that you created with the call, PAD\_\$CREATE\_FRAME. Programs that use frames often overwrite text at random points within the frame. You should periodically call PAD\_\$CLEAR\_FRAME to remove this discarded data. By doing so, you prevent data from accumulating in the transcript pad file. You also prevent the Display Manager from invoking the time-consuming frame-rewrite operation.

Clearing begins at the record indicated by the seek-key and continues to the end of the frame. If the first four bytes of the seek-key are 0, the entire frame is cleared. The seek-key is returned by STREAM\_\$PUT\_REC and STREAM\_\$PUT\_CHR. See the STREAM Calls section for more information.

PAD\_ \$CLOSE\_ FRAME

PAD\_ \$CLOSE\_ FRAME

Closes a frame, leaving its contents in the pad, and returns to line-oriented processing on the input pad.

#### FORMAT

PAD\_ \$CLOSE\_ FRAME (stream-id, status)

#### INPUT PARAMETERS

##### stream-id

Number of the stream on which the pad is open, in STREAM\_ \$ID\_ T format. This is a 2-byte integer.

#### OUTPUT PARAMETERS

##### status

Completion status, in STATUS\_ \$T format. This data type is 4-bytes long. See the PAD Data Types section for more information.

#### USAGE

After the frame is closed, you can view the frame by scrolling the transcript window backwards. Once the frame is closed, all frame operations except PAD\_ \$CREATE\_ FRAME are invalid.



**PAD\_\$COOKED**

Disables raw mode input or output to a pad.

**FORMAT**

PAD\_\$COOKED (stream-id, status)

**INPUT PARAMETERS****stream\_id**

Number of the stream on which the pad is open, in STREAM\_\$ID\_T format. This is a 2-byte integer.

**OUTPUT PARAMETERS****status**

Completion status, in STATUS\_\$T format. This data type is 4 bytes long. See the PAD Data Types section for more information.

**USAGE**

This call returns the pad to normal (cooked) processing if it is currently in raw mode due to a call to the PAD\_\$RAW procedure. PAD\_\$COOKED has no effect if called when the pad is not currently in raw mode. After you execute this procedure, the input window reappears and is empty.

PAD\_ \$CPR\_ ENABLE

PAD\_ \$CPR\_ ENABLE

Enables reporting of the keyboard cursor position for an input pad in raw mode. (You can only get keyboard cursor position reports on pads in raw mode).

## FORMAT

PAD\_ \$CPR\_ ENABLE (stream-id, report-cpr-type, status)

## INPUT PARAMETERS

### stream-id

Number of the stream on which the input pad is open, in STREAM\_ \$ID\_ T format. This is a 2-byte integer.

### report-cpr-type

Type of cursor position report. This is a 2-byte integer. Specify one of the following predefined values:

PAD\_ \$CPR\_ NONE

Requests no cursor position reports (the default).

PAD\_ \$CPR\_ CHANGE

Requests cursor position reports only when the cursor has moved through keystrokes since the last output call or the last position report.

PAD\_ \$CPR\_ ALL

Requests a cursor position report with every character.

PAD\_ \$CPR\_ PICK

Requests a cursor position report after the cursor has settled after being moved by the touchpad, bitpad, or mouse.

PAD\_ \$CPR\_ DRAW

Requests a cursor position report for all cursor positions during cursor movement from the touchpad, bitpad, or mouse.

PAD\_ \$CPR\_ PICK and PAD\_ \$CPR\_ DRAW also report new cursor positions resulting from Display Manager commands; for example, arrow keys, tabs, TR, TL, TB.

## OUTPUT PARAMETERS

### status

Completion status, in STATUS\_ \$T format. This data type is 4 bytes long. See the PAD Data Types section for more information.

## USAGE

You get cursor position reports in response to a `STREAM_$GET_REC` call, intermixed with raw character data. The Display Manager uses a single byte to represent a raw keystroke. It uses a 6-byte sequence to give you a cursor position report. The sequence looks like this:

- 1 byte -- `PAD_$CPR_FLAG`, indicating that the next 5 bytes are a cursor position report.
- 2 bytes -- The x coordinate of the cursor position.
- 2 bytes -- The y coordinate of the cursor position.
- 1 byte -- The raw keystroke or `PAD_$NO_KEY` if there is no keystroke accompanying this cursor position report, or `PAD_$LEFT_WINDOW` if the cursor moved outside the window.

The x and y coordinates are scaled according to the scaling factors in effect at the time of the `PAD_$CPR_ENABLE` call (see `PAD_$SET_SCALE`). The x and y coordinates are relative to the upper left corner of the frame. (If the cursor is not inside a frame, the x coordinate is relative to the start of the current line, and the y coordinate is meaningless.)

In raw mode, the Display Manager does not automatically echo typed keystrokes nor move the cursor. If your program requests `PAD_$CPR_ALL` but does not act to move the cursor (typically by displaying typed keystrokes), each keystroke produces a cursor position report, usually describing the same cursor position. If you don't intend to echo keyboard input, request `PAD_$CPR_CHANGE` instead to avoid redundant cursor position reports.

`PAD_$CPR_CHANGE` compares the present keyboard cursor with the last *output* cursor position. In raw mode, the position of the output cursor is under program control. Therefore, if your program does not move the output cursor to follow the input cursor (which you can move) you may receive a stream of cursor position reports, all showing the same position, as long as the keyboard cursor is not in the same position as the output cursor.

PAD\_\$CREATE

PAD\_\$CREATE

Creates a new pad and a window pane to view it.

## FORMAT

PAD\_\$CREATE (pathname, name-length, pane-type, related-stream-id, side,  
pane\_options, pane-size, pane-stream-id, status)

## INPUT PARAMETERS

### pathname

Pathname to a file to display in the window pane, in NAME\_\$PNAME\_T format. This is an array of up to 256 characters.

If the specified pathname refers to an existing file, the Display Manager positions the new window pane at the beginning of the file, and displays any existing data. If the given pathname does not refer to an existing file, a permanent file with that name is created. You usually use a null pathname when creating a transcript pad. You must specify a null pathname when creating an input pad.

### name-length

Length of the pathname in bytes. This is a 2-byte integer. A value of 0 creates a temporary file for the pad. You must specify 0 when creating an input pad.

### pane-type

The window pane type in PAD\_\$TYPE\_T format. This is a 2-byte integer. Specify one of the following predefined values:

PAD\_\$EDIT Creates a pad in which you can view and modify the associated file.

PAD\_\$INPUT Creates an input pad.

PAD\_\$READ\_EDIT  
Creates a pad in which you can view but not modify the associated file.

PAD\_\$TRANSCRIPT  
Creates a transcript pad.

### related-stream-id

The stream ID of a transcript pad, in STREAM\_\$ID\_T format. This is a 2-byte integer. The related-stream-id for an input window pane (PAD\_\$INPUT) must refer to an open transcript window pane that has no other input window pane associated with it.

**side**

The side of the transcript pad that the new window occupies, in PAD\_\$SIDE\_T format. This is a 2-byte integer. Specify one of the following predefined values:

PAD\_\$LEFT

PAD\_\$RIGHT

PAD\_\$TOP

PAD\_\$BOTTOM

You must specify PAD\_\$BOTTOM when creating an input window pane for a transcript window pane.

**pane-options**

Attributes of the pane. This is a 2-byte integer. In Pascal, specify any combination of the following set of predefined values:

PAD\_\$ABS\_SIZE

Specifies an absolute pane-size. If not given, the pane-size parameter is a relative value.

PAD\_\$INIT\_RAW

Indicates that a new input pad is initially in raw rather than cooked mode. This is for input pads only, it is invalid for any other pad types.

In FORTRAN, specify either 0, to indicate that the pane-size is relative, or give the sum of the desired options.

**pane-size**

Size of the pane. This is a 2-byte integer. A window pane always takes up one full side of the related window. The size refers only to the depth of the window.

You can express the pane size either as a percentage relative to the existing transcript window, or as an absolute value in terms of the current scale factor.

If you specify the pane-size as an absolute size, the Display Manager attempts to keep the window pane at that size. However, the window pane can never be larger than the related window, so that, if the related window shrinks below the size requested, the window pane also shrinks.

In addition, if you specify the pane size as an absolute size, the value given is multiplied by the current scale factors to yield raster units. The default scale factors are the current font size so that, unless you change the scale, you should express the pane-size in terms of lines or characters.

An input window pane will normally be one line deep, but can grow and shrink depending on how many lines of input are waiting for action. You should specify a size that accommodates this because the size parameter determines the *maximum* number of lines that the input window pane can occupy. The size of an input window pane can never be less than 1. (A common relative size is 20.)

PAD\_ \$CREATE.

## OUTPUT PARAMETERS

### pane-stream-id

Number of the stream on which the new window pane is open, in STREAM\_ \$ID \_T format.

### status

Completion status, in STATUS\_ \$T format. This data type is 4 bytes long. See the PAD Data Types for more information.

## USAGE

Use this call to create a new pad and window pane on a related stream. The related stream can be either the stream ID of a transcript pad that you previously created with a call to PAD\_ \$CREATE or PAD\_ \$CREATE\_ WINDOW. For transcript pads, the stream ID can be a standard output stream such as STREAM\_ \$STDOUT, or STREAM\_ \$ERRROUT.

You can create any number of window panes on top of the original transcript pad up to the maximum of 40 pads and 60 windows.

You must use PAD\_ \$CREATE to create an input pad for an existing transcript pad.

**PAD\_\$CREATE\_FRAME**

Creates a frame in a pad.

**FORMAT**

PAD\_\$CREATE\_FRAME (stream-id, width, height, status)

**INPUT PARAMETERS****stream-id**

Number of the stream on which the input pad is open, in STREAM\_\$ID\_T format. This is a 2-byte integer.

**width**

Width of the new frame in pixels. This is a 2-byte integer. Value can be up to 32767 raster units. Width is scaled according to the current scale factors.

**height**

Height of the new frame in pixels. This is a 2-byte integer. Value can be up to 32767 raster units. Height is scaled according to the current scale factors.

**OUTPUT PARAMETERS****status**

Completion status, in STATUS\_\$T format. This data type is 4 bytes long. See the PAD Data Types section for more information.

**USAGE**

Use this call to create a frame on an existing transcript pad. Because you can move the cursor anywhere within the frame, create frames when you want to have more control over the cursor position in a given area of the screen.

Your program must either close the frame with PAD\_\$CLOSE\_FRAME or delete the frame with PAD\_\$DELETE\_FRAME before exiting. (Note that you can review a closed frame by scrolling the transcript window backwards, but a deleted frame no longer exists.)

PAD\_\$CREATE\_ICON

PAD\_\$CREATE\_ICON

Creates a new pad and associated window in icon format.

## FORMAT

PAD\_\$CREATE\_ICON (pathname, name-length, type, unit, icon-pos, icon-char,  
window, stream-id, status)

## INPUT PARAMETERS

### pathname

Pathname to a file to display in the pad, in NAME\_\$PNAME\_T format. This is an array of up to 256 characters.

If the specified pathname refers to an existing file, the Display Manager positions the new window pane at the beginning of the file, and displays any existing data. If the given pathname does not refer to an existing file, a permanent file with that name is created. You usually create a null pathname when creating a transcript pad.

### name-length

Length of the pathname string. This is a 2-byte integer. A value of 0 creates a temporary file for the pad.

### type

Pad type in PAD\_\$TYPE\_T format. This is a 2-byte integer. Specify one of the following predefined values:

PAD\_\$TRANSCRIPT

Creates a transcript pad.

PAD\_\$EDIT Creates a pad in which you can view and modify the associated file.

PAD\_\$READ\_EDIT

Creates a pad in which you can view but not modify the associated file.

### unit

Display unit number associated with the stream-ID. This is a 2-byte integer. Usually, there is only one display per node so this value is often 1.

### icon-pos

X- and y-coordinates of the upper left corner of the icon window, in PAD\_\$POSITION\_T format. This data type is four bytes long. See the PAD Data Types section for more information.

### icon-char

Icon font character to be displayed in the icon window. This character must reside in the current icon font file. A null character value (") causes the Display Manager to select the default icon character for this pad type.



**window**

Window descriptor giving the position on the screen that the new window will occupy when expanded to full size (the icon window size is fixed by the the font character selected), in PAD\_\$WINDOW\_DESC\_T format. This data type is 8 bytes long. See the PAD Data Types section for more information.

The window specified is the usable part of the displayed window. The displayed window is larger by the size of the border and the legend. If you specify either the width or the height as zero, the window is created using the same rules as for Display Manager commands (see the *DOMAIN System Command Reference*).

**OUTPUT PARAMETERS****stream-id**

Number of the stream on which the new window is open, in STREAM\_\$ID\_T format. This is a 2-byte integer.

**status**

Completion status, in STATUS\_\$T format. This data type is 4 bytes long. See the PAD Data Types section for more information.

**USAGE**

Use this call to create a new pad and window in icon format. To change this window from icon format to a full-sized window, use PAD\_\$SELECT\_WINDOW. To change an existing window into icon format, use PAD\_\$MAKE\_ICON.

PAD\_\$CREATE\_WINDOW

PAD\_\$CREATE\_WINDOW

Creates a new pad and a window to view it.

## FORMAT

PAD\_\$CREATE\_WINDOW (pathname, name-length, pad-type, unit, window,  
stream-id, status)

## INPUT PARAMETERS

### pathname

Pathname to a file to display in the pad, in NAME\_\$PNAME\_T format. This is an array of up to 256 characters. When creating an edit or read/edit pad, this is the pathname of the permanent file for use as the pad. If a file with this name exists, the Display Manager positions the new window at the top of the pad. If such a file doesn't exist, a new file with that name is created. You usually use a null pathname when creating a transcript pad.

### name-length

Length of the pathname string. This is a 2-byte integer. When creating an edit or read/edit pad, a value of 0 creates a temporary file as the pad.

### pad-type

Pad type in PAD\_\$TYPE\_T format. This is a 2-byte integer. Specify one of the following predefined values:

PAD\_\$TRANSCRIPT

Creates a transcript pad.

PAD\_\$EDIT Creates a pad in which you can view and modify the associated file.

PAD\_\$READ\_EDIT

Creates a pad in which you can view but not modify the associated file.

### unit

Display unit number to use. This is a 2-byte integer. Usually there is only one node per display so this value is often 1.

### window

Window descriptor giving the position on the screen that the new window will occupy, in PAD\_\$WINDOW\_DESC\_T format. This data type is 8 bytes long. See the PAD Data Types section for more information.

The window specified is the usable part of the displayed window. The displayed window is larger by the size of the border and the legend. If you specify either the width or the height as zero, the window is created using the same rules as for Display Manager commands (see the *DOMAIN System Command Reference*).

## OUTPUT PARAMETERS

### stream-id

Number of the stream on which the new window is open, in STREAM\_\$ID\_T format. This is a 2-byte integer.

**status**

Completion status, in STATUS\_\$T format. This data type is 4 bytes long. See the PAD Data Types section for more information.

**USAGE**

Use this call to create a new pad and window to view it. Use PAD\_\$CREATE to create a new pad and window pane on an existing transcript pad. To create an input pad, you must use PAD\_\$CREATE.

PAD\_\$DEF\_PFK

PAD\_\$DEF\_PFK

Defines a program function key for use by a program.

## FORMAT

PAD\_\$DEF\_PFK (stream-id, key-name, definition, def-len, status)

## INPUT PARAMETERS

### stream-id

Number of the stream on which the pad is open, in STREAM\_\$ID\_T format. This is a 2-byte integer.

### key-name

Name of the key to be defined. This is a 4-byte character array. Use the key name exactly as it appears in the *DOMAIN System Command Reference*. Use uppercase letters (for example, F1) except when you are redefining a lowercase letter key (such as x). Do not use quotes in this character array (except to redefine the quote key).

### definition

Display Manager command you want executed whenever the specified key is pressed. This is an array of up to 128 characters.

### def-len

Length of the definition in bytes. This is a 2-byte integer. A value of 0 (zero) returns the key to its original definition.

## OUTPUT PARAMETERS

### status

Completion status, in STATUS\_\$T format. This data type is 4 bytes long. See the PAD Data Types section for more information.

## USAGE

PAD\_\$DEF\_PFK defines a program function key for use by a program. When you press a defined key, the definition string is entered as a Display Manager command.

Program function keys defined by PAD\_\$DEF\_PFK behave like keys defined through the Display Manager, except that the definition is only effective within windows viewing the associated pad.

Definitions remain in effect after the program finishes executing, but only within windows viewing the pad associated with the program.

The Display Manager command string you specify as the key definition is often the ES command, which contains a text string and lets the program function key simulate the typing of that text. You may specify the ER command, which introduces a two-digit hexadecimal number and feeds that value directly to the program when the user presses the key. The ER command essentially enables raw-mode input of the specified value, with no echoing or other processing by the Display Manager. The *DOMAIN System Command Reference* contains more details on these commands.

The rules for naming keys in PAD\_\$DEF\_PFK differ from the rules for naming keys in the KD (key definition) Display Manager command. That command implicitly converts letters to uppercase and allows the use of single quotes.

PAD\_\$DELETE\_FRAME

PAD\_\$DELETE\_FRAME

Deletes and clears the current frame.

#### FORMAT

PAD\_\$DELETE\_FRAME (stream-id, status)

#### INPUT PARAMETERS

##### stream-id

Number of the stream on which the pad is open, in STREAM\_\$ID\_T format. This is a 2-byte integer.

#### OUTPUT PARAMETERS

##### status

Completion status, in STATUS\_\$T format. This data type is 4 bytes long. See the PAD Data Types section for more information.

#### USAGE

PAD\_\$DELETE\_FRAME removes the current frame from the pad. After executing this procedure, the pad returns to line-oriented processing. You cannot perform further frame operations until you create another frame with a call to PAD\_\$CREATE\_FRAME.

**PAD\_\$DM\_CMD**

Executes a Display Manager command.

**FORMAT**

PAD\_\$DM\_CMD (stream-id, command, command-length, status)

**INPUT PARAMETERS****stream-id**

Number of the stream on which a pad is open, in STREAM\_\$ID\_T format. This is a 2-byte integer.

**command**

Display Manager command, in PAD\_\$STRING\_T format. This is an array of up to 256 characters.

**command-length**

Length of the command string in bytes. This is a 2-byte integer.

**OUTPUT PARAMETERS****status**

Completion status, in STATUS\_\$T format. This data type is 4 bytes long. See the PAD Data Types section for more information.

**USAGE**

Use this procedure with caution since it performs actions that you normally perform with the keyboard. Because of this, PAD\_\$DM\_CMD may produce unexpected results.

You can find a list of Display Manager commands in the *DOMAIN System Command Reference*.

PAD\_\$EDIT\_WAIT

PAD\_\$EDIT\_WAIT

Suspends program execution until you close an edit window pane, then converts the stream so that the program can access the new input.

## FORMAT

PAD\_\$EDIT\_WAIT (pane-stream-id, status)

## INPUT PARAMETERS

### pane-stream-id

Number of the stream on which the edit window is open, in STREAM\_\$ID\_T format. This is a 2-byte integer.

## OUTPUT PARAMETERS

### status

Completion status, in STATUS\_\$T format. This data type is 4 bytes long. See the PAD Data Types section for more information.

## USAGE

Your program suspends execution until you close the edit window pane with a CTRL/Y (PW; WC -Q command) or a CTRL/N (WC -Q command).

If you close the edit window pane with a CTRL/N, and the file did not exist before the edit window pane was created, PAD\_\$EDIT\_WAIT returns an error, usually indicating that the file was deleted while open.

You must use this procedure before reading a file edited through an edit window pane.



**PAD\_ \$FORCE\_ PROMPT**

Forces an unterminated string to be written as a prompt in an input window.

**FORMAT**

PAD\_ \$FORCE\_ PROMPT (stream-id, status)

**INPUT PARAMETERS****stream-id**

Number of the stream on which the input pad is open, in STREAM\_ \$ID\_ T format. This is a 2-byte integer.

**OUTPUT PARAMETERS****status**

Completion status, in STATUS\_ \$T format. This data type is 4 bytes long. See the PAD Data Types section for more information.

**USAGE**

This procedure forces the DM to display the last non-terminated string that was written to a transcript pad as the prompt in the related input pad. The string is not terminated with a PAD\_ \$NEW\_ LINE, which always forces the display of a string, and is usually intended as a prompt string.

This procedure is no-op if the stream id is that of a raw input window.

PAD\_\$ICON\_WAIT

PAD\_\$ICON\_WAIT

Waits until a window is expanded from an icon format to a full-window size or until the icon window moves.

## FORMAT

PAD\_\$ICON\_WAIT (stream-id, window-no, icon-moved, icon-pos, status)

## INPUT PARAMETERS

### stream-id

Number of the stream on which the pad is open, in STREAM\_\$ID\_T format. This is a 2-byte integer.

### window-no

Index into the window list returned by PAD\_\$INQ\_WINDOWS. This is a 2-byte integer. Window number one always refers to the first window created to view the pad.

## OUTPUT PARAMETERS

### icon-moved

A Boolean value indicating icon-window movement. It returns a value of TRUE if the icon window has moved.

### icon-pos

New position of the moved icon window in PAD\_\$POSITION\_T format. This data type is 4 bytes long. See the PAD\_\$ Data Types section for more information.

### status

Completion status, in STATUS\_\$T format. This data type is 4 bytes long. See the PAD Data Types section for more information.

## USAGE

This call may be used on any type of pad.

If the window is not currently in icon format, this call returns immediately.

**PAD\_\$INQ\_DISP\_TYPE**

Returns the type of display associated with the given stream ID.

**FORMAT**

PAD\_\$INQ\_DISP\_TYPE (stream-id, display-type, unit, status)

**INPUT PARAMETERS****stream-id**

Number of the stream associated with an input or transcript pad, in STREAM\_\$ID\_T format. This is a 2-byte integer.

**OUTPUT PARAMETERS****display-type**

Type of display associated with the specified stream ID, in PAD\_\$DISPLAY\_TYPE\_T format. This is a 2-byte integer. Returns one of the following predefined values:

PAD\_\$NONE No display

PAD\_\$BW\_15P  
Black and white portrait

PAD\_\$BW\_19L  
Black and white landscape

PAD\_\$COLOR\_DISPLAY  
Color display (1024 x 1024)

PAD\_\$800\_COLOR  
Color display with fewer pixels (1024 x 800)

PAD\_\$COLOR2\_DISPLAY  
Color display (1280x1024x8)

PAD\_\$COLOR3\_DISPLAY  
Color display (1024x800x8)

PAD\_\$COLOR4\_DISPLAY  
Color display (1024x800x4)

**unit**

Display unit number. This is a 2-byte integer. This parameter is reserved for future use, it will always have the value of 1.

**status**

Completion status, in STATUS\_\$T format. This data type is 4 bytes long. See the PAD Data Types section for more information.

PAD\_\$INQ\_DISP\_TYPE

## USAGE

PAD\_\$INQ\_DISP\_TYPE returns the display type and unit number associated with the stream ID. The display unit number can be used as an argument to PAD\_\$CREATE\_WINDOW.

## PAD\_\$INQ\_FONT

Returns information about the current font.

**FORMAT**

PAD\_\$INQ\_FONT (stream-id, font-width, font-height,  
font-name, font-size, font-len, status)

**INPUT PARAMETERS****stream-id**

The number of the stream on which the pad is open, in STREAM\_\$ID\_T format. This is a 2-byte integer.

**font-size**

The number of bytes available in the "font-name" string buffer. This is a 2-byte integer. PAD\_\$INQ\_FONT fills the "font-name" output parameter with this many characters of information. If you do not want to know the pathname, you can specify 0 (zero) as the value of "font-size."

**OUTPUT PARAMETERS****font-width**

Width of the font in raster units. This is a 2-byte integer. For fonts in which different characters have different widths, "font-width" describes the width of the space character.

**font-height**

Height of the font in raster units. This is a 2-byte integer. The height includes any interline spacing specified in the font file.

**font-name**

Full pathname of the font, up to the node entry directory (/), in PAD\_\$STRING\_T format. This is an array of up to 256 characters. The pathname is returned with the correct character case (i.e., upper-case characters in the pathname are returned as upper-case; lower-case as lower-case).

**font-len**

Length of the "font-file" pathname. This is a 2-byte integer. If this value is greater than the input parameter "font-size," the Display Manager truncates the returned pathname.

**status**

Completion status, in STATUS\_\$T format. This data type is 4 bytes long. See the PAD Data Types section for more information.

**USAGE**

Use this call to determine which font your program is currently using. Your program can use up to 100 different character fonts as long as you initially load all the fonts that you intend to use with PAD\_\$LOAD\_FONT. When you want your program to use a specific font, call PAD\_\$USE\_FONT to invoke a previously loaded font. Each time you want to change a loaded font, use PAD\_\$USE\_FONT.

PAD\_\$INQ\_FULL\_WINDOW

PAD\_\$INQ\_FULL\_WINDOW

Returns information about the entire window specified, including the border and legend.

#### FORMAT

PAD\_\$INQ\_FULL\_WINDOW (stream-id, window-no, window, status)

#### INPUT PARAMETERS

##### stream-id

Number of the stream on which the pad is open, in STREAM\_\$ID\_T format. This is a 2-byte integer.

##### window-no

Window number of the window open on the pad. This is a 2-byte integer. Window number one always refers to the first window created to view the pad.

#### OUTPUT PARAMETERS

##### window

Window descriptor giving the position on the screen that the window occupies, including the border and legend, in PAD\_\$WINDOW\_DESC\_T format. This data type is 8 bytes long. See the PAD Data Types section for more information.

The window gives the position of the top left corner, width and height of the window. The values appear in the following order: top, left, width, height. All values are expressed in raster units.

##### status

Completion status, in STATUS\_\$T format. This data type is 4 bytes long. See the PAD Data Types section for more information.

#### USAGE

Use this call to determine exactly how much screen space your window uses, including the border and legend. A call to PAD\_\$INQ\_WINDOWS returns similar information about the usable part of the display windows (not including the border and legend).

Note that if the specified stream-id and window-no refer to a window pane, the information returned is for the outermost containing window.

**PAD\_\$INQ\_ICON**

Returns information about a window in icon format.

**FORMAT**

PAD\_\$INQ\_ICON (stream-id, window-no, icon-pos, icon-char, status)

**INPUT PARAMETERS****stream-id**

Number of the stream on which the pad is open, in STREAM\_\$ID\_T format. This is a 2-byte integer.

**window-no**

Index into the window list returned by PAD\_\$INQ\_WINDOWS. This is a 2-byte integer. Window number one always refers to the first window created to view the pad.

**OUTPUT PARAMETERS****icon-pos**

Position of the icon, in PAD\_\$POSITION\_T format. This data type is 4 bytes long. See the PAD Data Types section for more information.

**icon-char**

Character currently displayed in the icon window.

**status**

Completion status, in STATUS\_\$T format. This data type is 4 bytes long. See the PAD Data Types section for more information.

**USAGE**

If the window is not currently in icon format, the information returned describes its previous icon status, if any, and its future icon status, should the Display Manager command ICON or the PAD\_\$MAKE\_ICON call be issued with the default setting for icon-pos and icon-char.

PAD\_\$INQ\_ICON\_FONT

PAD\_\$INQ\_ICON\_FONT

Returns information about the current icon font.

## FORMAT

PAD\_\$INQ\_ICON\_FONT (stream-id, window-no, font-name  
font-buf-size, font-len, status)

## INPUT PARAMETERS

### stream-id

Number of the stream on which the pad is open, in STREAM\_\$ID\_T format. This is a 2-byte integer.

### window-no

Window for which information is wanted. Window-no is an index into the window list returned by PAD\_\$INQ\_WINDOWS. This is a 2-byte integer. Window number one always refers to the first window created to view the pad.

### font-buf-size

Number of bytes available in the font-name buffer string. This is a 2-byte integer. PAD\_\$INQ\_FONT fills the output parameter, font-name, with this many characters of information.

## OUTPUT PARAMETERS

### font-name

Pathname the font from the node entry directory (/), in NAME\_\$PNAME\_T format. This is an array of up to 256 characters. The pathname is returned with the correct character case (i.e., upper-case characters in the pathname are returned as upper-case; lower-case as lower-case).

### font-len

Length of the font file pathname. This is a 2-byte integer. If this value is greater than the input parameter font-size, the Display Manager truncates the returned pathname to fit in the smaller number of characters.

### status

Completion status, in STATUS\_\$T format. This data type is 4 bytes long. See the PAD Data Types section for more information.

## USAGE

Use this call to get the pathname of the icon font in use. You can change the icon font in use with the call, PAD\_\$SET\_ICON\_FONT.

The default icon font file is /SYS/DM/FONTS/ICONS. You can create a new icon font file to contain your own icons by using the font editor EDFONT. See the *DOMAIN System Command Reference* for a complete description of EDFONT.



**PAD\_\$INQ\_KBD**

Returns information about the keyboard currently in use.

**FORMAT**

PAD\_\$INQ\_KBD (stream-id, buffer-size, kbd-suffix, length, status)

**INPUT PARAMETERS****stream-id**

Number of the stream associated with an input or transcript pad, in STREAM\_\$ID\_T format. This is a 2-byte integer.

**buffer-size**

Number of bytes available in the "kbd-suffix" string buffer. This is a 2-byte integer.

**OUTPUT PARAMETERS****kbd-suffix**

Suffix to be appended to Display Manager pathnames to locate a key definition file, in PAD\_\$STRING\_T format. This is an array of up to 256 characters. Suffixes used by standard DOMAIN software are:

Null string           Corresponds to the 880 keyboard.

Value of "2"           Corresponds to the low-profile keyboard.

Value of "3"           Corresponds to the low-profile keyboard with numeric keypad.

(Display Manager pathnames for key definitions are /SYS/DM/STD\_KEYES and USER\_DATA/KEY\_DEFS.)

**length**

Actual length of the string. This is a 2-byte integer. If the length parameter is greater than "kbd-suffix," it truncates "kbd-suffix."

**status**

Completion status, in STATUS\_\$T format. This data type is 4 bytes long. See the PAD Data Types section for more information.

**USAGE**

Use this call to determine which keyboard is in use. For example, you might want to set up program definition keys according to the type of keyboard in use.

PAD\_\$INQ\_POSITION

PAD\_\$INQ\_POSITION

Returns the position of the output cursor.

#### FORMAT

PAD\_\$INQ\_POSITION (stream-id, x, y, status)

#### INPUT PARAMETERS

##### stream-id

Number of the stream on which the pad is open, in STREAM\_\$ID\_T format. This is a 2-byte integer.

#### OUTPUT PARAMETERS

##### x

X position of the output cursor. This is a 2-byte integer.

##### y

Y position of the output cursor. This is a 2-byte integer.

##### status

Completion status, in STATUS\_\$T format. This data type is 4 bytes long. See the PAD Data Types section for more information.

#### USAGE

X and y are divided by the current scale factors.

If this procedure is executed when the cursor is inside a frame, x and y are relative to the upper left corner of the frame. If the cursor is not in a frame, x represents the position on the line and y is undefined.

**PAD\_\$INQ\_VIEW**

Returns information about the position of a window relative to a pad.

**FORMAT**

PAD\_\$INQ\_VIEW (stream-id, window-number, line, eof-linenum, x-offset,  
y-offset, status)

**INPUT PARAMETERS****stream-id**

Number of the stream associated with an input or transcript pad, in STREAM\_\$ID\_T format. This is a 2-byte integer.

**window-number**

Index into the window list returned by PAD\_\$INQ\_WINDOWS. This is a 2-byte integer. Window number one always refers to the first window created to view the pad.

**OUTPUT PARAMETERS****line**

Number of the line being viewed. This is a 4-byte integer.

**eof-linenum**

Last line or frame on the pad. This is a 4-byte integer.

**x-offset**

Distance the pad is horizontally scrolled. This is a 2-byte integer.

**y-offset**

Distance the pad is vertically scrolled. This is a 2-byte integer. Only frames can be vertically scrolled.

**status**

Completion status, in STATUS\_\$T format. This data type is 4 bytes long. See the PAD Data Types section for more information.

**USAGE**

Use this routine in conjunction with PAD\_\$SET\_VIEW to control the display of graphic images that are larger than the window. PAD\_\$INQ\_VIEW describes the pad element currently being viewed through the given window, usually a transcript pad element.

If the element currently in view is a frame, x-offset and y-offset describe how the window is positioned in relation to the frame. If you are viewing the current frame and not some previous part of the pad, the value of eof-linenum will be equal to the line parameter.

If the element currently in view is not a frame, the line parameter is the number of the top line in the window.

PAD\_\$INQ\_WINDOWS

PAD\_\$INQ\_WINDOWS

Returns information about windows viewing the current pad.

### FORMAT

PAD\_\$INQ\_WINDOWS (stream-id, windowlist, window-list-size,  
window-no, status)

### INPUT PARAMETERS

#### stream-id

Number of the stream on which the pad is open, in STREAM\_\$ID\_T format. This is a 2-byte integer.

#### window-list-size

Maximum number of windows on which information is desired. This is a 2-byte integer.

### OUTPUT PARAMETERS

#### windowlist

Information describing a window, in PAD\_\$WINDOW\_LIST\_T format. This data type is an array of up to 10 elements, each of which is in PAD\_\$WINDOW\_DESC\_T format (four 2-byte integers). See the PAD Data Types section for more information.

Windowlist indicates the top left corner and the width and height of each window open on the pad, up to wlistsize. The values appear in the following order: top, left, width, height. Top and left are expressed in raster units, but width and height are divided by the current scale factors.

#### window-no

The number of windows open on the pad. This is a 2-byte integer. Window number one always refers to the first window created to view the pad. Use this parameter in calls that require a window number.

#### status

Completion status, in STATUS\_\$T format. This data type is 4 bytes long. See the PAD Data Types section for more information.

**PAD\_\$IS\_ICON**

Checks if a window is in icon form.

**FORMAT**

is-icon = PAD\_\$IS\_ICON (stream-id, window-no, status)

**RETURN VALUE****is-icon**

A boolean value indicating whether or not the window associated with the stream is in icon form.

**INPUT PARAMETERS****stream-id**

Number of the stream on which the input pad is open, in STREAM\_\$ID\_T format. This is a 2-byte integer.

**window-number**

Index into the window list returned by PAD\_\$INQ\_WINDOWS. This is a 2-byte integer. Window number one always refers to the first window created to view the pad.

**OUTPUT PARAMETERS****status**

Completion status, in STATUS\_\$T format. This data type is 4 bytes long. See the PAD Data Types section for more information.

**USAGE**

The function returns TRUE if the window is an icon and FALSE if it is not.

PAD\_\$LOAD\_FONT

PAD\_\$LOAD\_FONT

Loads a character font.

## FORMAT

PAD\_\$LOAD\_FONT (stream-id, font-pathname, name-length, font-id, status)

## INPUT PARAMETERS

### stream-id

Number of the stream on which the pad is open, in STREAM\_\$ID\_T format. This is a 2-byte integer.

### font-pathname

Pathname of the file containing the character font, in PAD\_\$STRING\_T format. This is an array of up to 256 characters.

### name-length

Length of the pathname. This is a 2-byte integer.

## OUTPUT PARAMETERS

### font-id

Font identifier, to be used in later calls to PAD\_\$USE\_FONT. This is a 2-byte integer.

### status

Completion status, in STATUS\_\$T format. This data type is 4 bytes long. See the PAD Data Types section for more information.

## USAGE

Your program can use up to 100 different character fonts as long as you initially load all the fonts that you intend to use with PAD\_\$LOAD\_FONT. When you want your program to use a specific font, call PAD\_\$USE\_FONT to invoke a previously loaded font. Each time you want to change a loaded font, use PAD\_\$USE\_FONT. To determine which font your program is currently using, call PAD\_\$INQ\_FONT.

The Display Manager first attempts to find the font file using the pathname directly, with the normal defaults. If it fails to find the file, it searches in /SYS/DM/FONTS.

PAD\_\$LOAD\_FONT does not switch fonts. It merely loads the font into the invisible portion of display memory and returns a font ID. After loading the font, your program can call PAD\_\$USE\_FONT to use it.

You can load up to 100 fonts in a given pad.

**PAD\_\$LOCATE**

Returns the position of the keyboard cursor in response to a keystroke.

**FORMAT**

PAD\_\$LOCATE (stream-id, x, y, character, status)

**INPUT PARAMETERS****stream-id**

Number of the stream on which the pad is open, in STREAM\_\$ID\_T format. This is a 2-byte integer.

**OUTPUT PARAMETERS****x**

X position of the input cursor. This is a 2-byte integer.

**y**

Y position of the input cursor. This is a 2-byte integer.

**character**

Value of the key pressed.

**status**

Completion status, in STATUS\_\$T format. This data type is 4 bytes long. See the PAD Data Types section for more information.

**USAGE**

This procedure returns the cursor position only when a raw character is entered. If the pad is in raw mode, any keystroke will do. In cooked mode, the ER command must be used. This command is usually entered through a function key definition.

The keyboard cursor position must be within the transcript pad.

X and y are divided by the current scale factors.

PAD\_\$MAKE\_ICON

PAD\_\$MAKE\_ICON

Changes an existing window into icon format.

## FORMAT

PAD\_\$MAKE\_ICON (stream-id, window-no, icon-char, status)

## INPUT PARAMETERS

### stream-id

Number of the stream on which the pad is open, in STREAM\_\$ID\_T format. This is a 2-byte integer.

### window-no

Index into the window list returned by PAD\_\$INQ\_WINDOWS. This is a 2-byte integer. Window number one always refers to the first window created to view the pad.

### icon-char

Icon font character to be displayed in the icon window. This character must reside in the current icon font. A 0 (zero) causes the Display Manager to select the default icon character for this pad type.

## OUTPUT PARAMETERS

### status

Completion status, in STATUS\_\$T format. This data type is 4 bytes long. See the PAD Data Types section for more information.

## USAGE

This call changes an existing full-size window into icon format. (If the window is invisible at the time of the call, it first becomes visible and then becomes an icon.) To create a completely new pad and window in icon format, use PAD\_\$CREATE\_ICON.

If the window is already an icon, this call has no effect.

Specify the display position for the new icon using the PAD\_\$SET\_ICON\_POS routine before executing this call. If you do not do this, the Display Manager assigns a default icon position descriptor and font character.

The size of the icon window is not user-definable. It is determined automatically by the size of the font character specified.



**PAD\_\$MAKE\_INVISIBLE**

Makes a visible window invisible.

**FORMAT**

PAD\_\$MAKE\_INVISIBLE (stream-id, window-no, status)

**INPUT PARAMETERS****stream-id**

Number of the stream on which the pad is open, in STREAM\_\$ID\_T format. This is a 2-byte integer.

**window-no**

Index into the window list returned by PAD\_\$INQ\_WINDOWS. This is a 2-byte integer. Window number one always refers to the first window created to view the pad.

**OUTPUT PARAMETERS****status**

Completion status, in STATUS\_\$T format. This data type is 4 bytes long. See the PAD Data Types section for more information.

**USAGE**

The effect of this call is the same as if the window were completely obscured by other windows on the screen, except that no amount of pushing, popping, moving, or growing can make it reappear. Only a subsequent call to PAD\_\$SELECT\_WINDOW will restore it to visibility in its full-size format.

If the window is currently invisible, this call has no effect.

If the window is currently in icon format, it will first be made into a full-size window and then turned invisible.

PAD\_\$MOVE

PAD\_\$MOVE

Moves the output cursor.

## FORMAT

PAD\_\$MOVE (stream-id, rel-abs, x, y, status)

## INPUT PARAMETERS

### stream-id

Number of the stream on which the pad is open, in STREAM\_\$ID\_T format. This is a 2-byte integer.

### rel-abs

Indicates whether cursor movement is to be relative or absolute. This is a 2-byte integer in PAD\_\$REL\_ABS\_T format. Specify one of the following predefined values:

#### PAD\_\$RELATIVE

Movement is relative to the last cursor position. X and y denote positive or negative offsets to the current cursor position, scaled according to the current scale factors.

#### PAD\_\$ABSOLUTE

X and y are absolute, within the frame. X and y must be positive. Within a frame, movement is relative to the upper left corner of the frame. Outside a frame, x is relative to the left end of the current line and y is not used. In both cases, x and y are scaled according to the current scale factors.

### x

Change to the x-coordinate of the cursor position. This is a 2-byte signed integer.

### y

Change to the y-coordinate of the cursor position. This is a 2-byte signed integer.

## OUTPUT PARAMETERS

### status

Completion status, in STATUS\_\$T format. This data type is 4 bytes long. See the PAD Data Types section for more information.

## USAGE

PAD\_\$MOVE changes the position of the output cursor, which marks the place where the next program output will appear.

The cursor can move vertically only within a frame, not on a line. The Display Manager uses the y value only when a frame is active, and ignores it otherwise.

**PAD\_\$POP\_PUSH\_WINDOW**

Pops or pushes a window.

**FORMAT**

PAD\_\$POP\_PUSH\_WINDOW(stream-id, window-no, flag, status)

**INPUT PARAMETERS****stream-id**

Number of the stream on which the pad is open, in STREAM\_\$ID\_T format. This is a 2-byte integer.

**window-no**

The index into the window list returned by PAD\_\$INQ\_WINDOWS. This is a 2-byte integer. Window number one always refers to the first window created to view the pad.

**flag**

Indicates if the window is to be pushed or popped. This is a Boolean variable. A value of TRUE pops the specified window to the top of the screen, ensuring that no portion of the window is hidden by another window. A value of FALSE pushes the specified window to the bottom of the screen, allowing other windows to cover it wherever possible.

**OUTPUT PARAMETERS****status**

Completion status, in STATUS\_\$T format. This data type is 4 bytes long. See the PAD Data Types section for more information.

PAD\_ \$RAW

PAD\_ \$RAW

Places an input or transcript pad in raw mode.

## FORMAT

PAD\_ \$RAW (stream-id, status)

## INPUT PARAMETERS

### stream-id

Number of the stream on which the pad is open, in STREAM\_ \$ID\_ T format. This is a 2-byte integer. The stream-id given should refer to an input stream, usually standard input (STREAM\_ \$STDIN). PAD\_ \$RAW has no effect on output.

## OUTPUT PARAMETERS

### status

Completion status, in STATUS\_ \$T format. This data type is 4 bytes long. See the PAD Data Types section for more information.

## USAGE

PAD\_ \$RAW puts the pad into raw mode, and has no effect if it is already in raw mode. PAD\_ \$COOKED returns the pad to normal processing.

In raw mode, the Display Manager sends keyboard input directly to the program without echoing or processing it in any way. ASCII control characters are also sent to the program, but the Display Manager still handles its function keys.

The Display Manager immediately displays every character it receives, unless the window is in HOLD mode. If the window is in HOLD mode, new characters do not appear until the keyboard user scrolls the window or releases HOLD.

When it executes this procedure, the Display Manager clears the input pad and shrinks its window size to zero. The keyboard cursor moves to the current output cursor position in the transcript pad. While the pad is in raw mode, the keyboard and output cursors usually move together.

**NOTE:** A program using PAD\_ \$RAW *must* execute PAD\_ \$COOKED before termination. Most system utilities, including the Shell, will not work correctly in raw mode.

**PAD\_\$SELECT\_WINDOW**

Makes an invisible window visible and/or changes an icon-format window into a full-sized window.

**FORMAT**

PAD\_\$SELECT\_WINDOW (stream-id, window-no, status)

**INPUT PARAMETERS****stream-id**

Number of the stream on which the pad is open, in STREAM\_\$ID\_T format. This is a 2-byte integer.

**window-no**

Index into the window list returned by PAD\_\$INQ\_WINDOWS. This is a 2-byte integer. Window number one always refers to the first window created to view the pad.

**OUTPUT PARAMETERS****status**

Completion status, in STATUS\_\$T format. This data type is 4 bytes long. See the PAD Data Types section for more information.

**USAGE**

Use PAD\_\$MAKE\_INVISIBLE to make windows invisible.

If this call is used to expand an icon to full-size format, the position and dimensions of the large window are the same as those it had when it was last full size. If it was never full-size, its position and dimensions are those specified (or defaulted) when the icon was created (either by PAD\_\$CREATE\_ICON, or by the Display Manager commands CP, CV, CE, or CPB with the -I option specified).

PAD\_\$SET\_AUTO\_CLOSE

PAD\_\$SET\_AUTO\_CLOSE

Sets a window to close automatically when its pad closes.

## FORMAT

PAD\_\$SET\_AUTO\_CLOSE (stream-id, window-no, auto-close, status)

## INPUT PARAMETERS

### stream-id

Number of the stream on which the pad is open, in STREAM\_\$ID\_T format. This is a 2-byte integer.

### window-no

Index into the window list returned by PAD\_\$INQ\_WINDOWS. This is a 2-byte integer. Window number one always refers to the first window created to view the pad.

### auto-close

Indicates whether the window is to close automatically. This is a Boolean value. If TRUE, the window disappears when the pad onto which it opens is closed.

## OUTPUT PARAMETERS

### status

Completion status, in STATUS\_\$T format. This data type is 4 bytes long. See the PAD Data Types section for more information.

## USAGE

When a program first makes this call and then does a STREAM\_\$CLOSE, the window specified is closed and deleted from the screen. This is equivalent to specifying the Display Manager command WC -A for a window.

**PAD\_\$SET\_BORDER**

Adds a border to, or removes the border from, a full window.

**FORMAT**

PAD\_\$SET\_BORDER(stream-id, window-number, flag, status)

**INPUT PARAMETERS****stream-id**

Number of the stream on which the pad is open, in STREAM\_\$ID\_T format. This is a 2-byte integer.

**window-number**

Index into the window list returned by PAD\_\$INQ\_WINDOWS. This is a 2-byte integer. Window number one always refers to the first window created to view the pad.

**flag**

Indicates whether to add or remove a border. This is a Boolean variable. If TRUE, the window appears with a border around its edges and a legend at the top. If FALSE, any border and legend are removed from the window, making the window's usable area equal to the amount of space the window occupies on the screen.

**OUTPUT PARAMETERS****status**

Completion status, in STATUS\_\$T format. This data type is 4 bytes long. See the PAD Data Types section for more information.

**USAGE**

Use this procedure to remove or add a border to a full window that has no other panes associated with it. If you add a pane to a window from which the border was removed, the border will be redrawn. Likewise, if you change the input pad mode to cooked, the border will be redrawn.

To get a full window without any panes, you can either create a transcript pad and never make a PAD\_\$CREATE call to add panes, or create a transcript pane that covers the entire window. Another way to get a full window is to make an input pane invisible by using the PAD\_\$RAW call.

PAD\_\$SET\_FULL\_WINDOW

PAD\_\$SET\_FULL\_WINDOW

Moves a window or sets a window position for future use.

## FORMAT

PAD\_\$SET\_FULL\_WINDOW (stream-id, window-no, window, status)

## INPUT PARAMETERS

### stream-id

Number of the stream on which the pad is open, in STREAM\_\$ID\_T format. This is a 2-byte integer.

### window-no

Index into the window list returned by PAD\_\$INQ\_WINDOWS. This is a 2-byte integer. Window number one always refers to the first window created to view the pad.

### window

Window descriptor giving the position on the screen that the new window will occupy when expanded to a full-sized window, in PAD\_\$WINDOW\_DESC\_T format. This data type is 8 bytes long. See the PAD\_\$ Data Types section for more information.

The window specified is the entire window, including the border, legend, and usable part of the window. The call, PAD\_\$INQ\_FULL\_WINDOW returns information about the entire window.

## OUTPUT PARAMETERS

### status

Completion status, in STATUS\_\$T format. This data type is 4 bytes long. See the PAD Data Types section for more information.

## USAGE

If the window specified is currently in icon format or is invisible, this call establishes a full-size window position for future use (i.e., when your program calls PAD\_\$SELECT\_WINDOW to expand the icon into a full-size window, or you issue the Display Manager commands ICON or WI).

If the window specified is currently full-size, then the window is repositioned to the location given by window.



## PAD\_\$SET\_ICON\_FONT

Sets the current icon font to a specified font name.

**FORMAT**

PAD\_\$SET\_ICON\_FONT (stream-id, window-no, font-name,  
font-len, status)

**INPUT PARAMETERS****stream-id**

Number of the stream on which the pad is open, in STREAM\_\$ID\_T format. This is a 2-byte integer.

**window-no**

Window whose icon font you want to change. Window-no is an index into the window list returned by PAD\_\$INQ\_WINDOWS. This is a 2-byte integer. Window number one always refers to the first window created to view the pad.

**font-name**

Full pathname of the font, up to the node entry directory (/), in NAME\_\$PNAME\_T format. This is an array of up to 256 characters.

**font-len**

Length of the font file pathname. This is a 2-byte integer.

**OUTPUT PARAMETERS****status**

Completion status, in STATUS\_\$T format. This data type is 4 bytes long. See the PAD Data Types section for more information.

**USAGE**

Use this call to change the icon font in use. This call changes the font of the specified window only.

When a window is created either as a full window or an icon, the Display Manager assigns it an icon from "active icon font." The default active icon font is in /SYS/DM/FONTS/ICONS. You can specify another font to be the active icon font by using the FL command with the -I option.

You can create a new icon font file to contain your own icons by using the font editor EDFONT. See the *DOMAIN System Command Reference* for a complete description of EDFONT.

If the window is in icon format at the time of this call, the icon in the display changes to the new font immediately.

PAD\_\$SET\_ICON\_POS

PAD\_\$SET\_ICON\_POS

Moves an icon or sets an icon position for future use.

### FORMAT

PAD\_\$SET\_ICON\_POS (stream-id, window-no, icon-pos, icon-char, status)

### INPUT PARAMETERS

#### stream-id

Number of the stream on which the pad is open, in STREAM\_\$ID\_T format. This is a 2-byte integer.

#### window-no

Index into the window list returned by PAD\_\$INQ\_WINDOWS. This is a 2-byte integer. Window number one always refers to the original transcript pad.

#### icon-pos

New position (x and y coordinates) of the icon, in PAD\_\$POSITION\_T format. This data type is 4 bytes long. See the PAD Data Types section for more information.

#### icon-char

Character to be displayed in the icon window.

### OUTPUT PARAMETERS

#### status

Completion status, in STATUS\_\$T format. This data type is 4 bytes long. See the PAD Data Types section for more information.

### USAGE

If the window specified is currently in full-size format, this call establishes an icon position for future use (i.e., when your program calls PAD\_\$MAKE\_ICON to turn the window into icon format, or you use the Display Manager command ICON).

If the window specified is already in icon format, then the icon is repositioned to the location given by icon-pos, and the specified icon-char replaces the current one.

Compare this call to PAD\_\$SET\_FULL\_WINDOW, which performs the same operations for full-size windows.

The size of the icon window is not user-definable. It is determined automatically by the size of the font character specified.

**PAD\_\$SET\_SCALE**

Sets a scale factor for cursor operations.

**FORMAT**

PAD\_\$SET\_SCALE (stream-id, x-factor, y-factor, status)

**INPUT PARAMETERS****stream-id**

Number of the stream on which the pad is open, in STREAM\_\$ID\_T format. This is a 2-byte integer.

**x-factor**

Scale factor for the x-coordinate. This is a 2-byte integer.

**y-factor**

Scale factor for the y-coordinate. This is a 2-byte integer.

**OUTPUT PARAMETERS****status**

Completion status, in STATUS\_\$T format. This data type is 4 bytes long. See the PAD Data Types section for more information.

**USAGE**

Specify a scale factor of zero to use the scale of the current character font. This is the default.

Specify a nonzero scale factor to use that number as a multiplier for raster units. One raster unit is equal to one bit in the display.

The scale factor is used to convert between raster units and numbers supplied in routines such as PAD\_\$MOVE. When using the scale of the current font, you express dimensions in terms of characters and lines. In any case, the numbers you enter are multiplied by the scale factor to yield raster units, and raster units are divided by the scale factor before being returned.

The scale factor is used to process input or output for PAD\_\$CPR\_ENABLE, PAD\_\$CREATE\_FRAME, PAD\_\$INQ\_POSITION, PAD\_\$LOCATE, PAD\_\$MOVE, and PAD\_\$INQ\_WINDOWS. In PAD\_\$INQ\_WINDOWS, height and width are scaled, but top and left are not. PAD\_\$INQ\_FONT always returns dimensions in terms of raster units.

The scale factors set with this call apply to the specified stream until specifically reset, even after the calling program ends. Your program should not depend on the scale factors being correctly set, but should call PAD\_\$SET\_SCALE to explicitly set the scale factors as desired.

PAD\_\$SET\_TABS

PAD\_\$SET\_TABS

Sets tab stops within a pad.

## FORMAT

PAD\_\$SET\_TABS (stream-id, tab-stop-array, no-of-tabs, status)

## INPUT PARAMETERS

### stream-id

Number of the stream on which the pad is open, in STREAM\_\$ID\_T format. This is a 2-byte integer.

### tab-stop-array

Columns for tab stops. This is an array of up to 100 2-byte integers. Each element in the array contains a column number at which a tab stop will be set. Column numbers are scaled according to the PAD\_\$SET\_SCALE procedure.

For example, assume that the current vertical and horizontal scale factors are both equal to one. A three-element array containing the integers 100, 300, and 500 would specify tab stops at bit positions 100, 300, and 500 on the screen. Because the display contains approximately 100 bits per inch, these tab stops would be set about 1, 3, and 5 inches (2.54, 7.62, and 12.70 cm) from the left edge of the screen.

### no-of-tabs

Number of tab stops set. This is a 2-byte integer.

## OUTPUT PARAMETERS

### status

Completion status, in STATUS\_\$T format. This data type is 4 bytes long. See the PAD Data Types section for more information.

## USAGE

This procedure sets tabs only for the pad open on the specified stream. Tab stops for all other pads are unchanged.

The default tab setting has tabs every 4 columns.

**PAD\_\$SET\_VIEW**

Positions a window to establish a given view.

**FORMAT**

PAD\_\$SET\_VIEW (stream-id, window-no, line, x-offset,  
y-offset, status)

**INPUT PARAMETERS****stream-id**

Number of the stream associated with a transcript pad, in STREAM\_\$ID\_T format. This is a 2-byte integer.

**window-no**

Index into the window list returned by PAD\_\$INQ\_WINDOWS. This is a 2-byte integer. Window number one always refers to the first window created to view the pad.

**line**

Line number to view. This is a 4-byte integer.

**x-offset**

Distance to scroll the pad horizontally. This is a 2-byte integer.

**y-offset**

Distance to scroll the pad vertically (for frames only). This is a 2-byte integer.

**OUTPUT PARAMETERS****status**

Completion status, in STATUS\_\$T format. This data type is 4 bytes long. See the PAD Data Types section for more information.

**USAGE**

This routine repositions a window to establish a particular view of a transcript pad. Programs can call this routine after a call to PAD\_\$INQ\_VIEW and in conjunction with calls to PAD\_\$INQ\_WINDOWS to control the display of graphic images that are larger than the window.

PAD\_\$USE\_FONT

PAD\_\$USE\_FONT

Invokes a loaded font.

## FORMAT

PAD\_\$USE\_FONT (stream-id, font-id, status)

## INPUT PARAMETERS

### stream-id

Number of the stream on which the pad is open, in STREAM\_\$ID\_T format. This is a 2-byte integer.

### font-id

Font identifier returned by PAD\_\$LOAD\_FONT. This is a 2-byte integer.

## OUTPUT PARAMETERS

### status

Completion status, in STATUS\_\$T format. This data type is 4 bytes long. See the PAD Data Types section for more information.

## USAGE

Use this call to change your program's current character font.

Your program can use up to 100 different character fonts as long as you initially load all the fonts that you intend to use with PAD\_\$LOAD\_FONT. When you want your program to use a specific font, call PAD\_\$USE\_FONT to invoke a previously loaded font. Each time you want to change a loaded font, use PAD\_\$USE\_FONT. To determine which font your program is currently using, call PAD\_\$INQ\_FONT.

NOTE: Use PAD\_\$USE\_FONT only to *change* the current font in use. You will get erroneous results if the call specifies the font that is already currently in use.

**ERRORS**

STATUS\_ \$OK  
Successful completion.

PAD\_ \$2MNY\_ CLIENTS  
Operation illegal with more than one client process.

PAD\_ \$2MNY\_ INPUT\_ PADS  
Only one input pad per transcript.

PAD\_ \$BAD\_ KEY\_ NAME  
Key name not found.

PAD\_ \$EDIT\_ QUIT  
User quit (WC -Q) out of edit pane.

PAD\_ \$FONT\_ FILE\_ ERR  
Could not access font file.

PAD\_ \$ID\_ OOR  
Stream id out of range.

PAD\_ \$ILL\_ PARAM\_ COMB  
Conflict in PAD\_ \$CREATE call.

PAD\_ \$ILL\_ PTYPE  
Cannot do operation on this type of pad.

PAD\_ \$NO\_ SUCH\_ WINDOW  
Bad window number in INQ/SET\_ VIEW.

PAD\_ \$NO\_ WINDOW  
Window no longer exists.

PAD\_ \$NOT\_ ASCII  
Existing pad in PAD\_ \$CREATE is not ASCII.

PAD\_ \$NOT\_ INPUT  
Operation valid on input pads only.

PAD\_ \$NOT\_ RAW  
Operation requires pad be in raw mode.

PAD\_ \$NOT\_ TRANSCRIPT  
Operation valid on transcript pads only.

PAD\_ \$STREAM\_ NOT\_ OPEN  
No stream open on this SID.

PAD\_ \$STREAM\_ NOT\_ PAD  
Preferred stream is not a pad.

PAD\_ \$TOO\_ MANY\_ FONTS  
Too many fonts loaded in this pad.

PAD\_ \$VOOR  
Value out of range.





# PBUFS

## Table of Contents

PBUFS\_\$CREATE  
PBUFS\_\$OPEN  
Errors

PBUFS-2  
PBUFS-3  
PBUFS-4



## PBUFS

The PBUFS (Paste Buffer) programming calls maintain cut-and-paste buffers for data interchange between applications and the Display Manager. This section describes their error messages and call syntax. The PBUFS calls do not use unique data types. Refer to the Introduction at the beginning of this manual for a description of call syntax format.

PBUFS\_\$CREATE

PBUFS\_\$CREATE

Creates a paste buffer.

## FORMAT

PBUFS\_\$CREATE (buffer-name, type, stream-id, status)

## INPUT PARAMETERS

### buffer-name

Name of the paste buffer you want to create (not a pathname), in NAME\_\$NAME\_T format. This is an array of up to 32 characters. This array must be a full 32 bytes, padded with blanks. See the NAME Data Types section for more information.

### type

Indicates whether the paste buffer is to hold text or pictures. This is a Boolean value. TRUE designates a text buffer. FALSE designates a GMF buffer that can hold images.

## OUTPUT PARAMETERS

### stream-id

Number of a stream with which to refer to the new paste buffer, in STREAM\_\$ID\_T format. This is a 2-byte integer.

### status

Completion status, in STATUS\_\$T format. This data type is 4 bytes long. See the PAD Data Types section for more information.

## USAGE

This call creates a paste buffer of the specified name and type in the directory 'node\_data/paste\_buffers'. An error occurs if the named paste buffer already exists in /sys/node\_data/paste\_buffers.

The file has the temporary attribute, STREAM\_\$IRM\_TEMPORARY. The system will delete this file when you close the stream, unless you call STREAM\_\$REDEFINE to change the file's attributes first.

Calling PBUFS\_\$CREATE opens the stream for overwrite access (STREAM\_\$OVERWRITE).

You can call STREAM\_\$CREATE, specifying a pathname in /sys/node\_data/paste\_buffers to achieve the same effect.

**PBUFS\_\$OPEN**

Opens a pre-existing paste buffer.

**FORMAT**

PBUFS\_\$OPEN (buffer-name, type, stream-id, status)

**INPUT PARAMETERS****buffer-name**

Name of the paste buffer you want to open (not a pathname), in NAME\_\$NAME\_T format. This is an array of up to 32 characters. This array must be a full 32 bytes, padded with blanks. See the NAME Data Types section for more information.

**type**

Indicates whether the paste buffer is to hold text or pictures. This is a Boolean value. TRUE designates a text buffer. FALSE designates a GMF buffer that can hold images. The value you specify must match the value used when creating the paste buffer, or the paste buffer manager returns the completion status PBUFS\_\$WRONG\_TYPE.

**OUTPUT PARAMETERS****stream-id**

Number of the stream with which to refer to the paste buffer, in STREAM\_\$ID\_T format. This is a 2-byte integer.

**status**

Completion status, in STATUS\_\$T format. This data type is 4 bytes long. See the PAD Data Types section for more information.

**USAGE**

This call open a pre-existing paste buffer of the specified name and type in the directory 'node\_data/paste\_buffers'.

An error occurs if the named paste buffer does not already exist in /sys/node\_data/paste\_buffers. Use PBUFS\_\$CREATE to create a buffer.

You can call STREAM\_\$OPEN, on a file in /sys/node\_data/paste\_buffers, to achieve the same effect.

PBUFS ERRORS

**ERRORS**

PBUFS\_ \$WRONG\_ TYPE

The actual buffer type differs from the type specified.

# PFM

## Table of Contents

Data Types	PFM-2
PFM_\$CLEANUP	PFM-4
PFM_\$ENABLE	PFM-5
PFM_\$ERROR_TRAP	PFM-6
PFM_\$ESTABLISH_FAULT_HANDLER	PFM-7
PFM_\$INHIBIT	PFM-9
PFM_\$RELEASE_FAULT_HANDLER	PFM-10
PFM_\$RESET_CLEANUP	PFM-11
PFM_\$RLS_CLEANUP	PFM-12
PFM_\$SIGNAL	PFM-13
Errors	PFM-14





## PFM

The PFM (Process Fault Manager) programming calls control signals, faults, and exceptions for faults. This section describes their data types, call syntax, and error codes. Refer to the Introduction at the beginning of this manual for a description of data-type diagrams and call syntax format.

PFM DATA TYPES

CONSTANTS


PFM\_\$ALL\_FAULTS                    0                    Specified when establishing a handler to catch all faults.

DATA TYPES

PFM\_\$CLEANUP\_REC                    Cleanup routine information, 64 bytes long.

PFM\_\$FAULT\_FUNC\_P\_T                A 4-byte integer. A pointer to a fault handler function.

PFM\_\$FAULT\_REC\_T                    Parameter to fault handler function. The diagram below illustrates the PFM\_\$FAULT\_REC\_T data type:



Field Description:

pattern  
Reserved for PFM use.

status  
The returned status in STATUS\_\$T format.

PFM\_\$FH\_FUNC\_VAL\_T                A 2-byte integer. Specifies action to be taken when handler completes. One of the following predefined values:

PFM\_\$CONTINUE\_FAULT\_HANDLING  
Specifies that the fault be passed to next handler.

PFM\_\$RETURN\_TO\_FAULTING\_CODE  
Specifies that control be returned to the program.

PFM\_\$FH\_HANDLE\_T                    A 4-byte integer. Pointer to a fault handler.


PFM\_\$FH\_OPT\_SET\_T                    A 2-byte integer. Options for type of handler to establish. Any combination of the following predefined values:

PFM\_\$FH\_BACKSTOP  
 specifies that the handler should be called  
 after all other handlers.

PFM\_\$FH\_MULTI\_LEVEL  
 Specifies that handler applies to faults on its  
 program level, and all subordinate levels.

STATUS\_\$T

A status code. The diagram below illustrates the  
 STATUS\_\$T data type:



Field Description:

all  
 All 32 bits in the status code.

fail  
 The fail bit. If this bit is set, the error was not  
 within the scope of the module invoked, but  
 occurred within a lower-level module (bit 31).

subsys  
 The subsystem that encountered the error (bits  
 24 - 30).

modc  
 The module that encountered the error (bits 16 -  
 23).

code  
 A signed number that identifies the type of error  
 that occurred (bits 0 - 15).

PFM\_\$CLEANUP

PFM\_\$CLEANUP

Establishes a clean-up handler for faults.

## FORMAT

status = PFM\_\$CLEANUP (clean-up-record)

## RETURN VALUE

### status

Completion status, in STATUS\_\$T format. This data type is 4 bytes long. See the PFM Data Types section for more information.

When initially called to establish a clean-up handler, PFM\_\$CLEANUP returns the status PFM\_\$CLEANUP\_SET. After a fault occurs, PFM\_\$CLEANUP returns the status of the fault, or the status signaled by PFM\_\$SIGNAL or PFM\_\$ERROR\_TRAP.

## OUTPUT PARAMETERS

### clean-up-record

A record uniquely identifying the clean-up handler, in PFM\_\$CLEANUP\_REC format. This data type is 8 bytes long. See the PFM Data Types section for more information.

This parameter is passed as input to the PFM\_\$RLS\_CLEANUP and PFM\_\$RESET\_CLEANUP procedures in order to specify a particular handler. Your program cannot modify or copy this value.

## USAGE

PFM\_\$CLEANUP establishes a clean-up handler that is executed when a fault occurs. Clean-up handlers let the program "clean up" a task, possibly notifying you of the error condition and leaving any open files in a known and stable state.

You may establish more than one clean-up handler. Multiple cleanup handlers are executed consecutively, starting with the most recently established handler and continuing backward in time (LIFO). A built-in clean-up handler is always established when you invoke your program. This built-in handler is always called last. It closes any files that are still open and returns control to the invoking Shell.

The initial call to PFM\_\$CLEANUP establishes the clean-up handler and returns a status value of PFM\_\$CLEANUP\_SET. When a fault occurs, execution returns to the most recent PFM\_\$CLEANUP call. The clean-up handler associated with that call is then removed from the stack and executed.

**PFM\_ \$ENABLE**

Enables asynchronous faults.

**FORMAT**

PFM\_ \$ENABLE

**USAGE**

PFM\_ \$ENABLE enables asynchronous faults after they have been inhibited by a call to PFM\_ \$INHIBIT. PFM\_ \$ENABLE causes the operating system to pass asynchronous faults on to the program.

While faults are inhibited, the operating system holds at most one asynchronous fault. So, as soon as a PFM\_ \$ENABLE executes, your program receives one asynchronous fault. If more than one fault occurred while faults were inhibited, the program receives the *first* asynchronous fault.

Since a user cannot terminate a program while PFM\_ \$INHIBIT is in effect, it is good programming practice to inhibit asynchronous faults only during critical intervals, or enable faults occasionally to allow users to exit.

PFM\_\$ERROR\_TRAP

PFM\_\$ERROR\_TRAP

Simulates a fault with a given status code, storing traceback information.

### FORMAT

PFM\_\$ERROR\_TRAP (status)

### INPUT PARAMETERS

#### status

Error code, in STATUS\_\$T format. This data type is 4 bytes long. See the PFM Data Types section for more information.

### USAGE

Use this procedure to force an error exit with the specified status code, or in a fatal error situation where no status code otherwise returns. One possible use is in defining your own error condition.

This procedure differs from PFM\_\$SIGNAL in that traceback information is stored, so that it is possible to determine where the fault occurred.

**PFM\_\$ESTABLISH\_FAULT\_HANDLER**

Establishes a fault handler.

**FORMAT**

handler-id = PFM\_\$ESTABLISH\_FAULT\_HANDLER (target-status, options,  
function-pointer, status)

**RETURN VALUE****handler-id**

A value uniquely identifying the established handler, in PFM\_\$FH\_HANDLE\_T format. This is a 4-byte integer.

You pass this value to the PFM\_\$RELEASE\_FAULT\_HANDLER call when you want to release the handler.

**INPUT PARAMETERS****target-status**

A value specifying the type of fault that causes this handler to take effect. This is a 4-byte integer.

To establish a fault handler for all faults produced by a certain DOMAIN module, use any error status code returned by that module, with the fault code field set to 0. To establish a fault handler that handles all faults, use the constant PFM\_\$ALL\_FAULTS.

**options**

A value specifying the type of handler you want to establish, in PFM\_\$FH\_OPT\_SET\_T format. This is a 2-byte integer. Specify any combination of the following set of predefined values:

**PFM\_\$FH\_MULTI\_LEVEL**

To declare a multilevel fault handler that handles faults for its own program level and all subordinate levels.

**PFM\_\$FH\_BACKSTOP**

To establish a backstop fault handler that takes effect after all nonbackstop handlers have taken effect.

(In FORTRAN, you can combine these options by adding the constants.)

**function-pointer**

The address of the fault handler for the specified type(s) of faults, in PFM\_\$FAULT\_FUNC\_P\_T format. This is a 4-byte integer.

**OUTPUT PARAMETERS****status**

Completion status, in STATUS\_\$T format. This data type is 4 bytes long. See the PFM Data Types section for more information.

PFM\_\$ESTABLISH\_FAULT\_HANDLER

This call establishes a fault handler, making it take effect for all the faults of the specified type or types that occur after the time of the call.

The fault handler remains in effect until you release it using PFM\_\$RELEASE\_FAULT\_HANDLER or until the program ends.



**PFM\_ \$INHIBIT**

Inhibits asynchronous faults.

**FORMAT**

PFM\_ \$INHIBIT

**USAGE**

PFM\_ \$INHIBIT prevents asynchronous faults from being passed to the program. Use this call when an interval of your program cannot be interrupted, for example, when performing I/O. Use the complementary PFM\_ \$ENABLE call to re-enable asynchronous faults.

Asynchronous faults are produced from outside your program and are unrelated to anything within your program. They can occur at any point during your program's execution. A common example of an asynchronous fault is the Display Manager quit (DQ) command that occurs when someone types a CTRL/Q to stop a program.

Since a user cannot terminate a program while PFM\_ \$INHIBIT is in effect, it is good programming practice to inhibit asynchronous faults only during critical intervals.

While faults are inhibited, the operating system holds at most one asynchronous fault. So, as soon as a PFM\_ \$ENABLE executes, your program receives one asynchronous fault. If more than one fault occurred while faults were inhibited, the program receives the *first* asynchronous fault.

Inhibiting asynchronous faults has no effect on the processing of synchronous faults such as floating-point overflow errors, access violations, address errors, and so on.

PFM\_\$RELEASE\_FAULT\_HANDLER

PFM\_\$RELEASE\_FAULT\_HANDLER

Releases a fault handler.

### FORMAT

PFM\_\$RELEASE\_FAULT\_HANDLER (handler-id, status)

### INPUT PARAMETERS

#### handler-id

A value uniquely identifying the handler, in PFM\_\$FH\_HANDLE\_T format. This is a 4-byte integer.

A unique value is returned by PFM\_\$ESTABLISH\_FAULT\_HANDLER each time you establish a handler.

### OUTPUT PARAMETERS

#### status

Completion status, in STATUS\_\$T format. This data type is 4 bytes long. See the PFM Data Types section for more information.

### USAGE

This call causes the specified fault handler to cease having effect for faults that occur after the time of the call.

You should note that you cannot release a fault handler that is installed inside a fault handler.

To establish a fault handler, use the PFM\_\$ESTABLISH\_FAULT\_HANDLER call.

**PFM\_\$RESET\_CLEANUP**

Returns a clean-up handler to the top of the handler stack.

**FORMAT**

PFM\_\$RESET\_CLEANUP (clean-up-record, status)

**INPUT PARAMETERS****clean-up-record**

A record uniquely identifying the clean-up handler, in PFM\_\$CLEANUP\_REC format. This data type is 8 bytes long. See the PFM Data Types section for more information.

A unique record is returned by PFM\_\$CLEANUP each time a cleanup handler is established. The clean-up-record that is input must not have been altered or copied. If it has been, or if for some other reason the record is invalid, the procedure will fail with the status PFM\_\$INVALID\_CLEANUP\_REC.

**OUTPUT PARAMETERS****status**

Completion status, in STATUS\_\$T format. This data type is 4 bytes long. See the PFM Data Types section for more information.

**USAGE**

This procedure re-establishes the clean-up handler identified by the clean-up-record at the top of the stack, so that any subsequent errors invoke it first.

This procedure can only be used within a clean-up handler.

PFM\_\$RLS\_CLEANUP

PFM\_\$RLS\_CLEANUP

Releases a specified clean-up handler and any other clean-up handlers above it on the stack.

## FORMAT

PFM\_\$RLS\_CLEANUP (clean-up-record, status)

## INPUT PARAMETERS

### clean-up-record

A record uniquely identifying the clean-up handler, in PFM\_\$CLEANUP\_REC format. This data type is 8 bytes long. See the PFM Data Types section for more information.

A unique record is returned by PFM\_\$CLEANUP each time a clean-up handler is established. The clean-up-record that is input must not have been altered or copied. If it has been, or if for some other reason the record is invalid, the procedure will fail with the status PFM\_\$INVALID\_CLEANUP\_REC.

## OUTPUT PARAMETERS

### status

Completion status, in STATUS\_\$T format. This data type is 4 bytes long. See the PFM Data Types section for more information.

Possible status values are:

PFM\_\$INVALID\_CLEANUP\_REC

The clean-up-record has been altered or copied and is therefore invalid.

PFM\_\$BAD\_RLS\_ORDER

Program attempting to release a clean-up handler before releasing all handlers established after it. This status is only a warning; the handler is successfully released, and all handlers above it on the stack are also released.

## USAGE

PFM\_\$RLS\_CLEANUP releases the specified clean-up handler and all other clean-up handlers above it on the stack.

**PFM\_ \$SIGNAL**

Exits from the current procedure and signals a status for the clean-up handler on the top of the stack.

**FORMAT**

PFM\_ \$SIGNAL (status)

**INPUT PARAMETERS****status**

Status code, in STATUS\_ \$T format. This data type is 4 bytes long. See the PFM Data Types section for more information.

**USAGE**

PFM\_ \$SIGNAL can be called from within a clean-up handler or from normal code.

If invoked from within a clean-up handler, PFM\_ \$SIGNAL exits from the current clean-up handler and invokes the clean-up handler on the top of the stack, if there is one. If invoked from outside a clean-up handler, this routine invokes the top clean-up handler on the stack, with the status code given in the PFM\_ \$SIGNAL call.

Typically, PFM\_ \$SIGNAL is called at the end of one clean-up handler to invoke the next handler, and the status parameter is normally assigned the error status originally received from PFM\_ \$CLEANUP. When no more clean-up handlers from the current program are on the stack, PFM\_ \$SIGNAL causes the program to exit to the invoking program (which may be the Shell) with the status code set to the value given in the status parameter.

Traceback information (see the *DOMAIN System Command Reference*) is not stored when PFM\_ \$SIGNAL is called. When a fault occurs, however, the operating system automatically stores traceback information.

Unlike most subroutines, PFM\_ \$SIGNAL does not return to the place from which it was called.

PFM ERRORS

**ERRORS**

STATUS\_ \$OK

Successful completion.

PFM\_ \$BAD\_RLS\_ORDER

Cleanup handler released out of order.

PFM\_ \$CLEANUP\_NOT\_FOUND

Static cleanup handler not found.

PFM\_ \$CLEANUP\_SET

Cleanup handler established successfully.

PFM\_ \$CLEANUP\_SET\_SIGNALLED

PFM\_ \$CLEANUP\_SET was signalled.

PFM\_ \$FH\_NOT\_FOUND

Attempt to release non-existent fault handler.

PFM\_ \$FH\_WRONG\_LEVEL

Attempt to release fault handler at wrong level.

PFM\_ \$INVALID\_CLEANUP\_REC

Invalid clean-up record.

PFM\_ \$NO\_SPACE

No RWS space to create static clean-up handler.

# PGM

## Table of Contents

Data Types	PGM-2
PGM_\$DEL_ARG	PGM-7
PGM_\$EXIT	PGM-8
PGM_\$GET_ARG	PGM-9
PGM_\$GET_ARGS	PGM-11
PGM_\$GET_EC	PGM-12
PGM_\$GET_PUID	PGM-13
PGM_\$INVOKE	PGM-14
PGM_\$MAKE_ORPHAN	PGM-19
PGM_\$PROC_WAIT	PGM-20
PGM_\$SET_SEVERITY	PGM-21
Errors	PGM-23





## PGM

The PGM (Program Manager) programming calls load programs, start execution, and perform cleanup. This section describes their data types, call syntax, and error codes. Refer to the Introduction at the beginning of this manual for a description of data-type diagrams and call syntax format.


PGM DATA TYPES

CONSTANTS

PGM_\$ERROR	3	The error severity level.
PGM_\$FALSE	1	A test severity level.
PGM_\$INTERNAL_FATAL	5	The fatal severity level.
PGM_\$MAX_SEVERITY	15	The highest severity level.
PGM_\$OK	0	The success severity level.
PGM_\$OUTPUT_INVALID	4	A conditional severity level.
PGM_\$PROGRAM_FAULTED	6	The program fault severity level.
PGM_\$TRUE	0	A test severity level.
PGM_\$WARNING	2	The warning severity level.

DATA TYPES

EC_\$PTR_T	A 4-byte integer. Pointer to an eventcount.
EC2_\$EVENTCOUNT_T	User eventcount. The diagram below illustrates the EC2_\$EVENTCOUNT_T data type:



Field Description:

value  
Current EC value.

awaiters  
First process waiting.


PGM\_\$ARG

An argument returned by PGM\_\$GET\_ARGS. The diagram below illustrates the PGM\_\$ARG data type:

predefined  
type

byte:  
offset

field name



Field Description:

len  
Length of the argument.

chars  
The text of the argument, a character array of up to 128 elements.

PGM\_\$ARGV

A 128-element array of 4-byte integers. An array of pointers to returned arguments.

PGM\_\$ARGV\_PTR

A 4-byte integer. The address of a returned argument.

PGM\_\$CONNV

A 128-element array of 2-byte integers. An array of stream IDs.

PGM\_\$EC\_KEY

A 2-byte integer. Key specifying process eventcount. One of the following predefined values:

PGM\_\$CHILD\_PROC  
Currently the only valid key.

PGM\_\$MODE

A 2-byte integer. Specifies the mode in which to invoke a program. Any combination of the following predefined values:

PGM\_\$NAME

An array of up to 128 characters. The text of a retrieved argument.

PGM DATA TYPES

PGM\_\$OPTS


A 2-byte integer. Options for the mode in which to invoke a program. One of the following predefined values:

PGM\_\$WAIT  
 Specifies synchronous operation of the invoked program.

PGM\_\$BACKGROUND  
 Specifies parallel operation of the invoked process.

PGM\_\$PROC

Process handle record. The diagram below illustrates the PGM\_\$PROC data type:




Field Description:

p  
 The process pointer.

STATUS\_\$\$T

A status code. The diagram below illustrates the STATUS\_\$\$T data type:



Field Description:

**all**  
All 32 bits in the status code.

**fail**  
The fail bit. If this bit is set, the error was not within the scope of the module invoked, but occurred within a lower-level module (bit 31).

**subsys**  
The subsystem that encountered the error (bits 24 - 30).

**modc**  
The module that encountered the error (bits 16 - 23).

**code**  
A signed number that identifies the type of error that occurred (bits 0 - 15).

UID\_\$\$T

A type UID. The diagram below illustrates the UID\_\$\$T data type:

PGM DATA TYPES

predefined  
record

byte:  
offset

field name

0:

integer

high

4:

integer

low

Field Description:

high

The high four bytes of the UID.

high

The low four bytes of the UID.

**PGM\_\$DEL\_ARG**

Deletes a command line argument.

**FORMAT**

PGM\_\$DEL\_ARG (arg-number)

**INPUT PARAMETERS****arg-number**

Number indicating the argument to delete. This is a 2-byte integer.

**USAGE**

PGM\_\$DEL\_ARG deletes the specified argument from the argument vector whose address is returned by PGM\_GET\_ARGS. After execution of PGM\_\$DEL\_ARGS, the previously returned address refers to the newly changed argument vector.

Arguments in the argument vector are numbered 0 through n, where 0 is the program name, and n is the final argument. Because PGM\_\$DEL\_ARGS changes the argument vector, arguments following deleted arguments change in number. For example, say the argument vector contains six arguments (including the program name). After you delete the third argument, arguments 4, 5, and 6 must be referenced as arguments 3, 4, and 5.

PGM\_\$EXIT

PGM\_\$EXIT

Exits from a program to its caller.

## FORMAT

PGM\_\$EXIT

## USAGE

PGM\_\$EXIT can be used to exit from a program at any point and return to the program's caller.

PGM\_\$EXIT differs from a simple exit (for example, via FORTRAN's END statement) in that PGM\_\$EXIT is valid in a subroutine. Execution in a subroutine terminates the main program. FORTRAN's STOP statement, which can be used in main programs and subprograms, calls PGM\_\$EXIT.

When PGM\_\$EXIT is executed, any files left open by the program are closed, any storage acquired is released, and the inhibit count is reset to its value when the program was invoked.

PGM\_\$EXIT calls PFM\_\$SIGNAL with a status code equal to the last severity level set by a call to PGM\_\$SET\_SEVERITY. If no PGM\_\$SET\_SEVERITY calls have been made, the status code is PGM\_\$OK. PFM\_\$SIGNAL signals this severity to any established clean-up handlers, which normally execute in response to any status code other than PFM\_\$CLEANUP\_SET. Therefore, any established clean-up routines are normally executed after PGM\_\$EXIT is called.



**PGM\_\$GET\_ARG**

Returns one argument from the command line.

**FORMAT**

`arg-length = PGM_$GET_ARG (arg-number, argument, status, maxlen)`

**RETURN VALUE****arg-length**

Length, in bytes, of the returned argument. This is a 2-byte integer.

**INPUT PARAMETERS****arg-number**

Number of the argument to return. This is a 2-byte integer.

**OUTPUT PARAMETERS****argument**

String of length `arg-length`, containing the requested argument, in `PGM_$NAME` format. This is an array of up to 128 characters.

**status**

Completion status, in `STATUS_$T` format. This data type is 4 bytes long. See the PGM Data Types section for more information.

**INPUT PARAMETERS****maxlen**

Maximum length of the argument, in bytes. This is a 2-byte integer.

FORTTRAN automatically passes the length of a character string following the string itself. Therefore, to return a character string argument to a FORTTRAN program, omit the `maxlen` parameter. Use the following format for the call:

```
arg-length = PGM_$GET_ARG (arg-number, argument, status)
```

This format applies to character strings only. For an argument of any other type, use the standard call.

If the value of `maxlen` is less than the returned argument length, the program manager truncates the returned argument to `maxlen` bytes and returns the status `PGM_$ARG_TOO_BIG`.

PGM\_\$GET\_ARG

## USAGE

PGM\_\$GET\_ARG returns one argument from the program's caller. The argument is in character string format.

Argument numbers on the command line range from 0 to n. Argument 0 is the program name.

**PGM\_\$GET\_ARGS**

Returns the address of the argument vector.

**FORMAT**

PGM\_\$GET\_ARGS (argument-count, arg-vector-addr)

**OUTPUT PARAMETERS****argument-count**

Number of arguments in the argument vector. This is a 2-byte integer.

**arg-vector-addr**

Pointer to an argument vector; that is, an array of up to 128 pointers to PGM\_\$ARG.

**USAGE**

PGM\_\$GET\_ARGS returns the address of the argument vector.

The argument vector is an array of addresses pointing to the arguments. This array can be up to 128 elements.

The addresses are in PGM\_\$ARGV format. This is a 4-byte integer. See the PGM Data Types section for more information.

PGM\_\$GET\_EC

PGM\_\$GET\_EC

Gets an eventcount to wait for completion of a child process.

### FORMAT

PGM\_\$GET\_EC (process-handle, process-key, eventcount-pointer, status)

### INPUT PARAMETERS

#### process-handle

Process handle of the child process for which to wait, in PGM\_\$PROC format. This data type is 4 bytes long. See the PGM Data Types section for more information.

The process handle is returned by PGM\_\$INVOKE when you create a process. Note that the process handle is valid only when you invoke the program in default mode.

#### process-key

Key specifying which process eventcount the system should return, in PGM\_\$EC\_KEY format. This is a 2-byte integer.

Currently the only allowable value is PGM\_\$CHILD\_PROC.

### OUTPUT PARAMETERS

#### eventcount-pointer

The eventcount address to be obtained, in EC2\_\$PTR\_T format. This is a 4-byte integer.

EC2\_\$PTR\_T is a pointer to an EC2\_\$EVENTCOUNT\_T record. See the EC2 Data Types section for more information.

#### status

Completion status, in STATUS\_\$T format. This data type is 4 bytes long. See the PGM Data Types section for more information.

### USAGE

PGM\_\$GET\_EC returns a pointer to an eventcount that advances when a child process terminates. This eventcount address can be passed to EC2\_\$WAIT to wait for a specific child process to complete. You identify the child process by passing the process handle as an input parameter.

When a child process is created, the process eventcount value is 0. When a child process terminates, the process eventcount value is 1. To wait on a specific child process, you might use:

```
PGM_$PROC_EC (...gets process event count ....)
EC2_$WAIT    (...waits until eventcount is 1 ....)
```

See the Managing Programs Chapter of the *Programming With General System Calls* manual for more information.

**PGM\_\$GET\_PUID**

Gets the process UID of a process.

**FORMAT**

PGM\_\$GET\_PUID (process-handle, puid, status)

**INPUT PARAMETERS****process-handle**

Process handle of the child process for which you want a UID, in PGM\_\$PROC format. This data type is 4 bytes long. See the PGM Data Types section for more information.

The process handle is returned by PGM\_\$INVOKE when you create a process. Note that the process handle is valid only when you invoke the program in default mode.

**OUTPUT PARAMETERS****puid**

Process UID, in UID\_\$T format. This data type is 8 bytes long. See the PGM Data Types section for more information.

**status**

Completion status, in STATUS\_\$T format. This data type is 4 bytes long. See the PGM Data Types section for more information.

**USAGE**

PGM\_\$GET\_PUID, which returns the process UID of a child process.

PGM\_\$GET\_PUID, which is used in conjunction with other system calls. These calls are:

- PROC2\_\$GET\_INFO, which returns information about a process given a PUID.
- PROC2\_\$LIST, which returns a list of the PUIDs of all active user processes.
- PGM\_\$MAKE\_ORPHAN, which returns the PUID of the orphaned process.

PGM\_\$INVOKE

PGM\_\$INVOKE

Invokes a program.

## FORMAT

PGM\_\$INVOKE (pathname, namelength, arg-count, arg-vector, stream-count,  
connection-vector, mode, process-handle, status)

## INPUT PARAMETERS

### pathname

Pathname of the program to invoke, in NAME\_\$PNAME\_T format. This is an array of up to 256 characters.

The specified pathname must be an absolute pathname; the Shell's search rules do not apply.

### namelength

Length of the pathname, in bytes. This is a 2-byte integer.

### arg-count

Number of arguments to pass to the invoked program. This is a 2-byte integer.

This number corresponds to the number of elements in the argument vector.

### arg-vector

Array containing the addresses of the arguments to pass to the invoked program, in PGM\_\$ARGV format. This is an array of 4-byte integers.

A program can pass any number of arguments to a program it is invoking. Each argument must be preceded by a 2-byte integer indicating the number of bytes in the argument. The first argument must be the name of the program; the simple name, not the absolute pathname (that is, *date*, *not* //desperado/com/date). Note that if the invoked program calls PGM\_\$DEL\_ARG, the argument vector changes. See the description of PGM\_\$DEL\_ARG for details.

### stream-count

Number of streams to pass to the invoked program. This is a 2-byte integer.

You are permitted to pass up to 32 streams. In the invoked program these streams are numbered 0 to 31.

### connection-vector

Array containing stream IDs to pass to the invoked program, in PGM\_\$CONNV format. Each stream ID is a 2-byte integer, in STREAM\_\$ID\_T format. Up to 128 elements are permitted.

By default, every program is invoked with four streams, numbered 0 through 3. Stream 0 is standard input, stream 1 is standard output, Stream 2 is error input, stream 3 is error output.

Stream IDs refer to objects already opened by the calling program, using STREAM\_\$CREATE or STREAM\_\$OPEN. The first element in the connection-vector

array becomes stream 0 in the invoked program, the second element becomes stream 1, and so on.

You may leave "holes" in the connection vector by setting a stream ID equal to the predefined constant `STREAM_ $NO_ STREAM`.

#### mode

Mode in which to invoke the program, in `PGM_ $MODE` format. This is a 2-byte integer. Specify a null set, or one of the following predefined values:

`PGM_ $WAIT` The program executes as a separate program within the same process as the invoking program.

`PGM_ $BACK_ GROUND`  
The program executes as a separate process that runs to termination independently of the invoking process.

If you pass a null set (default), the program executes as a separate process that communicates its termination status to the invoking program. To specify a null set in C and FORTRAN, declare the variable and initialize it to 0.

## OUTPUT PARAMETERS

#### process-handle

Process handle of the process in which the invoked program runs, in `PGM_ $PROC` format. This data type is 4 bytes long. See the PGM Data Types section for more information.

The process handle is used as an input parameter in the `PGM_ $GET_ EC`, `PGM_ $PROC_ WAIT`, `PGM_ $GET_ PUID`, and `PGM_ $MAKE_ ORPHAN` calls to identify an invoked program.

Note that the process handle is valid only after creating a process in default mode. You will get an error (for example, 'reference to illegal address') if you attempt to use the process handle of a process created in background mode. The following calls use the process handle: `PGM_ $GET_ EC`, `PGM_ $GET_ PUID`, `PGM_ $MAKE_ ORPHAN` or `PGM_ $PROC_ WAIT`.

#### status

Completion status, in `STATUS_ $T` format. This data type is 4 bytes long. See the PGM Data Types section for more information.

Possible values are:

`STATUS_ $OK` Success status.

`PGM_ $BAD_ CONNV`  
Stream vector too large (>32).

Severity level values returned by the program:

`PGM_ $TRUE` Value of tested condition is true.

`PGM_ $FALSE` Value of tested condition is false.

PGM\_ \$INVOKE

PGM\_ \$WARNING

Unusual, but not fatal condition detected.

PGM\_ \$ERROR Syntactic or semantic errors in input; output is structurally sound.

PGM\_ \$INVALID\_ OUTPUT

Syntactic or semantic errors in input; output is not structurally sound.

PGM\_ \$INTERNAL\_ FATAL

Internal fatal error detected.

Any status returned by the invoked program.

Any status returned by modules that PGM\_ \$INVOKE calls.

## USAGE

PGM\_ \$INVOKE invokes a program in the specified mode, and passes that program any parameters that it needs. The addresses of arguments are passed to the invoked program by way of the arg-vector, which is an array of those addresses. The invoked program uses the system routines, PGM\_ \$GET\_ ARGS, PGM\_ \$GET\_ ARG, and PGM\_ \$DEL\_ ARG to access the arguments. See the documentation of those routines for details.

You can change standard input for the invoked program by opening the desired input file and passing its stream ID as the first element of the connection vector. The same is true for standard output, standard error input, and standard error output.

When the invoked program finishes executing, files it has opened are closed, storage it has acquired is released, and the inhibit count is the same as it was upon entry.

The behavior of an invoked program differs depending on the mode in which the program is invoked.

### Invoking a Program in Wait Mode

When you invoke a program this way, the invoking program executes the program and waits for it to complete before continuing.

A program invoked in wait mode calls PGM\_ \$SET\_ SEVERITY to indicate its completion status to the invoking program.

A program ends when one of the following takes place:

- A language defined termination statement is executed
- An unhandled fault occurs
- You call PGM\_ \$EXIT



Normal termination returns execution to the calling program. An unhandled error either terminates the program with an error status or invokes a clean-up handler. PGM\_\$EXIT invokes any established clean-up handlers, then exits to the calling program. Any severity levels set during program execution are returned in the status parameter.

### **Invoking a Program in Default Mode**

When you invoke a program specifying a null set, the invoking program creates a new process in which to run the program. The invoking process may wait for the child process to complete and determine its termination status by calling PGM\_\$PROC\_WAIT.

When a process invokes another process, the invoking process is referred to as the parent process and the invoked process is referred to as the child process. Executing a program in a child process is useful if you wish to perform concurrent processing or if your program requires a large amount of address space.

### **Waiting for a Child Process**

The PGM\_\$GET\_EC call permits you to get a process eventcount that is advanced when a specified process terminates. By using this call in conjunction with the system calls EC2\_\$READ and EC2\_\$WAIT, a parent process can wait for the completion of a child process (or a list of event counts).

### **Getting the Completion Status of a Child Process**

Once a child process has completed, examine its completion status. To obtain the completion status of a default mode process, call PGM\_\$PROC\_WAIT in the parent process. PGM\_\$PROC\_WAIT takes the process handle of the invoked program as an input parameter and returns its completion status. If the child process has not completed execution at the time of the PROC\_WAIT call, execution of the parent process suspends until a completion status is available.

A certain amount of resources in a parent process are used to keep track of a child process. When a call to PGM\_\$PROC\_WAIT is completed those resources are released. If you invoke a number of child processes without ever calling PROC\_WAIT, the parent process may run out of resources. If you are not interested in the completion status of the invoked program, invoke it using background mode.

### Invoking a Program in Background Mode

When you invoke a program specifying PGM\_\$BACK\_GROUND, the invoking program creates a new process in which to run the program. Background mode differs from default mode in that a background mode process runs completely independently of the parent. That is, there is no communication of the completion status.

Background mode is useful for performing processing that has no further dependence on the parent process. For example, a parent process may perform interactive data collection, invoke a program in a background process to manipulate the data, and then return to further data collection. This permits you to collect and manipulate the data concurrently.

Because a background mode process has no dependence on the parent, it is referred to as an orphan process. You can change a default child process into an orphan process by calling PGM\_\$MAKE\_ORPHAN.

Note that the process handle is valid only after creating a process in default mode. You will get an error (for example, 'reference to illegal address') if you attempt to use the process handle of a process created in background mode. The following calls use the process handle: PGM\_\$GET\_EC, PGM\_\$GET\_PUID, PGM\_\$MAKE\_ORPHAN or PGM\_\$PROC\_WAIT.

**PGM\_ \$MAKE\_ ORPHAN**

Changes a normal child process into an orphan process.

**FORMAT**

PGM\_ \$MAKE\_ ORPHAN (process-handle, puid, status)

**INPUT PARAMETERS****process-handle**

Process handle of the child process to orphan, in PGM\_ \$PROC format. This data type is 4 bytes long. See the PGM Data Types section for more information.

The process handle is returned by PGM\_ \$INVOKE when you create a process. Note that the process handle is valid only when you invoke the program in default mode.

**OUTPUT PARAMETERS****puid**

Process UID, in UID\_ \$T format. This data type is 8 bytes long. See the PGM Data Types section for more information.

**status**

Completion status, in STATUS\_ \$T format. This data type is 4 bytes long. See the PGM Data Types section for more information.

**USAGE**

PGM\_ \$MAKE\_ ORPHAN changes the specified child process into an orphan process.

An orphan process is one that is run in PGM\_ \$BACKGROUND mode. An orphan process runs independently of the parent process and no termination status is returned to the parent.

PGM\_\$PROC\_WAIT

PGM\_\$PROC\_WAIT

Waits for a process that has been created with PGM\_\$INVOKE to terminate and returns a completion status for the process.

#### FORMAT

PGM\_\$PROC\_WAIT (process-handle, status)

#### INPUT PARAMETERS

##### process-handle

Process handle of the child process for which to wait, in PGM\_\$PROC format. This data type is 4 bytes long. See the PGM Data Types section for more information.

The process handle is returned by PGM\_\$INVOKE when you create a process. Note that the process handle is valid only when you invoke the program in default mode.

#### OUTPUT PARAMETERS

##### status

The child process completion status, in STATUS\_\$T format. This data type is 4 bytes long. See the PGM Data Types section for more information.

#### USAGE

PGM\_\$PROC\_WAIT suspends the execution of a parent process until the completion of a specified child process. This call permits a child process to pass a completion status to the parent upon termination.

Using PGM\_INVOKE in default mode (the empty set) and then calling PGM\_\$PROC\_WAIT is equivalent to using PGM\_\$INVOKE in PGM\_\$WAIT mode.

**PGM\_\$SET\_SEVERITY**

Sets the severity level for a program.

**FORMAT**

PGM\_\$SET\_SEVERITY (severity-level)

**INPUT PARAMETERS****severity-level**

The severity level returned to the caller. This is a 2-byte integer. Specify only one of the following predefined values:

PGM\_\$OK        The program completed successfully and performed the requested action.

PGM\_\$TRUE     The program completed successfully; its purpose was to test a condition, and the value of that condition was TRUE.

PGM\_\$FALSE    The program completely successfully; its purpose was to test a condition, and the value of that condition was FALSE.

**PGM\_\$WARNING**

The program completed successfully and performed the requested action. However, an unusual (but nonfatal) condition was detected.

PGM\_\$ERROR    The program could not perform the requested action because of syntactic or semantic errors in the input. The output is structurally sound, however.

**PGM\_\$OUTPUT\_INVALID**

The program could not perform the requested action because of syntactic or semantic errors in the input, and the output is not structurally sound.

**PGM\_\$INTERNAL\_FATAL**

The program detected an internal fatal error and ceased processing. The state of the output is neither defined nor guaranteed.

**PGM\_\$PROGRAM\_FAULTED**

The program detected and handled a fault.

Severity levels are a subset of the general system status codes.

**USAGE**

Every program returns a severity level to its caller. By default, the severity level is PGM\_\$OK. Use PGM\_\$SET\_SEVERITY in the invoked program to change the level to another value.

C programmers wishing to use this call must specify the -entry option to /com/bind (in the UNIX environment) or the -e option to /bin/ld (in our operating system) to specify the start routine manually (normally "main"). The alternative to PGM\_\$SET\_SEVERITY is to use the exit()-- or return()--routine to return a program status.

PGM\_\$SET\_SEVERITY

The following are examples of appropriate changes to the severity level.

PGM\_\$TRUE or PGM\_\$FALSE would be returned by an "equal" program that compares its two arguments to see if they are equal.

PGM\_\$WARNING would be returned by DLF (DELETE\_FILE) if the file to be deleted did not exist.

PGM\_\$ERROR would be returned by a compiler if the input program contained an error that prevented a correct translation, but the output object module format was correct.

PGM\_\$OUTPUT\_INVALID would be returned by a compiler if an error in the input program caused the object module format to be invalid.

PGM\_\$INTERNAL\_FATAL would be returned if the program could not proceed because it detected that its data structures were corrupted.

PGM\_\$PROGRAM\_FAULTED would be returned if the program signaled a fault and wishes to inform the invoking program without resigalling the fault.

**ERRORS****STATUS\_ \$OK**

Successful completion.

**PGM\_ \$ERROR**

The program could not perform the requested action because of syntactic or semantic errors in the input. The output is structurally sound, however.

**PGM\_ \$FALSE**

The program completely successfully; its purpose was to test a condition, and the value of that condition was FALSE.

**PGM\_ \$INTERNAL\_FATAL**

The program detected an internal fatal error and ceased processing. The state of the output is neither defined nor guaranteed.

**PGM\_ \$OK**

The program completed successfully and performed the requested action.

**PGM\_ \$OUTPUT\_INVALID**

The program could not perform the requested action because of syntactic or semantic errors in the input, and the output is not structurally sound.

**PGM\_ \$PROGRAM\_FAULTE**

The program faulted.

**PGM\_ \$TRUE**

The program completed successfully; its purpose was to test a condition, and the value of that condition was TRUE.

**PGM\_ \$WARNING**

The program completed successfully and performed the requested action. However, an unusual (but non-fatal) condition was detected.





# PM

## Table of Contents

Data Types	PM-2
PM_\$GET_HOME_TXT	PM-3
PM_\$GET_SID_TXT	PM-4



## PM

The PM (Process Manager) programming calls create and delete processes. This section describes their data types and call syntax. The PM calls do not use produce unique error messages. Refer to the Introduction at the beginning of this manual for a description of data-type diagrams and call syntax format.

PM DATA TYPES


CONSTANTS

NAME\_\$PNAMLEN\_MAX            256            Maximum length of a pathname.

DATA TYPES

NAME\_\$PNAME\_T                            An array of up to NAME\_\$PNAMLEN\_MAX (256) characters.

STATUS\_\$T                                A status code. The diagram below illustrates the STATUS\_\$T data type:



Field Description:

all  
All 32 bits in the status code.

fail  
The fail bit. If this bit is set, the error was not within the scope of the module invoked, but occurred within a lower-level module (bit 31).

subsys  
The subsystem that encountered the error (bits 24 - 30).

modc  
The module that encountered the error (bits 16 - 23).

code  
A signed number that identifies the type of error that occurred (bits 0 - 15).

**PM\_\$GET\_HOME\_TXT**

Returns the home directory of the calling process as a string.

**FORMAT**

PM\_\$GET\_HOME\_TXT (maxlen, home, len)

**INPUT PARAMETERS****maxlen**

Maximum number of characters to be returned (at most, the size of the buffer you assign for home). This is a 2-byte positive integer. This parameter need not exceed 256.

**OUTPUT PARAMETERS****home**

Pathname of the home directory for the SID (log-in identifier) of this process. This is an array of up to 256 characters.

**len**

Number of characters returned in the home parameter. This is a 2-byte positive integer.

**USAGE**

The home directory is obtained from the network registry when you log in and is inherited by all your processes.

PM\_\$GET\_SID\_TXT

PM\_\$GET\_SID\_TXT

Returns the SID (log-in identifier) of the calling process as a string.

#### FORMAT

PM\_\$GET\_SID\_TXT (maxlen, sid, len)

#### INPUT PARAMETERS

##### maxlen

Maximum number of characters to be returned (at most, the size of the buffer you assign for home). This is a 2-byte positive integer. This parameter need not exceed 140.

#### OUTPUT PARAMETERS

##### sid

String containing the person, project, organization and node ID of the SID (log-in identifier) of this process, in the form:

person.group.project.nodeid

This is an array of up to 140 characters.

##### len

Number of characters returned in the log-in identifier. This is a 2-byte positive integer.

#### USAGE

Your SID is the full identifier obtained from the network registry when you log in and is inherited by all your processes.

# PRF

## Table of Contents

Data Types	PRF-2
PRF_\$CONFIG_FILE	PRF-5
PRF_\$DELETE_ENTRY	PRF-6
PRF_\$INIT	PRF-7
PRF_\$INQ_OPTION	PRF-8
PRF_\$NAME_PRINT	PRF-9
PRF_\$QUEUE_FILE	PRF-10
PRF_\$SET_OPTION	PRF-11
PRF_\$STREAM_PRINT	PRF-17
Errors	PRF-18





## PRF

The PRF (Print Library) programming calls provide application control of the print queue and the PRF options. This section describes their data types, call syntax, and error codes. Refer to the Introduction at the beginning of this manual for a description of data-type diagrams and call syntax format.

DATA TYPES


NUM\_VAR\_T

A print option expressed as a numeric value. The diagram below illustrates the NUM\_VAR\_T data type:

Total size: 4

byte: offset

field name



Field Description:

real32  
A 32-bit real number.

int32  
A positive 31-bit number.

pad\_int16  
Padding.

int16  
A positive 16-bit number.

pad\_int16\_n  
Padding.

int16\_n  
A positive or negative 16-bit number.

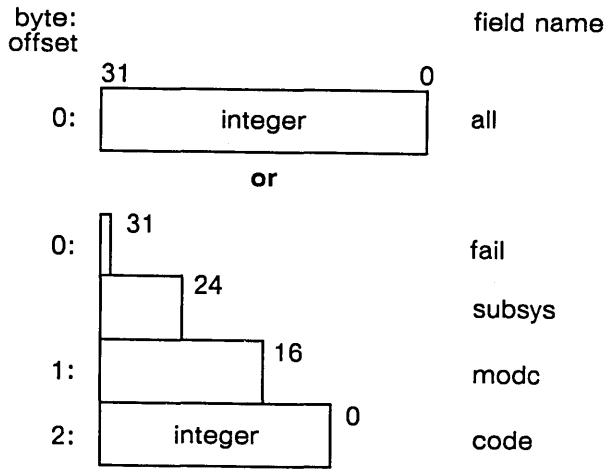
pad\_byte  
Padding.

int8  
A positive 8-bit integer.

int1  
A boolean (bit 0).

STATUS\_\$T

A status code. The diagram below illustrates the STATUS\_\$T data type:



Field Description:

all  
All 32 bits in the status code.

fail  
The fail bit. If this bit is set, the error was not within the scope of the module invoked, but occurred within a lower-level module (bit 31).

subsys  
The subsystem that encountered the error (bits 24 - 30).

modc  
The module that encountered the error (bits 16 - 23).

code  
A signed number that identifies the type of error that occurred (bits 0 - 15).

PRF DATA TYPES

UNIV\_PTR

A 4-byte integer. A pointer to allocated storage.



**PRF\_\$CONFIG\_FILE**

Sets the print options in the print request packet from a configuration file.

**FORMAT**

PRF\_\$CONFIG\_FILE (name, name-length, status)

**INPUT PARAMETERS****name**

The full pathname of the configuration file from which the print options are to be read.  
This an array of up to 256 characters.

**name-length**

The number of characters in the pathname. This is a 2-byte integer.

**OUTPUT PARAMETERS****status**

Completion status in STATUS\_\$T format. This data type is 4 bytes long. See the PRF Data Types section for more information.

**USAGE**

This procedure allows you to specify print options in a configuration file. The options are then read from the file into a print request packet. For a list of available options, refer to the description of the PRF\_\$SET\_OPTION call. The configuration file is formatted as in the following sample:

```
COPIES 1
MARGINS ON
PLOT ON
```

PRF\_\$DELETE\_ENTRY

PRF\_\$DELETE\_ENTRY

Deletes a print request from the print queue.

#### FORMAT

PRF\_\$DELETE\_ENTRY (queue-entry-name, queue-entry-name-length, status)

#### INPUT PARAMETERS

##### queue-entry-name

The full pathname of the print request packet. This an array of up to 256 characters.

##### queue-entry-name-length

The number of characters in the pathname of the print request packet. This is a 2-byte integer.

#### OUTPUT PARAMETERS

##### status

Completion status in STATUS\_\$T format. This data type is 4 bytes long. See the PRF Data Types section for more information.

#### USAGE

This procedure takes as input the print request packet name returned by PRF\_\$STREAM\_PRINT, PRF\_\$NAME\_PRINT, and PRF\_\$QUEUE\_PRINT and deletes the print request packet from the /sys/print/queue directory.

You should take care when making this call that you do not unintentionally delete your data file. If the delete option (refer to the list of options in the description of PRF\_\$SET\_OPTION) is ON, PRF\_\$DELETE\_ENTRY will also delete the data file. Thus, if you are using this call in conjunction with PRF\_\$QUEUE\_FILE, you will probably want to set the delete option to OFF since the PRF\_\$QUEUE\_FILE does not create additional copies of the data file.

**PRF\_\$INIT**

Sets print request packet options to default values.

**FORMAT**

PRF\_\$INIT (stream-id, status)

**INPUT PARAMETERS****stream-id**

The stream to write error and status messages to, in STREAM\_\$ID\_T format. This is a 2-byte integer. See the STREAM Data Types section for more information.

**OUTPUT PARAMETERS****status**

Completion status in STATUS\_\$T format. This data type is 4 bytes long. See the PRF Data Types section for more information.

**USAGE**

This procedure sets the print request packet options to their default values and sets the stream to be used for sending error and status information.

PRF\_\$INQ\_OPTION

PRF\_\$INQ\_OPTION

Returns the string and number representations of the print request packet option to the caller.

## FORMAT

PRF\_\$INQ\_OPTION (option, option-length, number-value, string,-value, string-length, status)

## INPUT/OUTPUT PARAMETERS

### option

Specifies the option to be used when printing the data. Refer to the description of the PRF\_\$SET\_OPTION for a list of available options.

### option-length

The number of characters in the specified option.

## OUTPUT PARAMETERS

### number-value

The numerical representation of the value of the option, in NUM\_VAR\_T FORMAT. See the PRF Data Types section for more information.

### string-value

The string representation of the value of the option. This an array of up to 256 characters.

### string-length

The number of characters in the string-value. This is a 2-byte integer.

### status

Completion status in STATUS\_\$T format. This data type is 4 bytes long. See the PRF Data Types section for more information.

## USAGE

This procedure returns to the caller the value (in string and number representation) of an option.



**PRF\_\$NAME\_PRINT**

Copies a data file to the /sys/print/spooler directory and generates print request packet.

**FORMAT**

PRF\_\$NAME\_PRINT (filename, filename-length, queue-name,  
queue-name-length, status)

**INPUT PARAMETERS****filename**

The full pathname of the data file. This an array of up to 256 characters.

**filename-length**

The number of characters in the pathname. This is a 2-byte integer.

**OUTPUT PARAMETERS****queue-name**

The full pathname of the print request packet. This an array of up to 256 characters.

**queue-name-length**

The number of characters in the pathname of the print request packet. This is a 2-byte integer.

**status**

Completion status in STATUS\_\$T format. This data type is 4 bytes long. See the PRF Data Types section for more information.

**USAGE**

This procedure copies the data file to the /sys/print/spooler directory and then generates a print request packet that contains the options describing how the file is to be printed.

PRF\_\$QUEUE\_FILE

PRF\_\$QUEUE\_FILE

Generates a print request packet in the /sys/print/queue directory.

### FORMAT

PRF\_\$QUEUE\_FILE (filename, filename-length, queue-name, queue-name-length,  
status)

### INPUT PARAMETERS

#### filename

The full pathname of the data file. This an array of up to 256 characters.

#### filename-length

The number of characters in the pathname. This is a 2-byte integer.

### OUTPUT PARAMETERS

#### queue-name

The full pathname of the print request packet. This an array of up to 256 characters.

#### queue-name-length

The number of characters in the pathname of the print request packet. This is a 2-byte integer.

#### status

Completion status in STATUS\_\$T format. This data type is 4 bytes long. See the PRF Data Types section for more information.

### USAGE

This procedure generates a print request packet that contains the parameters describing how the file is to be printed. Note that the data file specified as an input parameter is not copied to the /sys/print/spooler directory and will therefore be deleted after it has been printed unless you set the delete option to OFF (refer to the list of print options in the description of PRF\_\$SET\_OPTION).

**PRF\_\$SET\_OPTION**

Sets an option in the print request packet.

**FORMAT**

PRF\_\$SET\_OPTION (option, option-length, number-value, string-value, string-length, status)

**INPUT PARAMETERS****option**

Specifies the option to be used when printing the data. A description of the available options appears below under "USAGE".

**option-length**

The number of characters in the specified option.

**number-value**

The numerical representation of the value of the option, in NUM\_VAR\_T FORMAT. See the PRF Data Types section for more information.

**string-value**

The string representation of the value of the option. This an array of up to 256 characters.

**string-length**

The number of characters in the string-value. This is a 2-byte integer.

**string-encoding**

Indicates whether the value of the option is encoded as a string or number. This is a boolean value. If true, the option is encoded as a string; if false, a number.

**OUTPUT PARAMETERS****status**

Completion status in STATUS\_\$T format. This data type is 4 bytes long. See the PRF Data Types section for more information.

**USAGE**

This procedure sets a specified print option in the print request packet. The print server (PRSVR) uses the option to determine how to print the data. Following are the options that you can specify:

**SEA[RCH\_DIR] ON|OFF**

Search through all the directories of all the active processes on your node for the file(s) to be printed. This option is most useful in interactive mode, when the working directory of the process may be different from the working directory of the file to be printed. The default state is OFF.

**D[ELETE] ON|OFF**

Delete the data file after it has been printed. The default state is

on; i.e., if you do not specify this option, the print server (PRSVR) deletes the data file after it has been printed.

**COP[IES] n**

Print n copies of the file. If COP[IES] is specified, "n" is required. If this option is omitted altogether, one copy is printed by default.

**PR[INTER] name**

Specify the name of the printer that is to print the data file. This option is useful only if more than one printer is in use on the network. If you omit this option, the print server (PRSVR) uses the default printer name "P".

**S[ITE] entry\_dir**

Specify print queue (/SYS/PRINT) on an alternate node by giving that node's entry directory name. This option allows you to maintain more than one print queue directory. You may want to maintain separate queues for different organizations, or you may want two queues to provide redundancy in case of node failure.

**USER[\_NAME] name**

Specify the user name that will appear on the banner page of the printed file. The print server alarm facility also uses this name to determine who should be notified when printing is complete (see SIG below). This means that this name must be a valid login name (unless you don't care about sending an alarm). If you omit this option, the current login name is used.

**SIG[NAL] ALARM|OFF**

Request an alarm server signal when the file has finished printing. The default state is OFF.

**BAN[NER] ON|OFF**

Enable/disable banner page. The default is specified in the print server configuration file.

**CONFIG[\_FILE] [pathname]**

Specify a file containing print request packet options, one per line. Refer to the description of the PRF\_\$CONFIG\_FILE call for information concerning the format of the configuration file. If you omit the pathname, the print server executes the configuration file ~USER\_DATA/PRF.DB.

**PAPER\_SIZE A|B|LEGAL|A3|A4|A5|B4|B5**

Specify the page size to use on the LASER26 printer. If you omit this option, the default is the page size as specified in the print server configuration file.

**ORIENT[ATION] PORT[TRAIT]|LAND[SCAPE]**

Print the file according to the specified orientation. Portrait

specifies that text lines or the x axis of bitmaps will be parallel to the short edge of a page. Landscape specifies that text lines or the x axis of bitmaps will be perpendicular to the short leading edge. This option applies only to printers that include the PostScript interpreter (i.e., LASER26, LaserWriter, Genicom, and V80 printers with the PostScript decomposer software). If you omit this option, the default state is portrait. Specifying this option will override any auto-rotation performed on bitmaps.

**TEXT** Specify text mode for printing ASCII files. This is the default print mode.

**PLOT ON|OFF**

Specify plot mode. Include this option to print bitmap files created by a graphics metafile (GMF) manager or GPR or the CPSCR (COPY\_SCREEN) command. Turning PLOT off forces text mode.

**TRANSPARENT ON|OFF**

Specify that when the file is printed, the records of the file are passed directly to the printer driver routine with no processing by the print server. Turning TRANSPARENT off forces text mode.

**POST[SCRIPT] ON|OFF**

Specify that the printer uses a software PostScript interpreter; i.e., the specified printer is driven and licensed to run the print server (prsvr\_post) that contains a software version of the PostScript interpreter. If you specify OFF for this option, data is printed as text, plot, or transparent data. If ON, data is passed through the PostScript interpreter, as though the printer were an Apple LaserWriter or LASER26 printer. This option allows you to take full advantage of the PostScript language when printing text or bitmaps, or bypass the interpreter for higher throughput. The default state is off.

*The following options apply to text files only:*

**MARGINS ON|OFF**

Enable/disable the margin settings specified by the next four options. The default state is on; i.e., if you do not specify this option, the margin settings are those you have specified with the next four options.

**TOP n** Specify page top margin, in n (a real number) inches. The default is a value specified in the print server configuration file.

**BOT[TOM] n**

Specify page bottom margin, in n (a real number) inches. The default is a value specified in the print server configuration file.

**RIGHT n**

Specify page right margin, in n (a real number) inches. The default is 0 inches.

**LEFT n**

Specify page left margin, in n (a real number) inches. The default is 0 inches.

**HEADERS ON|OFF**

Enable/disable page headers and footers as specified by the next two options. If you omit the next two options, the default is as specified in the print server configuration file.

**HEAD[\_STRING] l-string/c-string/r-string**

Specify the contents of left, center, and right components of the page header generated by print server. Components may be empty strings. The following special characters return the values indicated when they appear in the header strings:

- @ = Escape character
- # = Current page number with 1 leading and 1 trailing space
- % = Current date
- ! = Filename
- & = Filename's last time, date modified
- \* = Insert a space in text string (literal spaces are not allowed)

Example: HEAD !/Page#/% will produce a header with the filename in the left component, the string "Page" followed by the current page number in the center component, and the current date in the right component. The default header is a string specified in the print server configuration file.

**FOOT[\_STRING] l-string/c-string/r-string**

Specify contents of page footers. The format is the same as for HEAD above. There is no default footer (null string).

**FTN ON|OFF**

Cause the print server to use FORTRAN forms control even if the file does not have the FORTRAN carriage control flag. Use of this option causes print server to interpret the first character of each line as a FORTRAN carriage control character (and not print it). This can be unfortunate if the file has ASCII carriage control, so be careful. The default state is OFF.

**WRAP ON|OFF**

Enable/disable automatic line wrapping. When enabled, the print server wraps any lines that exceed the right margin onto the next line. When disabled, the print server truncates lines that exceed the right margin.

**COL[UMNS] 1|2**

Specify that text will be printed in one or two columns. Currently restricted to use with LASER26, APPLE LASERWRITER, and printers that contain the PostScript interpreter. The default is one column.

*The following options are for use with printers supporting variable font and pitch sizes:*

**LPI n** Specify the line spacing in lines per inch. The default is six lines per inch.

**PITCH n**

Set the pitch (characters/inch) setting for printing the data file. The following pitch settings are available on the indicated printers:

Printronic: 10

Spinwriter: 12

Imagen: 8.5, 10, 12, 15, 17.1

GE: 9.6, 10, 10.6, 11.1, 12, 13.1, 13.8, 15, 16.7, 18

Versatec 12

**POINT n**

Set the point size for the font to be used. This is a real number in units of a point (1/72 inch). PostScript printers will accept any size from 1 to 100. Other printers are limited to the pitch sizes specified with the PITCH option.

**WEIGHT light|medium|bold**

Set the weight of the font to be used. This option is only valid for the GE printer type. The default is 'medium'.

**LQ ON|OFF**

Specify that the document is to be printed in 'letter quality' (ON) as opposed to 'draft' (OFF) mode. This option is only valid for the GE printer type. The default state is OFF.

*The following options apply to plot files:*

**RES[OLUTION] n**

Specify resolution of output plot in dots per inch. If you specify a resolution not available on the particular printer, the file is printed at the closest available resolution. The default resolution is specified in the PRSVR configuration file.

**WHITE[\_SPACE] n**

Specify amount of white space (in inches) to appear between multiple plots in one file. The default is three inches.

**BW[\_REV] ON|OFF**

Enable/disable black and white reversal for bitmaps. The default state is OFF.

**MAGN[IFICATION] n**

Specify bitmap magnification value. 'n' is an integer in the range -1 to 16. The default is 0.

- 1 : Selects auto-scaling to magnify the bitmap to fill the available page space.
- 0: Selects "one-to-one" scaling between the display and the printer for GMF bitmaps. (For GPR bitmaps, this translates to magnification 1.)
- 1-16 : Selects magnification by that amount. Portions of the magnified bitmap that exceed the printer page boundaries are clipped.



**PRF\_ \$STREAM\_ PRINT**

Copies data from a stream to the /sys/print/spooler directory and generates a print request packet in the /sys/print/queue directory.

**FORMAT**

PRF\_ \$STREAM\_ PRINT (stream-id, queue-name, name-length, status)

**INPUT PARAMETERS****stream-id**

The stream id of the data to be printed, in STREAM\_ \$ID\_ T format. This is a 2-byte integer. See the STREAM Data Types section for more information.

**OUTPUT PARAMETERS****queue-name**

The full pathname of the print request packet. This an array of up to 256 characters.

**name-length**

The number of characters in the pathname. This is a 2-byte integer.

**status**

Completion status in STATUS\_ \$T format. This data type is 4 bytes long. See the PRF Data Types section for more information.

**USAGE**

This procedure copies input data to the /sys/print/spooler directory. It then generates a print request packet that contains the options describing how the file is to be printed. This procedure will print data from the first record in the stream up to the EOF mark.

PRF ERRORS

ERRORS

STATUS\_ \$OK

Successful completion.

# PROC1

## Table of Contents

Data Types  
PROC1\_\$GET\_CPUP

PROC1-2  
PROC1-3



## PROC1


The PROC1 (Level 1 Process) programming call returns the CPU time used by a Level 1 process. This section describes its data types and call syntax. The PROC1 call does not produce unique error messages. Refer to the Introduction at the beginning of this manual for a description of data-type diagrams and call syntax format.

PROC1 DATA TYPES

DATA TYPES

TIME\_\$CLOCK\_T


Internal representation of time. The diagram below illustrates the TIME\_\$CLOCK\_T data type:



Field Description:

high  
High 32 bits of the clock.

low  
Low 16 bits of the clock.



Field Description:

high16  
High 16 bits of the clock.

low32  
Low 32 bits of the clock.

**PROC1\_\$GET\_CPU**

Returns the CPU time used by this process.

**FORMAT**

PROC1\_\$GET\_CPU (clock)

**OUTPUT PARAMETERS****clock**

The amount of CPU time used by this process since its creation, in **TIME\_\$CLOCK\_T** format. This data type is 6 bytes long. See the **TIME Data Types** section for more information.

**USAGE**

**PROC1\_\$GET\_CPU** returns the amount of CPU time that the calling process has used since its creation. The returned clock value has a resolution of 4 microseconds.

CPU time is the time during which the process is running in the CPU. This includes the time that the operating system is performing services for the process, but does not include the time that the process spends waiting for I/O transfers to complete.





# PROC2

## Table of Contents

Data Types

PROC2\_\$GET\_INFO

PROC2\_\$LIST

PROC2\_\$WHO\_AM\_I

Errors

PROC2-2

PROC2-5

PROC2-7

PROC2-8

PROC2-9



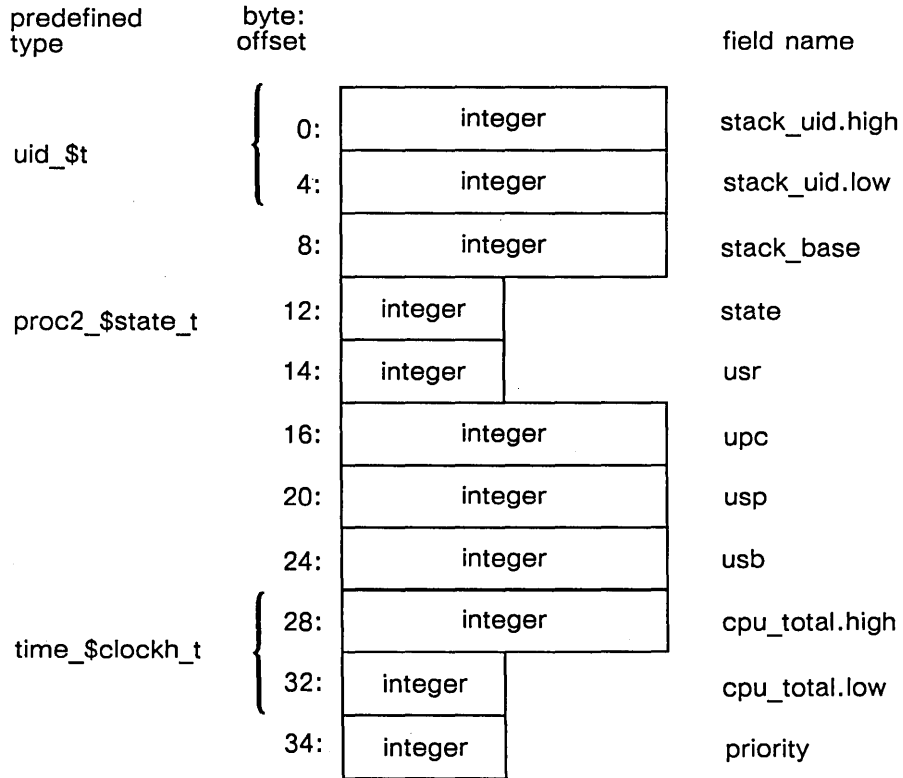
## PROC2

The PROC2 (Level 2 Process Manager) programming calls provide information about user processes. This section describes their data types, call syntax, and error codes. Refer to the Introduction at the beginning of this manual for a description of data-type diagrams and call syntax format.

DATA TYPES

PROC2\_\$INFO\_T

Process information record. The diagram below illustrates the PROC2\_\$INFO\_T data type:



Field Description:

stack\_uid  
 Uid of user stack.

stack\_base  
 Base address of user stack.

state  
 Process state - ready, waiting, etc..

usr  
 User status register.

upc  
 User program counter.

usp  
 User stack pointer.

usb  
 User stack base pointer (A6).

cpu\_total  
Cumulative cpu time used by process.

priority  
Process priority.

PROC2\_\$UID\_LIST\_T

An array of UIDs (in UID\_\$T format) of up to 24 elements.

PROC2\_\$STATE\_T

A 2-byte integer. State of a user process. Any combination of the following predefined values:

PROC2\_\$WAITING  
Process is waiting.


PROC2\_\$SUSPENDED  
Process is suspended.

PROC2\_\$SUSP\_PENDING  
Process suspension is pending.

PROC2\_\$BOUND  
Process is bound.

STATUS\_\$T

A status code. The diagram below illustrates the STATUS\_\$T data type:



Field Description:

all  
All 32 bits in the status code.

fail  
The fail bit. If this bit is set, the error was not within the scope of the module invoked, but occurred within a lower-level module (bit 31).

PROC2 DATA TYPES

subsys

The subsystem that encountered the error (bits 24 - 30).

modc

The module that encountered the error (bits 16 - 23).

code

A signed number that identifies the type of error that occurred (bits 0 - 15).

**PROC2\_\$GET\_INFO**

Returns information about a process.

**FORMAT**

PROC2\_\$GET\_INFO (process-uid, info, info-buf-length, status)

**INPUT PARAMETERS****process-uid**

The UID of the process for which you want information, in UID\_\$T format. This data type is 8 bytes long. See the PROC2 Data Types section for more information.

You can get process UIDs by calling PROC2\_\$WHO\_AM\_I and PROC2\_\$LIST.

If the process-uid in the call is the caller's own process, the only information returned is the stack UID and virtual address. If you want to find out the amount of CPU time used by the caller's process, use PROC1\_\$CPU\_TIME.

**info-buf-length**

Length of the information buffer allotted for returned information, in bytes. This is normally 36 bytes.

**OUTPUT PARAMETERS****info**

Information about the process, in PROC2\_\$INFO\_T format. This data type is 36 bytes long. See the PROC2 Data Types section for more information.

**status**

Completion status, in STATUS\_\$T format. This data type is 4 bytes long. See the PROC2 Data Types section for more information. Possible values are:

STATUS\_\$OK Completed successfully.

PROC2\_\$IS\_CURRENT

Specified calling process UID (success).

PROC2\_\$UID\_NOT\_FOUND

Specified UID is not on node.

**USAGE**

GET\_\$INFO returns information about a process when supplied with a process UID. The information returned consists of the following:

- The program state (ready, waiting, suspended, SUSP\_PENDING, bound).
- The User Status Register (USR).
- The User Program Counter (UPC).

PROC2\_\$GET\_INFO

- The user stack pointer (A7).
- The stack base pointer (A6).
- The amount of CPU time used.
- The CPU scheduling priority.



**PROC2\_\$LIST**

Returns a list of existing level 2 (user) processes in the caller's node.

**FORMAT**

PROC2\_\$LIST (uid-list, max-num-uids, number-uids)

**OUTPUT PARAMETERS****uid-list**

The UIDs of the active level 2 processes on the system, in PROC2\_\$UID\_LIST\_T format. This is a 24-element array of UIDs. Each UID is a 4-byte integer in UID\_\$T format.

**INPUT PARAMETERS****max-num-uids**

Maximum number of process UIDs to be returned. (At most, the size of the buffer you assign for uid-list. This is a 2-byte integer.

**OUTPUT PARAMETERS****number-uids**

Number of active level 2 processes on the node, even if that number is greater than max-num-uids. This is a 2-byte integer.

**USAGE**

The UIDs of all level 2 processes (user processes) on the caller's node, up to max-num-uids, are returned.

PROC2\_\$WHO\_AM\_I

PROC2\_\$WHO\_AM\_I

Returns the UID of the calling process.

#### FORMAT

PROC2\_\$WHO\_AM\_I (my-uid)

#### OUTPUT PARAMETERS

##### **my-uid**

The UID of the calling process, in UID\_\$T format. This data type is 8 bytes long. See the PROC2 Data Types section for more information.

#### USAGE

You can use a UID obtained through this call to find out information about your process through the PROC2\_\$GET\_INFO call.

**ERRORS**

STATUS\_ \$OK

Successful completion.

PROC2\_ \$BAD\_ STACK\_ BASE

Bad stack base.

PROC2\_ \$IS\_ CURRENT

Request is for current process.

PROC2\_ \$NOT\_ LEVEL\_ 2

Not a level two process.

PROC2\_ \$UID\_ NOT\_ FOUND

Process not found.



# RWS

## Table of Contents

Data Types	RWS-2
RWS_\$ALLOC	RWS-4
RWS_\$ALLOC_HEAP	RWS-5
RWS_\$ALLOC_HEAP_POOL	RWS-6
RWS_\$ALLOC_RW	RWS-9
RWS_\$ALLOC_RW_POOL	RWS-10
RWS_\$RELEASE_HEAP	RWS-13
RWS_\$RELEASE_HEAP_POOL	RWS-14
Errors	RWS-16



## RWS

The RWS (Read/Write Storage) programming allow programs to allocate "scratch" read/write storage. This section describes their data types, call syntax, and error codes. Refer to the Introduction at the beginning of this manual for a description of data-type diagrams and call syntax format.

RWS DATA TYPES

DATA TYPES

RWS\_\$POOL\_T

A 2-byte integer. Types of pools from which to allocate read/write or heap storage. One of the following predefined values:

RWS\_\$STD\_POOL

Standard pool makes storage accessible to calling process only.

RWS\_\$STREAM\_TM\_POOL


Stream pool makes storage accessible to calling program and to a program invoked with the UNIX EXEC system call.

RWS\_\$GLOBAL\_POOL

Global pool makes storage accessible to all processes.

STATUS\_\$T

A status code. The diagram below illustrates the STATUS\_\$T data type:



Field Description:

all

All 32 bits in the status code.

fail

The fail bit. If this bit is set, the error was not within the scope of the module invoked, but occurred within a lower-level module (bit 31).

subsys

The subsystem that encountered the error (bits 24 - 30).



modc

The module that encountered the error (bits 16 - 23).

code

A signed number that identifies the type of error that occurred (bits 0 - 15).

UNIV\_PTR

A 4-byte integer. A pointer to allocated storage.

RWS\_\$ALLOC

RWS\_\$ALLOC

Allocates read/write storage for C, FORTRAN or Pascal programs.

## FORMAT

RWS\_\$ALLOC (storage\_sz, storage\_ptr)

## INPUT PARAMETERS

**storage\_sz**

The number of bytes of storage needed. This is a 4-byte integer.

## OUTPUT PARAMETERS

**storage\_ptr**

The address of the new storage space, in UNIV\_PTR format. This is a 4-byte integer. A returned address of zero (NIL) means that RWS\_\$ALLOC could not allocate the desired storage.

## USAGE

RWS\_\$ALLOC allocates the specified number of bytes of read/write storage to the calling process and returns the address of the storage area.

This routine is useful for allocating different quantities of dynamic storage, depending on a run-time factor.

FORTRAN programmers: due to FORTRAN calling conventions, this is currently the only RWS call you can use to allocate read/write storage.

Pascal and C programmers can use other RWS calls to allocate read/write or heap (releaseable read/write) storage. See the calls, RWS\_\$ALLOC\_RW\_POOL and RWS\_\$ALLOC\_HEAP\_POOL for details.

C programmers might want to use the C library routine MALLOC to allocate storage.

**RWS\_\$ALLOC\_HEAP**

Allocates heap (releaseable read/write) storage for Pascal and C programs. *This call is obsolete and is documented here for maintenance purposes only. You should replace it with RWS\_\$ALLOC\_HEAP\_POOL.*

**FORMAT**

storage\_ptr = RWS\_\$ALLOC\_HEAP (storage\_sz)

**RETURN VALUE****storage\_ptr**

The address of the new storage space, in UNIV\_PTR format. This is a 4-byte integer. A returned address of zero (NIL) means that RWS\_\$ALLOC\_HEAP could not allocate the desired storage.

**INPUT PARAMETERS****storage\_sz**

The number of bytes of storage needed. This is a 4-byte integer.

**USAGE**

Note that RWS\_\$ALLOC\_HEAP\_POOL replaces this obsolete call, which we include for maintenance purposes only. For current and future development, use RWS\_\$ALLOC\_HEAP\_POOL.

RWS\_\$ALLOC\_HEAP allocates the specified number of bytes of releaseable read/write storage to the calling process and returns the address of the storage area. It allocates storage from the standard RWS pool, which makes the storage accessible to the calling program only. Use RWS\_\$RELEASE\_HEAP to release storage allocated with this call.

FORTTRAN programmers: due to FORTRAN calling conventions, RWS\_\$ALLOC is currently the only RWS call you can use to allocate read/write storage.

C programmers might want to use the C library routine MALLOC to allocate storage.

RWS\_\$ALLOC\_HEAP\_POOL

RWS\_\$ALLOC\_HEAP\_POOL

Allocates heap (releasable read/write) storage from a specified pool.

#### FORMAT

storage\_ptr = RWS\_\$ALLOC\_HEAP\_POOL(alloc\_pool, storage\_sz )

#### RETURN VALUE

**storage\_ptr**

The address of the new storage space, in UNIV\_PTR format. This is a 4-byte integer. A returned address of zero (NIL) means that RWS\_\$ALLOC\_HEAP\_POOL could not allocate the desired storage.

#### INPUT PARAMETERS

**alloc\_pool**

Pool from which the storage will be allocated, in RWS\_\$POOL\_T format. This is a 2-byte integer. Specify one of the following predefined values:

RWS\_\$GLOBAL\_POOL

Global pool makes storage accessible to all processes. Note that pointers are valid in *all* processes because they reserve the *identical* portion of address space.

RWS\_\$STD\_POOL

Standard pool makes storage accessible to the calling program only. Most programs use this type.

RWS\_\$STREAM\_TM\_POOL

Stream pool makes storage accessible to the calling program and to a program invoked with a UNIX EXEC system call. Use this type when your program needs to pass information across a UNIX EXEC system call.

**storage\_sz**

Number of bytes of storage needed. This is a 4-byte integer.

#### USAGE

RWS\_\$ALLOC\_HEAP\_POOL allocates a specified number of bytes of heap storage to the calling process and returns the address of the storage area.

When you no longer need the storage, call RWS\_\$RELEASE\_HEAP\_POOL to return the storage to the pool from which it was allocated.

Whether you allocate heap (releaseable read/write) storage with this call or read/write storage with RWS\_\$ALLOC\_RW\_POOL depends on how long you want to keep the storage. Once you allocate read/write storage, the storage exists until the program terminates. However, you can explicitly release heap storage once you have finished using it. The heap requires more system overhead initially to keep track of allocated storage. Read/write storage does not require any system overhead.

(Currently, the overhead for RWS\_\$ALLOC\_HEAP\_POOL is between 4 to 16 bytes -- The exact amount of call overhead is subject to change.)

Typically, you allocate heap storage if your program requires a substantial amount of storage for a limited time, or if you want to keep your working set as small as possible. You allocate read/write storage if you do not need to release storage before terminating a program, or if the amount of overhead for a heap is unacceptable.

When allocating heap or read/write storage, you control how your program accesses storage by specifying the type of storage pool to use:

- The standard pool (RWS\_\$STD\_POOL) permits access to the calling process only.
- The global pool (RWS\_\$GLOBAL\_POOL) permits access to all processes.
- The stream pool (RWS\_\$STREAM\_TM\_POOL) permits access to the calling program and a program invoked with a UNIX EXEC system call.

The global pool allows you to share information among processes. For example, you might want to create a global queue to pass messages between processes. Note that pointers are valid in *all* processes because all processes reserve an *identical* portion of address space.

The stream pool allows you to make storage accessible between a calling process and an overlay process. For example, the IOS manager uses a stream pool to pass an open stream to a program invoked with an EXEC call. It stores information about that stream in a stream pool.

The following table summarizes the aspects of each type of storage allocation.

#### Summary of Types of Storage Allocation

	Read/Write Storage	Heap Storage
<b>Standard Pool</b>	Storage kept until program exits or until it invokes a program with a UNIX EXEC system call.	Storage kept until you release it with RWS_\$RELEASE_HEAP, the program exits, or the program invokes a program with a UNIX EXEC call.
	No system overhead.	About 16 bytes of system overhead.
	Storage available to local process only.	
<b>Global Pool</b>	Storage kept until reboot.	Storage kept until you release it with RWS_\$RELEASE_HEAP or reboot.
	About 4 bytes of system overhead.	About 4 bytes of system overhead.
	Storage available to all processes.	

**Summary of Types of Storage Allocation, Cont.**

	<b>Read/Write Storage</b>	<b>Heap Storage</b>
<b>Stream Pool</b>	Storage kept until program exits.	Storage kept until you release it with RWS_\$RELEASE_HEAP.
	No system overhead.	About 16 bytes of system overhead.

Storage available to the local process or to a program invoked with a UNIX EXEC system call.

**NOTE:** Do not depend on the exact amount of system overhead used in RWS system calls. The amount of overhead is subject to change.

Note that this call replaces the obsolete RWS\_\$ALLOC\_HEAP call, which we include for maintenance purposes only. For current and future development, use RWS\_\$ALLOC\_HEAP\_POOL.

FORTTRAN programmers: due to FORTRAN calling conventions, RWS\_\$ALLOC is currently the only RWS call you can use to allocate read/write storage. C programmers might want to use the C library routine MALLOC to allocate storage.

**RWS\_\$ALLOC\_RW**

Allocates read/write storage for Pascal and C programs. *This call is obsolete and is documented here for maintenance purposes only. You should replace it with RWS\_\$ALLOC\_RW\_POOL.*

**FORMAT**

storage\_ptr = RWS\_\$ALLOC\_RW (storage\_sz)

**RETURN VALUE****storage\_ptr**

The address of the new storage space, in UNIV\_PTR format This is a 4-byte integer. A returned address of zero (NIL) means that RWS\_\$ALLOC\_RW could not allocate the desired storage.

**INPUT PARAMETERS****storage\_sz**

The number of bytes of storage needed. This is a 4-byte integer.

**USAGE**

Note that RWS\_\$ALLOC\_RW\_POOL replaces this obsolete call, which we include for maintenance purposes only. For current and future development, use RWS\_\$ALLOC\_RW\_POOL.

RWS\_\$ALLOC\_RW allocates the specified number of bytes of read/write storage to the calling process and returns the address of the storage area. It allocates storage from the standard RWS pool, which makes the storage accessible to the calling program only. This call does not require any system overhead.

FORTRAN programmers: due to FORTRAN calling conventions, RWS\_\$ALLOC is currently the only RWS call you can use to allocate read/write storage.

C programmers might want to use the C library routine MALLOC to allocate storage.

RWS\_\$ALLOC\_RW\_POOL

RWS\_\$ALLOC\_RW\_POOL

Allocates read/write storage from a specified pool.

#### FORMAT

storage\_ptr = RWS\_\$ALLOC\_RW\_POOL(alloc\_pool, storage\_sz )

#### RETURN VALUE

##### storage\_ptr

The address of the new storage space, in UNIV\_PTR format. This is a 4-byte integer. A returned address of zero (NIL) means that RWS\_\$ALLOC\_RW\_POOL could not allocate the desired storage.

#### INPUT PARAMETERS

##### alloc\_pool

Pool from which storage will be allocated, in RWS\_\$POOL\_T format. This is a 2-byte integer. Specify one of the following following predefined values:

##### RWS\_\$GLOBAL\_POOL

Global pool makes storage accessible to all processes. Note that pointers are valid in *all* processes because they reserve the *identical* portion of address space.

##### RWS\_\$STD\_POOL

Standard pool makes storage accessible to the calling program only. Most programs use this type.

##### RWS\_\$STREAM\_TM\_POOL

Stream pool makes storage accessible to the calling program and to a program invoked with a UNIX EXEC system call. Use this type when your program needs to pass information across a UNIX EXEC system call.

##### storage\_sz

Number of bytes of storage needed. This is a 4-byte integer.

#### USAGE

RWS\_\$ALLOC\_RW\_POOL allocates a specified number of bytes of read/write storage to the calling process and returns the address of the storage area.

Whether you allocate read/write storage with this call or heap (releaseable read/write) storage with RWS\_\$ALLOC\_HEAP\_POOL depends on how long you want to keep the storage. Once you allocate read/write storage, the storage exists until the program terminates. However, you can explicitly release heap storage once you have finished using it. The heap requires more system overhead initially, to keep track of allocated storage. Read/write storage does not require any system overhead.



Typically, you allocate read/write storage if you do not need to release storage before terminating a program, or if the amount of overhead for a heap is unacceptable. You allocate heap storage if your program requires a substantial amount of storage for a limited time, or if you want to keep your working set as small as possible.

When allocating read/write or heap storage, you control how your program accesses storage by specifying the type of storage pool to use:

- The standard pool (RWS\_\$STD\_POOL) permits access to the calling process only.
- The global pool (RWS\_\$GLOBAL\_POOL) permits access to all processes.
- The stream pool (RWS\_\$STREAM\_TM\_POOL) permits access to the calling program and a program invoked with a UNIX EXEC system call.

The global pool allows you to share information among processes. For example, you might want to create a global queue to pass messages between processes. Note that pointers are valid in *all* processes because all processes reserve an *identical* portion of address space.

The stream pool allows you to make storage accessible between a calling process and an overlay process. For example, the IOS manager uses a stream pool to pass an open stream to a program invoked with an EXEC call. It stores information about that stream in a stream pool.

The following table summarizes the aspects of each type of storage allocation.

#### Summary of Types of Storage Allocation

	Read/Write Storage	Heap Storage
<b>Standard Pool</b>	Storage kept until program exits or until it invokes a program with a UNIX EXEC system call.	Storage kept until you release it with RWS_\$RELEASE_HEAP, the program exits, or the program invokes a program with a UNIX EXEC call.
	No system overhead.	About 16 bytes of system overhead.
	Storage available to local process only.	
<b>Global Pool</b>	Storage kept until reboot.	Storage kept until you release it with RWS_\$RELEASE_HEAP or reboot.
	About 4 bytes of system overhead.	About 4 bytes of system overhead.
	Storage available to all processes.	

**Summary of Types of Storage Allocation, Cont.**

	<b>Read/Write Storage</b>	<b>Heap Storage</b>
<b>Stream Pool</b>	Storage kept until program exits.	Storage kept until you release it with RWS_\$RELEASE_HEAP.
	No system overhead.	About 16 bytes of system overhead.
	Storage available to the local process or to a program invoked with a UNIX EXEC system call.	

**NOTE:** Do not depend on the exact amount of system overhead used in RWS system calls. The amount of overhead is subject to change.

Note that this call replaces the obsolete RWS\_\$ALLOC\_RW call, which we include for maintenance purposes only. For current and future development, use RWS\_\$ALLOC\_RW\_POOL.

FORTRAN programmers: due to FORTRAN calling conventions, RWS\_\$ALLOC is currently the only RWS call you can use to allocate read/write storage. C programmers might want to use the C library routine MALLOC to allocate storage.

**RWS\_\$RELEASE\_HEAP**

Releases storage allocated by the RWS\_\$ALLOC\_HEAP call. *This call is obsolete and is documented here for maintenance purposes only. You should replace it with RWS\_\$ALLOC\_HEAP\_POOL.*

**FORMAT**

RWS\_\$RELEASE\_HEAP (storage\_ptr, status)

**INPUT PARAMETERS****storage\_ptr**

The address heap storage space, in UNIV\_PTR format. This is a 4-byte integer. This must be the pointer returned by a call to RWS\_\$ALLOC\_HEAP.

**OUTPUT PARAMETERS****status**

Completion status, in STATUS\_\$T format. This data type is 4 bytes long. See the RWS Data Types section for more information.

**USAGE**

Note that RWS\_\$RELEASE\_HEAP\_POOL replaces this less efficient call, which we include for maintenance purposes only. For current and future development, use RWS\_\$RELEASE\_HEAP\_POOL.

Use this call to release the storage that you previously allocated with RWS\_\$ALLOC\_HEAP.

RWS\_\$RELEASE\_HEAP\_POOL

RWS\_\$RELEASE\_HEAP\_POOL

Releases storage to the pool from which it was allocated.

## FORMAT

RWS\_\$RELEASE\_HEAP\_POOL (storage\_ptr, alloc\_pool, status)

## INPUT PARAMETERS

### storage\_ptr

Pointer to the address heap storage space, in UNIV\_PTR format. This is a 4-byte integer. This must be the pointer returned by a call to RWS\_\$ALLOC\_HEAP\_POOL.

### alloc\_pool

Pool where storage will be returned to, in RWS\_\$POOL\_T format. This is a 2-byte integer. Specify the same value you specified in the RWS\_\$ALLOC\_HEAP\_POOL call, which is one of the following predefined values:

RWS\_\$GLOBAL\_POOL

Global pool makes storage accessible to all processes. Note that pointers are valid in *all* processes because they reserve the *identical* portion of address space.

RWS\_\$STD\_POOL

Standard pool makes storage accessible to the calling program only. Most programs use this type.

RWS\_\$STREAM\_TM\_POOL

Stream pool makes storage accessible to the calling program and to a program invoked with a UNIX EXEC system call. Use this type when your program needs to pass information invoked with a UNIX EXEC system call.

## OUTPUT PARAMETERS

### status

Completion status, in STATUS\_\$T format. This data type is 4 bytes long. See the RWS Data Types section for more information.

## USAGE

Use RWS\_\$RELEASE\_HEAP\_POOL to release storage to the pool from which it was allocated. You allocate storage to a specific pool with the RWS\_\$ALLOC\_HEAP\_POOL call.

RWS\_\$ALLOC\_HEAP\_POOL dynamically allocates storage from one of the three types of storage pools, and returns a pointer to the new storage. When you no longer need the storage, you release it by passing the "storage\_ptr" and "alloc\_pool" to RWS\_\$RELEASE\_HEAP\_POOL. RWS\_\$RELEASE\_HEAP\_POOL returns the storage to the pool from which it was allocated.

Note that this call replaces the less efficient RWS\_\$RELEASE\_HEAP call, which we include for maintenance purposes only. For current and future development, use RWS\_\$RELEASE\_HEAP\_POOL.

RWS ERRORS

**ERRORS**

RWS\_ \$LEVEL\_ FAILURE

User program wrote over the storage where the system stored the program level information.

RWS\_ \$NOT\_ HEAP\_ ENTRY

Argument to RWS\_ \$RELEASE\_ HEAP did not refer to storage allocated with RWS\_ \$ALLOC\_ HEAP.

RWS\_ \$SCRIBBLED\_ OVER

User program wrote over the storage where the system stored the heap's process information.

RWS\_ \$WRONG\_ LEVEL

Attempted to release storage that was allocated by a program at a superior (lower) program level. This error can occur when using RWS\_ \$STD\_ POOL or RWS\_ \$STREAM\_ TM\_ POOL.

# SIO

## Table of Contents

Data Types  
SIO\_\$CONTROL  
SIO\_\$INQUIRE  
Errors

SIO-2  
SIO-7  
SIO-14  
SIO-20





## SIO

The SIO (Serial I/O) programming calls handle communications to devices across serial I/O lines. This section describes their data types, call syntax, and error codes. Refer to the Introduction at the beginning of this manual for a description of data-type diagrams and call syntax format.

SIO DATA TYPES

CONSTANTS

SIO_\$50	Baud rate.
SIO_\$75	Baud rate.
SIO_\$110	Baud rate.
SIO_\$134	Baud rate.
SIO_\$150	Baud rate.
SIO_\$300	Baud rate.
SIO_\$600	Baud rate.
SIO_\$1200	Baud rate.
SIO_\$2000	Baud rate.
SIO_\$2400	Baud rate.
SIO_\$3600	Baud rate.
SIO_\$4800	Baud rate.
SIO_\$7200	Baud rate.
SIO_\$9600	Baud rate.
SIO_\$19200	Baud rate.
SIO_\$EVEN_PARITY	Possible parity value.
SIO_\$MAX_LINE	Maximum number of SIO lines.
SIO_\$NO_PARITY	Possible parity value.
SIO_\$ODD_PARITY	Possible parity value.
SIO_\$STOP_1	Possible stop bit value.
SIO_\$STOP_1_POINT_5	Possible stop bit value.
SIO_\$STOP_2	Possible stop bit value.
SIO_\$5BPC	Bits per character value.
SIO_\$6BPC	Bits per character value.
SIO_\$7BPC	Bits per character value.
SIO_\$8BPC	Bits per character value.

## DATA TYPES

SIO\_\$ERR\_ENABLES\_T

A 2-byte integer. Determines which errors are enabled. Any combination of the following predefined values:

SIO\_\$CHECK\_PARITY

Check for received parity errors.

SIO\_\$CHECK\_FRAMING

Check for received framing errors.

SIO\_\$CHECK\_DCD\_CHANGE

Check for error when DCD line changes state.

SIO\_\$CHECK\_CTS\_CHANGE

Check for error when CTS line changes state.

SIO\_\$LINE\_T

A 2-byte integer. SIO line number. Possible values are integers from 0 through SIO\_\$MAX\_LINE (3).

SIO\_\$OPT\_T

A 2-byte integer. An SIO option. One of the following predefined values:

SIO\_\$ERASE

Set erase character.

SIO\_\$KILL

Set kill character.

SIO\_\$EOFCHR

Set EOF character.

SIO\_\$RAW

Transparent input and output.

SIO\_\$NO\_ECHO

Do not echo input.

SIO\_\$NO\_NL

Do not special case newlines.

SIO\_\$SPEED

Set bit rate.

SIO\_\$HOST\_SYNC

Use xoff/xon to synchronize with host.

SIO\_\$NLC\_DELAY

Constant delay for newlines.

SIO\_\$QUIT\_ENABLE

Enable quits from this line to calling process.

SIO DATA TYPES

SIO\_\$INPUT\_SYNC  
Respond xoff/xon on receive side.

SIO\_\$LINE  
Return line number (inquire only).

SIO\_\$RTS  
Set/clear RTS bit.

SIO\_\$DTR  
Set/clear DTR bit.

SIO\_\$DCD  
Read DTR bit (inquire only).

SIO\_\$DCD\_ENABLE  
Enable fault on DCD loss.

SIO\_\$CTS  
Read CTS bit (inquire only).

SIO\_\$CTS\_ENABLE  
Enable CTS gating of output.

SIO\_\$PARITY  
Control parity setting/processing.

SIO\_\$BITS\_PER\_CHAR  
Number of bits per character.

SIO\_\$STOP\_BITS  
Number of stop bits.

SIO\_\$ERR\_ENABLE  
Enable error reporting.

SIO\_\$SEND\_BREAK  
Establish break condition on line.

SIO\_\$QUITCHR  
Set quit character.

SIO\_\$BP\_ENABLE  
Enable bit pad processing on line.

SIO\_\$INT\_ENABLE  
Enable interrupts in this process.

SIO\_\$INTCHR  
Set interrupt character.

SIO\_\$SUSP\_ENABLE  
Enable process suspension character.

SIO\_\$SUSPCHR  
Set process suspension character.

SIO\_\$RAW\_NL  
 Display NL/CR on NL output in raw mode.

SIO\_\$UNUSED  
 Unused.

SIO\_\$HUP\_CLOSE  
 Set hangup-on-close.

SIO\_\$RTS\_ENABLE  
 Enable RTS flow control.

SIO\_\$SPEED\_FORCE  
 Set bit rate, even if disturbs partner channel.

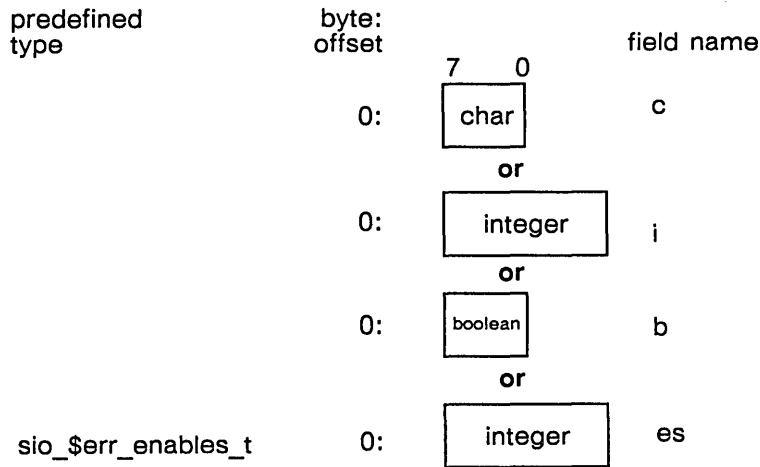
SIO\_\$FLUSH\_IN  
 Flush input buffer.

SIO\_\$FLUSH\_OUT  
 Flush output buffer.

SIO\_\$DRAIN\_OUT  
 Wait for output buffer to drain.

SIO\_\$VALUE\_T

Value corresponding to SIO options. The diagram below illustrates the SIO\_\$VALUE\_T data type:



Field Description:

c  
 A character value.

i  
 A 2-byte integer value.

b  
 A Boolean value.


SIO DATA TYPES

es

A set of enabled errors. This is a 2-byte field.

STATUS\_\$\$T

A status code. The diagram below illustrates the STATUS\_\$\$T data type:



Field Description:

all

All 32 bits in the status code.

fail

The fail bit. If this bit is set, the error was not within the scope of the module invoked, but occurred within a lower-level module (bit 31).

subsys

The subsystem that encountered the error (bits 24 - 30).

modc

The module that encountered the error (bits 16 - 23).

code

A signed number that identifies the type of error that occurred (bits 0 - 15).

## SIO\_\$CONTROL

Changes current settings of serial line options and values.

**FORMAT**

SIO\_\$CONTROL (stream-id, option, value, status)

**INPUT PARAMETERS****stream-id**

Stream ID of a stream attached to a serial line. This is a 2-byte integer.

The stream specified by stream-id must be attached to a serial line. Any other attachment results in an error.

**option**

The attribute that is to be set, in SIO\_\$OPT\_T format. This is a 2-byte integer. Specify only one of the following predefined values:

## SIO\_\$BITS\_PER\_CHAR

Sets the number of bits per character. This option takes a predefined 2-byte integer value. The default is SIO\_\$8BPC, which is 8 bits per character. Possible choices are: SIO\_\$xBPC, where x may be 5, 6, 7, or 8.

## SIO\_\$BP\_ENABLE

Enables/disables processing of bit pad input from a graphics tablet. This option takes a Boolean value. The default is FALSE (disabled).

When enabled, data received on the SIO line is not delivered through IOS\_\$GET. Instead, the SIO driver interrupt routine accumulates data and passes it a point at a time to the display driver. During this processing, subsequent points within plus or minus two in both x and y dimensions are ignored.

SIO\_\$CTS No meaning for this call.

## SIO\_\$CTS\_ENABLE

Sets whether the CTS\_ENABLE mode is on or off. This option takes a Boolean value. The default is FALSE (off).

Some devices use one of the RS-232 control lines for flow control instead of XON/XOFF. If such a line is wired to the CTS line on the connector and if SIO\_\$CTS\_ENABLE is TRUE, then transmission will be inhibited whenever CTS is FALSE.

SIO\_\$DCD No meaning for this call.

## SIO\_\$DCD\_ENABLE

Sets whether the DCD\_ENABLE mode is on or off. This option takes a Boolean value. The default is FALSE (off).

SIO\_\$CONTROL

If the connection is broken (i.e., the remote modem hangs up) the DCD line becomes FALSE. If SIO\_\$DCD\_ENABLE is TRUE then a fault with status FAULT\_\$STOP will occur at the time of the transition of DCD from TRUE to FALSE.

SIO\_\$DRAIN\_OUT

Causes the process to wait until all the characters in the output buffer have been transmitted before returning. This option takes a Boolean value. The default is FALSE (off).

SIO\_\$DTR

Sets the state of the outgoing DTR (Data Terminal Ready) line. This option takes a Boolean value. The default is TRUE (on).

On most modems the DTR line controls whether the modem will answer incoming calls (i.e., when DTR is TRUE.) When it is reset it causes the modem to hang up the phone line.

SIO\_\$EOFCHR

Sets the end-of-file character. This option takes a character value. The default is CTRL/Z.

SIO\_\$ERASE

Sets the erase character, which erases the character immediately before the current cursor position. This option takes a character value. The default is <BACKSPACE>.

SIO\_\$ERR\_ENABLE

Sets which kinds of errors can be reported in calls to IOS\_\$GET on this stream. This option takes a set of values, in SIO\_\$ERR\_ENABLES\_T format. Specify any combination of the following predefined 2-byte integer values:

SIO\_\$CHECK\_PARITY

Report received parity errors.

SIO\_\$CHECK\_FRAMING

Report received framing errors. This value is set by default.

SIO\_\$CHECK\_DCD\_CHANGE

Report an "error" when DCD line changes state.

SIO\_\$CHECK\_CTS\_CHANGE

Report an "error" when CTS line changes state.

SIO\_\$FLUSH\_IN

Causes the input buffer of an SIO line to be flushed. This option takes a Boolean value. The default is FALSE (off).

SIO\_\$FLUSH\_OUT

Causes the output buffer of an SIO line to be flushed. This option takes a Boolean value. The default is FALSE (off).

SIO\_\$HOST\_SYNC

Sets whether HOST\_SYNC mode is on or off. This option takes a Boolean value. The default is TRUE (on).



In `HOST_SYNC` mode, the node sends `XOFF (CTRL/S)` when its input buffer begins to fill, and `XON (CTRL/Q)` when its input buffer begins to empty again. This allows for synchronization of high-speed data transfer from computer to computer.

**SIO\_\$HUP\_CLOSE**

Causes the modem to be hung up on the last close (`IOS_CLOSE`) of the SIO line. The hangup is performed by dropping `DTR` for 3/4 second.

**SIO\_\$INPUT\_SYNC**

Sets whether the incoming sync mode is on or off. This option takes a Boolean value. The default is `FALSE` (off).

It is like `HOST_SYNC` except it controls processing of incoming `XON (CTRL/Q)` or `XOFF (CTRL/S)`. It works in raw or cooked mode.

**SIO\_\$INT\_ENABLE**

Enables/disables interrupts for the current process. This option takes a Boolean value. The default is `FALSE` (disabled).

**SIO\_\$INTCHR** Sets the process interrupt character. (This option is used primarily by `DOMAIN/IX`.) This option takes a character value. The default is `FALSE (CTRL/C)`.

**SIO\_\$KILL** Sets the kill character, which deletes characters from the cursor position to the end of the line. This option takes a character value. The default is `CTRL/X`.

**SIO\_\$LINE** No meaning for this call.

**SIO\_\$NLC\_DELAY**

Sets the value of a time delay to be used following transmission of a line feed character, to allow for carriage motion, scrolling time, and so on. This option takes a 2-byte integer value, specifying the number of milliseconds of delay. The default is zero.

**SIO\_\$NO\_ECHO**

Sets whether `NO_ECHO` mode is on or off. In `NO_ECHO` mode, input characters are not automatically echoed as output. This mode may be used to support a half-duplex connection. `NO_ECHO` mode is off by default.

**SIO\_\$NO\_NL** Sets whether `NO_NL` mode is on or off. This option takes a Boolean value. The default is `FALSE` (off).

Normally, newline characters (decimal 10) are transmitted as a carriage-return, line-feed. In `NO_NL` mode, the newline character is transmitted as is. This mode makes output transparent without going to raw input.

**SIO\_\$PARITY** Sets the state of parity detection or parity generation. This option takes a predefined 2-byte integer value. The default is `SIO_$NO_PARITY`.

Possible choices are: SIO\_\$ODD\_PARITY, SIO\_\$EVEN\_PARITY, and SIO\_\$NO\_PARITY.

If parity is enabled (whether odd or even) then one bit is added to each character. The parity bit is checked by the hardware on received characters and errors are reported, subject to the SIO\_\$ERR\_ENABLE option. If the number of bits per character is fewer than 8, then the parity bit is delivered with the data in raw mode and is stripped in cooked mode.

## SIO\_\$QUITCHR

Sets the quit character. This option takes a character value. The default is CTRL/|.

## SIO\_\$QUIT\_ENABLE

Sets whether THE quit\_enable mode is on or off. This option takes a Boolean value. The default is FALSE (off).

In quit\_enable mode, the node responds to CTRL/| and to the <BREAK> key, if any. The response is a quit fault in the process using SIO\_\$QUIT\_ENABLE. If SIO\_\$QUIT\_ENABLE is FALSE, then neither the <BREAK> key nor the CTRL/| sequence has any effect. In raw input mode, only the <BREAK> and CTRL/| sequence causes a quit fault.

## SIO\_\$RAW

Sets whether raw mode is on or off. This option takes a Boolean value. The default is FALSE (off). In raw mode, full 8-bit bytes are transmitted in both directions, without any interpretation. Each call (e.g., IOS\_\$GET and STREAM\_\$GET\_REC) that reads from the stream returns as many bytes as have been received since the last call.

When raw mode is turned on or off, any input that your program has received, but has not yet read, is flushed from the input buffer.

## SIO\_\$RAW\_NL

Sets whether NO\_NL mode is on or off in raw mode (i.e., when SIO\_\$RAW is TRUE). This option takes a Boolean value. The default is FALSE (off).

Normally, newline characters (decimal 10) are transmitted as a carriage-return, line-feed. In NO\_NL mode, the newline character is transmitted as is. This mode makes output transparent without going to raw input.

## SIO\_\$RTS

Sets the state of the RTS (Request to Send) line. This option takes a Boolean value. The default is TRUE (on).

The RTS line is an outgoing line.

## SIO\_\$RTS\_ENABLE

Enables/disables the RTS (Return to Send) Line. This is a Boolean value. The default is FALSE (off).

If TRUE, the operating system handles flow control. For this to work properly the CTS line must also be enabled.

#### SIO\_\$SEND\_BREAK

Causes a break condition on the line. This option takes a 2-byte integer value, specifying the duration of the break, in milliseconds. A reasonable value is 200.

SIO\_\$SPEED Sets the baud rate of the line. This option takes a predefined 2-byte integer value. The default is SIO\_\$9600.

Possible values are:

SIO\_\$50, SIO\_\$75, SIO\_\$110, SIO\_\$134, SIO\_\$150,  
SIO\_\$300, SIO\_\$600, SIO\_\$1200, SIO\_\$2000, SIO\_\$2400,  
SIO\_\$3600, SIO\_\$4800, SIO\_\$7200, SIO\_\$9600, SIO\_\$19200.

Some machine types are configured so that certain SIO lines are "partnered" with each other. If you use the SIO\_\$SPEED option to set a line's baud rate to a rate that is incompatible with the partner's rate, you receive the error status, SIO\_\$INCOMPATIBLE\_SPEED. However, you can override this error using the SIO\_\$SPEED\_FORCE option. If you use SIO\_\$SPEED\_FORCE to set the speed of a line and the new speed is incompatible with the partner, the speed of the partner is changed to 9600 baud. See the Usage section for details about partnered lines and incompatible speeds.

#### SIO\_\$SPEED\_FORCE

Sets the baud rate of the line even if the partner line's speed is incompatible. This option takes a predefined 2-byte integer value.

Possible values are the same as for SIO\_\$SPEED:

SIO\_\$50, SIO\_\$75, SIO\_\$110, SIO\_\$134, SIO\_\$150,  
SIO\_\$300, SIO\_\$600, SIO\_\$1200, SIO\_\$2000, SIO\_\$2400,  
SIO\_\$3600, SIO\_\$4800, SIO\_\$7200, SIO\_\$9600, SIO\_\$19200.

Some machine types are configured so that certain SIO lines are "partnered" with each other. If you use the SIO\_\$SPEED option to set a line's baud rate to a rate that is incompatible with the partner line's rate, you receive the error status, SIO\_\$INCOMPATIBLE\_SPEED. However, you can override this error using the SIO\_\$SPEED\_FORCE option. If you use SIO\_\$SPEED\_FORCE to set the speed of a line and the new speed is incompatible with the partner, the speed of the partner is changed to 9600 baud. See the Usage section for details about partnered lines and incompatible speeds.

#### SIO\_\$STOP\_BITS

Sets the number of stop bits. This option takes a predefined 2-byte integer value. The default is SIO\_\$STOP\_1. Possible values are: SIO\_\$STOP\_x where x may be 1, 1\_POINT\_5, or 2.

## SIO\_\$CONTROL

### SIO\_\$SUSP\_ENABLE

Enables/disables suspend faults for the current process. This option takes a Boolean value. The default is FALSE (disabled).

### SIO\_\$SUSPCHR

Sets the process suspend character. (This option is used primarily by DOMAIN/IX.) This option takes a character value. The default is CRTL/P.

### value

Each of the SIO\_\$CONTROL options accepts a corresponding value. For the character options, the value is simply the character. For the mode-setting options, the value is a Boolean (LOGICAL) data item. For most of the remaining options, the value is a 2-byte integer. In one case, you may specify a set of values. The type of value required for each option is described along with the option, above.

## OUTPUT PARAMETERS

### status

Completion status, in STATUS\_\$T format. This data type is 4 bytes long. See the SIO Data Types section for more information.

Possible values are:

STATUS\_\$OK The operation completed successfully.

### IOS\_\$NOT\_OPEN

No stream is open on the specified stream ID.

### SIO\_\$STREAM\_NOT\_SIO

The specified stream ID is not connected to a serial line.

### SIO\_\$BAD\_OPTION

The call specified an invalid option name.

### SIO\_\$INCOMPATIBLE\_SPEED

The specified speed is incompatible with the speed of the line's partner.

## USAGE

To poll the serial line for unread input, use IOS\_\$GET with the IOS\_\$COND\_OPT option.

The hardware configuration for some machine types is such that certain SIO lines are "partnered" with each other.

Below is a list of machine types and the SIO lines that are partnered on them.

<i>Machine Type</i>	<i>Partnered Lines</i>
DN400	No
DN420	Partners
DN600	
DN300	1,2
DSP80	1,2
DN460	0,1
DN660	2,3
DN550	1,2
DN3000	No Partners

A characteristic of partnered lines is that some baud rates are incompatible. The following lists show the baud rates that are incompatible for partnered lines.

<i>Incompatible Rates A</i>	<i>Incompatible Rates B</i>
SIO_ \$50 SIO_ \$7200	SIO_ \$75 SIO_ \$150 SIO_ \$2000 SIO_ \$19200

If one partner is set to a baud rate in the A list, attempting to set the other partner to a baud rate in the B list (using the SIO\_ \$SPEED option) will result in the error SIO\_ \$INCOMPATIBLE\_ SPEED. The same is true for the reverse (having a partnered line set to a rate in the B list and attempting to set its partner to a rate in the A list). Speeds other than those in the two lists are compatible.

You may force a line's baud rate to a rate that is incompatible to its partner, by using the SIO\_ \$SPEED\_ FORCE option. This will change the specified line's speed; however, it will also change the speed of the partnered line to SIO\_ \$9600 (which is always a compatible speed).

SIO\_\$INQUIRE

SIO\_\$INQUIRE

Obtains current settings of serial line options and values.

## FORMAT

SIO\_\$INQUIRE (stream-id, option, value, status)

## INPUT PARAMETERS

### stream-id

Stream-id of a stream attached to a serial line. This is a 2-byte integer.

### option

The attribute that is to be reported, in SIO\_\$OPT\_T format. This is a 2-byte integer. One of the following predefined values:

SIO\_\$BITS\_PER\_CHAR

Returns the number of bits per character. This option takes a predefined 2-byte integer value. The default is SIO\_\$8BPC, which is 8 bits per character. Possible choices are: SIO\_\$xBPC, where x may be 5, 6, 7, or 8.

SIO\_\$BP\_ENABLE

Enables/disables processing of bit pad input from a graphics tablet. This option takes a Boolean value. The default is FALSE (disabled).

When enabled, data received on the SIO line is not delivered through IOS\_\$GET. Instead, the SIO driver interrupt routine accumulates data and passes it a point at a time to the display driver. During this processing, subsequent points within plus or minus two in both x and y dimensions are ignored.

SIO\_\$CTS

Returns the state of the CTS (Clear to Send) line. This is an incoming line.

SIO\_\$CTS\_ENABLE

Returns whether CTS\_ENABLE mode is on or off. This option takes a Boolean value. The default is FALSE (off).

Some devices use one of the RS-232 control lines for flow control instead of XON/XOFF. If such a line is wired to the CTS line on the connector and if SIO\_\$CTS\_ENABLE is TRUE, then transmission will be inhibited whenever CTS is FALSE.

SIO\_\$DCD

Reports the state of the DCD (Data Carrier Detect) line. It is an incoming line, which usually means there is an active modem at the other end of the phone line.

SIO\_\$DCD\_ENABLE

Returns whether DCD\_ENABLE mode is on or off. This option takes a Boolean value. The default is FALSE (off).

If the connection is broken (i.e., the remote modem hangs up) the DCD line becomes FALSE. If SIO\_\$DCD\_ENABLE is TRUE then a fault with status FAULT\_\$STOP occurs at the time of the transition of DCD from TRUE to FALSE.

## SIO\_\$DRAIN\_OUT

No meaning for this call.

## SIO\_\$DTR

Returns the state of the outgoing DTR (Data Terminal Ready) line. This option takes a Boolean value. The default is TRUE (on).

On most modems the DTR line controls whether the modem answers incoming calls (i.e., when DTR is TRUE.) When it is reset it causes the modem to hang up the phone line.

SIO\_\$EOFCHR Returns the end-of-file character. This option takes a character value. The default is CTRL/Z.

## SIO\_\$ERASE

Returns the erase character, which erases the character immediately before the current cursor position. This option takes a character value. The default is <BACKSPACE>.

## SIO\_\$ERR\_ENABLE

Returns which kinds of errors can be reported in calls to IOS\_\$GET on this stream. This option takes a set of values, in SIO\_\$ERR\_ENABLES\_T format. Specify any combination of the following predefined 2-byte integer values:

## SIO\_\$CHECK\_PARITY

Report received parity errors.

## SIO\_\$CHECK\_FRAMING

Report received framing errors. This value is set by default.

## SIO\_\$CHECK\_DCD\_CHANGE

Report an "error" when DCD line changes state.

## SIO\_\$CHECK\_CTS\_CHANGE

Report an "error" when CTS line changes state.

## SIO\_\$FLUSH\_IN

No meaning for this call.

## SIO\_\$FLUSH\_OUT

No meaning for this call.

## SIO\_\$HOST\_SYNC

Returns whether the HOST\_SYNC mode is on or off. This option takes a Boolean value. The default is TRUE (on).

In HOST\_SYNC mode, the node sends XOFF (CTRL/S) when its input buffer begins to fill, and XON (CTRL/Q) when its input buffer begins to

## SIO\_\$INQUIRE

empty again. This allows for synchronization of high-speed data transfer from computer to computer.

### SIO\_\$HUP\_CLOSE

Causes the modem to be hung up on the last close (IOS\_\$CLOSE) of the SIO line. The hangup is performed by dropping DTR for 3/4 second. This option takes a Boolean value. The default is FALSE (off).

### SIO\_\$INPUT\_SYNC

Returns whether incoming sync mode is on or off. This option takes a Boolean value. The default is FALSE (off).

It is like HOST\_SYNC except it controls processing of {\_incoming\_} XON (CTRL/Q) or XOFF (CTRL/S). It works in raw or cooked mode.

### SIO\_\$INT\_ENABLE

Enables/disables interrupts for the current process. This option takes a Boolean value. The default is FALSE (disabled).

SIO\_\$INTCHR Returns the process interrupt character. (This option is used primarily by DOMAIN/IX.) This option takes a character value. The default is FALSE (CTRL/C).

SIO\_\$KILL Returns the kill character, which deletes characters from the cursor position to the end of the line. This option takes a character value. The default is CTRL/X.

SIO\_\$LINE Returns the serial line number corresponding to the stream ID. This option returns an integer value from 0 to 3.

### SIO\_\$NLC\_DELAY

Returns the value of a time delay to be used following transmission of a line feed character, to allow for carriage motion, scrolling time, and so on. This option takes a 2-byte integer value, specifying the number of milliseconds of delay. The default is zero.

### SIO\_\$NO\_ECHO

Returns whether the NO\_ECHO mode is on or off. In NO\_ECHO mode, input characters are not automatically echoed as output. This mode may be used to support a half-duplex connection. NO\_ECHO mode is off by default.

SIO\_\$NO\_NL Returns whether the NO\_NL mode is on or off. This option takes a Boolean value. The default is FALSE (off).

Normally, newline characters (decimal 10) are transmitted as a carriage-return, line-feed. In NO\_NL mode, the newline character is transmitted as is. This mode makes output transparent without going to raw input.

SIO\_\$PARITY Returns the state of parity detection or parity generation. This option takes a predefined 2-byte integer value. The default is



SIO\_\$NO\_PARITY. Possible choices are: SIO\_\$ODD\_PARITY, and SIO\_\$EVEN\_PARITY, and SIO\_\$NO\_PARITY.

If parity is enabled (whether odd or even) then one bit is added to each character. The parity bit is checked by the hardware on received characters and errors are reported, subject to the SIO\_\$ERR\_ENABLE option. If the number of bits per character is fewer than 8, then the parity bit is delivered with the data in raw mode and is stripped in cooked mode.

#### SIO\_\$QUITCHR

Returns the quit character. This option takes a character value. The default is CTRL/].

#### SIO\_\$QUIT\_ENABLE

Returns whether the QUIT\_ENABLE mode is on or off. This option takes a Boolean value. The default is FALSE (off).

In QUIT\_ENABLE mode, the node responds to CTRL/] and to the <BREAK> key, if any. The response is a quit fault in the process using SIO\_\$QUIT\_ENABLE. If SIO\_\$QUIT\_ENABLE is FALSE then neither the <BREAK> key nor the CTRL/] sequence has any effect. In raw input mode only the <BREAK>, and CTRL/] sequence causes a quit fault.

#### SIO\_\$RAW

Returns whether raw mode is on or off. This option takes a Boolean value. The default is FALSE (off). In raw mode, full 8-bit bytes are transmitted in both directions, without any interpretation. Each call (e.g., IOS\_\$GET and STREAM\_\$GET\_REC) that reads from the stream returns as many bytes as have been received since the last call.

When raw mode is turned on or off, any input that your program has received, but has not yet read, is flushed from the input buffer.

#### SIO\_\$RAW\_NL

Returns whether the NO\_NL mode is on or off in raw mode.(i.e., when SIO\_\$RAW is TRUE). This option takes a Boolean value. The default is FALSE (off).

Normally, newline characters (decimal 10) are transmitted as a carriage-return, line-feed. In NO\_NL mode, the newline character is transmitted as is. This mode makes output transparent without going to raw input.

#### SIO\_\$RTS

Returns the state of the RTS (Request to Send) line. This option takes a Boolean value. The default is TRUE (on).

The RTS line is an outgoing line.

#### SIO\_\$RTS\_ENABLE

Enables/disables the RTS (Return to Send) Line. This is a Boolean value. The default is FALSE (off).

## SIO\_ \$INQUIRE

If TRUE, the operating system handles flow control. For this to work properly the CTS line must also be enabled.

### SIO\_ \$SEND\_ BREAK

No meaning for this call.

**SIO\_ \$SPEED** Returns the baud rate of the line. This option takes a predefined 2-byte integer value. The default is SIO\_ \$9600.

Possible values are:

SIO\_ \$50, SIO\_ \$75, SIO\_ \$110, SIO\_ \$134, SIO\_ \$150,  
SIO\_ \$300, SIO\_ \$600, SIO\_ \$1200, SIO\_ \$2000, SIO\_ \$2400,  
SIO\_ \$3600, SIO\_ \$4800, SIO\_ \$7200, SIO\_ \$9600, SIO\_ \$19200.

### SIO\_ \$SPEED\_ FORCE

No meaning for this call.

### SIO\_ \$STOP\_ BITS

Returns the number of stop bits. This option takes a predefined 2-byte integer value. The default is SIO\_ \$STOP\_ 1. Possible values are: SIO\_ \$STOP\_ x where x may be 1, 1\_ POINT\_ 5, or 2.

### SIO\_ \$SUSP\_ ENABLE

Enables/disables suspend faults for the current process. This option takes a Boolean value. The default is FALSE (disabled).

### SIO\_ \$SUSPCHR

Returns the process suspend character. (This option is used primarily by DOMAIN/IX.) This option takes a character value. The default is CRTL/P.

## OUTPUT PARAMETERS

### value

Each of the SIO\_ \$INQUIRE options returns a corresponding value. For the character options, the value is simply the character. For the mode-setting options, the value is a Boolean (LOGICAL) data item. For most of the remaining options, the value is a 2-byte integer. In one case, a set of values may be returned. The type of value returned for each option is described along with the option, above.

### status

Completion status, in STATUS\_ \$T format. This data type is 4 bytes long. See the SIO Data Types section for more information.

Possible values are:

STATUS\_ \$OK The operation completed successfully.

### IOS\_ \$NOT\_ OPEN

No stream is open on the specified stream ID.

### SIO\_ \$STREAM\_ NOT\_ SIO

The specified stream ID is not connected to a serial line.

## SIO\_\$BAD\_OPTION

The call specified an invalid option name.

**USAGE**

The stream specified by stream ID must be attached to a serial line. Any other attachment results in an error.

When raw mode is turned on or off, any input that your program has received, but has not yet read, is flushed from the input buffer.

To poll the serial line for unread input, use IOS\_\$GET with the IOS\_\$COND\_OPT option.

SIO ERRORS

**ERRORS**

STATUS\_ \$OK

Successful completion.

SIO\_ \$BAD\_ OPTION

Bad option parameter.

SIO\_ \$ILLEGAL\_ STRID

Illegal stream ID.

SIO\_ \$INCOMPATIBLE\_ SPEED

Speed incompatible with partner SIO line.

# SMD

## Table of Contents

SMD_\$BLT_U	SMD-2
SMD_\$BORROW_DISPLAY_NC_U	SMD-4
SMD_\$BORROW_DISPLAY_U	SMD-5
SMD_\$CLEAR_KBD_CURSOR_U	SMD-6
SMD_\$CLEAR_WINDOW_U	SMD-7
SMD_\$COLOR	SMD-8
SMD_\$COND_EVENT_WAIT_U	SMD-9
SMD_\$COND_INPUT_U	SMD-10
SMD_\$DRAW_ABS_U	SMD-11
SMD_\$DRAW_BOX_U	SMD-12
SMD_\$DRAW_REL_U	SMD-13
SMD_\$EVENT_WAIT_U	SMD-14
SMD_\$GET_EC	SMD-15
SMD_\$INQ_DISP_TYPE	SMD-16
SMD_\$LOAD_FONT_FILE_U	SMD-17
SMD_\$LOAD_FONT_U	SMD-18
SMD_\$MAP_DISPLAY_U	SMD-19
SMD_\$MOVE_ABS_U	SMD-20
SMD_\$MOVE_KBD_CURSOR_U	SMD-21
SMD_\$MOVE_REL_U	SMD-22
SMD_\$OP_WAIT_U	SMD-23
SMD_\$RETURN_DISPLAY_U	SMD-24
SMD_\$SET_QUIT_CHAR	SMD-25
SMD_\$SET_TP_CURSOR	SMD-26
SMD_\$SOFT_SCROLL_U	SMD-27
SMD_\$STOP_TP_CURSOR	SMD-29
SMD_\$TP_DIRECT	SMD-30
SMD_\$TP_DISABLE	SMD-31
SMD_\$TP_ENABLE	SMD-32
SMD_\$UNLOAD_FONT_FILE_U	SMD-33
SMD_\$UNLOAD_FONT_U	SMD-34
SMD_\$UNMAP_DISPLAY_U	SMD-35
SMD_\$VECTOR_INIT_U	SMD-36
SMD_\$WRITE_STRING_U	SMD-37
Errors	SMD-39



## SMD

The SMD (Screen Manager) programming calls provide direct control over black-and-white displays. This section describes the call syntax and the error codes. Refer to the Introduction at the beginning of this manual for a description of call syntax format. Refer to the SMD insert files for data-type descriptions.

SMD\_\$BLT\_U

SMD\_\$BLT\_U

Starts a bit transfer from one area of display memory to another.

## FORMAT

SMD\_\$BLT\_U (registers, status)

## INPUT PARAMETERS

### registers

Values for the BLT register, in SMD\_\$BLT\_REGS\_T format. This is a thirteen-element array of 2-byte integers.

## OUTPUT PARAMETERS

### status

Completion status, in STATUS\_\$T format. This data type is 4 bytes long.

## USAGE

SMD\_\$BLT\_U starts the block transfer within display memory.

The BLT register contains values for CS; CD; source and destination Xs, Xe, Ys, and Ye; and the display mode.

In FORTRAN programs, specify the display mode as the sum of selected variables from Table SMD-1.

By default, the display driver waits for the BLT to complete before returning to the calling program. (This is a "busy" wait, meaning the CPU is active while waiting.) If the display mode includes IDONE, however, control returns to the caller immediately, and generates the SMD event SMD\_\$EVENT\_SCROLL\_BLT\_COMPLETE (see SMD\_\$EVENT\_WAIT\_U) when it finishes.

If the display mode includes IDONE, control returns to the calling program immediately, as noted above. However, if display memory is mapped into the process's address space, the program must not reference display memory or call any of the following vector-drawing routines, until the BLT completes:

SMD\_\$DRAW\_ABS\_U  
SMD\_\$DRAW\_REL\_U  
SMD\_\$MOVE\_ABS\_U  
SMD\_\$MOVE\_REL\_U



Table SMD-1. Display Mode Values

<i>Mnemonic</i>	<i>Meaning</i>
SMD_\$CLRMODE	Fill destination with a constant.
SMD_\$DECR	Source overlaps destination and x is being decremented, that is, destination is to right of source.
SMD_\$IDONE	Start BLT and immediately return control to calling program.
SMD_\$NONINTERLACE	Disable hardware interlacing.
SMD_\$BLT	Perform bit BLT operation. Required in all calls to SMD_\$BLT_U.

SMD\_\$BORROW\_DISPLAY\_NC\_U

SMD\_\$BORROW\_DISPLAY\_NC\_U

Requests use of the display driver and display memory without clearing the screen (black and white only).

## FORMAT

SMD\_\$BORROW\_DISPLAY\_NC\_U (unit, status)

## INPUT PARAMETERS

### unit

Unit number of the display to be used. This is a 2-byte integer. Currently, the only valid unit number is 1.

## OUTPUT PARAMETERS

### status

Completion status, in STATUS\_\$T format. This data type is 4 bytes long.

## USAGE

SMD\_\$BORROW\_DISPLAY\_NC\_U requests the use of the display memory and the display (SMD) driver routines. A program must execute this routine (or SMD\_\$BORROW\_DISPLAY\_U) before it can call any other display driver routines.

This procedure gains exclusive use of the display for the calling program. The display manager continues to operate for all other processes and pads, but the screen does not reflect its actions. Control of the screen returns to the display manager when SMD\_\$RETURN\_DISPLAY\_U is executed, or when you type CTRL/Q.

To gain access to display memory, call SMD\_\$MAP\_DISPLAY\_U.

**SMD\_\$BORROW\_DISPLAY\_U**

Requests use of the display driver and display memory, and clears the screen.

**FORMAT**

SMD\_\$BORROW\_DISPLAY\_U (unit, status)

**INPUT PARAMETERS****unit**

Unit number of the display to be used. This is a 2-byte integer. Currently, the only valid unit number is 1.

**OUTPUT PARAMETERS****status**

Completion status, in STATUS\_\$T format. This data type is 4 bytes long.

**USAGE**

SMD\_\$BORROW\_DISPLAY\_U requests the use of the display memory and the display driver routines. A program must execute this routine (or SMD\_\$BORROW\_DISPLAY\_NC\_U) before it can call any other display driver routines.

The display is cleared when this procedure is executed.

SMD\_\$BORROW\_DISPLAY\_NC\_U performs an identical borrowing operation, but does not clear the screen.

This procedure gains exclusive use of the Display for the calling program. The Display Manager continues to operate for all other processes and pads, but the screen does not reflect its actions. Control of the screen returns to the display manager when SMD\_\$RETURN\_DISPLAY\_U is executed, or when you type CTRL/Q.

To gain access to display memory, call SMD\_\$MAP\_DISPLAY\_U.

SMD\_\$CLEAR\_KBD\_CURSOR\_U

SMD\_\$CLEAR\_KBD\_CURSOR\_U

Clears the keyboard cursor from the display.

#### FORMAT

SMD\_\$CLEAR\_KBD\_CURSOR\_U (status)

#### OUTPUT PARAMETERS

##### status

Completion status, in STATUS\_\$T format. This data type is 4 bytes long.

#### USAGE

SMD\_\$CLEAR\_KBD\_CURSOR\_U disables the keyboard cursor and removes it from the display. To re-enable the cursor, call SMD\_\$MOVE\_KBD\_CURSOR\_U.

SMD\_\$CLEAR\_WINDOW\_U

Clears an area on the screen.

#### FORMAT

SMD\_\$CLEAR\_WINDOW\_U (boundaries, status)

#### INPUT PARAMETERS

##### boundaries

The x and y coordinates of the destination area to be cleared, in SMD\_\$WINDOW\_LIMITS\_T format. This data type is 8 bytes long. In FORTRAN, use a four-element array of 2-byte integers.

#### OUTPUT PARAMETERS

##### status

Completion status, in STATUS\_\$T format. This data type is 4 bytes long.

#### USAGE

SMD\_\$CLEAR\_WINDOW\_U clears the area of the screen within the boundaries.

This procedure returns control to its caller after the area is clear.

SMD\_\$COLOR

SMD\_\$COLOR

Sets the color of lines drawn on the display.

#### FORMAT

SMD\_\$COLOR (color)

#### INPUT PARAMETERS

##### color

Either SMD\_\$WHITE or SMD\_\$BLACK. This is a 2-byte integer.

#### USAGE

The color set with this call is used in all subsequent vector or box drawing calls executed by the program, until another SMD\_\$COLOR call is executed.

SMD\_\$WHITE makes subsequent vectors white or green. SMD\_\$BLACK makes subsequent vectors black.

This call does not change the color of the background.

**SMD\_\$COND\_EVENT\_WAIT\_U**

Checks an SMD eventcount, but does not wait.

**FORMAT**

SMD\_\$COND\_EVENT\_WAIT\_U (event-type, event-data, reserved, status)

**OUTPUT PARAMETERS****event-type**

The type of event that occurred. This is a 2-byte integer. Possible values are SMD\_\$INPUT, SMD\_\$SCROLL\_BLT\_COMPLETE, SMD\_\$TPAD\_DATA, SMD\_\$TPAD\_AND\_INPUT, and SMD\_\$NO\_EVENT.

**event\_data**

The data associated with the event, in SMD\_\$EVENT\_DATA\_T format. This is a 2-byte integer.

**reserved**

A 2-byte integer; reserved for future use.

**status**

Completion status, in STATUS\_\$T format. This data type is 4 bytes long.

**USAGE**

SMD\_\$COND\_EVENT\_WAIT\_U causes no suspension of the calling program. Programs can use this procedure to check for keyboard or touchpad input.

The only difference between SMD\_\$COND\_EVENT\_WAIT\_U and SMD\_\$EVENT\_WAIT\_U is that the first checks the eventcount but does not wait. The second waits for an event.

The SMD\_\$COND\_EVENT\_WAIT\_U routine returns one type of event (SMD\_\$NO\_EVENT) that SMD\_\$EVENT\_WAIT\_U does not. When this type is output, it means nothing happened.

SMD\_\$COND\_INPUT\_U

SMD\_\$COND\_INPUT\_U

Returns a character if one has been typed.

#### FORMAT

input-flag = SMD\_\$COND\_INPUT\_U(char)

#### RETURN VALUE

##### input-flag

Boolean (LOGICAL) value. Contains TRUE if a character has been typed and FALSE otherwise.

#### OUTPUT PARAMETERS

##### char

The character typed at the keyboard. This is a character variable.

#### USAGE

If a character has been typed at the keyboard, the value of this function is TRUE. The function returns the character and removes it from the keyboard input buffer.

If no characters have been typed, the value of the function is FALSE, and the returned char parameter is undefined.



**SMD\_\$DRAW\_ABS\_U**

Draws a vector given an absolute position.

**FORMAT**

SMD\_\$DRAW\_ABS\_U (column, line)

**INPUT PARAMETERS****column**

The number of the column to which the vector will be drawn. This is a 2-byte integer in the range 0 to 799.

**line**

The number of the line to which the vector will be drawn. This is a 2-byte integer in the range 0 to 1023.

**USAGE**

SMD\_\$DRAW\_ABS\_U draws a vector from the current position to the point specified by (column, line).

Call SMD\_\$VECTOR\_INIT\_U once to initialize the vector drawing package before using this procedure.

No error checking is performed on the arguments, for optimal performance. Incorrect program operation occurs if a column or line value is outside the specified range.

The current position is updated to (column, line) upon completion of this procedure. Use SMD\_\$MOVE\_ABS\_U or SMD\_\$MOVE\_REL\_U to set the position without drawing a vector.

SMD\_\$DRAW\_BOX\_U

SMD\_\$DRAW\_BOX\_U

Draws a box on the screen.

#### FORMAT

SMD\_\$DRAW\_BOX\_U (boundaries, status)

#### INPUT PARAMETERS

##### boundaries

The x and y coordinates of the box on the screen, in SMD\_\$WINDOW\_LIMITS\_T format. This data type is 8 bytes long. In FORTRAN this is a four-element array of 2-byte integers.

#### OUTPUT PARAMETERS

##### status

Completion status, in STATUS\_\$T format. This data type is 4 bytes long.

#### USAGE

SMD\_\$DRAW\_BOX\_U draws lines vertically and horizontally to connect the supplied endpoints.

The supplied values for Xs and Ys must be less than Xe and Ye, respectively.

SMD\_\$DRAW\_REL\_U

Draws a vector given a relative position.

#### FORMAT

SMD\_\$DRAW\_REL\_U (column, line)

#### INPUT PARAMETERS

##### column

The column number, relative to the current position, to which the vector will be drawn. This is a two-byte integer.

##### line

The line number, relative to the current position, to which the vector will be drawn. This is a two-byte integer.

#### USAGE

SMD\_\$DRAW\_REL\_U draws a vector from the current position to the point computed by adding the value of column to the current column and adding the value of line to the current line.

Call SMD\_\$VECTOR\_INIT\_U to initialize the vector drawing package before using this procedure.

When the value for column is added to the current position column number, the sum must be between 0 and 799 for portrait displays, or between 0 and 1023 for landscape displays. Similarly, when the value for line is added to the current position line number, the sum must be between 0 and 1023 for portrait displays, or 0 and 799 for landscape displays.

No error checking is performed on the arguments, for optimal performance. Incorrect program operation occurs if a computed column or line value is outside the specified range.

The current position is updated to the computed column and line values by this procedure. Use SMD\_\$MOVE\_REL\_U or SMD\_\$MOVE\_ABS\_U to set the position without drawing a vector.

SMD\_\$EVENT\_WAIT\_U

SMD\_\$EVENT\_WAIT\_U

Suspends the calling process until you type characters at the keyboard or until the current scroll or BLT is complete.

## FORMAT

SMD\_\$EVENT\_WAIT\_U (event-type, event-data, reserved, status)

## OUTPUT PARAMETERS

### event-type

The type of event that occurred. This is a 2-byte integer. Possible values are the following:

SMD\_\$INPUT  
SMD\_\$SCROLL\_BLT\_COMPLETE  
SMD\_\$TPAD\_DATA  
SMD\_\$TPAD\_AND\_INPUT

### event-data

The data associated with the event, in SMD\_\$EVENT\_DATA\_T format. This is a 2-byte integer.

### reserved

A 2-byte integer; reserved for future use.

### status

Completion status, in STATUS\_\$T format. This data type is 4 bytes long.

## USAGE

SMD\_\$EVENT\_WAIT\_U suspends the calling process until a display driver event occurs. Programs can use this procedure to read keyboard input.

An SMD\_\$INPUT event occurs when you type a character.

An SMD\_\$SCROLL\_BLT\_COMPLETE event occurs when a block transfer (BLT) is complete, or when the last block transfer required as part of a scrolling operation has been started.

The display driver notifies the calling program of SMD\_\$INPUT, SMD\_\$SCROLL\_BLT\_COMPLETE, SMD\_\$TPAD\_DATA, and SMD\_\$TPAD\_AND\_INPUT events.

**SMD\_\$GET\_EC**

Gets the eventcount address of the eventcount that will be advanced upon keyboard input or when a BLT is done.

**FORMAT**

SMD\_\$GET\_EC (smd-key, eventcount-pointer, status)

**INPUT PARAMETERS****smd-key**

This specifies which eventcount to obtain. It is in SMD\_\$EC\_KEY\_T format and may have a value of either SMD\_\$INPUT\_EC (for the keyboard) or SMD\_\$SCROLL\_BLT\_EC (for a user-initiated BLT.) This is a 2-byte integer.

**OUTPUT PARAMETERS****eventcount\_pointer**

The eventcount address to be obtained, in EC2\_\$PTR\_T format. EC2\_\$PTR\_T is a pointer to an EC2\_\$EVENTCOUNT\_T. This is a 4-byte integer.

**status**

Completion status, in STATUS\_\$T format. This data type is 4 bytes long.

**USAGE**

The SMD eventcounter is advanced whenever anything is entered via the keyboard and also whenever any user-initiated BLT is done.

SMD\_\$INQ\_DISP\_TYPE

SMD\_\$INQ\_DISP\_TYPE

Returns the type of the display physically attached to the given unit number.

## FORMAT

display\_type = SMD\_\$INQ\_DISP\_TYPE (unit)

## RETURN VALUE

### display\_type

The display configuration in `smd_$display_type_t` format. This is a 2-byte integer. It has one of the following predefined values:

SMD_\$NONE	No display
SMD_\$BW_15P	Black and white portrait
SMD_\$BW_19L	Black and white landscape
SMD_\$COLOR_DISPLAY	Color display (1024 x 1024)
SMD_\$800_COLOR	Color display with fewer pixels (1024 x 800)
SMD_\$COLOR2_DISPLAY	Color display (1280x1024x8)
SMD_\$COLOR3_DISPLAY	Color display (1024x800x8)
SMD_\$COLOR4_DISPLAY	Color display (1024x800x4)

## INPUT PARAMETERS

### unit

This parameter (a 2-byte integer) has three possible meanings, as follows:

1. The display unit, if the graphics routines are to operate in a borrowed display. Currently, the only valid display unit number for borrow-display mode is 1.
2. The stream identifier for the pad, if the graphics routines are to operate in frame or direct mode. Use `STREAM_$ID_T` format.
3. Any value, such as zero, if the graphics routines do not use the display.

## USAGE

Use this call to return the type of the display physically attached to the given unit number.

**SMD\_\$LOAD\_FONT\_FILE\_U**

Loads a font file.

**FORMAT**

SMD\_\$LOAD\_FONT\_FILE\_U (pathname, name-length, font-id, status)

**INPUT PARAMETERS****pathname**

Pathname, in NAME\_\$PNAME\_T format, of the file containing the font to be loaded.

**name-length**

The number of characters in the pathname. This is a 2-byte integer.

**OUTPUT PARAMETERS****font-id**

The internal identifier assigned to the font. Font-id is a 2-byte integer.

**status**

Completion status, in STATUS\_\$T format. This data type is 4 bytes long.

**USAGE**

SMD\_\$LOAD\_FONT\_FILE\_U loads the font in the named file into display memory and assigns an identifier to the font. Your program passes the font-id to SMD\_\$WRITE\_STRING to identify the font.

The images of all loaded fonts coexist in the invisible 28-K byte portion of display memory. This area is large enough for about eight small fonts.

If insufficient space is available in either display memory or internal tables to load the font, SMD\_\$LOAD\_FONT\_FILE\_U returns an error. In this case, your program must unload one or more font files to create space for the new font.

To unload fonts loaded with this routine, use SMD\_\$UNLOAD\_FONT\_FILE\_U.

Fonts loaded with this routine are no longer usable when the program exits or aborts. They are not, however, unloaded from display memory.

SMD\_\$LOAD\_FONT\_U

SMD\_\$LOAD\_FONT\_U

Loads a font into display memory and returns a font-id.

## FORMAT

font-id = SMD\_\$LOAD\_FONT\_U (table-ptr, status)

## RETURN VALUE

**font-id**

The internal identifier assigned to the font. This is a two-byte integer.

## INPUT PARAMETERS

**table-ptr**

Address of the font table. This is a 4-byte integer.

## OUTPUT PARAMETERS

**status**

Completion status, in STATUS\_\$T format. This data type is 4 bytes long.

## USAGE

This function returns an integer font-id. Use the font-id to identify the font in calls to SMD\_\$WRITE\_STRING\_U.

FORTRAN programs can use the IADDR function to get an address. Pascal programs can use the ADDR function.

The MS\_\$MAP call can be used to map a font file prior to using this call.

The display driver loads the font into any available space in the invisible 28K bytes of display memory. This area is large enough for about eight small fonts.

If insufficient space remains in invisible display memory or for internal tables, an error occurs. Your program must then unload one or more fonts to make room for the new one.

To unload the font, use SMD\_\$UNLOAD\_FONT\_U.



**SMD\_\$MAP\_DISPLAY\_U**

Maps display memory into the process' address space.

**FORMAT**

SMD\_\$MAP\_DISPLAY\_U (display-address, status)

**OUTPUT PARAMETERS****display-address**

The address of the first byte of the display memory. Display memory is mapped starting at this address, for the next 128-K bytes. This value is in SMD\_\$DISPLAY\_MEMORY\_PTR\_T format. This is a 4-byte integer.

**status**

Completion status, in STATUS\_\$T format. This data type is 4 bytes long.

**USAGE**

SMD\_\$MAP\_DISPLAY\_U creates an association between the display memory and 128-K bytes of the calling process' address space. Following this call, your program can directly access the display memory by references to the mapped portion of the address space.

Be careful not to access display memory while a bit BLT is underway. Doing so causes the offending program to abort with a hardware bus error fault. To avoid this problem, do not include IDONE in calls to SMD\_\$BLT\_U or use SMD\_\$EVENT\_WAIT\_U to wait for BLT completion.

To unmap the display memory, use SMD\_\$UNMAP\_DISPLAY\_U. Display memory is automatically unmapped when SMD\_\$RETURN\_DISPLAY\_U is executed, or when you type CTRL/Q to exit from the program.

SMD\_\$MAP\_DISPLAY\_U returns an error status if the display has not been borrowed, or if the display is already mapped into the calling process' address space.

SMD\_\$MOVE\_ABS\_U

SMD\_\$MOVE\_ABS\_U

Sets the current position for vector drawing.

### FORMAT

SMD\_\$MOVE\_ABS\_U (column, line)

### INPUT PARAMETERS

#### column

The number of the column to which the position is set. This is a 2-byte integer in the range 0 through 799.

#### line

The number of the line to which the position is set. This is a 2-byte integer in the range 0 through 1023.

### USAGE

SMD\_\$MOVE\_ABS\_U sets the current position, from which the next vector will be drawn.

Call SMD\_\$VECTOR\_INIT\_U to initialize the vector drawing package before using this procedure.

No error checking is performed on the arguments, for optimal performance. Incorrect program operation occurs if a column or line value is outside the specified range.

**SMD\_\$MOVE\_KBD\_CURSOR\_U**

Moves the keyboard cursor to a specified position.

**FORMAT**

SMD\_\$MOVE\_KBD\_CURSOR\_U (position, status)

**INPUT PARAMETERS****position**

X and y coordinates, in SMD\_\$POS\_T format, for the cursor position. This data type is 4 bytes long.

**OUTPUT PARAMETERS****status**

Completion status, in STATUS\_\$T format. This data type is 4 bytes long.

**USAGE**

SMD\_\$MOVE\_KBD\_CURSOR\_U moves the keyboard cursor to a new position.

If the cursor was previously removed from the display (via SMD\_\$CLEAR\_KBD\_CURSOR\_U), this call re-enables it.

In the position parameter, valid line values are 0-1023 for portrait displays, or 0 to 799 for landscape displays, and valid column values are 0-799 for portrait displays, or 0 1023 for landscape displays. The values represent the position of the lower left point of the cursor.

The keyboard cursor is 8 bits wide and 13 bits high.

SMD\_\$MOVE\_REL\_U

SMD\_\$MOVE\_REL\_U

Sets the current position for vector drawing.

## FORMAT

SMD\_\$MOVE\_REL\_U (column, line)

## INPUT PARAMETERS

### column

The column number, relative to the current position, from which the new column position is computed. This is a 2-byte integer.

### line

The line number, relative to the current position, from which the new line position is computed. This is a 2-byte integer.

## USAGE

SMD\_\$MOVE\_REL\_U computes a new position based upon the current value and the supplied arguments.

Call SMD\_\$VECTOR\_INIT\_U to initialize the vector drawing package before using this procedure.

When the value for column is added to the current column number, the sum must be between 0 and 799 for portrait displays, or between 0 and 1023 for landscape displays. Similarly, when the value for line is added to the current line number, the sum must be between 0 and 1023 for portrait displays and between 0 and 799 for landscape displays.

No error checking is performed on the arguments for optimal performance. Incorrect program operation occurs if the computed column or line value is outside the specified range.

SMD\_\$OP\_WAIT\_U

Waits for the current scroll or BLT operation to complete.

**FORMAT**

SMD\_\$OP\_WAIT\_U

**USAGE**

SMD\_\$OP\_WAIT\_U waits for completion of the current scroll or BLT operation. When this routine returns, the program can safely reference display memory.

If no scroll or BLT operation is underway, this routine returns immediately.

SMD\_\$RETURN\_DISPLAY\_U

SMD\_\$RETURN\_DISPLAY\_U

Returns control of the display to the Display Manager.

#### FORMAT

SMD\_\$RETURN\_DISPLAY\_U (unit, status)

#### INPUT PARAMETERS

##### unit

Unit number of the display to be returned. This is a 2-byte integer. Currently, the only valid unit number is 1.

#### OUTPUT PARAMETERS

##### status

Completion status, in STATUS\_\$T format. This data type is 4 bytes long.

#### USAGE

SMD\_\$RETURN\_DISPLAY\_U returns control of the display to the Display Manager. After executing this procedure, the calling program can no longer use display driver calls.

If the display was mapped into the process' address space, this procedure unmaps it.

After execution of this procedure, the Display Manager updates the display to reflect all input, output, and scrolling operations that occurred for all pads and processes while the screen was under direct program control.

If SMD\_\$BORROW\_DISPLAY\_U has not yet been successfully executed, an error status is returned.

**SMD\_\$SET\_QUIT\_CHAR**

Defines the quit character.

**FORMAT**

SMD\_\$SET\_QUIT\_CHAR (character, status)

**INPUT PARAMETERS**

**character**

A single character that is the new quit character. This is a character variable.

**OUTPUT PARAMETERS**

**status**

Completion status, in STATUS\_\$T format. This data type is 4 bytes long.

**USAGE**

The default quit character is CTRL/Q.

SMD\_\$SET\_TP\_CURSOR

SMD\_\$SET\_TP\_CURSOR

Changes normal cursor to touchpad cursor and moves it to the indicated position.

## FORMAT

SMD\_\$SET\_TP\_CURSOR (unit, position, buttons)

## INPUT PARAMETERS

### unit

A 2-byte integer indicating which display unit to use. Currently, the only valid unit number is 1.

### position

The new position of the cursor, in SMD\_\$POS\_T format. This data type is 4 bytes long.

### buttons

A 2-byte integer that specifies a combination of buttons according to which bits are set. If the bit is set, the corresponding button is down. Following are the bit settings and their meanings:

Bit Position	Button
-----	-----
0	lefthand button
1	middle button
2	righthand button
3	fourth button

Given the following assignments:

```
unit := 1;           { must always be 1 }
position.line := 312;
position.col := 17;
buttons := 3;       { the lefthand button is down }
```

the corresponding call would be:

```
smd_$set_tp_cursor (unit, position, buttons)
```

## USAGE

This call is for use by programs that process data from a locator device other than the touch pad, such as a tablet. SMD\_\$SET\_TP\_CURSOR should not be used with the touch pad.

If the keyboard cursor is currently displayed and the touch pad is enabled (see SMD\_\$TP\_ENABLE and SMD\_\$MOVE\_KBD\_CURSOR\_U) executing SMD\_\$SET\_TP\_CURSOR removes the keyboard cursor, and displays the touchpad cursor at the location denoted by position.



**SMD\_\$\$SOFT\_SCROLL\_U**

Starts horizontal or vertical scrolling on the screen, 2 raster lines at a time.

**FORMAT**

SMD\_\$\$SOFT\_SCROLL\_U (boundaries, direction, displacement, status)

**INPUT PARAMETERS****boundaries**

The x and y coordinates for the edges of the area to be scrolled, in SMD\_\$\$WINDOW\_LIMITS\_T format. This data type is 8 bytes long. In FORTRAN this is a four-element array of 2-byte integers.

**direction**

Direction in which to scroll in SMD\_\$\$DIRECTION\_T format. This is a 2-byte integer. Possible values are: SMD\_\$\$UP, SMD\_\$\$DOWN, SMD\_\$\$LEFT, and SMD\_\$\$RIGHT.

**displacement**

Number of horizontal or vertical raster lines to scroll. This is a 2-byte integer.

**OUTPUT PARAMETERS****status**

Completion status, in STATUS\_\$\$T format. This data type is 4 bytes long.

**USAGE**

SMD\_\$\$SOFT\_SCROLL\_U scrolls two raster lines at a time, until the number of lines scrolled is equal to the displacement.

Scrolling takes place only within the specified area. The rest of the display does not change.

This procedure starts the scrolling operation, then returns control to the calling program.

Scrolling replicates the two raster lines on the boundary opposite to the direction of scrolling (for example, if scrolling up, the two raster lines on the bottom of the scrolled area), and does not clear them. Therefore, if you want the scrolling operation to produce blank lines, you must make certain the lines that are replicated are blank. You can clear a section of the display explicitly using SMD\_\$\$CLEAR\_WINDOW.

Lines scrolled beyond the stated boundaries are lost.

## SMD\_ \$SOFT\_SCROLL\_U

Because scrolling occurs two lines at a time, the display driver can "interlace" other tasks with soft scrolling. Thus, the program can call most other display driver routines while scrolling is underway. The program must not, however, reference display memory, and it must not call any of the following routines:

```
SMD_ $DRAW_ABS_U
SMD_ $DRAW_BOX_U
SMD_ $DRAW_REL_U
SMD_ $MOVE_ABS_U
SMD_ $MOVE_REL_U
SMD_ $RETURN_DISPLAY_U
SMD_ $SOFT_SCROLL_U
```

The display driver waits for scrolling to complete before executing one of these procedures. Calls to SMD\_ \$BLT\_U are executed only if the display mode does not include IDONE. If the display mode value includes IDONE, the driver waits for the current scrolling operation to complete before starting the BLT. Attempting a BLT with any part of its source or destination in the scrolled area is not recommended.

The program must not reference display memory while scrolling is underway. The program can call SMD\_ \$EVENT\_WAIT\_U to find out when the completion of scrolling is imminent, and can then prepare data for another display operation. After preparing the data, the program can call SMD\_ \$OP\_WAIT to wait until references to display memory are safe.

**SMD\_\$STOP\_TP\_CURSOR**

Turns off the touch pad cursor and puts back the blinking cursor, if the blinking cursor would otherwise be displayed.

**FORMAT**

SMD\_\$STOP\_TP\_CURSOR (unit)

**INPUT PARAMETERS****unit**

A 2-byte integer indicating which display unit to use. Currently, the only valid unit number is 1.

**USAGE**

This call is for use only by programs that process data from a locator device other than the touchpad. SMD\_\$SET\_TP\_CURSOR should not be used with the touchpad.

If the touchpad cursor is currently displayed, it is replaced with the blinking keyboard cursor.

SMD\_\$TP\_DIRECT

SMD\_\$TP\_DIRECT

Controls whether locator device data directly controls the touch pad cursor or is sent to the user program.

## FORMAT

SMD\_\$TP\_DIRECT (on-off, status)

## INPUT PARAMETERS

### on-off

A Boolean (logical) variable indicating whether to send the locator device directly to the program (TRUE) or to the internal display driver routine that controls the touchpad cursor (FALSE).

## OUTPUT PARAMETERS

### status

Completion status, in STATUS\_\$T format. This data type is 4 bytes long.

## USAGE

If the value of on-off is FALSE, (the initial state) locator device (for example, touchpad) data causes the touchpad cursor to move, as long as the touchpad cursor and keyboard cursor are enabled (see SMD\_\$TP\_DISABLE). Locator device data are delivered through SMD\_\$EVENT\_WAIT as SMD\_\$TPAD\_AND\_INPUT, but only when keystrokes are also delivered.

If the on-off parameter is TRUE, the touchpad cursor is not displayed in response to locator device data, and data is delivered in a continuous stream through SMD\_\$EVENT\_WAIT as SMD\_\$TPAD\_DATA.

If the on-off parameter is TRUE, locator data is always delivered, regardless of whether or not the touchpad cursor or keyboard cursor is enabled.

**SMD\_\$TP\_DISABLE**

Prevents locator device data from moving the touchpad cursor.

**FORMAT**

SMD\_\$TP\_DISABLE (status)

**OUTPUT PARAMETERS****status**

Completion status, in STATUS\_\$T format. This data type is 4 bytes long.

**USAGE**

This call is for programs that modify display memory directly while the keyboard cursor is displayed. The touchpad cursor can interfere with display memory modifications because its location is unknown to the user program. SMD\_\$TP\_DISABLE prevents display of the touchpad cursor.

This call also prevents display of the keyboard cursor if the touchpad has moved it and the user program has not yet been given the new position through a SMD\_\$EVENT\_WAIT\_U call.

The SMD vector drawing routines modify display memory directly. Therefore, SMD\_\$TP\_DISABLE should be called before calling the vector routines, if the keyboard cursor is displayed at the same time the vectors are drawn.

SMD routines that modify display memory directly, other than the vector drawing routines (SMD\_\$DRAW\_REL\_U and SMD\_\$DRAW\_ABS\_U), automatically disable the touchpad cursor when they begin executing and re-enable the cursor when they finish.

In many cases the keyboard cursor is cleared (removed from the display using SMD\_\$CLEAR\_KBD\_CURSOR\_U) before display modifications are made. This call is not needed in such cases.

The touchpad cursor is initially disabled.

SMD\_\$TP\_ENABLE

SMD\_\$TP\_ENABLE

Allows the touch pad cursor to be displayed and moved around the screen.

#### FORMAT

SMD\_\$TP\_ENABLE (status)

#### OUTPUT PARAMETERS

##### status

Completion status, in STATUS\_\$T format. This data type is 4 bytes long.

#### USAGE

This call enables display of the touchpad cursor. The touchpad cursor can be disabled by calling SMD\_\$TP\_DISABLE.

The touchpad cursor is initially disabled.

In order for the touchpad or other locator device to move the cursor, three conditions must be satisfied:

1. SMD\_\$MOVE\_KBD\_CURSOR must be called to display the cursor in the first place.
2. SMD\_\$TP\_ENABLE must be called to allow the locator device to affect the cursor.
3. On\_off, an input parameter of SMD\_\$TP\_DIRECT, must be FALSE. (See SMD\_\$TP\_DIRECT.) Your program must make explicit calls to satisfy 1 and 2. The third condition is satisfied by default.

SMD\_\$UNLOAD\_FONT\_FILE\_U

Unloads a font file from display memory.

**FORMAT**

SMD\_\$UNLOAD\_FONT\_FILE\_U (font-id, status)

**INPUT PARAMETERS**

**font-id**

The internal identifier assigned to the font to be unloaded. This 2-byte value is returned by SMD\_\$LOAD\_FONT\_FILE\_U.

**OUTPUT PARAMETERS**

**status**

Completion status, in STATUS\_\$T format. This data type is 4 bytes long.

**USAGE**

SMD\_\$UNLOAD\_FONT\_FILE\_U unloads a font that was loaded with SMD\_\$LOAD\_FONT\_FILE\_U. Following this call, the font is no longer usable in calls to SMD\_\$WRITE\_STRING.

An error is returned if the font is not loaded in display memory, or if the associated font file is not mapped into the process' address space.

SMD\_\$UNLOAD\_FONT\_U

SMD\_\$UNLOAD\_FONT\_U

Unloads a font from display memory.

#### FORMAT

SMD\_\$UNLOAD\_FONT\_U (font-id, status)

#### INPUT PARAMETERS

##### font-id

The internal identifier assigned to the font to be unloaded. This 2-byte value is returned by SMD\_\$LOAD\_FONT\_U.

#### OUTPUT PARAMETERS

##### status

Completion status, in STATUS\_\$T format. This data type is 4 bytes long.

#### USAGE

This procedure unloads the specified font, making it unavailable for use. The program must reload the font before using it again.

Use SMD\_\$UNLOAD\_FONT\_U for fonts loaded with SMD\_\$LOAD\_FONT\_U. Use SMD\_\$UNLOAD\_FONT\_FILE\_U for font files loaded with SMD\_\$LOAD\_FONT\_FILE\_U.



SMD\_\$UNMAP\_DISPLAY\_U

Unmaps display memory from the process' address space.

**FORMAT**

SMD\_\$UNMAP\_DISPLAY\_U (status)

**OUTPUT PARAMETERS**

**status**

Completion status, in STATUS\_\$T format. This data type is 4 bytes long.

**USAGE**

SMD\_\$UNMAP\_DISPLAY\_U unmaps the display memory from the calling process' address space. Following this call, the 128K-byte portion of the address space onto which the display memory was mapped is no longer usable.

An error status is returned if the display has not been borrowed or if the display memory is not mapped when this call is made.

SMD\_\$VECTOR\_INIT\_U

SMD\_\$VECTOR\_INIT\_U

Initializes the vector-drawing routines.

## FORMAT

SMD\_\$VECTOR\_INIT\_U (display-address)

## INPUT PARAMETERS

### display-address

Starting address of the display memory in the program's address space. This is a 4-byte integer. This value is returned by SMD\_\$MAP\_DISPLAY\_U.

## USAGE

SMD\_\$VECTOR\_INIT\_U initializes the vector-drawing routines supplied with the display driver. These routines are named SMD\_\$DRAW\_ABS\_U, SMD\_\$DRAW\_REL\_U, SMD\_\$MOVE\_ABS\_U, and SMD\_\$MOVE\_REL\_U. You must use SMD\_\$VECTOR\_INIT\_U once before calling any vector-drawing routines.

The vector-drawing routines operate incorrectly if the value of display-address differs from that returned by SMD\_\$MAP\_DISPLAY\_U.

The current position is set to line 0 and column 0. If you call a vector drawing routine before calling a position moving routine, the display driver draws a vector from the current position. All position-moving and vector-drawing routines update the current position.

**SMD\_\$WRITE\_STRING\_U**

Displays a string of text on the screen.

**FORMAT**

SMD\_\$WRITE\_STRING\_U (position, font\_id, string, length, waitflag, status)

**INPUT PARAMETERS****position**

Row and column positions for the first character of the string, in SMD\_\$POS\_T format. This data type is 4 bytes long.

**font\_id**

Internal font identifier, returned by SMD\_\$LOAD\_FONT\_U.

**string**

String of ASCII text to be displayed. This is an array of up to 120 characters.

**length**

Length, in bytes, of the string to be displayed. This is a 2-byte integer.

**waitflag**

Boolean (logical) value indicating whether to wait for scrolling to complete before displaying the string.

**OUTPUT PARAMETERS****status**

Completion status, in STATUS\_\$T format. This data type is 4 bytes long.

**USAGE**

This procedure displays a character string on the screen.

The ordinal (ASCII) value of each character in the input string is used as an index into the font denoted by font-id. For instance, the character "A" in the input string causes the output of the 60th character in the font. If no 60th character is defined in the font, the character "A" is ignored. No error occurs and no graphic character is displayed. The cursor is moved to the right an amount equal to the size of one space character.

The position parameter defines the base from which the first character is written. For the standard font, valid line values are 0 to 792 and valid column values are 12 to 1019 for a portrait display; for a landscape display, valid line values are 0 to 1016 and valid column values are 12 to 795.

The ordinal (ASCII) values of the characters in the string must be in the range 0 through 127.

Set the value of waitflag to TRUE if the string is to be displayed within an area that is being scrolled (via SMD\_\$SOFT\_SCROLL\_U) and FALSE otherwise.

SMD\_\$WRITE\_STRING\_U

SMD\_\$WRITE\_STRING does not clear characters from the portion of the display where the string is written. Since characters may not fill the entire character box, pieces of previous characters may appear along with the string you wish to write. SMD\_\$CLEAR\_WINDOW can be used to clear a section of the display.

**ERRORS**

- SMD\_ \$ACCESS\_DENIED  
Display borrow request denied by screen manager.
- SMD\_ \$ALREADY\_ACQUIRED  
Display already acquired.
- SMD\_ \$ALREADY\_BORROWED  
Display already borrowed by this process.
- SMD\_ \$ALREADY\_MAPPED  
Display memory is already mapped.
- SMD\_ \$BORROW\_ERROR  
Error borrowing display from screen manager.
- SMD\_ \$CANT\_BORROW\_BOTH  
Cannot borrow both displays simultaneously.
- SMD\_ \$DISP\_ACQD  
Pad/stream operations not allowed while display acquired.
- SMD\_ \$DISPLAY\_IN\_USE  
Unable to borrow: display in use.
- SMD\_ \$DISPLAY\_MAP\_ERROR  
Error-mapping display memory.
- SMD\_ \$FONT\_NOT\_LOADED  
Specified font is not loaded.
- SMD\_ \$FONT\_NOT\_MAPPED  
Font associated with specified ID is not mapped.
- SMD\_ \$FONT\_TABLE\_FULL  
Internal font table is full.
- SMD\_ \$FONT\_TOO\_LARGE  
Font too large.
- SMD\_ \$HDM\_FULL  
Hidden display memory is full.
- SMD\_ \$HDMT\_UNLOAD\_ERR  
Error unloading internal (HDMT) table.
- SMD\_ \$ILLEGAL\_CALLER  
Invalid use of display driver procedure.
- SMD\_ \$ILLEGAL\_DIRECTION  
Invalid direction from SM.
- SMD\_ \$ILLEGAL\_UNIT  
Invalid display unit number.
- SMD\_ \$INVALID\_BLT\_COORD  
Invalid screen coordinates in BLT request.

SMD ERRORS

- SMD\_\$INVALID\_BLT\_CTL  
Invalid BLT control register.
- SMD\_\$INVALID\_BLT\_MODE  
Invalid BLT mode register.
- SMD\_\$INVALID\_BLTD\_INT  
Invalid BLT-done interrupt.
- SMD\_\$INVALID\_BUFFER\_SIZE  
Invalid buffer size.
- SMD\_\$INVALID\_CRSR\_NUMBER  
Invalid cursor number.
- SMD\_\$INVALID\_DIRECTION  
Invalid direction argument.
- SMD\_\$INVALID\_DISPLACEMENT  
Invalid scroll displacement argument.
- SMD\_\$INVALID\_IR\_STATE  
Invalid interrupt routine state.
- SMD\_\$INVALID\_KEY  
Invalid eventcount key.
- SMD\_\$INVALID\_LENGTH  
Invalid length argument.
- SMD\_\$INVALID\_POS  
Invalid position argument.
- SMD\_\$INVALID\_WID  
Invalid DM window ID.
- SMD\_\$INVALID\_WINDOW  
Invalid window limits argument.
- SMD\_\$NO\_MORE\_WIDS  
No more direct mode window IDs are available.
- SMD\_\$NOT\_BORROWED  
Cannot return: display not borrowed.
- SMD\_\$NOT\_IMPLEMENTED  
Nonconforming and main memory BLTs are not implemented.
- SMD\_\$NOT\_MAPPED  
Display memory is not mapped.
- SMD\_\$NOT\_ON\_COLOR  
Operation not implemented on color display.
- SMD\_\$PROCESS\_NOT\_FOUND  
Process not found.
- SMD\_\$PROTOCOL\_VIOL  
Internal protocol violation.

SMD\_\$QUIT\_WHILE\_WAITING

Quit while waiting.

SMD\_\$RETURN\_ERROR

Error returning display to screen manager.

SMD\_\$TOO\_MANY\_PAGES

Too many pages to be wired.

SMD\_\$UNEXP\_BLT\_INUSE

Unexpected BLT in use.

SMD\_\$UNSUPPORTED\_FONT\_VE

Unsupported font version number.

SMD\_\$WAIT\_QUIT

Quit while waiting.

SMD\_\$WINDOW\_OBSCURED

Acquire denied because window is obscured.

STATUS\_\$OK

Successful completion.





# STREAM

## Table of Contents

Data Types	STREAM-2
STREAM_\$CLOSE	STREAM-13
STREAM_\$CREATE	STREAM-14
STREAM_\$CREATE_BIN	STREAM-17
STREAM_\$CREATE_HERE	STREAM-20
STREAM_\$DELETE	STREAM-23
STREAM_\$FORCE_WRITE_FILE	STREAM-24
STREAM_\$GET_BUF	STREAM-25
STREAM_\$GET_CONDITIONAL	STREAM-27
STREAM_\$GET_EC	STREAM-29
STREAM_\$GET_PRIOR_REC	STREAM-30
STREAM_\$GET_REC	STREAM-32
STREAM_\$INQUIRE	STREAM-34
STREAM_\$OPEN	STREAM-41
STREAM_\$PUT_CHR	STREAM-43
STREAM_\$PUT_REC	STREAM-45
STREAM_\$REDEFINE	STREAM-47
STREAM_\$REPLACE	STREAM-52
STREAM_\$SEEK	STREAM-54
STREAM_\$TRUNCATE	STREAM-56
Errors	STREAM-57



# STREAM

The STREAM (Streams) programming calls perform device-independent I/O. This section describes their data types, call syntax, and error codes. Refer to the Introduction at the beginning of this manual for a description of data-type diagrams and call syntax format.

## STREAM DATA TYPES

### CONSTANTS

STREAM_\$MAX	127	Maximum number of possible stream IDs.
STREAM_\$NO_STREAM		Place-holder for stream ID when passing streams.
STREAM_\$SUBS		Subsys component of status return denoting streams.
STREAM_\$DIR_ENTTYPE_FILE		The file-type value for the enttype field of the DIR_ENTRY_T record.
STREAM_\$DIR_ENTTYPE_LINK		The link-type value for the enttype field of the DIR_ENTRY_T record.
STREAM_\$DIR_ENTRY_SIZE		Size of a directory entry record.

The following are mnemonic definitions used to specify attributes in the inquire and redefine input mask. Attributes followed by an asterisk (\*) are attributes to which a stream must be open for information to be returned on an inquire. Attributes followed by a plus sign (+) cannot be redefined.

MNEMONIC	Bit	Explanation
STREAM_\$STRID	0	Stream ID. +
STREAM_\$OBJ_NAME	1	Object name.
STREAM_\$OBJ_NAMLEN	1	Object name length.
STREAM_\$REC_LGTH	2	Record length.
STREAM_\$TEMPORARY	3	Temporary or permanent. *
STREAM_\$EXPLICIT_TYPE	4	Explicit record type.
STREAM_\$AB_FLAG	5	ASCII or binary file.
STREAM_\$EXPLICIT_ML	6	Explicit move mode. *
STREAM_\$CC	7	Carriage control.
STREAM_\$REC_TYPE	8	Record type.
STREAM_\$CONC	9	Object concurrency.
STREAM_\$OCONC	10	Concurrency at open. *
STREAM_\$OPOS	11	Access type. *
STREAM_\$PRE_EXIST	12	Pre-existing object. +
STREAM_\$HDR_LGTH	13	Header length. +
STREAM_\$FILE_LGTH	14	File length. +

STREAM_\$SEEK_KEY	15	Seek key. * +
STREAM_\$CUR_REC_LEN	16	Current record length. * +
STREAM_\$CUR_REL_REC_NO	17	Current relative record number. * +
STREAM_\$BLKS_USED	18	Number of blocks used. +
STREAM_\$DTU	19	Date and time last used. +
STREAM_\$DTM	20	Date and time last modified. +
STREAM_\$SPARSE	21	Sparsely written file. * +
STREAM_\$OTYPE	22	Object type.
STREAM_\$CLOSE_ON_EXEC	23	Close stream on DOMAIN/IX Exec call.
STREAM_\$NDELAY	24	Forced STREAM_\$GET_CONDITIONAL.
STREAM_\$APPEND_MODE	25	File in append mode.
STREAM_\$FORCED_LOCATE	26	Force locate mode.

**DATA TYPES**

STREAM\_\$PARM1\_T                   A 2-byte integer. Specifies the type of data on which the seek is being performed. One of the following predefined values:

STREAM\_\$KEY  
Seek with key returned earlier by stream manager.

STREAM\_\$REC  
Record-oriented seek.

STREAM\_\$CHR  
Character-oriented seek.

STREAM\_\$EOF  
Seek to the end-of-file.

STREAM\_\$PARM2\_T                   A 2-byte integer. Specifies the type of seek being performed. One of the following predefined values:

STREAM\_\$RELATIVE  
Seek relative to current position.

STREAM\_\$ABSOLUTE  
Seek relative to BOF or EOF.

STREAM DATA TYPES

STREAM\_\$OPOS\_T

A 2-byte integer. Specifies the access type of an object on open/create. One of the following predefined values:

STREAM\_\$READ  
Open/create for read-only access.

STREAM\_\$WRITE  
Create (new) for write access.

STREAM\_\$OVERWRITE  
Write access; truncate file to BOF if it already exists.

STREAM\_\$UPDATE  
Write access; file may already exist; position to start of file on open.

STREAM\_\$APPEND  
Write access; if file already exists, position to EOF on open.

STREAM\_\$MAKE\_BACKUP  
Create new file: rename existing file to .BAK on close.

STREAM\_\$OMODE\_T

A 2-byte integer. Specifies the concurrency at open of an object. One of the following predefined values:

STREAM\_\$NO\_CONC\_WRITE  
Allows no concurrent writers to open file while this stream is open.

STREAM\_\$CONTROLLED\_SHARING  
Currently the same as NO\_CONC\_WRITE.

STREAM\_\$REGULATED  
Allows unrestricted reading and writing of the file.

STREAM\_\$FCONC\_T

A 2-byte integer. Specifies the object concurrency. One of the following predefined values:

STREAM\_\$N\_OR\_1  
Allows N readers or 1 writer in file concurrently.

STREAM\_\$N\_AND\_1  
Allows N readers AND up to 1 writer concurrently.

STREAM\_\$N\_AND\_N

Allows any number of writers or readers concurrently.

STREAM\_\$STRICT\_N\_OR\_1

Disallows multiple writers even when they are in a process family.

STREAM\_\$CC\_T

A 2-byte integer. Specifies the type of carriage control employed in the object. One of the following predefined values:

STREAM\_\$CC\_T

ASCII ("Apollo standard") carriage control.

STREAM\_\$F&&\_CC

Fortran-77 standard (column 1) carriage control.

STREAM\_\$RTYPE\_T

A 2-byte integer. Specifies the record structure of the object. One of the following predefined values:

STREAM\_\$V1

Variable length records with count fields.

STREAM\_\$F2

Fixed-length records.

STREAM\_\$UNDEF

No record structure in data.

STREAM\_\$IR\_OPT

A 2-byte integer. Specifies method for accessing attribute record. One of the following predefined values:

STREAM\_\$USE\_STRID

Use the stream-id to access the attribute record.

STREAM\_\$NAME\_CONDITIONAL

Inquire is about the filename; only return information if file is open.

STREAM\_\$NAME\_UNCONDITIONAL

Use the filename to access the attribute block.

STREAM\_\$INQUIRE\_MASK\_T

A 2-byte integer. Attributes to inquire. Specify any combination of the mnemonic constants for object attributes.

STREAM\_\$REDEF\_MASK\_T

A 2-byte integer. Attributes to redefine. Specify any combination of the valid mnemonic constants for object attributes.

STREAM DATA TYPES

STREAM\_\$EC\_KEY\_T

A 2-byte integer. Specifies an eventcount to get.  
One of the following predefined values:

STREAM\_\$GETREC\_EC\_KEY  
Stream eventcount key.

STREAM\_\$EDIT\_WAIT\_EC\_KEY  
Edit pad eventcount key.

STREAM\_\$OPEN\_OPTIONS\_T

A 2-byte integer. Options available at open time.  
One of the following predefined values:

STREAM\_\$NO\_DELAY  
Do not wait for I/O (currently applies only to  
opening pipes).

STREAM\_\$OPEN\_OPTIONS\_SET\_T

A 2-byte integer. Options available at open time  
with the STREAM\_\$OPEN\_OPT call. Currently  
the only option available is:

STREAM\_\$NO\_DELAY  
Do not wait for I/O (currently applies only to  
opening pipes).

STREAM\_\$IR\_REC\_T

Attribute record for INQUIRE and REDEFINE  
calls. The streams chapter of the *Programming  
With General System Calls* manual describes how  
to use the attribute record. The diagram below  
illustrates the STREAM\_\$IR\_REC\_T data type:

STREAM\_\$ID\_T

A 2-byte integer. Open stream identifier.



STREAM\_ \$SK\_ T

Seek key returned on most stream calls. The diagram below illustrates the STREAM\_ \$SK\_ T data type:

predefined type	byte: offset	field name
	0:	rec_ adr
	4:	byte_ adr
	8:	flags

Field Description:

rec\_ adr  
The address of the record sought.

byte\_ adr  
The address of the byte sought.

flags  
Flags containing seek information.

predefined type	byte: offset	field name
	0:	offset

Field Description:

offset  
The offset of the record or character sought.

STREAM DATA TYPES

predefined type	byte: offset		field name
stream_\$id_t	0:	integer	strid
	2:	integer	obj_namlen
	4:	integer	rec_lgth
	8:	integer	flags1*
	10:	integer	flags2*
	12:	integer	unused
	14:	integer	hdr_lgth
	16:	integer	file_lgth
stream_\$sk_t	20:	integer	seek_key.rec_adr
	24:	integer	seek_key.byte_adr
	28:	integer	seek_key.flags
	32:	integer	cur_rec_len
	36:	integer	cur_rel_rec_no
	40:	integer	blks_used
time_\$clockh_t	44:	integer	dtu
time_\$clockh_t	48:	integer	dtm
uid_\$t	52:	integer	otype.high
	56:	integer	otype.low
name_\$name_t	60:	integer	flags3*
	62:	integer	flags4*
	64:	integer	unused
	68:	char	obj_name
	~	~	
	n:	char	

\*see below for field names

Field Description:

strid  
The stream ID of the object.

obj\_namlen  
The length of the object's name.

rec\_lgth

The length of the longest record in the object.

flags1

A bit mask containing predefined or Boolean values indicating object attributes. The following table lists the bit numbers within the mask, the record field names, and a short description of each attribute:

*Bit # Field Name Description*

Bit 0	temporary	Temporary or permanent object
Bit 1	explicit_type	Explicit fixed-length records
Bit 2	ab_flag	ASCII or binary file
Bit 3	explicit_ml	Explicit move mode
Bit 4-5	unused1	
Bit 6	cc	Type of carriage control
Bit 7-9	unused2	
Bit 10-11	rec_type	Record type
Bit 12-13	unused3	
Bit 14-15	conc	Object concurrency

flags2

A bit mask containing predefined or Boolean values indicating object attributes. The following table lists the bit numbers within the mask, the record field names, and a short description of each attribute:

*Bit # Field Name Description*

Bit 0-1	unused4	
Bit 2-3	oconc	Concurrency at open
Bit 4-5	unused5	
Bit 6-8	opos	Access type
Bit 9	pre_exist	Pre-existing object
Bit 10-15	unused6	

hdr\_lgth

The length of the object header.

file\_lgth  
The length of the file.

seek\_key  
The current seek-key.

cur\_rec\_len  
The length of the current record.

cur\_rel\_rec\_no  
The current record number relative to BOF.

blks\_used  
The number of blocks occupied by the file.

dtu  
The date and time of the last use of the object.

dtm  
The date and time of the modification of the object.

otype  
Specifies the type of the object.

flags3  
A bit mask containing predefined or Boolean values indicating object attributes. The following table lists the bit numbers within the mask, the record field names, and a short description of each attribute:

*Bit # Field Name Description*

Bit 0	sparse	File may contain "holes."
Bit 1-15	unused7	

flags4  
A bit mask containing predefined or Boolean values indicating object attributes. The following table lists the bit numbers within the mask, the record field names, and a short description of each attribute:

*Bit # Field Name Description*

Bit 0-12	unused8	
Bit 13	close_on_exec	Close stream on UNIX Exec call


STREAM DATA TYPES

Bit 14 ndelay Forced  
 STREAM\_\$GET\_CONDITIONAL  
 Bit 25 append\_mode File in  
 append mode  
 Bit 26 forced\_locate Force  
 locate mode

obj\_name  
 The name of the object.

STREAM\_\$DIR\_ENTRY\_T

The directory entry returned by  
 STREAM\_\$GET\_REC. The diagram below  
 illustrates the STREAM\_\$DIR\_ENTRY\_T data  
 type:



Field Description:

enttype  
 Type of the directory entry. Either  
 NAME\_\$FILE or NAME\_\$LINK.

entlen  
 Length of the directory entry name.

entname  
 Name of the directory entry.

unusedn  
 Reserved for future use by Apollo.

STREAM\_\$FORCE\_WRITE\_OPTIONS\_T

A 2-byte integer. Options available for force  
 writing to disk. Any combination of the following  
 predefined values:

STREAM\_\$FW\_FILE  
 Specifies that a file should be force-written.


STREAM DATA TYPES

STREAM\_\$FW\_DIR

Specifies that the directory of the file should be force-written.

STATUS\_\$T

A status code. The diagram below illustrates the STATUS\_\$T data type:



Field Description:

all

All 32 bits in the status code.

fail

The fail bit. If this bit is set, the error was not within the scope of the module invoked, but occurred within a lower-level module (bit 31).

subsys

The subsystem that encountered the error (bits 24 - 30).

modc

The module that encountered the error (bits 16 - 23).

code

A signed number that identifies the type of error that occurred (bits 0 - 15).

**STREAM\_\$CLOSE**

Closes a stream.

**FORMAT**

STREAM\_\$CLOSE (stream-id, status)

**INPUT PARAMETERS****stream-id**

Number of the stream to be closed, in STREAM\_\$ID\_T format. This is a 2-byte integer.

The number used for stream identification becomes available for reuse. If the object is open on more than one stream, STREAM\_\$CLOSE closes only the specified stream.

**OUTPUT PARAMETERS****status**

Completion status, in STATUS\_\$T format. This is a 4-byte integer.

**USAGE**

STREAM\_\$CLOSE closes the stream, so that you can no longer use the stream-id to operate on the object. Closing a stream to an object releases any locks maintained for the current user and may thus make the object available to other users.

If the stream is a disk file opened for any type of write access (STREAM\_\$WRITE, STREAM\_\$OVERWRITE, STREAM\_\$APPEND, or STREAM\_\$UPDATE), STREAM\_\$CLOSE updates its header, reflecting any changes made to the file while it was open, and indicating the date and time of last use and modification.

A program can close only the streams it has opened, and those opened by programs it has invoked (that is, opened at lower levels). Trying to close a stream opened at a higher level produces an error status code.

Closing a temporary object deletes it if no other process is using it.

STREAM\_\$CREATE

STREAM\_\$CREATE

Creates an object (if the object does not already exist) and opens a stream to it.

## FORMAT

STREAM\_\$CREATE (pathname, name-length, access, concurrency, stream-id, status)

## INPUT PARAMETERS

### pathname

Name of the object to be created, in NAME\_\$PNAME\_T format. This is a character array of up to 256 elements.

### name-length

Length of the pathname, in bytes. This is a 2-byte integer. To create a temporary object, specify a length of 0.

### access

Type of access requested, in STREAM\_\$OPOS\_T format. Possible values are:

STREAM\_\$APPEND

Permits adding data to the end of an object. The stream pointer points to the end of the object (EOF).

STREAM\_\$MAKE\_BACKUP

Creates a temporary file, with the same type and attributes as the file specified in the pathname. This access is used to create a backup file. (See below for a detailed description.)

STREAM\_\$OVERWRITE

Permits replacing the entire contents of an object. The stream pointer is positioned at the start of the object data and data is truncated.

STREAM\_\$READ

Permits reading data from an existing object.

STREAM\_\$UPDATE

Permits replacing selected portions of the contents of an object. The stream pointer is positioned at the start of the object data, just past the header if it has one.

STREAM\_\$WRITE

Permits writing data to a new object. If writing is attempted on an existing object, an error status is returned.

If you specify the access option STREAM\_\$WRITE, the pathname must refer to a new object; otherwise, an error status is returned.

If you specify the access option STREAM\_\$MAKE\_BACKUP, a new, unnamed temporary file is created by this call, which has the same type and other attributes as the file given by the pathname (if it exists). The new file is created on the same volume (i.e. the same node) as the file given by the pathname. The file given by the pathname is not



opened or modified in this case, but is examined to extract its attributes. Even though the existing file is not modified, it is conceptually being replaced, so this operation requires write access to the file.

The application then writes the new file, and when it is closed (by the STREAM\_\$CLOSE call) the name of the file given at create time is changed to pathname.BAK. The new (formerly unnamed temporary) file gets the old name, and becomes permanent.

If the ".bak" file already exists, it is deleted. (The caller must have either D or P rights to delete the file.) If the ".bak" file is locked at the time STREAM\_\$CLOSE is called, it is deleted when it is unlocked.

If the pathname mentioned in the create call does not exist, then an ordinary STREAM\_\$CREATE is done, as though the access option had been STREAM\_\$WRITE instead of STREAM\_\$MAKE\_BACKUP.

#### **concurrency**

Requested concurrency at open, in STREAM\_\$OMODE\_T format. Possible values are:

STREAM\_\$CONTROLLED\_SHARING  
No concurrent writing.

STREAM\_\$NO\_CONC\_WRITE  
No concurrent writing.

STREAM\_\$UNREGULATED  
Unregulated read and write access.

#### **OUTPUT PARAMETERS**

##### **stream-id**

Number of the stream on which the object is open, in STREAM\_\$ID\_T format. This is a 2-byte integer.

##### **status**

Completion status, in STATUS\_\$T format. This is a 4-byte integer.

#### **USAGE**

If the pathname specifies an object that does not exist, the stream manager creates a new UASC disk file with that pathname and opens a stream to it.

If the pathname specifies an existing object, the stream manager opens a stream to it for overwrite, update, or append access. If write access is specified for an existing object, an error status is returned.

STREAM\_\$CREATE can open existing objects of any type, but can create only disk files.

## STREAM\_\$CREATE

Default attributes for a new disk file are listed below:

<i>Attribute</i>	<i>Default Value</i>
Data type	ASCII
Record type	UASC file format
Location	Lowest level in the directory pathname. If no pathname is specified, assume the current working directory.
Concurrency	No default.
Object concurrency	One writer or any number of readers.
Carriage control	DOMAIN standard.

If the object already exists, its attributes remain the same when it is opened. For streams to serial lines, however, "cooked" input mode and NO\_WAIT are always in effect when the stream is opened. To change the object's attributes, call STREAM\_\$REDEFINE (or SIO\_\$CONTROL) before writing the object.

Both STREAM\_\$CREATE and STREAM\_\$OPEN open a stream to an object. However, STREAM\_\$CREATE creates the object if it does not exist, whereas STREAM\_\$OPEN returns an error if the object does not exist.

**STREAM\_ \$CREATE\_ BIN**

Creates a binary record-structured file (if the file does not already exist) and opens a stream to it.

**FORMAT**

STREAM\_ \$CREATE\_ BIN (pathname, name-length, access, concurrency, stream-id, status)

**INPUT PARAMETERS****pathname**

Name of the object to be created, in NAME\_ \$PNAME\_ T format. This is a character array of up to 256 elements.

**name-length**

Length of the pathname, in bytes. This is a 2-byte integer. To create a temporary object, specify a length of 0.

**access**

Type of access requested, in STREAM\_ \$OPOS\_ T format. Possible values are:

**STREAM\_ \$APPEND**

Permits adding data to the end of an object. The stream pointer points to the end of the object (EOF).

**STREAM\_ \$MAKE\_ BACKUP**

Creates a temporary file, with the same type and attributes as the file specified in the pathname. This access is used to create a backup file. (See below for a detailed description.)

**STREAM\_ \$OVERWRITE**

Permits replacing the entire contents of an object. The stream pointer is positioned at the start of the object data and data is truncated.

**STREAM\_ \$READ**

Permits reading data from an existing object.

**STREAM\_ \$UPDATE**

Permits replacing selected portions of the contents of an object. The stream pointer is positioned at the start of the object data, just past the header if it has one.

**STREAM\_ \$WRITE**

Permits writing data to a new object. If writing is attempted on an existing object, an error status is returned.

If you specify the access option STREAM\_ \$WRITE, the pathname must refer to a new object; otherwise, an error status is returned.

If you specify the access option STREAM\_ \$MAKE\_ BACKUP, a new, unnamed temporary file is created by this call, which has the same type and other attributes as the

## STREAM\_\$CREATE\_BIN

file given by the pathname (if it exists). The new file is created on the same volume (i.e. the same node) as the file given by the pathname. The file given by the pathname is not opened or modified in this case, but is examined to extract its attributes. Even though the existing file is not modified, it is conceptually being replaced, so this operation requires write access to the file.

The application then writes the new file, and when it is closed (by the STREAM\_\$CLOSE call) the name of the file given at create time is changed to pathname.BAK. The new (formerly unnamed temporary) file gets the old name, and becomes permanent.

If the ".bak" file already exists, it is deleted. (The caller must have either D or P rights to delete the file.) If the ".bak" file is locked at the time STREAM\_\$CLOSE is called, it is deleted when it is unlocked.

If the pathname mentioned in the create call does not exist, then an ordinary STREAM\_\$CREATE\_BIN is done, as though the access option had been STREAM\_\$WRITE instead of STREAM\_\$MAKE\_BACKUP.

### concurrency

Requested concurrency at open, in STREAM\_\$OMODE\_T format. Possible values are:

STREAM\_\$CONTROLLED\_SHARING  
No concurrent writing.

STREAM\_\$NO\_CONC\_WRITE  
No concurrent writing.

STREAM\_\$UNREGULATED  
Unregulated read and write access.

## OUTPUT PARAMETERS

### stream-id

Number of the stream on which the object is open, in STREAM\_\$ID\_T format. This is a 2-byte integer.

### status

Completion status, in STATUS\_\$T format. This is a 4-byte integer.

## USAGE

If the named object does not exist, the stream manager creates a binary file of fixed length records and opens a stream to it.

If the named object already exists, the file attributes remain the same and the stream manager opens a stream to it for overwrite, update, or append access. If write access is specified for an existing object, an error status is returned.

To change the file's attributes, call STREAM\_\$REDEFINE (or SIO\_\$CONTROL) before writing the object.

Both STREAM\_\$CREATE\_BIN and STREAM\_\$OPEN open a stream to a file. However, STREAM\_\$CREATE\_BIN creates the file if it does not exist, whereas STREAM\_\$OPEN returns an error if the file does not exist.

STREAM\_\$CREATE\_BIN differs from STREAM\_\$CREATE in that  
STREAM\_\$CREATE creates a UASC file by default.

STREAM\_\$CREATE\_HERE

STREAM\_\$CREATE\_HERE

Creates an object at the specified location and opens a stream to it.

## FORMAT

STREAM\_\$CREATE\_HERE (pathname, name-length, access, concurrency,  
loclen, locname, stream-id, status)

## INPUT PARAMETERS

### pathname

Name of the object to be created, in NAME\_\$PNAME\_T format. This is a character array of up to 256 elements.

### name-length

Length of the pathname, in bytes. This is a 2-byte integer. To create a temporary object, specify a length of 0.

### access

Type of access requested, in STREAM\_\$OPOS\_T format. Possible values are:

STREAM\_\$APPEND

Permits adding data to the end of an object. The stream pointer points to the end of the object (EOF).

STREAM\_\$MAKE\_BACKUP

Creates a temporary file, with the same type and attributes as the file specified in the pathname. This access is used to create a backup file. (See below for a detailed description.)

STREAM\_\$OVERWRITE

Permits replacing the entire contents of an object. The stream pointer is positioned at the start of the object data and data is truncated.

STREAM\_\$READ

Permits reading data from an existing object.

STREAM\_\$UPDATE

Permits replacing selected portions of the contents of an object. The stream pointer is positioned at the start of the object data, just past the header if it has one.

STREAM\_\$WRITE

Permits writing data to a new object. If writing is attempted on an existing object, an error status is returned.

If you specify the access option STREAM\_\$WRITE, the pathname must refer to a new object; otherwise, an error status is returned.

If you specify the access option STREAM\_\$MAKE\_BACKUP, a new, unnamed temporary file is created by this call, which has the same type and other attributes as the file given by the pathname (if it exists). The new file is created on the same volume (i.e.

the same node) as the file given by the pathname. The file given by the pathname is not opened or modified in this case, but is examined to extract its attributes. Even though the existing file is not modified, it is conceptually being replaced, so this operation requires write access to the file.

The application then writes the new file, and when it is closed (by the STREAM\_ \$CLOSE call) the name of the file given at create time is changed to pathname.BAK. The new (formerly unnamed temporary) file gets the old name, and becomes permanent.

If the ".bak" file already exists, it is deleted. (The caller must have either D or P rights to delete the file.) If the ".bak" file is locked at the time STREAM\_ \$CLOSE is called, it is deleted when it is unlocked.

If the pathname mentioned in the create call does not exist, then an ordinary STREAM\_ \$CREATE\_ HERE is done, as though the access option had been STREAM\_ \$WRITE instead of STREAM\_ \$MAKE\_ BACKUP.

#### concurrency

Requested concurrency at open, in STREAM\_ \$OMODE\_ T format. Possible values are:

STREAM\_ \$CONTROLLED\_ SHARING  
No concurrent writing.

STREAM\_ \$NO\_ CONC\_ WRITE  
No concurrent writing.

STREAM\_ \$UNREGULATED  
Unregulated read and write access.

#### loclen

Length of locname, in bytes. This is a 2-byte integer.

#### locname

Location at which to create the object, in NAME\_ \$PNAME\_ T format. This is a character array of up to 256 elements.

The location can be a tree name or a leaf name.

### OUTPUT PARAMETERS

#### stream-id

Number of the stream on which the object is open, in STREAM\_ \$ID\_ T format. This is a 2-byte integer.

#### status

Completion status, in STATUS\_ \$T format. This is a 4-byte integer.

### USAGE

This call creates an object at a specified location. It is especially useful for creating a temporary file on the same logical volume as an existing object.

If the pathname specifies an object that does not exist, the locname must specify the parent

## STREAM\_\$CREATE\_HERE

directory for the new object. If the pathname specifies an existing object, the locname and loclen are ignored.

If both the object and the location pathnames are valid, the stream manager opens a stream to the object for overwrite, update, or append access. If write access is specified for an existing object, an error status is returned.

STREAM\_\$CREATE\_HERE can open existing objects of any type, but can create only disk files.

Both STREAM\_\$CREATE\_HERE and STREAM\_\$OPEN open a stream to an object. STREAM\_\$CREATE\_HERE, like STREAM\_\$CREATE, creates the object if it does not exist, whereas STREAM\_\$OPEN returns an error if the object does not exist.



**STREAM\_ \$DELETE**

Deletes an object and closes the associated stream.

**FORMAT**

STREAM\_ \$DELETE (stream-id, status)

**INPUT PARAMETERS****stream-id**

Number of a stream on which the object is open, in STREAM\_ \$ID\_ T format. This is a 2-byte integer.

**OUTPUT PARAMETERS****status**

Completion status, in STATUS\_ \$T format. This is a 4-byte integer.

An error occurs if the stream is open for read access only.

An error status is returned if the stream is not open.

**USAGE**

STREAM\_ \$DELETE deletes the object, then closes the specified stream.

If the object cannot be deleted, an error occurs and all streams associated with the object remain open. Input pads cannot be deleted.

If the object is open on more than one stream, STREAM\_ \$DELETE deletes the object causing "object deleted" errors when other streams try to read or write the object.

Files or pads are deleted immediately, even if several processes have opened the object.

For serial lines and magnetic tape files, this call operates exactly like STREAM\_ \$CLOSE.

STREAM\_\$FORCE\_WRITE\_FILE

STREAM\_\$FORCE\_WRITE\_FILE

Forcibly writes a disk file open on the given stream.

## FORMAT

STREAM\_\$FORCE\_WRITE\_FILE (stream-id, options, status)

## INPUT PARAMETERS

### stream-id

The number of the stream on which the disk file is open, in STREAM\_\$ID\_T format. This is a 2-byte integer.

### options

The object types to be force-written, in STREAM\_\$FORCE\_WRITE\_OPTIONS\_T format. Possible values are:

STREAM\_\$FW\_FILE

Forces the file to disk.

STREAM\_\$FW\_DIR

Forces the file's directory to disk.

## OUTPUT PARAMETERS

### status

Completion status, in STATUS\_\$T format. This is a 4-byte integer.

## USAGE

Programs can call STREAM\_\$FORCE\_WRITE\_FILE to ensure that the file is stored on disk before it is closed.

If a program is handling a critical file, it can call this routine on the open stream to force-write the file's directory, thereby ensuring that the file pointer is saved in the directory.

**STREAM\_\$\_GET\_BUF**

Reads data from an object into a specified buffer.

**FORMAT**

STREAM\_\$\_GET\_BUF (stream-id, bufptr, buflen, retptr, retlen, seek-key, status)

**INPUT PARAMETERS****stream-id**

Number of a stream on which the object is open, in STREAM\_\$\_ID\_T format. This is a 2-byte integer.

**bufptr**

Pointer to the buffer into which the data may be read, in UNIV\_PTR format. This is a 4-byte integer.

To obtain a value for bufptr, FORTRAN programs can use the IADDR function. Pascal programs can use the ADDR function. The buffer can be aligned on a byte boundary; therefore, the value of bufptr can be odd.

**buflen**

Number of bytes of data to be read. This is a 4-byte integer.

**OUTPUT PARAMETERS****retptr**

Pointer to the data returned, in UNIV\_PTR format. This is a 4-byte integer.

Address the returned data only by using retptr. The stream manager may use "locate mode," in which it doesn't copy the desired data to the location indicated by bufptr. FORTRAN programs that call the stream manager in locate mode should use the "pointer variable" FORTRAN extension.

The value of retptr is meaningful only until execution of the next stream call on this stream.

**retlen**

Number of bytes of data returned. This is a 4-byte integer.

**seek-key**

Unique key identifying the location of the data read, in STREAM\_\$\_SK\_T format. This is a three-element array of 4-byte integers.

To obtain a seek-key value for the current stream position, call STREAM\_\$\_GET\_BUF with a buflen of 0.

If the returned status is nonzero, the seek-key may not be useful.

**status**

Completion status, in STATUS\_\$\_T format. This is a 4-byte integer.

STREAM\_\$GET\_BUF

## USAGE

For UASC files, STREAM\_\$GET\_BUF returns the requested number of bytes, including newline characters. UASC file records are delimited by the line-feed character (16#0A).

For non-UASC files, STREAM\_\$GET\_BUF functions the same as STREAM\_\$GET\_REC. That is, for fixed- or variable-length records STREAM\_\$GET\_BUF returns one record, and for nonrecord-structured files STREAM\_\$GET\_BUF returns the requested number of characters.

FORTTRAN programs using this procedure in locate mode should use the pointer variable FORTTRAN extension. Otherwise, call STREAM\_\$REDEFINE to set move mode before using this procedure.

**STREAM\_\$GET\_CONDITIONAL**

Reads a record if the record is available; otherwise, it sets the returned record length to zero.

**FORMAT**

STREAM\_\$GET\_CONDITIONAL (stream-id, bufptr, buflen, retptr, retlen,  
seek-key, status)

**INPUT PARAMETERS****stream-id**

Number of a stream on which the object is open, in STREAM\_\$ID\_T format. This is a 2-byte integer.

**bufptr**

Pointer to the buffer into which the data may be read, in UNIV\_PTR format. This is a 4-byte integer.

To obtain a value for bufptr, FORTRAN programs can use the IADDR function. Pascal programs can use the ADDR function. The buffer can be aligned on a byte boundary; therefore, the value of bufptr can be odd.

**buflen**

Number of bytes of data requested. This is a 4-byte integer.

If the number of bytes remaining in the record is less than buflen, STREAM\_\$GET\_CONDITIONAL returns the remainder of the record. If the number of bytes remaining in the record is greater than buflen, the stream manager reads enough data to fill the buffer and returns a negative value in retlen. The absolute value of retlen is the number of bytes remaining in the record.

**OUTPUT PARAMETERS****retptr**

Pointer to the data returned, in UNIV\_PTR format. This is a 4-byte integer.

Address the returned data only by using retptr. The stream manager may use "locate mode," in which it doesn't copy the desired data to the location indicated by bufptr. FORTRAN programs that call the stream manager in locate mode should use the "pointer variable" FORTRAN extension.

Records are aligned on word boundaries. Therefore, if the procedure reads an entire record, the value of retptr will be word-aligned and positive. The value of retptr is meaningful only until execution of the next stream call for this stream.

STREAM\_\$GET\_CONDITIONAL

**retlen**

Number of bytes of data actually returned. This is a 4-byte integer.

If the call returned any data, retlen has a value equal to the requested number of bytes. It has a value of 0 if the call returned no data.

**seek-key**

Unique key identifying the location of the data returned, in STREAM\_\$SK\_T format. This is a three-element array of 4-byte integers.

If the returned status is nonzero, the seek-key is not useful.

**status**

Completion status, in STATUS\_\$T format. This is a 4-byte integer.

**USAGE**

STREAM\_\$GET\_CONDITIONAL performs read operations on streams such as SIO lines, input pads, and mailboxes, for which data may not yet be available at the time of the call. Under these conditions, STREAM\_\$GET\_REC waits for data. STREAM\_\$GET\_CONDITIONAL never waits. If data is not immediately available, it returns a length of zero.

This call is commonly used in conjunction with STREAM\_\$GET\_EC and EC2\_\$WAIT.

Since data from ordinary files is always available, this call is equivalent to STREAM\_\$GET\_REC for files.

No error occurs if the stream manager cannot find data at the current stream position, unless the current position is known to be at EOF. In this case, a zero is returned in retlen.

**STREAM\_\$GET\_EC**

Gets the eventcount address of the eventcount to be advanced upon any activity within the specified stream.

**FORMAT**

STREAM\_\$GET\_EC (stream-id, stream-key, eventcount-pointer, status)

**INPUT PARAMETERS****stream-id**

The stream ID, in STREAM\_\$ID\_T format. This is a 2-byte integer.

**stream-key**

The key that specifies which eventcount to get, in STREAM\_\$EC\_KEY\_T format. The only value allowed is STREAM\_\$GETREC\_EC\_KEY.

**OUTPUT PARAMETERS****eventcount-pointer**

The eventcount address to be obtained, in EC2\_\$PTR\_T format. EC2\_\$PTR\_T is a pointer to an EC2\_\$EVENTCOUNT\_T array.

**status**

Completion status, in STATUS\_\$T format. This is a 4-byte integer.

**USAGE**

The eventcount is advanced whenever data becomes available through the stream. This call is valid for all streams, including those open to files, pads, mailboxes, and COMMENTS.

If the input pad is in raw mode then an event is counted after each single character stroke; if the keyboard is in cooked mode then an event is counted after each carriage return.

See the description of EC2\_\$WAIT for a description of eventcount data structures. See *Programming With General System Calls* for a discussion of eventcounts.

STREAM\_\$GET\_PRIOR\_REC

STREAM\_\$GET\_PRIOR\_REC

Reads the previous record.

## FORMAT

STREAM\_\$GET\_PRIOR\_REC (stream-id, bufptr, buflen, retptr, retlen,  
seek-key, status)

## INPUT PARAMETERS

### stream-id

Number of a stream on which the object is open, in STREAM\_\$ID\_T format. This is a 2-byte integer.

### bufptr

Pointer to the buffer into which the data may be read, in UNIV\_PTR format. This is a 4-byte integer.

To obtain a value for bufptr, FORTRAN programs can use the IADDR function. Pascal programs can use the ADDR function. The buffer can be aligned on a byte boundary; therefore, the value of bufptr can be odd.

### buflen

Number of bytes of data requested. This is a 4-byte integer.

If the number of bytes remaining in the record is less than buflen, STREAM\_\$GET\_PRIOR\_REC returns the remainder of the record. The value returned in retlen is the number of bytes actually read. If the number of bytes remaining in the record is greater than buflen, the stream manager reads enough data to fill the buffer and returns a negative value in retlen. The absolute value of the returned retlen is the number of bytes remaining in the record.

## OUTPUT PARAMETERS

### retptr

Pointer to the data returned, in UNIV\_PTR format. This is a 4-byte integer.

Address the returned data only by using retptr. The stream manager may use "locate mode," in which it doesn't copy the desired data to the location indicated by bufptr. FORTRAN programs that call the stream manager in locate mode should use the "pointer variable" FORTRAN extension.

Records are aligned on word boundaries. Therefore, if the procedure reads an entire record, the value of retptr will be word-aligned and positive. The value of retptr is meaningful only until execution of the next stream call on this stream.

### retlen

Number of bytes of data returned. This is a 4-byte integer.

If the number of bytes remaining in the record is less than buflen, STREAM\_\$GET\_PRIOR\_REC returns the remainder of the record. The value of retlen is the number of bytes actually read.



If the number of bytes remaining in the record is greater than buflen, the stream manager reads enough data to fill the buffer and returns a negative value in retlen. The absolute value of the returned retlen is the number of unread bytes remaining in the record.

**seek-key**

Unique key identifying the location of the data returned, in STREAM\_\$SK\_T format. This is a three-element array of 4-byte integers.

The seek-key identifies the beginning of the returned data, as it does for STREAM\_\$GET\_REC. Use it in STREAM\_\$SEEK calls followed by STREAM\_\$GET\_REC (not STREAM\_\$GET\_PRIOR\_REC) calls.

If the returned status is nonzero, the seek-key is not useful.

**status**

Completion status, in STATUS\_\$T format. This is a 4-byte integer.

**USAGE**

STREAM\_\$GET\_PRIOR\_REC reads the previous record from the object. The object must be open on the specified stream.

This call operates on file types UASC, PAD, HDR\_UNDEF, and REC with record format F2. It will not work on REC files with record format V1, or on nonfile-type streams such as SIO lines, input pads, or mailboxes. STREAM\_\$GET\_PRIOR\_REC works as follows:

**UASC and REC files**

If the seek-key is positioned at the beginning of a record, it is positioned to the previous record. If the seek-key is in the middle of a record, then its position is not changed.

**HDR\_UNDEF files**

The seek-key is repositioned by subtracting the caller's buffer size from the current position.

After these respective actions are taken, an ordinary STREAM\_\$GET\_REC operation is done.

FORTRAN programs using this procedure in locate mode should use the pointer variable FORTRAN extension. Otherwise, call STREAM\_\$REDEFINE to set move mode before using this procedure.

An error occurs if the stream manager cannot find a record at the current stream position - for example, if the current position is beyond EOF or at BOF (beginning of file).

STREAM\_\$GET\_REC

STREAM\_\$GET\_REC

Reads the next sequential record from an object.

## FORMAT

STREAM\_\$GET\_REC (stream-id, bufptr, buflen, retptr, retlen, seek-key, status)

## INPUT PARAMETERS

### stream-id

Number of a stream on which the object is open, in STREAM\_\$ID\_T format. This is a 2-byte integer.

### bufptr

Pointer to the buffer into which the record may be read, in UNIV\_PTR format. This is a 4-byte integer.

To obtain a value for bufptr, FORTRAN programs can use the IADDR function. Pascal programs can use the ADDR function. The buffer can be aligned on a byte boundary; therefore, the value of bufptr can be odd.

### buflen

Number of bytes of data to be read. This is a 4-byte integer.

If the number of bytes remaining in the record is less than buflen, STREAM\_\$GET\_REC returns the remainder of the record. If the number of bytes remaining in the record is greater than buflen, the stream manager reads enough data to fill the buffer and returns a negative value in retlen. The absolute value of retlen is the number of bytes remaining in the record.

## OUTPUT PARAMETERS

### retptr

Pointer to the data returned, in UNIV\_PTR format. This is a 4-byte integer.

Address the returned data only by using retptr. The stream manager may use "locate mode," in which it doesn't copy the desired data to the location indicated by bufptr. FORTRAN programs that call the stream manager in locate mode should use the "pointer variable" FORTRAN extension.

Records are aligned on word boundaries. Therefore, if the procedure reads an entire record, the value of retptr will be word-aligned and positive. The value of retptr is meaningful only until execution of the next stream call for this stream.

### retlen

Number of bytes of data returned. This is a 4-byte integer.

If the number of bytes remaining in the record is less than buflen, STREAM\_\$GET\_PRIOR\_REC returns the remainder of the record. The value of retlen is the number of bytes actually read.

If the number of bytes remaining in the record is greater than buflen, the stream manager

reads enough data to fill the buffer and returns a negative value in retlen. The absolute value of the returned retlen is the number of unread bytes remaining in the record.

**seek-key**

Unique key identifying the location of the data read, in STREAM\_\$SK\_T format. This is a three-element array of 4-byte integers.

To obtain a seek-key value for the current stream position, call STREAM\_\$GET\_REC with a buflen of 0.

Use the seek-key value in STREAM\_\$SEEK calls followed by STREAM\_\$GET\_REC calls.

If the returned status is nonzero, the seek-key is not useful.

**status**

Completion status, in STATUS\_\$T format. This is a 4-byte integer.

**USAGE**

STREAM\_\$GET\_REC returns at most the requested number of bytes of the next sequential record in the object.

FORTTRAN programs using this procedure in locate mode should use the pointer variable FORTTRAN extension. Otherwise, call STREAM\_\$REDEFINE to set move mode before using this procedure.

An error occurs if the stream manager cannot find a record at the current stream position -- for example, if the current position is at EOF.

STREAM\_\$INQUIRE

STREAM\_\$INQUIRE

Returns information about an object.

## FORMAT

STREAM\_\$INQUIRE (input-mask, inquiry-type, attributes, error-mask, status)

## INPUT PARAMETERS

### input-mask

Integer bit mask indicating the attributes for which information is requested, in STREAM\_\$INQUIRE\_MASK\_T format. This is a 4-byte integer.

The following lists the predefined symbols for each bit position in the mask.

Bit 0	STREAM_\$STRID	Stream-id
Bit 1	STREAM_\$OBJ_NAME	Object name
Bit 1	STREAM_\$OBJ_NAMLEN	Object name length
Bit 2	STREAM_\$REC_LGTH	Record length
Bit 3	STREAM_\$TEMPORARY	Temporary or permanent *
Bit 4	STREAM_\$EXPLICIT_TYPE	Explicit record type
Bit 5	STREAM_\$AB_FLAG	ASCII or binary file
Bit 6	STREAM_\$EXPLICIT_ML	Explicit move mode *
Bit 7	STREAM_\$CC	Carriage control
Bit 8	STREAM_\$REC_TYPE	Record type
Bit 9	STREAM_\$CONC	Object concurrency (not implemented)
Bit 10	STREAM_\$COCONC	Concurrency at open *
Bit 11	STREAM_\$OPOS	Access type *
Bit 12	STREAM_\$PRE_EXIST	Pre-existing object
Bit 13	STREAM_\$HDR_LGTH	Header length
Bit 14	STREAM_\$FILE_LENGTH	File length
Bit 15	STREAM_\$SEEK_KEY	Seek key *
Bit 16	STREAM_\$CUR_REC_LEN	Current record length *
Bit 17	STREAM_\$CUR_REL_REC_NO	Current relative record number *
Bit 18	STREAM_\$BLKS_USED	Number of blocks used
Bit 19	STREAM_\$DTU	Date and time last used
Bit 20	STREAM_\$DTM	Date and time last modified
Bit 21	STREAM_\$SPARSE	Sparsely written file *
Bit 22	STREAM_\$OTYPE	Object type
Bit 23	STREAM_\$CLOSE_ON_EXEC	Close stream on UNIX Exec call
Bit 24	STREAM_\$NDELAY	Forced STREAM_\$GET_CONDITIONAL
Bit 25	STREAM_\$APPEND_MODE	File in append mode
Bit 26	STREAM_\$FORCED_LOCATE	Force locate mode

\* Attributes to which a stream must be open for information to be returned.

Pascal and C programs specify these predefined values as members of a set. FORTRAN programs must add the desired values to each other to result in a correct input-mask value.

**inquiry-type**

Type of inquiry, in STREAM\_\$IR\_OPT format. Possible values are:

**STREAM\_\$USE\_STRID**

Specifies an inquiry by stream ID. On input, the attribute record must contain the stream-id to which the request applies. On output, the attribute record contains the requested information, if the stream is open. If the stream is not open, an error is returned.

**STREAM\_\$NAME\_CONDITIONAL**

Specifies an inquiry by name to be executed only if a stream is open to the object. On input, the attribute record must contain the object's pathname and name-length. On output, the attribute record contains the requested information if a stream is open. If no stream is open to the object, an error is returned.

**STREAM\_\$NAME\_UNCONDITIONAL**

Specifies an unconditional inquiry by name. On input, the attribute record must contain the object's pathname and name-length. On output, the attribute record contains the requested information whether or not a stream is open.

**INPUT/OUTPUT PARAMETERS****attributes**

Record containing attribute information, in STREAM\_\$IR\_REC\_T format. On input, this record contains a pathname or stream-id that identifies the object. On output, this record contains the returned information.

For serial lines, STREAM\_\$INQUIRE returns the default values.

The attribute parameter is able to convey a large amount of information by passing it in many fields of one record (including a number of bit fields). These record fields are listed below along with their size and a brief explanation of the information they transmit.

Stream-id	STREAM_\$ID_T. A 2-byte integer. Specified on input in conjunction with the STREAM_\$USE_STRID inquiry-type.
Name-length	a 2-byte integer. Name-length of the object. Specified on input in conjunction with the STREAM_\$NAME_CONDITIONAL and STREAM_\$NAME_UNCONDITIONAL inquiry-type.
Record length	A longword. If the object has variable-length records, the length of the first record is returned; if fixed-length records, the fixed record length is returned.

STREAM\_\$INQUIRE

Flag1

A field of 16 bits containing Boolean and enumerated values. Each Boolean value occupies one bit in the flag. Enumerated types may occupy more than one bit depending on the number of possible values. The following table lists the bit number(s), the corresponding attribute, the data type of the attribute, and a brief explanation of possible values.

<i>Bit Number</i>	<i>Attribute</i>	<i>Data Type</i>	<i>Explanation</i>
Bit 0	Temporary	Boolean	TRUE if the object is temporary.
Bit 1	Explicit type	Boolean	TRUE if record type is explicit.
Bit 2	ASCII/Binary	Boolean	TRUE if data is ASCII, otherwise it is binary
Bit 3	Force move mode	Boolean	TRUE if move mode is used (only applies to open streams).
Bits 4,5	Unused 1		
Bit 6	Carriage control	STREAM_\$CC_T	Either STREAM_\$F77_CC (FORTRAN) or STREAM_\$APOLLO_CC (DOMAIN).
Bits 7-9	Unused 2		
Bits 10,11	Record Type	STREAM_\$RTYPE_T	Either STREAM_\$F2 (fixed length), STREAM_\$V1 (variable-length), or STREAM_\$UNDEF (undefined).
Bits 12,13	Unused 3		
Bits 14,15	Object Concurrency	STREAM_\$FCONC_T	Not implemented. Always is STREAM_\$N_AND_N.

Flag2 A word. This is a field of bits containing Boolean and enumerated values. Each Boolean value occupies one bit in the flag. Enumerated types may occupy more than one bit depending on the number of possible values. The following table lists the bit number(s), the corresponding attribute, the data type of the attribute, and a brief explanation of possible values.

<i>Bit Number</i>	<i>Attribute</i>	<i>Data Type</i>	<i>Explanation</i>
Bits 16,17	Unused 4		
Bits 18,19	Concurrency at open	STREAM_OMODE_T	Either STREAM_\$UNREGULATED, STREAM_\$NO_CONC_WRITE, or STREAM_\$CONTROLLED_SHARING. Only returned from opened streams.
Bits 20,21	Unused 5		
Bits 2-24	Access Type	STREAM_\$OPOS_T	Either STREAM_\$READ, STREAM_\$WRITE, STREAM_\$UPDATE, STREAM_\$APPEND, or STREAM_\$OVERWRITE. Only returned from opened streams.
Bit 25	Pre-existing	Boolean	TRUE if object already exists.
Bits 26-31	Unused 5		

Unused 6

Header length A 2-byte integer. Length of streams header.

File length A 4-byte integer. Total file length, in bytes (including header).

Seek key STREAM\_\$SK\_T. A three-element INTEGER\*4 array. Current stream position. This attribute is only returned from opened streams.

Current record length

A 4-byte integer. Size of current record. This attribute is only returned from opened streams.

Current relative record number

A 4-byte integer. Only applies to files with fixed-length records. This attribute is only returned from opened streams.

Blocks used

A 4-byte integer. Number of disk blocks currently used for the file. Only applies to pads and disk files.

Date/Time Used

TIME\_\$CLOCKH\_T. A 4-byte integer. Date and time of last use. Only applies to pads and disk files.

STREAM\_\$INQUIRE

Date/Time Modified

TIME\_\$CLOCKH\_T. A 4-byte integer. Date and time of last modification. Only applies to pads and disk files.

Object type

UID\_\$T. A two element INTEGER\*4 array. Type UID of the object. The following table lists valid UID types.

UASC_\$UID	UASC file
RECORDS_\$UID	Record-structured file
HDR_\$UNDEF_\$UID	Nonrecord-structured file
OBJECT_FILE_\$UID	Object module file
PAD_\$UID	Saved transcript pad
INPUT_PAD_\$UID	Input pad
SIO_\$UID	Serial line descriptor file
MBX_\$UID	Mailbox object
MT_\$UID	Magnetic tape descriptor file

Sparse flag

Boolean. When TRUE, file allocation may have "holes."



**Flag3** A word. This is a field of bits containing Boolean and enumerated values. Each Boolean value occupies one bit in the flag. Enumerated types may occupy more than one bit depending on the number of possible values. The following table lists the bit number(s), the corresponding attribute, the data type of the attribute, and a brief explanation of possible values.

<i>Bit Number</i>	<i>Attribute</i>	<i>Data Type</i>	<i>Explanation</i>
Bits 16-27	Unused 7		
Bit 29	Close on exec	Boolean	TRUE causes stream to be closed upon an AUX exec call.
Bit 30	No delay mode	Boolean	When TRUE, any system call that reads data from a stream will act like STREAM_\$GET_CONDITIONAL (returns if data not available).
Bit 31	Append mode	Boolean	When TRUE, any call to STREAM_\$PUT_REC or STREAM_\$PUT_CHR does a seek to EOF before writing any data.
Bit 28	Force locate mode	Boolean	Normally, if the force move mode bit (bit 6) is not set, streams may use either move mode or locate mode. If this bit is set, streams will only use locate mode, the caller need not supply a buffer. This option can only be set for file-type streams (UASC, REC, HDR_UNDEF, and CASE_HM).

Unused 8 A 4-byte integer.

**Name** NAME\_\$PNAME\_T. A character array of up to 256 elements. The name of the object. Specified on input with STREAM\_\$NAME\_CONDITIONAL and STREAM\_\$NAME\_UNCONDITIONAL inquiry-types.

The array specified as the attribute parameter need not be large enough for every field, but just sufficient to span the required fields. For example, to inquire on explicit move mode, only a six-element INTEGER\*2 array is required (for a FORTRAN program), because the necessary flag is in FLAG1.

Accessing attribute record bit fields is discussed in detail in the *Programming With General System Calls*.

## OUTPUT PARAMETERS

### **error-mask**

Integer bit-mask indicating the requested fields that could not be returned, in STREAM\_\$INQUIRE\_MASK\_T format. This is a 4-byte integer.

## STREAM\_\$INQUIRE

This procedure may complete with partial success if it can return some, but not all, of the requested attributes. If an attribute is unavailable, STREAM\_\$INQUIRE sets the corresponding bit in the error mask and continues to inquire about other attributes. In cases of partial success, the returned status code is nonzero. The program must check the error mask to find out where the error occurred.

### **status**

Completion status, in STATUS\_\$T format. This is a 4-byte integer.

### **USAGE**

STREAM\_\$INQUIRE returns the attributes of the object specified in the mask. To receive the information you must specify either the stream ID of the object or the object name and name-length. However, some attributes such as access type apply only to objects to which a stream is open. These parameters are marked with asterisks in the inquiry mask parameter description above.

The stream position does not change as a result of this call.

**STREAM\_\$OPEN**

Opens a stream to an existing object.

**FORMAT**

STREAM\_\$OPEN (pathname, name-length, access, concurrency, stream-id, status)

**INPUT PARAMETERS****pathname**

Name of the object to be opened, in NAME\_\$PNAME\_T format. This is a character array of up to 256 elements.

**name-length**

Length of the pathname. This is a 2-byte integer.

**access**

Type of access requested, in STREAM\_\$OPOS\_T format. Possible values are:

**STREAM\_\$APPEND**

Permits adding data to the end of an object. The stream pointer points to the end of the object (EOF).

**STREAM\_\$MAKE\_BACKUP**

Not for use with STREAM\_\$OPEN. (See below for a detailed description.)

**STREAM\_\$OVERWRITE**

Permits replacing the entire contents of an object. The stream pointer is positioned at the start of the object data and data is truncated.

**STREAM\_\$READ**

Permits reading data from an existing object.

**STREAM\_\$UPDATE**

Permits replacing selected portions of the contents of an object. The stream pointer is positioned at the start of the object data, just past the header if it has one.

**STREAM\_\$WRITE**

Same as STREAM\_\$UPDATE.

**concurrency**

Requested concurrency at open, in STREAM\_\$OMODE\_T format. Possible values are:

**STREAM\_\$CONTROLLED\_SHARING**

Does not allow concurrent read and write access.

**STREAM\_\$NO\_CONC\_WRITE**

Does not allow concurrent read and write access.

**STREAM\_\$UNREGULATED**

Allows concurrent read and write access.

STREAM\_\$OPEN

## OUTPUT PARAMETERS

### **stream-id**

Number of the stream on which the object was opened, in STREAM\_\$ID\_T format. This is a 2-byte integer.

### **status**

Completion status, in STATUS\_\$T format. This is a 4-byte integer.

## USAGE

This routine opens a stream to the named object and assigns access and concurrency types. It returns the stream ID to be used in subsequent stream activity with the object.

An error occurs if the object does not exist.

STREAM\_\$OPEN does not return information about the object's characteristics. Use STREAM\_\$INQUIRE to obtain that information.

If the object already exists, its attributes remain the same when it is opened. For streams to SIO lines, however, "cooked" input mode is always in effect when the stream is opened. To change the object's attributes, call STREAM\_\$REDEFINE (or SIO\_\$CONTROL) before writing to the object.

**STREAM\_\$PUT\_CHR**

Writes data to an object without terminating the current record, if one exists.

**FORMAT**

STREAM\_\$PUT\_CHR (stream-id, bufptr, buflen, seek-key status)

**INPUT PARAMETERS****stream-id**

Number of a stream on which an object is open, in STREAM\_\$ID\_T format. This is a 2-byte integer.

**bufptr**

Pointer to the data to be written, in UNIV\_PTR format. This is a 4-byte integer.

FORTRAN programs can use IADDR to obtain the buffer address for the bufptr parameter. Pascal programs can use ADDR. Alternately, programs in either language can use pointer variables.

**buflen**

Number of bytes of data to be written. This is a 4-byte integer.

**OUTPUT PARAMETERS****seek-key**

Unique key identifying the location of the data written, in STREAM\_\$SK\_T format. This is a three-element array of 4-byte integers.

The seek key allows random access to the output data by a subsequent STREAM\_\$SEEK call.

**status**

Completion status, in STATUS\_\$T format. This is a 4-byte integer.

**USAGE**

STREAM\_\$PUT\_CHR writes the specified number of bytes to the object, but does not terminate a record.

Use this procedure to write data in nonrecord-structured files or to compose records piece by piece. Be sure to call STREAM\_\$REDEFINE to set STREAM\_\$UNDEF as the record type, and HDR\_UNDEF\_\$UID as the object type, before writing any output.

Records of fixed-length format automatically change to variable-length if this write operation extends the current record beyond the length of existing records. In this case, no error occurs. For files with explicit fixed-length records, an error occurs if this write operation extends the current record beyond the fixed-record size.

For files with variable-length records, no record length checking is performed. Therefore, take care not to alter the count field of the following record.

## STREAM\_\$PUT\_CHR

Record size can increase as a result of this call, but cannot decrease. For instance, after you overwrite the first 20 bytes of a 32-byte record, the last 12 bytes still contain the original data, and the count field remains the same. To terminate a record and update its count field, use STREAM\_\$PUT\_REC.

STREAM\_\$PUT\_CHR and STREAM\_\$PUT\_REC operate identically when applied to SIO lines, UASC files, and keyboards.

**STREAM\_\$PUT\_REC**

Writes data to an object and terminates the current record, if one exists.

**FORMAT**

STREAM\_\$PUT\_REC (stream-id, bufptr, buflen, seek-key, status)

**INPUT PARAMETERS****stream-id**

Number of a stream on which the object is open, in STREAM\_\$ID\_T format. This is a 2-byte integer.

**bufptr**

Pointer to the data to be written, in UNIV\_PTR format. This is a 4-byte integer.

FORTRAN programs can use IADDR to obtain the buffer address for the bufptr parameter. Pascal programs can use ADDR. Alternately, programs in either language can use pointer variables.

**buflen**

Number of bytes of data to be written. This is a 4-byte integer.

For files with explicit fixed-length records, an error occurs if the total record length is not equal to the fixed-record length.

If you specify a buflen of zero, STREAM\_\$PUT\_REC simply terminates the current record and updates its count field.

**OUTPUT PARAMETERS****seek-key**

Unique key identifying the location of the data in the object, in STREAM\_\$SK\_T format. This is a three-element array of 4-byte integers.

The seek key allows random access to the output data by a subsequent STREAM\_\$SEEK call.

**status**

Completion status, in STATUS\_\$T format. This is a 4-byte integer.

**USAGE**

STREAM\_\$PUT\_REC queues data for output to the object. It does not necessarily write the data on the device. Device writes may be performed asynchronously.

Records of default format (implicit fixed-length) automatically change to variable-length if the new record differs in length from any existing records. No error occurs.

Existing data in variable-length records are overwritten if the stream position is not at EOF. No error occurs.

STREAM\_\$PUT\_REC

STREAM\_\$PUT\_REC never inserts newline characters in the object on its own. You must do this yourself if you want newlines to appear in the object.

STREAM\_\$PUT\_CHR and STREAM\_\$PUT\_REC operate identically when applied to SIO lines, UASC files, and keyboards.



**STREAM\_ \$REDEFINE**

Changes one or more attributes of an object that is open on a stream.

**FORMAT**

STREAM\_ \$REDEFINE (stream-id, input-mask, attributes, error-mask, status)

**INPUT PARAMETERS****stream-id**

Number of a stream on which the object is open, in STREAM\_ \$ID\_ T format. This is a 2-byte integer.

**input-mask**

An integer bit mask showing which attributes you want to redefine, in STREAM\_ \$REDEF\_ MASK\_ T format. This is a 4-byte integer.

Bits 4 through 9 are the ones most commonly changed. The following lists the predefined symbols for each bit position in the mask.

Bit 1	STREAM_ \$OBJ_ NAME	Object name
Bit 1	STREAM_ \$OBJ_ NAMLEN	Object name length
Bit 2	STREAM_ \$REC_ LGTH	Record length
Bit 3	STREAM_ \$TEMPORARY	Temporary or permanent
Bit 4	STREAM_ \$EXPLICIT_ TYPE	Explicit record type
Bit 5	STREAM_ \$AB_ FLAG	ASCII or binary file
Bit 6	STREAM_ \$EXPLICIT_ ML	Explicit move mode
Bit 7	STREAM_ \$CC	Carriage control
Bit 8	STREAM_ \$REC_ TYPE	Record type
Bit 9	STREAM_ \$CONC	Object concurrency (not implemented)
Bit 10	STREAM_ \$OCONC	Concurrency at open
Bit 11	STREAM_ \$OPOS	Access type
Bit 21	STREAM_ \$SPARSE	Sparsely written file
Bit 22	STREAM_ \$OTYPE	Object type
Bit 23	STREAM_ \$CLOSE_ ON_ EXEC	Close stream on DOMAIN/IX Exec call
Bit 24	STREAM_ \$NDELAY	Forced STREAM_ \$GET_ CONDITIONAL
Bit 25	STREAM_ \$APPEND_ MODE	File in append mode
Bit 26	STREAM_ \$FORCED_ LOCATE	Force locate mode

Pascal and C programs specify these predefined values as members of a set. FORTRAN programs must add the desired values to each other to result in a correct input-mask value.

Note that some bit numbers are missing (0, 12 - 20). This is because STREAM\_ \$INQUIRE and STREAM\_ \$REDEFINE use the same attribute record, however certain attributes that can be inquired upon cannot be redefined.

**attributes**

Record containing new values for attributes, in STREAM\_ \$IR\_ REC\_ T format.

The attribute parameter is able to specify redefinition of a large number of attributes by passing information in many fields of one record (including a number of bit fields). These record fields are listed below along with their size and a brief explanation of the information they transmit.

STREAM\_ \$REDEFINE

- Stream-id            STREAM\_ \$ID\_ T. A 2-byte integer. Not redefinable.
- Name-length        a 2-byte integer. Name-length of the object. Specified on input in conjunction with the STREAM\_ \$NAME\_ CONDITIONAL and STREAM\_ \$NAME\_ UNCONDITIONAL inquiry-type.
- Record length      A longword. If the object has variable-length records, the length of the longest record is returned; if fixed-length records, the fixed record length is returned.
- Flag1              A field of 16 bits containing Boolean and enumerated values. Each Boolean value occupies one bit in the flag. Enumerated types may occupy more than one bit depending on the number of possible values. The following table lists the bit number(s), the corresponding attribute, the data type of the attribute, and a brief explanation of possible values.

<i>Bit Number</i>	<i>Attribute</i>	<i>Data Type</i>	<i>Explanation</i>
Bit 0	Temporary	Boolean	TRUE if the object is temporary.
Bit 1	Explicit type	Boolean	TRUE if record type is explicit.
Bit 2	ASCII/Binary	Boolean	TRUE if data is ASCII, otherwise it is binary.
Bit 3	Force move mode	Boolean	TRUE if move mode is used. (only applies to open streams)
Bits 4,5	Unused 1		
Bit 6	Carriage control	STREAM_ \$CC_ T	Either STREAM_ \$F77_ CC (FORTRAN) or STREAM_ \$APOLLO_ CC (DOMAIN).
Bits 7-9	Unused 2		
Bits 10,11	Record Type	STREAM_ \$RTYPE_ T	Either STREAM_ \$F2 (fixed length), STREAM_ \$V1 (variable-length), or STREAM_ \$UNDEF (undefined).
Bits 12,13	Unused 3		
Bits 14,15	Object Concurrency	STREAM_ \$FCONC_ T	Not implemented. Always STREAM_ \$N_ AND_ N.

- Flag2              A word. This is a field of bits containing Boolean and enumerated values. Each Boolean value occupies one bit in the flag. Enumerated types may occupy more than one bit depending on the number of possible values. The following table lists the bit number(s), the corresponding attribute, the data type of the attribute, and a brief explanation of possible values.

<i>Bit Number</i>	<i>Attribute</i>	<i>Data type</i>	<i>Explanation</i>
Bits 16,17	Unused 4		
Bits 18,19	Concurrency at open	STREAM_ OMODE_ T	Either STREAM_ \$UNREGULATED, STREAM_ \$NO_ CONC_ WRITE, or STREAM_ \$CONTROLLED_ SHARING. Only returned from opened streams.
Bits 20,21	Unused 5		
Bits 2 -24	Access Type	STREAM_ \$OPOS_ T	Either STREAM_ \$READ, STREAM_ \$WRITE, STREAM_ \$UPDATE, STREAM_ \$APPEND, or STREAM_ \$OVERWRITE. Only returned from opened streams.
Bit 25	Pre-existing	Boolean	Not redefinable.
Bits 26-31	Unused 5		

Unused 6

Header length     A 2-byte integer. Not redefinable.

File length        A 4-byte integer. Not redefinable.

Seek key            STREAM\_ \$SK\_ T. A three-element INTEGER\*4 array. Not redefinable.

Current record length  
                    A 4-byte integer. Not redefinable.

Current relative record number  
                    A 4-byte integer. Not redefinable.

Blocks used        A 4-byte integer. Not redefinable.

Date/Time Used    TIME\_ \$CLOCKH\_ T. A 4-byte integer. Not redefinable.

Date/Time Modified  
                    TIME\_ \$CLOCKH\_ T. A 4-byte integer. Not redefinable.

.Object type        UID\_ \$T. A two element INTEGER\*4 array. Type UID of the object. The following table lists valid UID types.

UASC_ \$UID	UASC file
RECORDS_ \$UID	Record-structured file
HDR_ \$UNDEF_ \$UID	Nonrecord-structured file
OBJECT_ FILE_ \$UID	Object module file
PAD_ \$UID	Saved transcript pad
INPUT_ PAD_ \$UID	Input pad
SIO_ \$UID	Serial line descriptor file
MBX_ \$UID	Mailbox object
MT_ \$UID	Magnetic tape descriptor file

STREAM\_ \$REDEFINE

- Sparse flag Boolean. When TRUE, file allocation may have "holes". See the *Programming With General System Calls* manual.
- Flag3 A word. This is a field of bits containing Boolean and enumerated values. Each Boolean value occupies one bit in the flag. Enumerated types may occupy more than one bit depending on the number of possible values. The following table lists the bit number(s), the corresponding attribute, the data type of the attribute, and a brief explanation of possible values.

Bit Number	Attribute	Data Type	Explanation
Bits 16-27	Unused 7		
Bit 29	Close on exec	Boolean	TRUE causes stream to be closed upon an DOMAIN/IX exec call.
Bit 30	No delay mode	Boolean	When TRUE, any system call that reads data from a stream will act like STREAM_\$GET_CONDITIONAL (returns if data not available).
Bit 31	Append mode	Boolean	When TRUE, any call to STREAM_\$PUT_REC or STREAM_\$PUT_CHR does a seek to EOF before writing any data.
Bit 28	Force locate mode	Boolean	Normally, if the force move mode bit (bit 6) is not set, streams may use either move mode or locate mode. If this bit is set, streams will only use locate mode; the caller need not supply a buffer. This option can only be set for file-type streams (UASC, REC, HDR_UNDEF, and CASE_HM).

Unused 8 A 4-byte integer.

Name NAME\_ \$PNAME\_ T. A character array of up to 256 elements. The name of the object. Specified on input with STREAM\_ \$NAME\_ CONDITIONAL and STREAM\_ \$NAME\_ UNCONDITIONAL inquiry types.

The array specified as the attribute parameter need not be large enough for every field, but just sufficient to span the required fields. For example, to redefine explicit move mode, only a six-element INTEGER\*2 array is required (for a FORTRAN program), because the necessary flag is in FLAG1.

Accessing attribute record bit fields is discussed in detail in the *Programming With General System Calls* handbook.

## OUTPUT PARAMETERS

### **error-mask**

An integer bit-mask indicating any requested changes that were not made, in STREAM\_ \$REDEF\_ MASK\_ T format. This is a 4-byte integer.

This procedure may complete with partial success if it can redefine some, but not all, of the requested attributes. If an attribute is not changed, STREAM\_ \$REDEFINE sets the corresponding bit in the error mask and continues to redefine other attributes. In cases of partial success, the returned status code is nonzero. The program must check the error mask to find out where the error occurred.

### **status**

Completion status, in STATUS\_ \$T format. This is a 4-byte integer.

## USAGE

STREAM\_ \$REDEFINE changes one or more attributes of an object to which you have a stream open. Wherever bits are set in the input mask, STREAM\_ \$REDEFINE tries to copy information from the corresponding fields of the attribute record to the object. Wherever bits are not set in the input mask, the corresponding attributes of the object do not change.

FORTRAN programs that use the stream manager to read files must call STREAM\_ \$REDEFINE to request explicit move mode.

You can use STREAM\_ \$REDEFINE only on streams with write access. However, you can use it to change read access to write access.

You cannot use STREAM\_ \$REDEFINE to change the stream position (use STREAM\_ \$SEEK instead), to change the object's length (use STREAM\_ \$TRUNCATE instead), or to change a serial line's attributes (use SIO\_ \$CONTROL).

STREAM\_ \$REPLACE

STREAM\_ \$REPLACE

Writes data to an object without changing the length of the current record.

## FORMAT

STREAM\_ \$REPLACE (stream-id, bufptr, buflen, seek-key, status)

## INPUT PARAMETERS

### stream-id

Number of a stream on which the object is open, in STREAM\_ \$ID\_ T format. This is a 2-byte integer.

### bufptr

Pointer to the data to be written, in UNIV\_ PTR format. This is a 4-byte integer.

FORTRAN programs can use IADDR to obtain the buffer address for the bufptr parameter. Pascal programs can use ADDR. Alternately, programs in either language can use pointer variables.

### buflen

Number of bytes of data to be written. This is a 4-byte integer.

## OUTPUT PARAMETERS

### seek-key

Unique key identifying the location of the output data in the object, in STREAM\_ \$SK\_ T format. This is a three-element array of 4-byte integers.

The seek key allows random access to the output data by a subsequent STREAM\_ \$SEEK call.

### status

Completion status, in STATUS\_ \$T format. This is a 4-byte integer.

## USAGE

STREAM\_ \$REPLACE replaces a record in the object. Call STREAM\_ \$PUT\_ REC to add records to an object.

For record-structured objects, this call terminates the current record. The length of the current record must be exactly the same as the length of the record it replaces. An error occurs if the record lengths are different. You can use STREAM\_ \$PUT\_ REC and STREAM\_ \$PUT\_ CHR to overwrite existing data in a file or pad with no record length checking.

For nonrecord-structured objects, this call writes the specified number of characters. No record-length errors can occur.

Like STREAM\_ \$PUT\_ REC, STREAM\_ \$REPLACE queues data for output to the object. It does not necessarily write the data on the device. Device writes may be performed asynchronously.

STREAM\_ \$SEEK

STREAM\_ \$SEEK

Moves the stream position.

## FORMAT

STREAM\_ \$SEEK (stream-id, seek-base, seek-type,  
{seek-key|signed-offset}, status)

## INPUT PARAMETERS

### stream-id

Number of a stream on which the object is open, in STREAM\_ \$ID\_ T format. This is a 2-byte integer.

### seek-base

Type of data on which the seek is based, in STREAM\_ \$PARM1\_ T format. Possible values are:

STREAM\_ \$CHR

Character-based seek.

STREAM\_ \$EOF

End-of-file based seek. (Any offset specified with STREAM\_ \$EOF is ignored.)

STREAM\_ \$KEY

Keyed value-based seek.

STREAM\_ \$REC

Record-based seek.

### seek-type

Value defining the relationship between the seek-base and the seek-key or signed-offset, in STREAM\_ \$PARM2\_ T format. Possible values are:

STREAM\_ \$RELATIVE

Moves the stream position relative to the current stream marker. Relative positioning is only valid for character (STREAM\_ \$CHR) and record (STREAM\_ \$REC) based seeks. Specifying either STREAM\_ \$EOF or STREAM\_ \$KEY with STREAM\_ \$RELATIVE results in an error status. A positive offset moves the stream position towards EOF and a negative offset moves the stream position towards the beginning of the object. Relative seeks start at 0; that is, an offset of 0 denotes the current position.

STREAM\_ \$ABSOLUTE

Seeks for an absolute position in the object. In absolute seeks, all four seek bases are valid. STREAM\_ \$EOF and STREAM\_ \$KEY can only be used in absolute seeks. Absolute character and record-based seeks start from the beginning of the object (past the header) if the offset is positive. If the offset is negative, the seek starts at EOF. Absolute seeks start at 1. An error occurs if you specify an offset of 0 for an absolute seek.



**seek-key**

Unique value identifying the data sought, in STREAM\_ \$SK\_ T format. This is a three-element array of 4-byte integers.

**signed-offset**

Offset to be used in calculating the new stream position. This is a 4-byte integer.

**OUTPUT PARAMETERS****status**

Completion status, in STATUS\_ \$T format. This is a 4-byte integer.

**USAGE**

STREAM\_ \$SEEK moves the stream marker to a specific location, or to an offset from a known location. It does not move any data.

An error occurs if the object does not support random access.

An error occurs if the program attempts to move the stream position beyond the beginning or end of the file.

Character-based seeks in record-structured objects cannot move the stream position beyond the current record.

Record-based seeks apply only to objects with fixed-length records. For objects with variable-length records, save the seek keys returned when the object is written, then perform keyed seeks.

For UASC files, specifying STREAM\_ \$REC simulates a fixed-length record seek by finding the length of the first record in the file and using that as the record length for all records. STREAM\_ \$CHR-seeks work like they do in nonrecord-structured files.

The stream marker must be aligned on a record boundary when you specify STREAM\_ \$REC. If alignment is incorrect, an error occurs. Similarly, if positioning is relative to EOF (that is, a negative offset), the current EOF must be on a record boundary.

The following examples illustrate the difference between absolute and relative seeks. The first example is an absolute seek for the 16th character in the object:

```
STREAM_ $SEEK(stream_id, STREAM_ $CHR, STREAM_ $ABSOLUTE,
              16, status)
```

The second example positions the stream marker seven records closer to the beginning of the object than it was before the call.

```
STREAM_ $SEEK(stream_id, STREAM_ $REC, STREAM_ $RELATIVE,
              -7, status)
```

STREAM\_\$TRUNCATE

STREAM\_\$TRUNCATE

Writes EOF at the current stream position.

## FORMAT

STREAM\_\$TRUNCATE (stream-id, status)

## INPUT PARAMETERS

### stream-id

Number of a stream on which the object is open, in STREAM\_\$ID\_T format. This is a 2-byte integer.

## OUTPUT PARAMETERS

### status

Completion status, in STATUS\_\$T format. This is a 4-byte integer.

## USAGE

STREAM\_\$TRUNCATE decreases the value of the file length attribute to match the stream pointer's current position. (Writing data to a stream that lengthens the object implicitly increases this attribute's value.) This sets EOF to the stream pointer's position, effectively deleting any data in the object past the stream pointer. If the stream position is already at EOF, truncating the object has no effect.

You can only truncate disk files and pads that the Display Manager is not using. Trying to truncate any other type of object returns an error status code.

Truncating an object does not close the stream.

**ERRORS**

STATUS\_\$OK

Successful completion.

STREAM\_\$ALREADY\_EXISTS

STREAM\_\$WRITE specified on STREAM\_\$CREATE.

STREAM\_\$BAD\_CHAR\_SEEK

Attempted character seek before start of current (variable length) record.

STREAM\_\$BAD\_COUNT\_FIELD\_IN\_FILE

Count field for current record is bad.

STREAM\_\$BAD\_FILE\_HDR

File header is no good (CRC error).

STREAM\_\$BAD\_LOCATION

Bad location parameter in create call.

STREAM\_\$BAD\_OPEN\_xp

OPEN\_XP must reference a stream that is already open on this node.

STREAM\_\$BAD\_POS\_ON\_REC\_SEEK

Relative record seek is not legal unless the reference point is on record boundary.

STREAM\_\$BAD\_RELATED\_PAD

PAD\_\$CREATE attempted with an invalid or unopened related pad.

STREAM\_\$BAD\_SHARED\_CURSOR\_REFCNT

Reference count on a shared file cursor went below zero.

STREAM\_\$BOF\_ERR

Attempted seek beyond BOF; e.g., offset=0.

STREAM\_\$CANT\_DELETE\_OLD\_NAME

WARNING : New name added but old cannot be deleted.

STREAM\_\$CANT\_SWITCH

Too many mapped objects to perform switch.

STREAM\_\$CLOSE\_ANOMALY

WARNING : Close successful but name of (temporary) object on this stream no longer references the same object.

STREAM\_\$CONCURRENCY\_VIOLATION

Requested access violates concurrency constraints.

STREAM\_\$DEVICE\_MUST\_BE\_LOCAL

Cannot open stream to remote device.

STREAM\_\$DIR\_NOT\_FOUND

Could not find directory in pathname on create.

STREAM\_\$END\_OF\_FILE

End of file.

STREAM\_\$EOF\_PAD\_PUT\_ERR

PUT\_REC legal only at EOF on pads; EOF has moved.

STREAM ERRORS

STREAM\_\$FILE\_TROUBLE\_WARNING  
WARNING: (SALVAGER) File trouble bit set in VTOCE.

STREAM\_\$FROM\_STRID\_NOT\_OPEN  
From stream is not open on switch request.

STREAM\_\$ID\_OOR  
Stream ID out-of-range (invalid).

STREAM\_\$ILL\_FORCED\_LOCATE  
Forced locate is only legal for disk files.

STREAM\_\$ILLEGAL\_NAME\_REDEFINE  
Attempted name change requires copying file.

STREAM\_\$ILLEGAL\_OBJ\_TYPE  
Cannot open a stream for this type of object.

STREAM\_\$ILLEGAL\_OPERATION  
This operation is illegal on named stream.

STREAM\_\$ILLEGAL\_PAD\_CLOSE  
Illegal to close transcript pad before related input pad.

STREAM\_\$ILLEGAL\_PAD\_CREATE\_TYPE  
PAD\_CREATE illegal with this type of object.

STREAM\_\$ILLEGAL\_PARAM\_COMB  
Illegal parameter combination for this operation.

STREAM\_\$ILLEGAL\_W\_VAR\_LGTH\_RECS  
Operation illegal with variable length records.

STREAM\_\$INQUIRE\_TYPE\_ERR  
Inquire (by name) about object that cannot be opened on a stream because of its type.

STREAM\_\$INQUIRE\_WARNING  
WARNING : Inquire-by-name is returning data only on first of multiple streams on which object is currently open.

STREAM\_\$INSUFF\_MEMORY  
Not enough virtual memory.

STREAM\_\$INSUFFICIENT\_RIGHTS  
Insufficient rights for requested access to object.

STREAM\_\$INTERNAL\_FATAL\_ERR  
Internal fatal error on table reverification.

STREAM\_\$INTERNAL\_MM\_ERR  
Internal fatal error in stream memory management (windowing).

STREAM\_\$INVALID\_DATA  
Bad data in call to VT\_\$PUT.

STREAM\_\$NAME\_CONFUSION  
Object already open under another name on another stream.

STREAM\_\$NAME\_NOT\_FOUND  
Name not found.

STREAM\_ \$NAME\_ REQD  
STREAM\_ \$OPEN without a name is illegal.

STREAM\_ \$NEED\_ MOVE\_ MODE  
Forced locate is set and could not do it.

STREAM\_ \$NEVER\_ CLOSED  
System (or process) crash prevented complete close.

STREAM\_ \$NO\_ AVAIL\_ TARGET  
No available target stream to switch to.

STREAM\_ \$NO\_ MORE\_ STREAMS  
No more streams.

STREAM\_ \$NO\_ RIGHTS  
No rights to access object.

STREAM\_ \$NO\_ SUCH\_ VERSION  
Specified DSEE version does not exist.

STREAM\_ \$NO\_ TABLE\_ SPACE  
Table space error; cover stream table exhausted.

STREAM\_ \$NOT\_ OPEN  
Operation attempted on unopened stream.

STREAM\_ \$NOT\_ THRU\_ LINK  
Cannot create file though link.

STREAM\_ \$OBJ\_ DELETED  
File has been deleted.

STREAM\_ \$OBJECT\_ NOT\_ FOUND  
Object associated with this name not found (may not exist).

STREAM\_ \$OBJECT\_ READ\_ ONLY  
Cannot open this object for writing.

STREAM\_ \$OUT\_ OF\_ SHARED\_ CURSORS  
Per-mode shared file cursor pool is exhausted.

STREAM\_ \$PART\_ REC\_ WARN  
WARNING : Partial record at the end of a file with fixed length records.

STREAM\_ \$PERM\_ FILE\_ NEEDS\_ NAME  
Only temporary files may be unnamed.

STREAM\_ \$PUT\_ BAD\_ REC\_ LEN  
Attempted put of wrong length record.

STREAM\_ \$READ\_ ONLY\_ ERR  
Attempted write to read-only stream.

STREAM\_ \$REDEFINE\_ PAD\_ ERR  
Cannot redefine this attribute of a pad.

STREAM\_ \$REPLACE\_ LGTH\_ ERR  
Attempted record length change on replace request.

STREAM ERRORS

STREAM\_ \$RESOURCE\_LOCK\_ERR

Unable to lock resources required to process request.

STREAM\_ \$SIO\_NOT\_LOCAL

SIO object not in /DEV.

STREAM\_ \$SOMETHING\_FAILED

Partial or complete failure of inquire or redefine (ERR\_MASK is nonempty).

STREAM\_ \$STREAM\_NOT\_FOUND

No stream found in conditional inquire.

STREAM\_ \$XP\_BUF\_TOO\_SMALL

Buffer supplied to GET\_XP too small.

# TIME

## Table of Contents

Data Types	TIME-2
TIME_\$CLOCK	TIME-4
TIME_\$GET_EC	TIME-5
TIME_\$WAIT	TIME-6
Errors	TIME-7





## TIME

The TIME (Timer) programming calls manage the real-time interval clock. This section describes their data types, call syntax, and error codes. Refer to the Introduction at the beginning of this manual for a description of data-type diagrams and call syntax format.


TIME DATA TYPES

CONSTANTS

- TIME\_\$ABSOLUTE                      Specifies absolute time.
- TIME\_\$CLOCKH\_KEY                  Eventcount key value.
- TIME\_\$RELATIVE                     Specifys relative time.

DATA TYPES

- EC2\_\$PTR\_T                          A 4-byte integer. Address of an eventcount.
- STATUS\_\$T                          A status code. The diagram below illustrates the STATUS\_\$T data type:



Field Description:

**all**  
All 32 bits in the status code.

**fail**  
The fail bit. If this bit is set, the error was not within the scope of the module invoked, but occurred within a lower-level module (bit 31).

**subsys**  
The subsystem that encountered the error (bits 24 - 30).


**modc**  
The module that encountered the error (bits 16 - 23).

code

A signed number that identifies the type of error that occurred (bits 0 - 15).

TIME\_\$CLOCK\_T


Internal representation of time. The diagram below illustrates the TIME\_\$CLOCK\_T data type:



Field Description:

high  
High 32 bits of the clock.

low  
Low 16 bits of the clock.



Field Description:

high16  
High 16 bits of the clock.

low32  
Low 32 bits of the clock.

TIME\_\$KEY\_T

A 2-byte integer. An event count key. One of the following predefined values:

TIME\_\$CLOCKH\_KEY  
Only permissible value.

TIME\_\$REL\_ABS\_T

A 2-byte integer. An indicator of type of time. One of the following predefined values:

TIME\_\$RELATIVE  
Relative time.

TIME\_\$ABSOLUTE  
Absolute time.

TIME\_\$CLOCK

TIME\_\$CLOCK

Returns the current UTC time.

#### FORMAT

TIME\_\$CLOCK (clock)

#### OUTPUT PARAMETERS

##### clock

The current Coordinated Universal Time, in TIME\_\$CLOCK\_T format. This data type is 6 bytes long. See the CAL Data Types section for more information.

#### USAGE

TIME\_\$CLOCK returns the current time of day, in UTC format. It is represented as the number of 4-microsecond periods that have elapsed since January 1, 1980 at 00:00.

To get the local time, use CAL\_\$GET\_LOCAL\_TIME instead of this procedure. To compute the local time from the value returned by TIME\_\$CLOCK, use CAL\_\$APPLY\_LOCAL\_OFFSET.

**TIME\_\$GET\_EC**

Gets the address of the time eventcount, which is advanced about every 0.25 second.

**FORMAT**

TIME\_\$GET\_EC (time-key, eventcount-pointer, status)

**INPUT PARAMETERS****time-key**

The key specifying which time eventcount the system should return, in TIME\_\$KEY\_T format. This is a 2-byte integer.

Currently the only allowable value is TIME\_\$CLOCKH\_KEY.

**OUTPUT PARAMETERS****eventcount-pointer**

The eventcount address to be obtained, in EC2\_\$PTR\_T format. This is a 4-byte integer.

EC2\_\$PTR\_T is a pointer to an EC2\_\$EVENTCOUNT\_T array. See the EC2 Data Types section for more information.

**status**

Completion status, in STATUS\_\$T format. This data type is 4 bytes long. See the TIME\_\$ Data Types section for more information.

**USAGE**

EC2\_\$PTR\_T is a pointer to an EC2\_\$EVENTCOUNT\_T array. See the EC2 Data Types section for a description of eventcount data structures.

TIME\_\$GET\_EC outputs an eventcount that gets advanced about every 1/4 second. Thus, it can be passed to EC2\_\$WAIT to wait for a specific interval of time to elapse.

The interval between successive advances of the eventcount is nominally 262,144 microseconds (about 0.25 second). The exact interval changes slightly with system load.

For a ten-second wait, you might use:

```
TIME_$GET_EC (...gets time_eventcount ...)
EC2_$READ   (...gets current_eventcount_value ...)
EC2_$WAIT   (...current_eventcount + 40, time_eventcount...)
```

See Eventcounts Chapter of the *Programming With General System Calls* manual for more information.

TIME\_\$WAIT

TIME\_\$WAIT

Suspends the calling process for a specified time.

## FORMAT

TIME\_\$WAIT (rel-abs, clock, status)

## INPUT PARAMETERS

### rel-abs

Type of clock value supplied, in TIME\_\$REL\_ABS\_T format. This is a 2-byte integer. Specify only one of the following predefined values:

TIME\_\$RELATIVE

Clock specifies the number of 4-microsecond periods to wait before resuming process execution.

TIME\_\$ABSOLUTE

Clock contains the UTC system time for which to wait before resuming process execution.

### clock

The relative or absolute time for which to wait before resuming process execution, in TIME\_\$CLOCK\_T format. This data type is 6 bytes long. See the CAL Data Types section for more information.

Note that if you specify TIME\_\$ABSOLUTE in the rel\_abs parameter, TIME\_\$WAIT expects a UTC time. (You can remove a local time offset with the CAL\_\$REMOVE\_LOCAL\_OFFSET call.)

## OUTPUT PARAMETERS

### status

Completion status, in STATUS\_\$T format. This data type is 4 bytes long. See the CAL Data Types section for more information.

## USAGE

TIME\_\$WAIT suspends the calling process until a relative time elapses or an absolute time occurs.

A nonzero status is returned if the operating system has insufficient internal table space to process the request.

**ERRORS**

STATUS\_\$OK

Successful completion.

TIME\_\$BAD\_INT

Bad timer interrupt.

TIME\_\$BAD\_KEY

Bad key to TIME\_\$GET\_EC.

TIME\_\$NO\_Q\_ENTRY

Error from TIME\_\$ADVANCE.

TIME\_\$NOT\_FOUND

Entry to be canceled not found.

TIME\_\$WAIT\_QUIT

Wait interrupted by quit fault.

C

C

C

C

C



# TONE

## Table of Contents

TONE\_\$TIME

TONE-2



## TONE

This section describes the call syntax for the TONE programming call. The TONE call does not use special data types or produce unique error messages. Refer to the Introduction at the beginning of this manual for a description of call syntax format.

TONE\_\$TIME

TONE\_\$TIME

Makes a tone. The tone remains on for the time indicated in the call.

**FORMAT**

TONE\_\$TIME (time)

**INPUT PARAMETERS**

**time**

Length of the tone, in TIME\_\$CLOCK\_T format. This data type is a 48-bit integer value.

**USAGE**

Only DOMAIN nodes shipped after April 19, 1982 contain a working speaker.

# TPAD

## Table of Contents

### Data Types

TPAD\_\$INQ\_DTYPE  
TPAD\_\$INQUIRE  
TPAD\_\$RE\_RANGE  
TPAD\_\$SET\_CURSOR  
TPAD\_\$SET\_MODE

TPAD-2  
TPAD-4  
TPAD-5  
TPAD-6  
TPAD-7  
TPAD-8



## TPAD

The TPAD (Touchpad) programming calls maintain the touchpad and mouse. This section describes their data types and call syntax. The TPAD calls do not produce unique error messages. Refer to the Introduction at the beginning of this manual for a description of data-type diagrams and call syntax format.

## TPAD DATA TYPES

### DATA TYPES

#### TPAD\_\$MODE\_T

A 2-byte integer. Cursor mode operations for the touchpad and mouse. They establish how movements of the finger affect the cursor position on the display. Note the only meaningful mode for a mouse is TPAD\_\$RELATIVE. One of the following predefined values:

##### TPAD\_\$ABSOLUTE

Touchpad corresponds directly to the display. When a finger is placed on the touchpad, the cursor jumps to the corresponding position on the screen.

##### TPAD\_\$RELATIVE

Cursor movement is relative to the current position. It moves only when a finger moves across the pad, it does not move when a finger is placed on touchpad.


##### TPAD\_\$REL\_ABS

Cursor moves when finger is placed on the touchpad and when a finger moves across the pad. It does not move if a finger is lifted and replaced quickly.



STATUS\_\$\$T

A status code. The diagram below illustrates the STATUS\_\$\$T data type:



Field Description:

all

All 32 bits in the status code.

fail

The fail bit. If this bit is set, the error was not within the scope of the module invoked, but occurred within a lower-level module (bit 31).

subsys

The subsystem that encountered the error (bits 24 - 30).

modc

The module that encountered the error (bits 16 - 23).

code

A signed number that identifies the type of error that occurred (bits 0 - 15).

TPAD\_\$INQ\_DTYPE

TPAD\_\$INQ\_DTYPE

Returns the name of the last type of locating device used.

#### FORMAT

device = TPAD\_\$INQ\_DTYPE

#### INPUT PARAMETERS

None.

#### OUTPUT PARAMETERS

None.

#### RETURN VALUE

##### device

Value indicating the last locating device used to provide input, in TPAD\_\$DEV\_TYPE\_T format. This data type is two bytes long. Specify one of the following predefined values:

TPAD\_\$UNKNOWN

TPAD\_\$HAVE\_TOUCHPAD

TPAD\_\$HAVE\_MOUSE

TPAD\_\$HAVE\_BITPAD

#### USAGE

If no locator input has been detected since the node was last booted, this call returns TPAD\_\$UNKNOWN.

**TPAD\_\$INQUIRE**

Returns information about the current touchpad mode.

**FORMAT**

TPAD\_\$INQUIRE (mode, x-scale, y-scale, hysteresis, origin)

**INPUT PARAMETERS**

None.

**OUTPUT PARAMETERS****mode**

Cursor mode, in TPAD\_\$MODE\_T format. This is a 2-byte integer. Specify one of the following predefined values:

TPAD\_\$ABSOLUTE

TPAD\_\$RELATIVE

TPAD\_\$REL\_ABS

**x-scale**

Scale factor in the x dimension. This is a 2-byte integer.

**y-scale**

Scale factor in the y dimension. This is a 2-byte integer.

**hysteresis**

Hysteresis factor, in pixels. This is a 2-byte integer.

This hysteresis factor prevents the touchpad manager from responding to any minor movements you make unintentionally. This value defines a "box" around your finger. The touchpad manager does not move the cursor if your finger stays within this box.

If your finger moves beyond the box, the touchpad manager subtracts the hysteresis value from the distance moved, and moves the cursor the remaining distance. The default factor is 5.

**origin**

The point of origin for x and y in SMD\_\$POS\_T format. This data type is 4 bytes long.

**USAGE**

Use this call to save the touchpad mode for later restoration, or to change one aspect of the mode without changing any other aspects. For example, you can use the output from this call as the input to the TPAD\_\$SET\_MODE call.

TPAD\_\$RE\_RANGE

TPAD\_\$RE\_RANGE

Re-establishes the touchpad raw data range.

**FORMAT**

TPAD\_\$RE\_RANGE ( )

**INPUT PARAMETERS**

None.

**OUTPUT PARAMETERS**

None.

**USAGE**

This call re-establishes the touchpad raw data range over the next 1000 data points. This is also done for you at system boot. See the section on Touchpad Modes in *Programming with General System Calls* for a description of the touch pad raw data range.

## TPAD\_\$SET\_CURSOR

Re-establishes the touchpad origin in relative mode. The call to TPAD\_\$SET\_CURSOR can occur at any time and affects subsequent touchpad inputs.

## FORMAT

TPAD\_\$SET\_CURSOR (origin)

## INPUT PARAMETERS

**origin**

A screen position that will be the origin for subsequent data points from the touchpad in relative mode or in absolute/relative mode, in SMD\_\$POS\_T format. This data type is 4 bytes long.

## OUTPUT PARAMETERS

None.

## USAGE

The system *remembers* the last cursor position delivered by a locator device. When a new data point comes from the mouse, or from the touchpad or bitpad in relative mode, a **displacement** is computed and applied to the last locator position. The TPAD\_\$SET\_CURSOR call makes the system *forget* the last locator position, and use the value passed in the call instead. The next locator data will then start from this new position instead of its former position.

You will rarely need to make this call, as GPR and the display manager make the call at appropriate times.

The origin is automatically re-established when you take your finger from the touchpad for more than one-eighth of a second. One effect of this is that the cursor typically doesn't move the next time you touch the pad in relative mode, unless you explicitly call TPAD\_\$SET\_CURSOR before that next touch.

This call has meaning for relative and absolute/relative mode only. In absolute/relative mode, when you first touch the pad, the pad inputs coordinates in the absolute mode. To have effect, the call to TPAD\_\$SET\_CURSOR must occur after this first touch, but during the relative part of this use of the touchpad (that is, before you lift your finger for more than one-half second.)

TPAD\_\$SET\_MODE

TPAD\_\$SET\_MODE

Sets the touch pad mode.

## FORMAT

TPAD\_\$SET\_MODE (mode, x-scale, y-scale, hysteresis, origin)

## INPUT PARAMETERS

### mode

Cursor mode, in TPAD\_\$MODE\_T format. This is a 2-byte integer. Specify one of the following predefined values:

TPAD\_\$ABSOLUTE

TPAD\_\$RELATIVE

TPAD\_\$REL\_ABS

### x-scale

Scale factor in the x dimension. This is a 2-byte integer.

### y-scale

Scale factor in the y dimension. This is a 2-byte integer.

### hysteresis

Hysteresis factor, in pixels. This is a 2-byte integer.

The hysteresis factor prevents the touchpad manager from responding to any minor movements you make unintentionally. This value defines a "box" around your finger. The touchpad manager does not move the cursor if your finger stays within this box.

If your finger moves beyond the box, the touchpad manager subtracts the hysteresis value from the distance moved, and moves the cursor the remaining distance. The default factor is 5.

### origin

The point of origin for x and y in SMD\_\$POS\_T format. This data type is 4 bytes long.

## OUTPUT PARAMETERS

None.

## USAGE

Use this call to set to mode, scale factors, and hysteresis factors of locator devices. You can also change the origin for relative or absolute/relative mode. This call applies to the touchpad, mouse, and bit pad locator devices. Note that the mouse uses only the scale and hysteresis factors and ignores the other mode settings, since it is an inherently relative device.

# VEC

## Table of Contents

VEC_\$ADD_CONSTANT	VEC-2
VEC_\$ADD_CONSTANT_I	VEC-3
VEC_\$ADD_VECTOR	VEC-4
VEC_\$ADD_VECTOR_I	VEC-5
VEC_\$COPY	VEC-7
VEC_\$COPY_I	VEC-8
VEC_\$DOT	VEC-9
VEC_\$DOT_I	VEC-10
VEC_\$DP_SP	VEC-11
VEC_\$DP_SP_I	VEC-12
VEC_\$INIT	VEC-13
VEC_\$MAT_MULT	VEC-14
VEC_\$MAT_MULTN	VEC-15
VEC_\$MAX	VEC-16
VEC_\$MAX_I	VEC-17
VEC_\$MULT_ADD	VEC-18
VEC_\$MULT_ADD_I	VEC-19
VEC_\$MULT_CONSTANT	VEC-21
VEC_\$MULT_CONSTANT_I	VEC-22
VEC_\$POSTMULT	VEC-23
VEC_\$POSTMULTN	VEC-24
VEC_\$PREMULT	VEC-25
VEC_\$PREMULTN	VEC-26
VEC_\$SP_DP	VEC-27
VEC_\$SP_DP_I	VEC-28
VEC_\$SUB	VEC-29
VEC_\$SUB_I	VEC-30
VEC_\$SUM	VEC-31
VEC_SUM_I	VEC-32
VEC_\$SWAP	VEC-33
VEC_\$SWAP_I	VEC-34
VEC_\$ZERO	VEC-35
VEC_\$ZERO_I	VEC-36





## VEC

The VEC (Vector Library) programming calls perform floating point and integer vector matrix arithmetic. This section describes their call syntax. The VEC calls do not use special data types or produce unique error messages. Refer to the Introduction at the beginning of this manual for a description of call syntax format.

The majority of the calls in this section have four versions: single-precision floating-point, double-precision floating-point, 16-bit integer (INTEGER\*2), and 32-bit integer (INTEGER\*4). The names of all single-precision vector routines are in the form VEC\_\$Iname. Double-precision routines are named VEC\_\$Dname. 16-bit integer routines are named VEC\_\$Iname16. 32-bit integer routines are named VEC\_\$Iname. For example, VEC\_\$DOT and VEC\_\$DDOT are single- and double-precision versions of DOT product routines. VEC\_\$IDOT and VEC\_\$IDOT16 are the 32-bit and 16-bit versions, respectively.

Each type of routine generally requires parameters of the same type. For the double-precision routines, all floating-point parameters are double-precision; for the single-precision routines, all floating-point parameters must be single precision; for the integer procedures and functions, the parameters and returned values are integers, etc.

Routine names that end in I denote "incremental" routines, which step through vectors at increments other than 1.

When calling any of the vector routines, make sure that the indices you pass are valid. For best performance, these routines do **not** check index values for validity; hence, passing a value of zero can cause a variety of errors.

**NOTE:** All matrices are assumed to be stored in FORTRAN (column-major) order. Because Pascal and C store matrix elements in row-major order, you may need to transpose or otherwise rearrange elements when calling vector routines from C or Pascal programs.

VEC\_\$ADD\_CONSTANT

VEC\_\$ADD\_CONSTANT

Adds a constant to a vector.

### FORMAT

VEC\_\$ADD\_CONSTANT (start\_vec, length, constant, result\_vec)  
VEC\_\$DADD\_CONSTANT (start\_vec, length, constant, result\_vec)  
VEC\_\$IADD\_CONSTANT (start\_vec, length, constant, result\_vec)  
VEC\_\$IADD\_CONSTANT16 (start\_vec, length, constant, result\_vec)

### INPUT PARAMETERS

**start\_vec**

Floating-point or integer vector to which the value will be added.

**length**

Number of elements to add. This is a 4-byte integer.

**constant**

Value to be added to each element of start\_vec.

### OUTPUT PARAMETERS

**result\_vec**

Floating-point or integer vector containing the sum.

### USAGE

These routines add a constant to a vector, returning the result in a second vector. The routines perform the following operation:

```
DO 10 I + 1,LENGTH
    RESULT_VEC(I) = CONSTANT + START_VEC(I)
10 CONTINUE
```

## VEC\_\$ADD\_CONSTANT\_I

Adds a constant to a vector, stepping through the vector by increments.

**FORMAT**

```
VEC_$ADD_CONSTANT_I (start_vec, inc1, length, constant, result_vec, inc2)
VEC_$DADD_CONSTANT_I (start_vec, inc1, length, constant, result_vec, inc2)
VEC_$IADD_CONSTANT_I (start_vec, inc1, length, constant, result_vec, inc2)
VEC_$IADD_CONSTANT16_I (start_vec, inc1, length, constant, result_vec, inc2)
```

**INPUT PARAMETERS****start\_vec**

Floating-point or integer vector to which the value will be added.

**inc1**

Increment for the index of start\_vec. This is a 4-byte integer.

**length**

Number of elements to add. This is a 4-byte integer.

**constant**

Value to be added to each element of start\_vec.

**OUTPUT PARAMETERS****result\_vec**

Floating-point or integer vector containing the sum.

**INPUT PARAMETERS****inc2**

Increment for the index of result\_vec. This is a 4-byte integer.

**USAGE**

These routines add a constant to a vector, stepping through their elements by user-specified increments. The routines perform the following operation:

```
J=1
K=1
DO 10 I = 1,LENGTH
  RESULT_VEC(K) = CONSTANT+START_VEC(J)
  K = K+INC2
  J = J+INC1
10 CONTINUE
```

VEC\_\$ADD\_VECTOR

VEC\_\$ADD\_VECTOR

Adds two vector.

### FORMAT

VEC \$ADD\_VECTOR (start\_vec, add\_vec, length, result\_vec)  
VEC \$DADD\_VECTOR (start\_vec, add\_vec, length, result\_vec)  
VEC \$IADD\_VECTOR (start\_vec, add\_vec, length, result\_vec)  
VEC \$IADD\_VECTOR16 (start\_vec, add\_vec, length, result\_vec)

### INPUT PARAMETERS

**start\_vec**

Floating-point or integer vector to which the value will be added.

**add\_vec**

Floating-point or integer vector to be added.

**length**

Number of elements to add. This is a 4-byte integer.

### OUTPUT PARAMETERS

**result\_vec**

Floating-point or integer vector containing the sum.

### USAGE

These routines add two vectors, returning the sum in a third vector. The routines perform the following operation:

```
DO 10 I = 1,LENGTH
  RESULT_VEC(I) = START_VEC(I) + ADD_VEC(I)
10 CONTINUE
```

**VEC\_\$ADD\_VECTOR\_I**

Adds two vectors, stepping through them by increments.

**FORMAT**

```
VEC $ADD_VECTOR_I(start_vec, inc1, add_vec, inc2, length, result_vec, inc3)
VEC $DADD_VECTOR_I(start_vec, inc1, add_vec, inc2, length, result_vec,
                    inc3)
VEC $IADD_VECTOR_I(start_vec, inc1, add_vec, inc2, length, result_vec,
                    inc3)
VEC $IADD_VECTOR16_I(start_vec, inc1, add_vec, inc2, length, result_vec,
                     inc3)
```

**INPUT PARAMETERS****start\_vec**

Floating-point or integer vector to which the value will be added.

**inc1**

Increment for the index of start\_vec. This is a 4-byte integer.

**add\_vec**

Floating-point or integer vector to be added.

**inc2**

Increment for the index of add\_vec. This is a 4-byte integer.

**length**

Length of the vector. This is a 4-byte integer.

**OUTPUT PARAMETERS****result\_vec**

Floating-point or integer vector containing the sum.

**INPUT PARAMETERS****inc3**

Increment for the index of result\_vec. This is a 4-byte integer.

VEC\_\$ADD\_VECTOR\_I

## USAGE

These routines add two vectors, stepping through their elements by user-specified increments. The routines perform the following operation:

```
J=1
K=1
DO 10 I = 1,LENGTH
  VEC3(J) = VEC1(K) + VEC2(L)
  J = J + INC3
  K = K + INC1
  L = L + INC2
10 CONTINUE
```

**VEC\_\$COPY**

Copies elements from one vector to another.

**FORMAT**

```
VEC_$COPY (start_vec, result_vec, length)
VEC_$DCOPY (start_vec, result_vec, length)
VEC_$ICOPY (start_vec, result_vec, length)
VEC_$ICOPY16 (start_vec, result_vec, length)
```

**INPUT PARAMETERS****start\_vec**

Floating-point or integer vector from which elements will be copied.

**length**

Number of elements to copy. This is a 4-byte integer.

**OUTPUT PARAMETERS****result\_vec**

Floating-point or integer vector to which elements will be copied.

**USAGE**

These routines copy elements from one vector to another. The routines perform the following operation:

```
DO 10 I = 1,LENGTH
  RESULT_VEC(I) = START_VEC(I)
10 CONTINUE
```

VEC\_\$COPY\_I

VEC\_\$COPY\_I

Copies elements from one vector to another, stepping through the vectors by increments.

### FORMAT

VEC\_\$COPY\_I (vec1, inc1, vec2, inc2, length)  
VEC\_\$DCOPY\_I (vec1, inc1, vec2, inc2, length)  
VEC\_\$ICOPY\_I (vec1, inc1, vec2, inc2, length)  
VEC\_\$ICOPY16\_I (vec1, inc1, vec2, inc2, length)

### INPUT PARAMETERS

**vec1**

Floating-point or integer vector from which elements will be copied.

**inc1**

Increment for the index of vec1. This is a 4-byte integer.

**inc2**

Increment for the index of vec2. This is a 4-byte integer.

**length**

Number of elements to copy. This is a 4-byte integer.

### OUTPUT PARAMETERS

**vec2**

Floating-point or integer vector to which elements will be copied.

### USAGE

These routines copy a vector but use an increment to step through the vector. VEC\_\$COPY\_I moves 32 bits regardless of data type and VEC\_\$DCOPY\_I moves 64 bits. The routines perform the following operation:

```
      J = 1  
      K = 1  
      DO 10 I = 1,LENGTH  
        VEC2(J) = VEC1(K)  
        J = J+INC2  
        K = K+INC1  
10 CONTINUE
```



**VEC\_\$DOT**

Calculates the dot product of two vectors.

**FORMAT**

```
result = VEC_$DOT (vec1, vec2, length)
result = VEC_$DDOT (vec1, vec2, length)
result = VEC_$IDOT (vec1, vec2, length)
result = VEC_$IDOT16 (vec1, vec2, length)
```

**RETURN VALUE****result**

Floating-point or integer dot product.

**INPUT PARAMETERS****vec1, vec2**

Floating-point or integer vectors.

**length**

Number of elements to add. This is a 4-byte integer.

**USAGE**

These routines calculate the dot product of two vectors. The routines perform the following operation:

```
DO 10 I = 1,LENGTH
  TEMP = TEMP + VEC1(I) * VEC2(I)
10 CONTINUE
VEC_$DOT = TEMP
```

VEC\_\$DOT\_I

VEC\_\$DOT\_I

Calculates the dot product of two vectors, stepping through the vectors by increments.

### FORMAT

```
result = VEC_$DOT_I (vec1, inc1, vec2, inc2, length)
result = VEC_$DDOT_I (vec1, inc1, vec2, inc2, length)
result = VEC_$IDOT_I (vec1, inc1, vec2, inc2, length)
result = VEC_$IDOT16_I (vec1, inc1, vec2, inc2, length)
```

### RETURN VALUE

#### result

Floating-point or integer dot product.

### INPUT PARAMETERS

#### vec1

Floating-point or integer vector.

#### inc1

Increment for the index of vec1. This is a 4-byte integer.

#### vec2

Floating-point or integer vector.

#### inc2

Increment for the index of vec1. This is a 4-byte integer.

#### length

Number of elements to add. This is a 4-byte integer.

### USAGE

These routines calculate the dot product of two vectors, stepping through the vectors with a user-supplied increment. The routines perform the following operation:

```
      J = 1
      K = 1
      DO 10 I = 1,LENGTH
        TEMP = TEMP + VEC1(J) * VEC2(K)
        J = J + INC1
        K = K + INC2
10 CONTINUE
      VEC_$DOT = TEMP
```

**VEC\_\$DP\_SP**

Copies a double-precision vector to a single precision vector.

**FORMAT**

VEC\_\$DP\_SP (dp\_vec, sp\_vec, length)

**INPUT PARAMETERS****dp\_vec**

Floating-point double-precision vector.

**OUTPUT PARAMETERS****sp\_vec**

Floating-point single-precision resultant vector.

**INPUT PARAMETERS****length**

Number of elements to copy. This is a 4-byte integer.

**USAGE**

VEC\_\$DP\_SP copies a double-precision vector to a single-precision resultant vector. The routine performs the following operation:

```
DO 10 I=1, LENGTH
  SP_VEC(I) = SNGL(DP_VEC(I))
10 CONTINUE
```

VEC\_\$DP\_SP\_I

VEC\_\$DP\_SP\_I

Copies a double-precision vector to a single-precision vector, stepping through the vectors incrementally.

### FORMAT

VEC\_\$DP\_SP\_I (dp\_vec, inc1, sp\_vec, inc2, length)

### INPUT PARAMETERS

**dp\_vec**

Double-precision floating-point vector.

**inc1**

Increment for the index of dp\_vec. This is a 4-byte integer.

### OUTPUT PARAMETERS

**sp\_vec**

Single-precision floating-point resultant vector.

### INPUT PARAMETERS

**inc2**

Increment for the index of sp\_vec. This is a 4-byte integer.

**length**

Number of elements to copy. This is a 4-byte integer.

### USAGE

VEC\_\$DP\_SP\_I copies a double precision vector to a single precision resultant vector, stepping through the vectors by user-supplied increments. The routine performs the following operation:

```
J=1
K=1
DO 10 I=1, LENGTH
  SP_VEC(K) = SNGL (DP_VEC(J))
  J = J+INC1
  K = K+INC2
10 CONTINUE
```

**VEC\_\$INIT**

Initializes a vector with a constant.

**FORMAT**

VEC\_\$INIT (vector, length, constant)  
VEC\_\$DINIT (vector, length, constant)  
VEC\_\$IINIT (vector, length, constant)  
VEC\_\$IINIT16 (vector, length, constant)

**INPUT/OUTPUT PARAMETERS****vector**

Upon input Floating-point or integer vector to be initialized.

Upon output Floating-point or integer vector which has been initialized to a constant.

**INPUT PARAMETERS****length**

Number of elements to initialize. This is a 4-byte integer.

**constant**

Floating-point or integer constant to be assigned to elements of vector.

**USAGE**

These routines perform the following operation:

```
DO 10 I = 1, LENGTH  
    VECTOR (I) = CONSTANT  
10 CONTINUE
```

VEC\_\$MAT\_MULT

VEC\_\$MAT\_MULT

Multiplies two 4 x 4 matrices, returning the result in a 4 x 4 matrix.

#### FORMAT

VEC\_\$MAT\_MULT (matrix1, matrix2, out\_matrix)  
VEC\_\$DMAT\_MULT (matrix1, matrix2, out\_matrix)  
VEC\_\$IMAT\_MULT (matrix1, matrix2, out\_matrix)  
VEC\_\$IMAT\_MULT16 (matrix1, matrix2, out\_matrix)

#### INPUT PARAMETERS

##### matrix1

A 4 x 4 floating-point or integer matrix to be multiplied.

##### matrix2

A 4 x 4 floating-point or integer matrix to be multiplied.

#### OUTPUT PARAMETERS

##### out\_matrix

Floating-point or integer matrix containing the product.

#### USAGE

These routines multiply two 4 x 4 matrices, returning the result as a third 4 x 4 matrix. They are intended for use in graphics applications, and perform the following operation:

```
DO 10 J = 1,4  
DO 10 I = 1,4  
CALL VEC_$POST_MULT (MATRIX1, MATRIX2(J, I), OUT_MATRIX(J, I))  
10 CONTINUE
```

## VEC\_\$MAT\_MULTN

Multiplies two variably-dimensioned matrices, returning the result in a third matrix.

**FORMAT**

VEC\_\$MAT\_MULTN (matrix1, matrix2, m, n, s, out\_matrix)  
VEC\_\$DMAT\_MULTN (matrix1, matrix2, m, n, s, out\_matrix)  
VEC\_\$IMAT\_MULTN (matrix1, matrix2, m, n, s, out\_matrix)  
VEC\_\$IMAT\_MULTN16 (matrix1, matrix2, m, n, s, out\_matrix)

**INPUT PARAMETERS****matrix1**

First floating-point or integer matrix of dimensions m x n to be multiplied.

**matrix2**

Second floating-point or integer matrix of dimensions n x s to be multiplied.

**m, n, s**

The various matrix dimensions. These are 4-byte integers.

**OUTPUT PARAMETERS****out\_matrix**

Floating-point or integer matrix of dimensions m x s containing the product.

**USAGE**

These routines multiply two matrices with dimensions specified by you, returning the result as a third matrix. Note that the matrices are assumed to be stored in FORTRAN (column-major) order. These routines perform the following operation:

```
DO 10 I = 1,M
DO 10 J = 1,S
  OUT_MATRIX(I,J) = 0.0
  DO 10 K = 1,N
    OUT_MATRIX(I,J) = OUT_MATRIX(I,J) + MATRIX1(I,K) * MATRIX2(K,J)
  10 CONTINUE
```

VEC\_\$MAX

VEC\_\$MAX

Finds the element with the greatest maximum absolute value in a vector and returns its value and location.

#### FORMAT

VEC\_\$MAX (start\_vec, length, result, result\_loc)  
VEC\_\$DMAX (start\_vec, length, result, result\_loc)  
VEC\_\$IMAX (start\_vec, length, result, result\_loc)  
VEC\_\$IMAX16 (start\_vec, length, result, result\_loc)

#### INPUT PARAMETERS

**start\_vec**

Floating-point or integer vector to be searched.

**length**

Number of elements to examine. This is a 4-byte integer.

#### OUTPUT PARAMETERS

**result**

Floating-point or integer maximum absolute value of all the elements searched.

**result\_loc**

Location of value within the vector. This is a 4-byte integer.

#### USAGE

These routines search through a vector and return the maximum absolute value found and its location within the vector. The routines perform the following:

```
RESULT = ABS (START_VEC(1))
RESULT_LOC = 1
DO 10 I=2, LENGTH
  IF (ABS(START_VEC(I)) .GT. RESULT) THEN
    RESULT_LOC = I
    RESULT = ABS(START_VEC(I))
  ENDIF
10 CONTINUE
```



## VEC\_\$MAX\_I

Searches a vector for the maximum absolute value, stepping through the vector by increments greater than 1.

**FORMAT**

```
VEC_$MAX_I (start_vec, inc, length, result, result_loc)
VEC_$DMAX_I (start_vec, inc, length, result, result_loc)
VEC_$IMAX_I (start_vec, inc, length, result, result_loc)
VEC_$IMAX16_I (start_vec, inc, length, result, result_loc)
```

**INPUT PARAMETERS****start\_vec**

Floating-point or integer vector to be searched.

**inc**

Increment used to step through the start\_vec. This is a 4-byte integer.

**length**

Number of elements to examine. This is a 4-byte integer.

**OUTPUT PARAMETERS****result**

Floating-point or integer maximum absolute value of all the elements searched.

**result\_loc**

Location of value within given vector. This is a 4-byte integer.

**USAGE**

These routines search through a vector by a positive increment and return the greatest absolute value found and its location within the vector. The routines perform the following:

```
RESULT = ABS(START_VEC(1))
RESULT_LOC = 1
J = 1+INC1
DO 10 I=2,LENGTH
  IF (ABS(START_VEC(J)) .G.T.RESULT) THEN
    RESULT_LOC = I
    RESULT = ABS(START_VEC(J))
  ENDIF
  J = J+INC1
10 CONTINUE
```

VEC\_\$MULT\_ADD

VEC\_\$MULT\_ADD

Multiplies a vector by a constant and adds the result to a second vector.

#### FORMAT

VEC\_\$MULT\_ADD (add\_vec, mult\_vec, length, constant, result\_vec)  
VEC\_\$DMULT\_ADD (add\_vec, mult\_vec, length, constant, result\_vec)  
VEC\_\$IMULT\_ADD (add\_vec, mult\_vec, length, constant, result\_vec)  
VEC\_\$IMULT\_ADD16 (add\_vec, mult\_vec, length, constant, result\_vec)

#### INPUT PARAMETERS

**add\_vec**

Floating-point or integer vector to be added.

**mult\_vec**

Floating-point or integer vector to be multiplied by the constant.

**length**

Number of elements to add. This is a 4-byte integer.

**constant**

Floating-point or integer constant to be multiplied by mult\_vec.

#### OUTPUT PARAMETERS

**result\_vec**

Floating-point or integer vector containing the sum.

#### USAGE

These routines multiply one vector (input as mult\_vec) by a constant, and add the result to a second vector (input as add\_vec). The result is returned in a third vector. The routines perform the following operation:

```
      DO 10 I = 1, LENGTH  
        RESULT_VEC(I) = ADD_VEC(I) + CONSTANT*MULT_VEC(I)  
10 CONTINUE
```

## VEC\_\$MULT\_ADD\_I

Multiplies a vector by a constant and adds the result to a second vector, stepping through both vectors and the result by increments.

## FORMAT

```
VEC_$MULT_ADD_I (add_vec, inc1, mult_vec, inc2,
                length, constant, result_vec, inc3)
VEC_$DMULT_ADD_I (add_vec, inc1, mult_vec, inc2,
                 length, constant, result_vec, inc3)
VEC_$IMULT_ADD_I (add_vec, inc1, mult_vec, inc2,
                 length, constant, result_vec, inc3)
VEC_$IMULT_ADD16_I (add_vec, inc1, mult_vec, inc2,
                   length, constant, result_vec, inc3)
```

## INPUT PARAMETERS

**add\_vec**

Floating-point or integer vector to be added.

**inc1**

Increment for the index of **add\_vec**. This is a 4-byte integer.

**mult\_vec**

Floating-point or integer vector to be multiplied by the constant.

**inc2**

Increment for the index of **mult\_vec**. This is a 4-byte integer.

**length**

Number of elements on which to operate. This is a 4-byte integer.

**constant**

Floating-point or integer constant to be multiplied by **mult\_vec**.

## OUTPUT PARAMETERS

**result\_vec**

Floating-point or integer vector containing the sum.

## INPUT PARAMETERS

**inc3**

Increment for the index of **result\_vec**. This is a 4-byte integer.

VEC\_\$MULT\_ADD\_I

## USAGE

These routines multiply one vector by a constant and add the result to a second vector. The result is returned in a third vector. The indices to all three vectors are incremented by user-specified values. The routines perform the following operation:

```
J = 1
K = 1
L = 1
DO 10 I = 1,LENGTH
  RESULT(J) = ADD_VEC(K) + CONSTANT*MULT_VEC(L)
  J = J + INC3
  K = K + INC1
  L = L + INC2
10 CONTINUE
```

## VEC\_\$MULT\_CONSTANT

Multiplies a vector by a scalar constant and returns the result in a second vector.

**FORMAT**

VEC\_\$MULT\_CONSTANT (mult\_vec, length, constant, result\_vec)  
VEC\_\$DMULT\_CONSTANT (mult\_vec, length, constant, result\_vec)  
VEC\_\$IMULT\_CONSTANT (mult\_vec, length, constant, result\_vec)  
VEC\_\$IMULT\_CONSTANT16 (mult\_vec, length, constant, result\_vec)

**INPUT PARAMETERS****mult\_vec**

Floating-point or integer vector to be multiplied.

**length**

Number of elements to multiply. This is a 4-byte integer.

**constant**

Floating-point or integer constant to multiply by mult\_vec.

**OUTPUT PARAMETERS****result\_vec**

Floating-point or integer vector containing the product.

**USAGE**

These routines multiply one vector by a scalar constant, returning the result in a second vector. The routines perform the following operation:

```
      DO 10 I = 1,LENGTH  
        RESULT_VEC(I) = CONSTANT * MULT_VEC(I)  
10 CONTINUE
```

VEC\_\$MULT\_CONSTANT\_I

VEC\_\$MULT\_CONSTANT\_I

Multiplies a vector by a scalar constant, returns the result in a second vector, and steps through the vectors by increments.

### FORMAT

VEC\_\$MULT\_CONSTANT\_I (mult\_vec, inc1, length, constant, result\_vec, inc2)  
VEC\_\$DMULT\_CONSTANT\_I (mult\_vec, inc1, length, constant, result\_vec, inc2)  
VEC\_\$IMULT\_CONSTANT\_I (mult\_vec, inc1, length, constant, result\_vec, inc2)  
VEC\_\$IMULT\_CONSTANT16\_I(mult\_vec, inc1, length, constant, result\_vec, inc2)

### INPUT PARAMETERS

**mult\_vec**

Floating-point or integer vector from which data will be copied.

**inc1**

Increment for the index of mult\_vec. This is a 4-byte integer.

**length**

Number of elements on which to operate. This is a 4-byte integer.

**constant**

Floating-point or integer constant to multiply by mult\_vec.

### OUTPUT PARAMETERS

**result\_vec**

Floating-point or integer vector containing the result.

### INPUT PARAMETERS

**inc2**

Increment for the index of result\_vec. This is a 4-byte integer.

### USAGE

These routines multiply elements of a vector by a scalar constant and store the result in a second vector. The routines step through both vectors by increments. The routines perform the following:

```
J = 1
K = 1
DO 10 I=1,LENGTH
    RESULT(J) = MULT_VEC(K) * CONSTANT
    J = J+INC2
    K = K+INC1
10 CONTINUE
```

**VEC\_\$POSTMULT**

Postmultiplies a 4 x 4 matrix by a 4 x 1 column vector, returning the result in a second vector.

**FORMAT**

VEC\_\$POSTMULT (matrix, col\_vec, result\_vec)  
VEC\_\$DPOSTMULT (matrix, col\_vec, result\_vec)  
VEC\_\$IPOSTMULT (matrix, col\_vec, result\_vec)  
VEC\_\$IPOSTMULT16 (matrix, col\_vec, result\_vec)

**INPUT PARAMETERS****matrix**

A 4 x 4 floating-point or integer matrix to be postmultiplied.

**col\_vec**

A 4 x 1 floating-point or integer column vector.

**OUTPUT PARAMETERS****result\_vec**

Floating-point or integer vector containing the product.

**USAGE**

These routines postmultiply a 4 x 4 matrix by a 4 x 1 column vector, and return a 4 x 1 column vector. They are intended for use in graphics applications, and perform the following operation:

```
DO 10 J = 1,4
  RESULT_VEC(J) = 0.0
DO 10 I = 1,4
  RESULT_VEC(J) = RESULT_VEC(J) + COL_VEC(I)*MATRIX(J,I)
10 CONTINUE
```

VEC\_\$POSTMULTN

VEC\_\$POSTMULTN

Postmultiplies a variably-dimensioned matrix by an n x 1 vector, returning the result in a second vector.

#### FORMAT

VEC\_\$POSTMULTN (matrix, col\_vec, m, n, result\_vec)  
VEC\_\$DPOSTMULTN (matrix, col\_vec, m, n, result\_vec)  
VEC\_\$IPOSTMULTN (matrix, col\_vec, m, n, result\_vec)  
VEC\_\$IPOSTMULTN16 (matrix, col\_vec, m, n, result\_vec)

#### INPUT PARAMETERS

##### matrix

An m x n floating-point or integer matrix to be postmultiplied.

##### col\_vec

An n x 1 floating-point or integer column vector.

##### m, n

Dimensions of the matrices. These are 4-byte integers.

#### OUTPUT PARAMETERS

##### result\_vec

An m x 1 floating-point or integer vector containing the product.

#### USAGE

These routines postmultiply a m x n matrix by a n x 1 column vector, and return an m x 1 column vector. They perform the following operation:

```
DO 10 I = 1,M
  RESULT_VEC(I) = 0.0
DO 10 J = 1,N
  RESULT_VEC(I) = RESULT_VEC(I) + COL_VEC(J) * MATRIX(I,J)
10 CONTINUE
```



**VEC\_\$PREMULT**

Premultiplies a 4 x 4 matrix by a 1 x 4 row vector, returning the result in a second vector.

**FORMAT**

VEC\_\$PREMULT (row\_vec, matrix, result\_vec)  
VEC\_\$DPREMULT (row\_vec, matrix, result\_vec)  
VEC\_\$IPREMULT (row\_vec, matrix, result\_vec)  
VEC\_\$IPREMULT16 (row\_vec, matrix, result\_vec)

**INPUT PARAMETERS****row\_vec**

A 1 x 4 floating-point or integer row vector.

**matrix**

A 4 x 4 floating-point or integer matrix to be premultiplied.

**OUTPUT PARAMETERS****result\_vec**

Floating-point or integer vector containing the product.

**USAGE**

These routines premultiply a 4 x 4 matrix by a 1 x 4 row vector, and return a 1 x 4 row vector. They perform the following operation:

```
DO 10 I = 1,4
  RESULT_VEC(I) = 0.0
DO 10 J = 1,4
  RESULT_VEC(I) = RESULT_VEC(I) + ROW_VEC(J) * MATRIX(J,I)
10 CONTINUE
```

VEC\_\$PREMULTN

VEC\_\$PREMULTN

Premultiplies a variably-dimensioned matrix by a 1 x n row vector, returning the result in a second vector.

### FORMAT

VEC\_\$PREMULTN (row\_vec, matrix, m, n, result\_vec)  
VEC\_\$DPREMULTN (row\_vec, matrix, m, n, result\_vec)  
VEC\_\$IPREMULTN (row\_vec, matrix, m, n, result\_vec)  
VEC\_\$IPREMULTN16 (row\_vec, matrix, m, n, result\_vec)

### INPUT PARAMETERS

**row\_vec**

A 1 x m floating-point or integer row vector.

**matrix**

An m x n floating-point or integer matrix to be premultiplied.

**m, n**

Dimensions of the matrices. These are 4-byte integers.

### OUTPUT PARAMETERS

**result\_vec**

One by n floating-point or integer vector containing the product.

### USAGE

These routines premultiply a variably-dimensioned matrix by a 1 x m row vector, and return a 1 x n row vector. They perform the following operation:

```
DO 10 I = 1,N
  RESULT_VEC(I) = 0.0
  DO 10 J = 1,M
    RESULT_VEC(I) = RESULT_VEC(I) + ROW_VEC(J) * MATRIX(J,I)
  10 CONTINUE
```

**VEC\_\$\$SP\_DP**

Copies a single-precision vector to a double-precision vector.

**FORMAT**

VEC\_\$\$SP\_DP (sp\_vec, dp\_vec, length)

**INPUT PARAMETERS**

**sp\_vec**

Single-precision floating-point vector.

**OUTPUT PARAMETERS**

**dp\_vec**

Double-precision floating-point resultant vector.

**INPUT PARAMETERS**

**length**

Number of elements to copy. This is a 4-byte integer.

**USAGE**

VEC\_\$\$SP\_DP copies a single-precision vector to a double-precision resultant vector. The routine performs the following:

```
      DO 10 I=1, LENGTH
        DP_VEC(I) = DBLE (SP_VEC(I))
      10 CONTINUE
```

VEC\_\$SP\_DP\_I

VEC\_\$SP\_DP\_I

Copies a single-precision vector to a double-precision vector, stepping through the vectors incrementally.

### FORMAT

VEC\_\$SP\_DP\_I (sp\_vec, inc1, dp\_vec, inc2, length)

### INPUT PARAMETERS

**sp\_vec**

Floating-point, single-precision vector.

**inc1**

Increment for the index of sp\_vec. This is a 4-byte integer.

### OUTPUT PARAMETERS

**dp\_vec**

Floating-point double-precision resultant vector.

**inc2**

Increment for the index of dp\_vec. This is a 4-byte integer.

### INPUT PARAMETERS

**length**

Number of elements to copy. This is a 4-byte integer.

### USAGE

VEC\_\$SP\_DP\_I copies a single-precision vector to a double-precision resultant vector, stepping through the vectors by user-supplied increments. It does the following:

```
J = 1
K = 1
DO 10 I=1, LENGTH
    DP_VEC(K) = DBLE (SP_VEC(J))
    J = J+INC1
    K = K+INC2
10 CONTINUE
```

**VEC\_\$SUB**

Subtracts one vector from another.

**FORMAT**

VEC\_\$SUB (start\_vec, sub\_vec, length, result\_vec)  
VEC\_\$DSUB (start\_vec, sub\_vec, length, result\_vec)  
VEC\_\$ISUB (start\_vec, sub\_vec, length, result\_vec)  
VEC\_\$ISUB16 (start\_vec, sub\_vec, length, result\_vec)

**INPUT PARAMETERS****start\_vec**

Floating-point or integer vector from which the values will be subtracted.

**sub\_vec**

Floating-point or integer vector to be subtracted.

**length**

Number of elements to subtract. This is a 4-byte integer.

**OUTPUT PARAMETERS****result\_vec**

Floating-point or integer vector containing the difference.

**USAGE**

These routines subtract one vector from another, returning the result in a third vector. The routines perform the following operation:

```
DO 10 I = 1,LENGTH
  RESULT_VEC(I) = START_VEC(I) - SUB_VEC(I)
10 CONTINUE
```

VEC\_\$SUB\_I

VEC\_\$SUB\_I

Subtracts one vector from another, stepping through the vectors by increments.

#### FORMAT

VEC\_\$SUB\_I (start\_vec, inc1, sub\_vec, inc2, length, result\_vec, inc3)  
VEC\_\$DSUB\_I (start\_vec, inc1, sub\_vec, inc2, length, result\_vec, inc3)  
VEC\_\$ISUB\_I (start\_vec, inc1, sub\_vec, inc2, length, result\_vec, inc3)  
VEC\_\$ISUB16\_I (start\_vec, inc1, sub\_vec, inc2, length, result\_vec, inc3)

#### INPUT PARAMETERS

**start\_vec**

Floating-point or integer vector from which the values will be subtracted.

**inc1**

Increment for the index of start\_vec. This is a 4-byte integer.

**sub\_vec**

Floating-point or integer vector to be subtracted.

**inc2**

Increment for the index of sub\_vec. This is a 4-byte integer.

**length**

Number of elements to subtract. This is a 4-byte integer.

#### OUTPUT PARAMETERS

**result\_vec**

Floating-point or integer vector containing the difference.

#### INPUT PARAMETERS

**inc3**

Increment for the index of result\_vec. This is a 4-byte integer.

#### USAGE

These routines subtract one vector from another, returning the result in a third vector. The indices to all three vectors are incremented by user-specified values. The routines perform the following operation:

```
DO 10 I = 1,LENGTH
  RESULT_VEC(J) = START_VEC(K) - SUB_VEC(L)
  J = J + INC3
  K = K + INC1
  L = L + INC2
10 CONTINUE
```

**VEC\_\$\$SUM**

Sums the elements of a vector.

**FORMAT**

```
sum = VEC_$$SUM (vec, length)
sum = VEC_$$DSUM (vec, length)
sum = VEC_$$ISUM (vec, length)
sum = VEC_$$ISUM16 (vec, length)
```

**RETURN VALUE****sum**

Floating-point or integer sum of first "length" elements of vec.

**INPUT PARAMETERS****vec**

Floating-point or integer vector to be summed.

**length**

Number of elements to sum. This is a 4-byte integer.

**USAGE**

These routines sum the elements of a vector. The routines perform the following operation:

```
DO 10 I = 1,LENGTH
    SUM = SUM + VEC(I)
10 CONTINUE
```

VEC\_SUM\_I

VEC\_SUM\_I

Sums the elements of a vector, stepping through the vector by increments.

#### FORMAT

```
sum = VEC_$SUM_I (vec, inc, length)
sum = VEC_$DSUM_I (vec, inc, length)
sum = VEC_$ISUM_I (vec, inc, length)
sum = VEC_$ISUM16_I (vec, inc, length)
```

#### RETURN VALUE

**sum**

Floating-point or integer sum of the elements.

#### INPUT PARAMETERS

**vec**

Floating-point or integer vector to be summed.

**inc**

Increment for the index of vec. This is a 4-byte integer.

**length**

Number of elements to sum. This is a 4-byte integer.

#### USAGE

These functions step through a vector incrementally, summing its elements. The sum is returned as the value of the function. The routines do the following:

```
      J = 1
      SUM = 0.0
      DO 10 I=1,LENGTH
          SUM = SUM+VEC(J)
          J = J+INC1
10 CONTINUE
```



**VEC\_\$SWAP**

Swaps the elements of two vectors.

**FORMAT**

VEC\_\$SWAP (vec1, vec2, length)  
VEC\_\$DSWAP (vec1, vec2, length)  
VEC\_\$ISWAP (vec1, vec2, length)  
VEC\_\$ISWAP16 (vec1, vec2, length)

**INPUT/OUTPUT PARAMETERS****vec1, vec2**

Floating-point or integer vectors to be swapped.

**INPUT PARAMETERS****length**

Number of elements to swap. This is a 4-byte integer.

**USAGE**

These routines swap the elements of two vectors. They perform the following operation:

```
DO 10 I = 1,LENGTH
  TEMP = VEC1(I)
  VEC1(I) = VEC2(I)
  VEC2(I) = TEMP
10 CONTINUE
```

VEC\_\$SWAP\_I

VEC\_\$SWAP\_I

Swaps the elements of two vectors, stepping through the vectors by increments.

#### FORMAT

VEC\_\$SWAP\_I (vec1, inc1, vec2, inc2, length)  
VEC\_\$DSWAP\_I (vec1, inc1, vec2, inc2, length)  
VEC\_\$ISWAP\_I (vec1, inc1, vec2, inc2, length)  
VEC\_\$ISWAP16\_I (vec1, inc1, vec2, inc2, length)

#### INPUT/OUTPUT PARAMETERS

**vec1**

Floating-point or integer vector to be swapped.

#### INPUT PARAMETERS

**inc1**

Increment for the index of vec1. This is a 4-byte integer.

#### INPUT/OUTPUT PARAMETERS

**vec2**

Floating-point or integer vector to be swapped.

#### INPUT PARAMETERS

**inc2**

Increment for the index of vec2. This is a 4-byte integer.

**length**

Number of elements to swap. This is a 4-byte integer.

#### USAGE

These routines step through two vectors by increments, swapping their elements. They perform the following operation:

```
J = 1
K = 1
DO 10 I=1,N
    TEMP = VEC1(J)
    VEC1(J) = VEC2(K)
    VEC2(K) = TEMP
    J = J + INC1
    K = K + INC2
10 CONTINUE
```

## VEC\_\$ZERO

Zeros a vector.

**FORMAT**

VEC\_\$ZERO (vector, length)  
VEC\_\$DZERO (vector, length)  
VEC\_\$IZERO (vector, length)  
VEC\_\$IZERO16 (vector, length)

**INPUT/OUTPUT PARAMETERS****vector**

Floating-point or integer vector to be zeroed.

**INPUT PARAMETERS****length**

Number of elements to zero. This is a 4-byte integer.

**USAGE**

These routines zero the elements of a vector. They perform the following operation:

```
DO 10 I = 1,LENGTH  
  VEC(I) = 0.0  
10 CONTINUE
```

VEC\_\$ZERO\_I

VEC\_\$ZERO\_I

Zeros a vector by increments.

#### FORMAT

VEC\_\$ZERO\_I (vector, inc, length)  
VEC\_\$DZERO\_I (vector, inc, length)  
VEC\_\$IZERO\_I (vector, inc, length)  
VEC\_\$IZERO16\_I (vector, inc, length)

#### INPUT/OUTPUT PARAMETERS

##### vector

Floating-point or integer vector to be zeroed.

#### INPUT PARAMETERS

##### inc

Increment for the index of vector. This is a 4-byte integer.

##### length

Number of elements to zero. This is a 4-byte integer.

#### USAGE

These routines step through a vector by increments, zeroing its elements. VEC\_\$ZERO\_I zeros 32 bits regardless of data type and VEC\_\$DZERO\_I zeros 64 bits. The routines perform the following:

```
      J = 1
      DO 10 I=1,LENGTH
          VEC(J) = 0.0
          J = J+INC1
10 CONTINUE
```

# VFMT

## Table of Contents

Data Types  
VFMT\_\$DECODE  
VFMT\_\$ENCODE  
VFMT\_\$READ  
VFMT\_\$RS  
VFMT\_\$WRITE  
VFMT\_\$WS  
Errors

VFMT-2  
VFMT-3  
VFMT-5  
VFMT-7  
VFMT-8  
VFMT-8  
VFMT-10  
VFMT-11  
VFMT-12



## VFMT


The VFMT (Variable Text Format) programming calls convert program data from one format to another. This section describes their data types, call syntax, and error codes. Refer to the Introduction at the beginning of this manual for a description of data-type diagrams and call syntax format.

VFMT DATA TYPES

DATA TYPES

STATUS\_ \$T

A status code. The diagram below illustrates the STATUS\_ \$T data type:



Field Description:

all  
All 32 bits in the status code.

fail  
The fail bit. If this bit is set, the error was not within the scope of the module invoked, but occurred within a lower-level module (bit 31).

subsys  
The subsystem that encountered the error (bits 24 - 30).

modc  
The module that encountered the error (bits 16 - 23).

code  
A signed number that identifies the type of error that occurred (bits 0 - 15).

VFMT\_ \$STRING\_ T

An array of up to 200 characters. Access control string.



**VFMT\_ \$DECODE**

Decodes data from a text buffer and writes the decoded data into program variables.

**FORMAT**

return-value = VFMT\_ \$DECODE{2|5|10} (control-string, text-buffer, size,  
count, status, a1, a2, ... a10)

**RETURN VALUE****return-value**

Position in the text buffer of the last decoded character. This is a 2-byte integer.

The return-value indicates the position in the text buffer of the last decoded character. The first buffer position is 1.

**INPUT PARAMETERS****control-string**

Character string giving instructions for decoding the input data. See the VFMT chapter of the *Programming With General System Calls* manual for details about how to construct control strings.

**text-buffer**

Buffer containing data to be decoded, in VFMT\_ \$STRING\_ T format. This is an array of up to 200 characters.

**size**

Number of bytes of data in the text buffer. This is a 2-byte integer.

**OUTPUT PARAMETERS****count**

Number of fields successfully decoded. This is a 2-byte integer.

**status**

Completion status, in STATUS\_ \$T format. This data type is 4 bytes long. See the VFMT Data Types section for more information.

**INPUT/OUTPUT PARAMETERS****a1, a2, ... a10**

Up to ten variables containing decoded data. The number required depends on the number immediately following VFMT\_ \$DECODE. The number of variables cannot exceed this number. In Pascal programs, you must specify exactly this number of variables, using dummy variables if necessary.

If you are decoding ASCII text strings, you must provide two variables for each text string: a character string to contain the decoded string, and a 2-byte integer variable to contain the length of the decoded string, which is returned by VFMT.

## VFMT\_\$DECODE

The returned length of the output string depends on what you specify in the control string. VFMT determines the output string length using the M directive or, if M is omitted, using the previous value of the string length variable. In the latter way you may use this integer variable to input a maximum string length.

If you specify a maximum string length with the M directive, and the length of the string decodes is less than the maximum, VFMT returns the actual length (the lesser value) in the string-length variable.

## USAGE

This is actually a description of three separate DOMAIN system calls: VFMT\_\$DECODE2, VFMT\_\$DECODE5, and VFMT\_\$DECODE10. The number immediately following VFMT\_\$DECODE indicates the maximum number of variables the call can handle. The number must follow immediately, with no embedded spaces.

**VFMT\_ \$ENCODE**

Encodes and writes data into a text buffer.

**FORMAT**

VFMT\_ \$ENCODE[2|5|10] (control-string, text-buffer, capacity, size,  
a1, a2 , ... a20)

**INPUT PARAMETERS****control-string**

Character string giving instructions for encoding the output data. See the VFMT chapter of the *Programming With General System Calls* manual for details about how to construct control strings.

**OUTPUT PARAMETERS****text-buffer**

Buffer to contain the encoded data, in VFMT\_ \$STRING\_ T format. This is an array of up to 200 characters.

**INPUT PARAMETERS****capacity**

Maximum number of characters that may be placed in the text-buffer. This is a 2-byte integer.

**OUTPUT PARAMETERS****size**

Number of characters placed in the buffer. This is a 2-byte integer.

**INPUT PARAMETERS****a1, a2 , ... a20**

Up to twenty variables containing data for encoding. The number of arguments required depends on the number immediately following VFMT\_ \$ENCODE, if any (see Usage description below). If you append a number to the call, then the number of variables cannot exceed this number. In Pascal programs, you must specify exactly this number of variables, using dummy variables if necessary.

If you are encoding ASCII text strings, you must provide two variables for each text string: a character string containing the string, and a 2-byte integer variable containing the length of the string.

**USAGE**

This is actually a description of four separate DOMAIN system calls: VFMT\_ \$ENCODE, VFMT\_ \$ENCODE2, VFMT\_ \$ENCODE5, and VFMT\_ \$ENCODE10. The number

VFMT\_\$ENCODE

immediately following VFMT\_\$ENCODE (if any) indicates the maximum number of variables the call can handle. The number must follow immediately, with no embedded spaces. If no number is specified, you may use a variable number of arguments up to a maximum of 20.

Any individual VFMT\_\$ENCODE[2|5|10] call can write out a maximum of 512 bytes per call. If you attempt to write more than that limit, it is truncated.

**VFMT\_\$READ**

Reads character data from standard input and decodes them into variables.

**FORMAT**

VFMT\_\$READ{2|5|10} (control-string, count, status, a1, a2, ... a10)

**INPUT PARAMETERS****control-string**

A character string giving instructions for decoding the input data. See the VFMT chapter of the *Programming With General System Calls* manual for details about how to construct control strings.

**OUTPUT PARAMETERS****count**

Number of fields successfully decoded. This is a 2-byte integer.

**status**

Completion status, in STATUS\_\$T format. This data type is 4 bytes long. See the VFMT Data Types section for more information.

**INPUT/OUTPUT PARAMETERS****a1, a2, ... a10**

Up to ten variables containing decoded data. The number required depends on the number immediately following VFMT\_\$READ. The number of variables cannot exceed this number. In Pascal programs, you must specify exactly this number of variables, using dummy variables if necessary.

If you are decoding ASCII text strings, you must provide two variables for each text string: a character string to contain the decoded string, and a 2-byte integer variable to contain the length of the decoded string, which is returned by VFMT.

The returned length of the output string depends on what you specify in the control string. VFMT determines the output string length using the M directive or, if M is omitted, using the previous value of the string length variable. In the latter way you may use this integer variable to input a maximum string length.

If you specify a maximum string length with the M directive, and the length of the string decodes is less than the maximum, VFMT returns the actual length (the lesser value) in the string-length variable.

**USAGE**

This is actually a description of three separate DOMAIN system calls: VFMT\_\$READ2, VFMT\_\$READ5, and VFMT\_\$READ10. The number immediately following VFMT\_\$READ indicates the maximum number of variables the call can handle. The number must follow immediately, with no embedded spaces.

VFMT\_ \$RS

VFMT\_ \$RS

Reads character data from a stream and decodes them into variables.

## FORMAT

VFMT\_ \$RS{2|5|10} (stream-id, control-string, count, status, a1, a2, ...a10)

## INPUT PARAMETERS

### stream-id

Number of the stream from which data are read, in STREAM\_ \$ID\_ T format. This is a 2-byte integer.

### control-string

A character string giving instructions for decoding the input data. See the VFMT chapter of the *Programming With General System Calls* manual for details about how to construct control strings.

## OUTPUT PARAMETERS

### count

Number of fields successfully decoded. This is a 2-byte integer.

### status

Completion status, in STATUS\_ \$T format. This data type is 4 bytes long. See the VFMT Data Types section for more information.

## INPUT/OUTPUT PARAMETERS

### a1, a2, ... a10

Up to ten variables containing decoded data. The number required depends on the number immediately following VFMT\_ \$RS. The number of variables cannot exceed this number. In Pascal programs, you must specify exactly this number of variables, using dummy variables if necessary.

If you are decoding ASCII text strings, you must provide two variables for each text string: a character string to contain the decoded string, and a 2-byte integer variable to contain the length of the decoded string, which is returned by VFMT.

The returned length of the output string depends on what you specify in the control string. VFMT determines the output string length using the M directive or, if M is omitted, using the previous value of the string length variable. In the latter way you may use this integer variable to input a maximum string length.

If you specify a maximum string length with the M directive, and the length of the string decoded is less than the maximum, VFMT returns the actual length (the lesser value) in the string-length variable.

## USAGE

This is actually a description of three separate DOMAIN system calls: VFMT\_ \$RS2,

VFMT\_\$RS5, and VFMT\_\$RS10. The number immediately following VFMT\_\$RS indicates the maximum number of variables the call can handle. The number must follow immediately, with no embedded spaces.

VFMT\_\$WRITE

VFMT\_\$WRITE

Encodes data and writes them to standard output.

## FORMAT

VFMT\_\$WRITE[2|5|10] (control-string, a1, a2 , ... a20)

## INPUT PARAMETERS

### control-string

A character string containing the control information for encoding. See the VFMT chapter of the *Programming With General System Calls* manual for details about how to construct control strings.

### a1, a2 , ... a20

Up to twenty variables containing data for encoding. The number of arguments required depends on the number immediately following VFMT\_\$WRITE, if any (see Usage description below). If you append a number to the call, then the number of variables cannot exceed this number. In Pascal programs, you must specify exactly this number of variables, using dummy variables if necessary.

If you are encoding ASCII text strings, you must provide two variables for each text string: a character string containing the string, and a 2-byte integer variable containing the length of the string.

## USAGE

This is actually a description of four separate DOMAIN system calls: VFMT\_\$WRITE, VFMT\_\$WRITE2, VFMT\_\$WRITE5, and VFMT\_\$WRITE10. The number immediately following VFMT\_\$WRITE (if any) indicates the maximum number of variables the call can handle. The number must follow immediately, with no embedded spaces. If no number is specified, you may use a variable number of arguments up to a maximum of 20.

Any individual VFMT\_\$WRITE[2|5|10] call can write out a maximum of 512 bytes per call. If you attempt to write more than that, it is truncated.



**VFMT\_\$WS**

Encodes data and writes them to a stream.

**FORMAT**

VFMT\_\$WS[2|5|10] (stream-id, control-string, a1, a2 , ... a20)

**INPUT PARAMETERS****stream-id**

The number of the stream to which data are written, in STREAM\_\$ID\_T format. This is a 2-byte integer.

**control-string**

A character string containing the control information for encoding. See the VFMT chapter of the *Programming With General System Calls* manual for details about how to construct control strings.

**a1, a2 , ... a20**

Up to twenty variables containing data for encoding. The number of arguments required depends on the number immediately following VFMT\_\$WS, if any (see Usage description below). If you append a number to the call, then the number of variables cannot exceed this number. In Pascal programs, if you append a number to the call, then you must specify exactly this number of variables, using dummy variables if necessary.

If you are encoding ASCII text strings, you must provide two variables for each text string: a character string containing the string, and a 2-byte integer variable containing the length of the string.

**USAGE**

This is actually a description of four separate DOMAIN system calls: VFMT\_\$WS, VFMT\_\$WS2, VFMT\_\$WS5, and VFMT\_\$WS10. The number immediately following VFMT\_\$WS (if any) indicates the maximum number of variables the call can handle. The number must follow immediately, with no embedded spaces. If no number is specified, you may use a variable number of arguments up to a maximum of 20.

Any individual VFMT\_\$WS[2|5|10] call can write out a maximum of 512 bytes per call.

VFMT ERRORS

**ERRORS**

VFMT\_ \$UNTERMINATED\_CTL\_STRING

Unterminated control string.

VFMT\_ \$INVALIED\_CTL\_STRING

Invalid control string.

VFMT\_ \$TOO\_FEW\_ARGS

Too few arguments supplied for read/decode.

VFMT\_ \$FW\_REQUIRED

Field width missing on "(" designator.

VFMT\_ \$EOS

Encountered end of string where more text was expected.

VFMT\_ \$NULL\_TOKEN

Encountered null token where numeric token was expected.

VFMT\_ \$NONNUMERIC\_CHAR

Non-numeric character found where numeric was expected.

VFMT\_ \$SIGN\_NOT\_ALLOWED

Sign encountered in unsigned field.

VFMT\_ \$VALUE\_TOO\_LARGE

Value out of range in text string.

VFMT\_ \$NONMATCHING\_CHAR

Character in text string does not match control string.

VFMT\_ \$NONMATCHING\_DELIMITER

Terminator in text string does not match specified terminator.

STATUS\_ \$OK

Successful completion.