



External Development Guide

IDL Version 7.1

May 2009 Edition

Copyright © ITT Visual Information Solutions
All Rights Reserved

Restricted Rights Notice

The IDL®, IDL Advanced Math and Stats™, ENVI®, and ENVI Zoom™ software programs and the accompanying procedures, functions, and documentation described herein are sold under license agreement. Their use, duplication, and disclosure are subject to the restrictions stated in the license agreement. ITT Visual Information Solutions reserves the right to make changes to this document at any time and without notice.

Limitation of Warranty

ITT Visual Information Solutions makes no warranties, either express or implied, as to any matter not expressly set forth in the license agreement, including without limitation the condition of the software, merchantability, or fitness for any particular purpose.

ITT Visual Information Solutions shall not be liable for any direct, consequential, or other damages suffered by the Licensee or any others resulting from use of the software packages or their documentation.

Permission to Reproduce this Manual

If you are a licensed user of these products, ITT Visual Information Solutions grants you a limited, nontransferable license to reproduce this particular document provided such copies are for your use only and are not sold or distributed to third parties. All such copies must contain the title page and this notice page in their entirety.

Export Control Information

The software and associated documentation are subject to U.S. export controls including the United States Export Administration Regulations. The recipient is responsible for ensuring compliance with all applicable U.S. export control laws and regulations. These laws include restrictions on destinations, end users, and end use.

Acknowledgments

ENVI® and IDL® are registered trademarks of ITT Corporation, registered in the United States Patent and Trademark Office. ION™, ION Script™, ION Java™, and ENVI Zoom™ are trademarks of ITT Visual Information Solutions.

ESRI®, ArcGIS®, ArcView®, and ArcInfo® are registered trademarks of ESRI.

Portions of this work are Copyright © 2008 ESRI. All rights reserved.

Numerical Recipes™ is a trademark of Numerical Recipes Software. Numerical Recipes routines are used by permission.

GRG2™ is a trademark of Windward Technologies, Inc. The GRG2 software for nonlinear optimization is used by permission.

NCSA Hierarchical Data Format (HDF) Software Library and Utilities. Copyright © 1988-2001, The Board of Trustees of the University of Illinois. All rights reserved.

NCSA HDF5 (Hierarchical Data Format 5) Software Library and Utilities. Copyright © 1998-2002, by the Board of Trustees of the University of Illinois. All rights reserved.

CDF Library. Copyright © 2002, National Space Science Data Center, NASA/Goddard Space Flight Center.

NetCDF Library. Copyright © 1993-1999, University Corporation for Atmospheric Research/Unidata.

HDF EOS Library. Copyright © 1996, Hughes and Applied Research Corporation.

SMACC. Copyright © 2000-2004, Spectral Sciences, Inc. and ITT Visual Information Solutions. All rights reserved.

This software is based in part on the work of the Independent JPEG Group.

Portions of this software are copyrighted by DataDirect Technologies, © 1991-2003.

BandMax®. Copyright © 2003, The Galileo Group Inc.

Portions of this computer program are copyright © 1995-1999, LizardTech, Inc. All rights reserved. MrSID is protected by U.S. Patent No. 5,710,835. Foreign Patents Pending.

Portions of this software were developed using Unisearch's Kakadu software, for which ITT has a commercial license. Kakadu Software. Copyright © 2001. The University of New South Wales, UNSW, Sydney NSW 2052, Australia, and Unisearch Ltd, Australia.

This product includes software developed by the Apache Software Foundation (www.apache.org/).

MODTRAN is licensed from the United States of America under U.S. Patent No. 5,315,513 and U.S. Patent No. 5,884,226.

QUAC and FLAASH are licensed from Spectral Sciences, Inc. under U.S. Patent No. 6,909,815 and U.S. Patent No. 7,046,859 B2.

Portions of this software are copyrighted by Merge Technologies Incorporated.

Support Vector Machine (SVM) is based on the LIBSVM library written by Chih-Chung Chang and Chih-Jen Lin (www.csie.ntu.edu.tw/~cjlin/libsvm/), adapted by ITT Visual Information Solutions for remote sensing image supervised classification purposes.

IDL Wavelet Toolkit Copyright © 2002, Christopher Torrence.

IMSL is a trademark of Visual Numerics, Inc. Copyright © 1970-2006 by Visual Numerics, Inc. All Rights Reserved.

Other trademarks and registered trademarks are the property of the respective trademark holders.



Contents

Chapter 1	
External Development Overview	11
About This Manual	12
Supported Inter-Language Communication Techniques in IDL	13
Dynamic Linking Terms and Concepts	20
When Is It Appropriate to Combine External Code with IDL?	22
Skills Required to Combine External Code with IDL	23
IDL Organization	27
External Definitions	29
Interpreting Logical Boolean Values	30
Compilation and Linking Details	31
Recommended Reading	32

Part I: Techniques That Do Not Use IDL's Internal API

Chapter 2	
Using SPAWN and Pipes	37
Chapter 3	
Using CALL_EXTERNAL	43
The CALL_EXTERNAL Function	44
Passing Parameters	54
Using Auto Glue	56
Basic C Examples	58
Wrapper Routines	62
Passing String Data	64
Passing Array Data	68
Passing Structures	70
Fortran Examples	72
Chapter 4	
Remote Procedure Calls	77
IDL and Remote Procedure Calls	78
Using IDL as an RPC Server	79
Client Variables	80
Linking to the Client Library	81
Compatibility with Older IDL Code	83
The IDL RPC Library	85
RPC Examples	110

Part II: IDL's Internal API

Chapter 5	
IDL Internals:	
Types	113
Type Codes	114
Mapping of Basic Types	116
IDL_MEMINT and IDL_FILEINT Types	119

Chapter 6

IDL Internals:

Keyword Processing	121
IDL and Keyword Processing	122
Creating Routines that Accept Keywords	123
Overview Of IDL Keyword Processing	124
The IDL_KW_PAR Structure	126
The IDL_KW_ARR_DESC_R Structure	129
Keyword Processing Options	130
The KW_RESULT Structure	132
Processing Keywords	133
Cleaning Up	136
Keyword Examples	137
The Pre-IDL 5.5 Keyword API	144

Chapter 7

IDL Internals:

Variables	151
IDL and Internal Variables	152
The IDL_VARIABLE Structure	153
Scalar Variables	156
Array Variables	157
Structure Variables	159
Heap Variables	164
Temporary Variables	165
Creating an Array from Existing Data	172
Getting Dynamic Memory	174
Accessing Variable Data	176
Copying Variables	177
Storing Scalar Values	178
Obtaining the Name of a Variable	180
Looking Up Main Program Variables	181
Looking Up Variables in Current Scope	182

Chapter 8**IDL Internals:****String Processing 183**

String Processing and IDL 184

Accessing IDL_STRING Values 185

Copying Strings 186

Deleting Strings 187

Setting an IDL_STRING Value 188

Obtaining a String of a Given Length 189

Chapter 9**IDL Internals:****Error Handling 191**

Message Blocks 192

Issuing Error Messages 195

Looking Up A Message Code by Name 201

Checking Arguments 202

Chapter 10**IDL Internals:****Type Conversion 205**

Converting Arguments to C Scalars 206

General Type Conversion 207

Converting to Specific Types 208

Chapter 11**IDL Internals:****UNIX Signals 209**

IDL and Signals 210

Signal Handlers 213

Establishing a Signal Handler 214

Removing a Signal Handler 215

UNIX Signal Masks 216

Chapter 12**IDL Internals:****Timers 221**

IDL and Timers 222

Making Timer Requests 223

Canceling Asynchronous Timer Requests	225
Blocking UNIX Timers	226

Chapter 13

IDL Internals: Files and Input/Output 229

IDL and Input/Output Files	230
File Information	232
Opening Files	236
Closing Files	239
Preventing File Closing	240
Checking File Status	241
Allocating and Freeing File Units	243
Detecting End of File	245
Flushing Buffered Data	246
Reading a Single Character	247
Output of IDL Variables	248
Adding to the Journal File	249

Chapter 14

IDL Internals: Miscellaneous 251

Dynamic Memory	252
Exit Handlers	255
User Interrupts	256
Functions for Returning System Variables	257
Terminal Information	258
Ensuring UNIX TTY State	260
Type Information	261
User Information	263
Constants	264
Macros	265

Part III: Techniques That Use IDL's Internal API

Chapter 15

Adding System Routines 269

IDL and System Routines	270
The System Routine Interface	272
Example: Hello World	273

Example: Doing a Little More (MULT2)	274
Example: A Complete Numerical Routine Example (FZ_ROOTS2)	277
Example: An Example Using Routine Design Iteration (RSUM)	286
Registering Routines	296
Enabling and Disabling System Routines	299
LINKIMAGE	307
Dynamically Loadable Modules	309

Chapter 16

Callable IDL 323

Calling IDL as a Subroutine	324
When is Callable IDL Appropriate?	325
Licensing Issues and Callable IDL	328
Using Callable IDL	329
Initialization	331
Diverting IDL Output	337
Executing IDL Statements	339
Runtime IDL and Embedded IDL	340
Cleanup	341
Issues and Examples: UNIX	342
Issues and Examples: Microsoft Windows	358

Chapter 17

Adding External Widgets to IDL 369

IDL and External Widgets	370
WIDGET_STUB	371
WIDGET_CONTROL/WIDGET_STUB	372
Functions for Use with Stub Widgets	374
Internal Callback Functions	377
UNIX WIDGET_STUB Example: WIDGET_ARROWB	379

Appendix A

Obsolete Internal Interfaces 385

Interfaces Obsolete in IDL 6.3	386
Interfaces Obsolete in IDL 5.5	388
Interfaces Obsolete in IDL 5.2.1	401
Simplified Routine Invocation	404

Obsolete Error Handling API	411
Index	413



Chapter 1

External Development Overview

This chapter discusses the following topics:

About This Manual	12	Skills Required to Combine External Code with IDL	23
Supported Inter-Language Communication Techniques in IDL	13	IDL Organization	27
Dynamic Linking Terms and Concepts	20	External Definitions	29
When Is It Appropriate to Combine External Code with IDL?	22	Interpreting Logical Boolean Values	30
		Compilation and Linking Details	31
		Recommended Reading	32

About This Manual

The *External Development Guide* describes options for using code not written in the IDL language alongside IDL itself. It is divided into three parts:

Part I: Techniques That Do Not Use IDL's Internal API

This section discusses techniques that allow IDL to work together with programs written in other programming languages, using IDL's "public" interfaces. Little or no familiarity with IDL's internal interfaces is required. For many users, the techniques in this section will solve most problems that require IDL to use — or be used by — other programs. Topics covered in Part I include:

- Letting IDL programs interact with other programs via pipes.
- Incorporating COM objects and ActiveX controls into IDL programs.
- Incorporating Java objects into IDL programs.
- Using IDL as a Remote Procedure Call server on a UNIX system.
- Calling routines written in other programming languages from within IDL using the `CALL_EXTERNAL` function.

Part II: IDL's Internal API

This section describes IDL's internal implementation in enough detail to allow you to write an IDL system routine in another compiled programming language (usually C) and link it with IDL.

Part III: Techniques That Use IDL's Internal API

This section describes the process of combining IDL with code written in another programming language. Topics covered in Part III include:

- Creating a system routine using the interface described in Part II and linking that routine into IDL at runtime.
- Calling IDL as a subroutine from another program ("Callable IDL").
- Adding user-defined widgets to IDL widget applications.

Supported Inter-Language Communication Techniques in IDL

IDL supports a number of different techniques for communicating with the operating system and programs written in other languages. These methods are described, in brief, below.

Options are presented in approximate order of increasing complexity. We recommend that you favor the simpler options at the head of this list over the more complex ones that follow if they are capable of solving your problem.

It can be difficult to choose the best option — there is a certain amount of overlap between their abilities. We highlight the advantages and disadvantages of each method as well as make recommendations to help you decide which approach to take. By comparing this list with the requirements of the problem you are trying to solve, you should be able to quickly determine the best solution.

Translate into IDL

Advantages

All the benefits of using a high level, interpreted, array oriented environment with high levels of platform independence.

Disadvantages

Not always possible.

Recommendation

Writing in IDL is the easiest path. If you have existing code in another language that is simple enough to translate to IDL, this is the best way to go. You should investigate the other options if the existing code is sufficiently complex, has desirable performance advantages, or is the reference implementation of some standardized package. Another good reason for considering the techniques described in this book is if you wish to access IDL abilities from a large program written in some other language.

SPAWN

The simplest (but most limited) way to access programs external to IDL is to use the SPAWN procedure. Calling SPAWN spawns a child process that executes a specified command. The output from SPAWN can be captured in an IDL string variable. In addition, IDL can communicate with a child process through a bi-directional pipe using SPAWN. More information about SPAWN can be found in [Chapter 2, “Using SPAWN and Pipes”](#) or in the documentation for “[SPAWN](#)” (*IDL Reference Guide*).

Advantages

- Simplicity
- Allows use of existing standalone programs.
- Data can be sent to and returned by the program via a pipe, making sophisticated inter-program communication possible quickly and easily.

Disadvantages

- Can be a slow when transferring large datasets.
- Programs may not have a useful user interface.

Recommendation

SPAWN is the easiest form of interprocess communication supported by IDL and allows accessing operating system commands directly.

Microsoft COM and ActiveX

IDL supports the inclusion of COM objects and ActiveX controls within IDL applications running on Microsoft Windows systems by encapsulating the object or control in an IDL object. Full access to the COM object or ActiveX control’s methods is available in this manner, allowing you to incorporate features not available in IDL into IDL programs. For more information, see [Chapter 2, “Overview: COM and ActiveX in IDL”](#) (*IDL Connectivity Bridges*).

Advantages

- Integrates easily with an important interprocess communication mechanism under Microsoft Windows.
- May support a higher level interface than the function call interfaces supported by the remaining options.

Disadvantages

- Only supported under Microsoft Windows.

Recommendation

Incorporate COM objects or ActiveX controls into your Windows-only IDL application if doing so provides functionality you cannot easily duplicate in IDL.

Use the IDL ActiveX control if you are writing a Windows-only application in a language that supports ActiveX and you wish to use IDL to perform computation or graphics within a framework established by this other application.

Sun Java

IDL also supports the inclusion of Java objects within IDL applications by encapsulating the object or control in an IDL object. Full access to the Java object is available in this manner, allowing you to incorporate features not available in IDL into IDL programs. For more information, see [Chapter 5, “Using Java Objects in IDL”](#) (*IDL Connectivity Bridges*).

Advantages

- Integrates easily with all types of Java code.
- Can easily leverage existing Java objects into IDL.

Disadvantages

- Only supported under Microsoft Windows, Linux, Solaris, and Macintosh platforms supported in IDL.

Recommendation

Incorporate Java objects into your IDL application if doing so provides functionality you cannot easily duplicate in IDL.

UNIX Remote Procedure Calls (RPCs)

UNIX platforms can use Remote Procedure Calls (RPCs) to facilitate communication between IDL and other programs. IDL is run as an RPC server and your own program is run as a client. IDL’s RPC functionality is documented in [Chapter 4, “Remote Procedure Calls”](#).

Advantages

- Code executes in a process other than the one running IDL, possibly on another machine, providing robustness and protection in a distributed framework.
- API is similar to that employed by Callable IDL, making it reasonable to switch from one to the other.
- Possibility of overlapped execution on a multi-processor system.

Disadvantages

- Complexity of managing RPC servers.
- Bandwidth limitations of network for moving large amounts of data.
- Only supported under UNIX.

Recommendation

Use RPC if you are coding in a distributed UNIX-only environment and the amount of data being moved is reasonable on your network. `CALL_EXTERNAL` might be more appropriate for especially simple tasks, or if the external code is not easily converted into an RPC server, or you lack RPC experience and knowledge.

CALL_EXTERNAL

IDL's `CALL_EXTERNAL` function loads and calls routines contained in shareable object libraries. IDL and the called routine share the same memory and data space. `CALL_EXTERNAL` is much easier to use than either system routines (`LINKIMAGE`, `DLMs`) or Callable IDL and is often the best (and simplest) way to communicate with other programs. `CALL_EXTERNAL` is also supported on all IDL platforms.

While many of the topics in this book can enhance your understanding of `CALL_EXTERNAL`, specific documentation and examples can be found in [Chapter 3, “Using `CALL_EXTERNAL`”](#) and the documentation for “`CALL_EXTERNAL`” (*IDL Reference Guide*).

Advantages

- Allows calling arbitrary code written in other languages.
- Requires little or no understanding of IDL internals.

Disadvantages

- Errors in coding can easily corrupt the IDL program.
- Requires understanding of system programming, compiler, and linker.
- Data must be passed to and from IDL in precisely the correct type and size or memory corruption and program errors will result.
- System and hardware dependent, requiring different binaries for each target system.

Recommendation

Use `CALL_EXTERNAL` to call code written for general use in another language (that is, without knowledge of IDL internals). For safety, you should call your `CALL_EXTERNAL` functions within special IDL procedures or functions that do error checking of the inputs and return values. In this way, you can reduce the risks of corruption and give your callers an appropriate IDL-like interface to the new functionality. If you use this method to incorporate external code into IDL, We highly recommend that you also use the `MAKE_DLL` procedure and the `AUTO_GLUE` keyword to `CALL_EXTERNAL`.

If you lack knowledge of IDL internals, `CALL_EXTERNAL` is the best way to add external code quickly. Programmers who do understand IDL internals will often write a system routine instead to gain flexibility and full integration into IDL.

IDL System Routine (LINKIMAGE, DLMS)

It is possible to write system routines for IDL using a compiled language such as C. Such routines are written to have the standard IDL calling interface, and are dynamically linked, as with `CALL_EXTERNAL`. They are more difficult to write, but more flexible and powerful. System routines provide access to variables and other objects inside of IDL.

This book contains the information necessary to successfully add your own code to IDL as a system routine. Especially important is [Chapter 15, “Adding System Routines”](#). Additional information about system routines can be found in [Chapter 3, “Using `CALL_EXTERNAL`”](#) and in the documentation for “`LINKIMAGE`” (*IDL Reference Guide*).

Advantages

- This is the most fully integrated option. It allows you to write IDL system routines that are indistinguishable from those written by ITT Visual Information Solutions.
- In use, system routines are very robust and fault tolerant.
- Allows direct access to IDL user variables and other important data structures.

Disadvantages

- All the disadvantages of `CALL_EXTERNAL`.
- Requires in-depth understanding of IDL internals, discussed in Part II of this manual.

Recommendation

Use system routines if you require the highest level of integration of your code into the IDL system. UNIX users with RPC experience should consider using RPCs to get the benefits of distributed processing. If your task is sufficiently simple or you do not have the desire or time to learn IDL internals, `CALL_EXTERNAL` is an efficient way to get the job done.

Callable IDL

IDL is packaged in a shareable form that allows other programs to call IDL as a subroutine. This shareable portion of IDL can be linked into your own programs. This use of IDL is referred to as “Callable IDL” to distinguish it from the more usual case of calling your code *from* IDL via `CALL_EXTERNAL` or as a system routine (`LINKIMAGE`, `DLM`).

This book contains the information necessary to successfully call IDL from your own code.

Advantages

- Supported on all systems.
- Allows extremely low level access to IDL.

Disadvantages

- All the disadvantages of `CALL_EXTERNAL` or IDL system routines.

- IDL imposes some limitations on programming techniques that your program can use.

Recommendation

Most platforms offer a specialized method to call other programs that might be more appropriate. Windows users should consider the ActiveX control or COM component. UNIX users should consider using the IDL RPC server. If these options are not appropriate for your task and you wish to call IDL from another program, then use Callable IDL.

Dynamic Linking Terms and Concepts

All systems on which IDL runs support the concept of dynamic linking. Dynamic linking consists of compiling and linking code into a form which is loadable by programs at run time as well as link time. The ability to load them at run time is what distinguishes them from ordinary object files. Various operating systems have different names for such loadable code:

- UNIX: Sharable Libraries
- Windows: Dynamic Link Libraries (DLL)

In this manual, we will call such files *sharable libraries* in order to have a consistent and uniform way to refer to them. It should be understood that this is a generic usage that applies equally to all of these systems. Sharable libraries contain functions that can be called by any program that loads them. Often, you must specify special compiler and linker options to build a sharable library. On many systems, the linker gives you control over which functions and data (often referred to as *symbols*) are visible from the outside (public symbols) and which are hidden (private symbols). Such control over the interface presented by a sharable library can be very useful. Your system documentation discusses these options and explains how to build a sharable library.

Dynamic linking is the enabling technology for many of the techniques discussed in this manual. If you intend to use any of these techniques, you should first be sure to study your system documentation on this topic.

CALL_EXTERNAL

CALL_EXTERNAL uses dynamic linking to call functions written in other languages from IDL.

LINKIMAGE and Dynamically Loadable Modules (DLMS)

These mechanisms use dynamic linking to add external code that supports the standard IDL system routine interface to IDL as system routines.

Callable IDL

Most of IDL is built as a sharable library. The actual IDL program that implements the standard interactive IDL program links to this library and uses it to do its work. Since IDL is a sharable library, it can be called by other programs.

Remote Procedure Calls (RPCs)

The IDL RPC server is a program that links to the IDL sharable library. The IDL RPC client side library is also a sharable library. Your RPC client program links against it to obtain access to the IDL RPC system.

When Is It Appropriate to Combine External Code with IDL?

IDL is an interactive program that runs across numerous operating systems and hardware platforms. The IDL user enjoys a large amount of portability across these platforms because IDL provides access to system abilities at a relatively high level of abstraction. The large majority of IDL users have no need to understand its inner workings or to link their own code into it.

There are, however, reasons to combine external code with IDL:

- Many sites have an existing investment in other code that they would prefer to use from IDL rather than incurring the cost of rewriting it in the IDL language.
- It is often best to use the reference implementation of a software package rather than re-implement it in another language, risk adding incorrect behaviors to it, and incur the ongoing maintenance costs of supporting it.
- IDL may be largely suitable for a given task, requiring only the addition of an operation that cannot be performed efficiently in the IDL language.

A programmer who is considering adding compiled code to IDL should understand the following caveats:

- We attempt to keep the interfaces described in this document stable, and we endeavor to minimize gratuitous change. However, we reserve the right to make any changes required by the future evolution of the system. Code linked with IDL is more likely to require updates and changes to work with new releases of IDL than programs written in the IDL language.
- The act of linking compiled code to IDL is inherently less portable than use of IDL at the user level.
- Troubleshooting and debugging such applications can be very difficult. With standard IDL, malfunctions in the program are clearly the fault of ITT Visual Information Solutions, and given a reproducible bug report, we attempt to fix them promptly. A program that combines IDL with other code makes it difficult to unambiguously determine where the problem lies. The level of support ITT Visual Information Solutions can provide in such troubleshooting is minimal. The programmer is responsible for locating the source of the difficulty. If the problem is in IDL, a simple program demonstrating the problem must be provided before we can address the issue.

Skills Required to Combine External Code with IDL

There is a large difference between the level at which a typical user sees IDL compared to that of the internals programmer. To the user, IDL is an easy-to-use, array-oriented language that combines numerical and graphical abilities, and runs on many platforms. Internally, IDL is a large C language program that includes a compiler, an interpreter, graphics, mathematical computation, user interface, and a large amount of operating system-dependent code.

The amount of knowledge required to effectively write internals code for IDL can come as a surprise to the user who is only familiar with IDL's external face. To be successful, the programmer must have experience and proficiency in many of the following areas:

Microsoft COM

To incorporate a COM object into your IDL program, you should be familiar with COM interfaces in general and the interface of the object you are using in particular.

Microsoft ActiveX

To incorporate an ActiveX control into your IDL widget application, you should be familiar with COM interfaces in general and the interface of the control you are using in particular.

Sun Java

To incorporate a Java object into your IDL program, you should be familiar with Java object classes in general and the methods and data members of the object you are using in particular.

UNIX RPC

To use IDL as an RPC server, a knowledge of Sun RPC (Also known as ONC RPC) is required. Sun RPC is the fundamental enabling technology that underlies the popular NFS (Network File System) software available on all UNIX systems, and as such, is universally available on UNIX. The system documentation on this subject should be sufficient.

ANSI C

IDL is written in ANSI C. To understand the data structures and routines described in this document, you must have a complete understanding of this language.

System C Compiler, Linker, and Libraries

In order to successfully integrate IDL with your code, you must fully understand the compilation tools being used as well as those used to build IDL and how they might interact. IDL is built with the standard C compiler used (and usually supplied) by the vendor of each platform to ensure full compatibility with all system components.

Inter-language Calling Conventions (C++, Fortran, ...)

It is possible to link IDL directly with code written in compiled languages other than C although the details differ depending on the machine, language, and compiler used. It is the programmer's responsibility to understand the inter-language calling conventions and rules for the target environment—there are too many possibilities for ITT Visual Information Solutions to actively document them all. ANSI C is a standard system programming language on all systems supported by IDL, so it is usually straightforward to combine it with code written in other compiled languages. You need to understand:

- The conventions used to pass parameters to functions in both languages. For example, C uses call-by-value while Fortran uses call-by-reference. It is easy to compensate for such conventions, but they must be taken into account.
- Any systematic name changes applied by the compilers. For example, some compilers add underscores at the beginning or end of names of functions and global data.
- Any run-time initialization that must be performed. On many systems, the real initial entry point for the program is not `main()`, but a different function that performs some initialization work and then calls your `main()` function. Usually these issues have been addressed by the system vendor, who has a large interest in allowing such inter-language usage:
 - If you call IDL from a program written in a language other than C, has the necessary initialization occurred?
 - If you use IDL to call code written in a language other than C, do you need to take steps to initialize the runtime system for that language?
 - Are the two runtime systems compatible?

Alternatives to direct linking (Microsoft COM or Active X) exist on some systems that simplify the details of inter-language linking.

C++

We are often asked if IDL can call C++ code. Compatibility with C has always been a strong design goal for C++, and C++ is largely a superset of the C language. It certainly is possible to combine IDL with C++ code. Callable IDL is especially simple, as all you need to do is to include the `idl_export.h` header file in your C++ code and then call the necessary IDL functions directly. Calling C++ code from IDL (`CALL_EXTERNAL`, System Routines) is also possible, but there are some issues you should be aware of:

- As a C program, IDL is not able to directly call C++ methods, or use other object-oriented features of the C++ language. To use these C++ features, you must supply a function with C linkage (using an extern “C” specification) for IDL to call. That routine, which is written in C++ is then able to use the C++ features.
- IDL does not initialize any necessary C++ runtime code. Your system may require such code to be executed before your C++ code can run. Consult your system documentation for details. (Please be aware that this information can be difficult to find; locating it may require some detective work on your part.)

Fortran

Issues to be aware of when combining IDL with Fortran:

- The primary issue surrounding the calling of Fortran code from IDL is one of understanding the calling conventions of the two languages. C passes everything by value, and supplies an operator that lets you explicitly take the address of a memory object. Fortran passes everything by reference (by address). Difficulties in calling FORTRAN from C usually come down to handling this issue correctly. Some people find it helpful to write a C wrapper function to call their Fortran code, and then have IDL call the wrapper. This is generally not necessary, but may be convenient.
- IDL is a C program, and as such, does not initialize any necessary Fortran runtime code. Your system may require such code to be executed before your Fortran code can run. In particular, Fortran code that does its own input output often requires such startup code to be executed. Consult your system documentation for details. One common strategy that can minimize this sort of problem is to use IDL’s I/O facilities to do I/O, and have your Fortran code limit itself to computation.

Operating System Features and Conventions

With the exception of purely numerical code, the programmer must usually fully understand the target operating system environment in which IDL is running in order to write code to link with it.

Microsoft Windows

You must be an experienced Windows programmer with an understanding of Windows APIs and DLLs.

UNIX

You should understand system calls, signals, processes, standard C libraries, and possibly even X Windows depending on the scope of the code being linked.

IDL Organization

In order to properly write code to be linked with IDL, it is necessary to understand a little about its internal operation. This section is intended to give just enough background to understand the material that follows. Traditional interpreted languages work according to the following algorithm:

```
while (statements remaining) {
    Get next statement.
    Perform lexical analysis and parse statement.
    Execute statement.
}
```

This description is accurate at a conceptual level, and most early interpreters did their work in exactly this way due to its simplicity. However, this scheme is inefficient because:

- The meaning of each statement is determined by the relatively expensive operations of lexical analysis, parsing, and semantic analysis each and every time the statement is encountered.
- Since each statement is considered in isolation, any statement that requires jumping to a different location in the program will require an expensive search for the target location. Usually, this search starts at the top of the file and moves forward until the target is found.

To avoid these problems, the IDL system uses a two-step process in which compilation and interpretation are separate. The core of the system is the interpreter. The interpreter implements a simple, stack-based postfix language, in which each instruction corresponds to a primitive of the IDL language. This internal form is a compact binary version of the IDL language routine. Routines written in the IDL language are compiled into this internal form by the IDL compiler when the `.RUN` executive command is issued, or when any other command requires a new routine to be executed. Once the IDL routine is compiled, the original version is ignored, and all references to the routine are to the compiled version. Some of the advantages of this organization are:

- The expensive compilation process is only performed once, no matter how often the resulting code is executed.
- Statements are not considered in isolation, so the compiler keeps track of the information required to make jumping to a new location in the program fast.
- The binary internal form is much faster to interpret than the original form.

- The internal form is compact, leading to better use of main memory, and allowing more code to fit in any memory cache the computer might be using.

The Interpreter Stack

The primary data structure in the interpreter is the stack. The stack contains pointers to variables, which are implemented by **IDL_VARIABLE** structures (see “[The IDL_VARIABLE Structure](#)” on page 153). Pointers to **IDL_VARIABLES** are referred to as **IDL_VPTRs**. Most interpreter instructions work by removing a predefined number of elements from the stack, performing their function, and then pushing the **IDL_VPTR** to the resulting **IDL_VARIABLE** back onto the stack. The removed items are the arguments to the instruction, and the new element represents the result. In this sense, the IDL interpreter is no different from any other postfix language interpreter. When an IDL routine is compiled, the compiler checks the number of arguments passed to each system routine against the minimum and maximum number specified in an internal table of routines, and signals an error if an invalid number of arguments is specified.

At execution time, the interpreter instructions that execute system procedures and functions operate as follows:

1. Look up the requested routine in the internal table of routines.
2. Execute the routine that implements the desired routine.
3. Remove the arguments from the stack.
4. If the routine was a function, push its result onto the stack.

Thus, the compiler checks for the proper number of arguments, and the interpreter does all the work related to pushing and popping elements from the stack. The called function need only worry about executing its operation and providing a result.

External Definitions

The file `idl_export.h`, found in the `external/include` subdirectory of the IDL distribution, supplies all the IDL-specific definitions required to write code for inclusion with IDL. As such, this file defines the interface between IDL and your code. It will be worth your while to examine this file, reading the comments and getting a general idea of what is available. If you are not writing in C, you will have to translate the definitions in this file to suit the language you are using.

Warning

`idl_export.h` contains some declarations which are necessary to the compilation process, but which are still considered private to ITT Visual Information Solutions. Such declarations are likely to be changed in the future and should not be depended on. In particular, many of the structure data types discussed in this document have more fields than are discussed here—such fields should not be used. For this reason, you should always include `idl_export.h` rather than entering the type definitions from this document. This will also protect you from changes to these data structures in future releases of IDL. Anything in `idl_export.h` that is not explicitly discussed in this document should not be relied upon.

The following two lines should be included near the top of every C program file that is to become part of IDL:

```
#include <stdio.h>
#include "idl_export.h"
```

Interpreting Logical Boolean Values

IDL is written in the C programming language, and this manual therefore discusses C language functions and data structures from the IDL program. In this documentation, you will see references to logical (boolean) arguments and results referred to in any of the following forms: True, False, TRUE, FALSE, IDL_TRUE, IDL_FALSE, and possibly other permutations on these. In all cases, the meaning of true and false in this manual correspond to those of the C programming language: A zero (0) value is interpreted as “false”, and a non-zero value is “true”.

When reading this manual, please be aware of the following points:

- Unless otherwise specified, the actual word used when discussing logical values is not important (i.e. true, True, TRUE, and IDL_TRUE) all mean the same thing.
- Internally, IDL uses the IDL_TRUE and IDL_FALSE macros described in “[Macros](#)” on page 265, for hard-wired logical constants. These macros have the values 1, and 0 respectively. This convention is nothing more than reflection of the need for a consistent standard within our code, and a desire to keep IDL names within a standard namespace to avoid collisions with user selected names. Otherwise, any of those other alternative names might have been used with equally good results.
- We don’t use the IDL_TRUE and IDL_FALSE convention in the text of this book because it would be unnecessarily awkward, preferring the more natural True/TRUE and False/FALSE.
- The convention for truth values in the IDL Language differ from those used in the C language. It is important to keep the language being used in mind when reading code to avoid drawing incorrect conclusions about its meaning.

Compilation and Linking Details

Once you've written your code, you need to compile it and link it into IDL before it can be run. Information on how to do this is available in the various subdirectories of the `external` subdirectory of the IDL distribution. References to files that are useful in specific situations are contained in this book.

In addition:

- The IDL `MAKE_DLL` procedure, documented in the *IDL Reference Manual*, provides a portable high level mechanism for building sharable libraries from code written in the C programming language.
- The IDL `!MAKE_DLL` system variable is used by the `MAKE_DLL` procedure to construct C compiler and linker commands appropriate for the target platform. If you do not use `MAKE_DLL` to compile and link your code, you may find the value of `!MAKE_DLL.CC` and `!MAKE_DLL.LD` helpful in determining which options to specify to your compiler and linker, in conjunction with your system and compiler documentation. For the C language, the options in `!MAKE_DLL` should be very close to what you need. For other languages, the `!MAKE_DLL` options should still be helpful in determining which options to use, as on most systems, all the language compilers accept similar options.
- The UNIX IDL distribution has a `bin` subdirectory that contains platform specific directories that in turn hold the actual IDL binary and related files. Included with these files is a `Makefile` that shows how to build IDL from the shareable libraries present in the directory. The link line in this makefile should be used as a starting point when linking your code with Callable IDL—simply omit `main.o` and include your own object files, containing your own main program.
- A more detailed description of the issues involved in compiling and linking your code can be found in this book under “[Compiling Programs That Call IDL](#)” on page 342.

Recommended Reading

There are many books written on the topics discussed in the previous section. The following list includes books we have found to be the most useful over the years in the development and maintenance of IDL. There are thousands of books not mentioned here. Some of them are also excellent. The absence of a book from this list should not be taken as a negative recommendation.

The C Language

Kernighan, Brian W. and Dennis M. Ritchie. *The C Programming Language, Second Edition*. Englewood Cliffs, New Jersey: Prentice Hall, 1988. ISBN 0-13-110370-9. This is the original C language reference, and is essential reading for this subject.

In addition, you should study the vendor supplied documentation for your compiler.

Microsoft Windows

The following books will be useful to anyone building IDL system routines or applications that call IDL in the Microsoft Windows environment.

Petzold, Charles. *Programming Windows, The Definitive Guide to the Win32 API*, Microsoft Press, 1998. ISBN 157231995X (Supersedes: *Programming Windows 95*).

Richter, Jeffrey. *Programming Applications for Microsoft Windows*. Microsoft Press, 1999. ISBN 1572319968 (Supersedes: *Advanced Windows, Third Edition*).

The Microsoft Developer Network (MSDN) supplies essential documentation for programming in the Windows environment. This documentation is part of the Visual C++ environment. More information on the MSDN is available at <http://msdn.microsoft.com>.

Sun Java

Flanagan, David. *Java in a Nutshell, Fourth Edition*, O'Reilly & Associates, March 2002. ISBN 0596002831. This book provides an accelerated introduction to the Java language and key APIs.

In addition, you should study the Java tutorials and documentation provided on the Sun's Java website (<http://www.java.sun.com>).

UNIX

Stevens, W. Richard. *Advanced Programming in the UNIX Environment*. Reading, Massachusetts: Addison Wesley, 1992. ISBN 0-201-56317-7. This is the definitive

reference for UNIX system programmers. It covers all the important UNIX concepts and covers the major UNIX variants in complete detail.

Rochkind, Marc J. *Advanced UNIX Programming (Second Edition)*. Boston: Addison-Wesley Professional, 2004. ISBN 0-13-141154-3. This volume is also extremely well written and does an excellent job of explaining and motivating the fundamental UNIX concepts that underlie the UNIX system calls.


The vendor-supplied documentation and manual pages should be used in combination with the books listed above.

X Windows

The X Windows series by O'Reilly & Associates contains all the information needed to program for the X Window system. There are several volumes—the ones you will need depend on the type of programming you are doing.

Scheifler, Robert W. and James Gettys. *X Window System*. Digital Press. This is purely a reference manual, as opposed to the O'Reilly books which contain a large amount of tutorial as well as reference information. This book is primarily useful for those using XLIB to draw graphics into Motif Draw Widgets and for those who need to understand the base layers of X Windows. Motif programmers may not require this information since Motif hides many of these details.

There are many other X Windows books on the market with varying levels of quality and usefulness. Note that most X Windows books are updated with each version of the system. (X Version 11, Release 6 is the current version at this printing.)



***Part I: Techniques
That Do Not Use IDL's
Internal API***



Chapter 2

Using SPAWN and Pipes

IDL's SPAWN procedure spawns a child process to execute a command or series of commands. General use of [SPAWN](#) is described in detail in the *IDL Reference Guide*. This section describes how to use SPAWN to communicate with the spawned child process using operating system pipes.

By default, calls to the SPAWN procedure cause the IDL process to wait until the child process has finished before continuing, with output sent to the standard output or captured in an IDL variable. Alternatively, IDL can attach a bidirectional pipe to the standard input and output of the child process, and then continue without waiting for the child process to finish. The pipe created in this manner appears in the IDL process as a normal logical file unit.

Once a process has been started in this way, the normal IDL input/output facilities can be used to communicate with it. The ability to use a child process in this manner allows you to solve specialized problems using other languages and to take advantage of existing programs.

In order to start such a process, use the UNIT keyword to SPAWN to specify a named variable in which the logical file unit number will be stored. Once the child process has done its work, use the FREE_LUN procedure to close the pipe and delete the process.

When using a child process in this manner, it is important to understand the following points:

- Closing the file unit causes the child process to be killed. Therefore, do not close the unit until the child process completes its work.
- A pipe is simply a buffer maintained by the operating system with an interface that makes it appear as a file to the programs using it. It has a fixed length and can therefore become completely filled. When this happens, the operating system puts the process that is filling the pipe to sleep until the process at the other end consumes the buffered data. The use of a bidirectional pipe can lead to deadlock situations in which both processes are waiting for the other. This can happen if the parent and child processes do not synchronize their reading and writing activities.
- Most C programs use the input/output facilities provided by the Standard C Library (*stdio*). In situations where IDL and the child process are carrying on a running dialog (as opposed to a single transaction), the normal buffering performed by *stdio* on the output file can cause communications to hang. We recommend calling the *stdio setbuf()* function as the first statement of the child program to eliminate such buffering.

```
(void) setbuf(stdout, (char *) 0);
```

It is important that this statement occur before any output operation is executed; otherwise, it may not have any effect.

Example: Communicating with a Child Process via an Operating System Pipe

The C program shown in the following example (*test_pipe.c*) accepts floating-point values from its standard input and returns their average on the standard output. In actual practice, such a trivial program would never be used from IDL, since it is simpler and more efficient to perform the calculation within IDL itself. The example does, however, serve to illustrate a method by which significant programs can be called from IDL.

In the interest of brevity, some error checking that would normally be included in such a program has been omitted. For example, a real program would need to check

the non-zero return values from `fread(3)` and `fwrite(3)` to ensure that the desired amount of data was actually transferred.

The code for this example can be found in the `spawn` subdirectory of the external directory of the IDL distribution. Instructions for building it can be found in the README file located in that directory.

```

1  #include <stdio.h>
2  #include <stdlib.h>
3  #include <errno.h>
4  #include <string.h>
5
6  int main(int argc, char **argv)
7  {
8      float *data, total = 0.0;
9      char *err_str;
10     int i, n;
11
12     /* Make sure the output is not buffered */
13     setbuf(stdout, (char *) 0);
14
15     /* Find out how many points */
16     if (!fread(&n, sizeof(n), 1, stdin)) goto error;
17
18     /* Get memory for the array */
19     if (!(data = (float *) malloc(n * sizeof(*data)))) goto error;
20
21     /* Read the data */
22     if (!fread(data, sizeof(*data), n, stdin)) goto error;
23
24     /* Calculate the average */
25     for (i=0; i < n; i++) total += data[i];
26     total /= (float) n;
27
28     /* Return the answer */
29     if (!fwrite(&total, sizeof(*data), 1, stdout)) goto error;
30     return 0;                /* Success */
31
32 error:
33     err_str = strerror(errno);
34     if (!err_str) err_str = "<unknown error>";
35     fprintf(stderr, "test_pipe: %s\n", err_str);
36     return 1;                /* Failure */
37 }

```

Table 2-1: `test_pipe.c`

This program performs the following steps:

1. Reads a long integer that tells how many data points to expect, because it is desirable to be able to average an arbitrary number of points.
2. Obtains dynamic memory via the *malloc()* function, and reads the data into it.
3. Calculates the average of the points.
4. Returns the answer as a single floating-point value.

Since the amount of input and output for this program is explicitly known and because it reads all of its input at the beginning and writes all of its results at the end, a deadlock situation cannot occur.

The following IDL statements use *test_pipe* to determine the average of the values 0 to 9:

IDL	<pre> 1 PRO test_pipe 2 3 ; Start test_pipe. The use of the NOSHELL keyword is not 4 ; necessary, but serves to speed up the start-up process. 5 SPAWN, 'test_pipe', UNIT=UNIT, /NOSHELL 6 7 ; Send the number of points followed by the actual data. 8 WRITEU, UNIT, 10L, FINDGEN(10) 9 10 ; Read the answer. 11 READU, UNIT, ANSWER 12 13 ; Announce the result. 14 PRINT, 'Average = ', ANSWER 15 16 ; Close the pipe, delete the child process, and deallocate the 17 ; logical file unit. 18 FREE_LUN, UNIT 19 END </pre>
------------	--

Table 2-2: pro test_pipe

Executing the IDL TEST_PIPE procedure gives the result:

```
Average =          4.50000
```

This mechanism provides the IDL user a simple and efficient way to augment IDL with code written in other languages such as C or Fortran. It is, however, not as efficient as writing the required operation entirely in IDL. The actual cost depends primarily on the amount of data being transferred.

For example, the above example can be performed entirely in IDL using a simple statement such as the following:

```
PRINT, 'Average = ', TOTAL(FINDGEN(10))/10.0
```




Chapter 3

Using

CALL_EXTERNAL

This chapter discusses the following topics:

The CALL_EXTERNAL Function	44	Passing String Data	64
Passing Parameters	54	Passing Array Data	68
Using Auto Glue	56	Passing Structures	70
Basic C Examples	58	Fortran Examples	72
Wrapper Routines	62		

The CALL_EXTERNAL Function

IDL allows you to integrate programs written in other languages with your IDL code, either by calling a compiled function from an IDL program or by linking a compiled function into IDL's internal system routine table:

- The CALL_EXTERNAL function allows you to call external functions (written in C/C++ or Fortran, for example) from your IDL programs. You should be comfortable writing and building programs in the external language being used, but significant knowledge of IDL's internals beyond basic type mapping between the languages is generally not necessary.
- An alternative to CALL_EXTERNAL is to write an IDL system routine and merge it with IDL at runtime. Routines merged in this fashion are added to IDL's internal system routine table and are available in the same manner as IDL built-in routines. This technique is discussed in [Chapter 15, "Adding System Routines"](#). To write a system routine, you will need to understand the IDL internals discussed in later sections of this book.

This chapter covers the basics of using CALL_EXTERNAL from IDL, then discusses platform-specific options for the UNIX and Windows versions of IDL. It can be helpful to refer to the documentation for "[CALL_EXTERNAL](#)" (*IDL Reference Guide*) when reading this material.

The CALL_EXTERNAL function loads and calls routines contained in shareable object libraries. Arguments passed to IDL are passed to this external code, and returned data from the external code is automatically presented as the result from CALL_EXTERNAL as an IDL variable. IDL and the called routine share the same process address space. Because of this, CALL_EXTERNAL avoids the overhead of process creation of the SPAWN routine. In addition, the shareable object library is only loaded the first time it is referenced, saving overhead on subsequent calls.

CALL_EXTERNAL is much easier to use than writing a system routine. Unlike a system routine, however, CALL_EXTERNAL does not check the type or number of parameters. Programming errors in the external routine are likely to result in corrupted data (either in the routine or in IDL) or to cause IDL to crash. See "[Common CALL_EXTERNAL Pitfalls](#)" on page 51 for help in avoiding some of the more common mistakes.

Example Code in the IDL Distribution

This chapter contains examples of CALL_EXTERNAL use. All of the code for these examples, along with additional examples, can be found in the `call_external`

subdirectory of the `external` directory of the IDL distribution. The C language examples use the `MAKE_DLL` procedure, and can therefore be easily run on any platform supported by IDL. To build the sharable library containing the external C code and then run all of the provided examples, execute the following IDL statements:

```
PUSHD, FILEPATH(' ', SUBDIRECTORY=['external', 'call_external', 'C'])
ALL_CALLEXT_EXAMPLES
POPD
```

Additional information on these examples, including details on running the individual examples, can be found in the README file located in that directory.

CALL_EXTERNAL Compared to UNIX Child Process

In many situations, a UNIX IDL user has a choice of using the `SPAWN` procedure to start a child process that executes external code and communicates with IDL via a pipe connecting the two processes. The advantages of this approach are:

- Simplicity.
- The processes do not share address space, and are therefore protected from each other's mistakes.

The advantages of `CALL_EXTERNAL` are:

- IDL and the called routine share the same memory and data space. Although this can be a disadvantage (as noted above) there are times where sharing address space is advantageous. For example, large data can be easily and cheaply shared in this manner.
- `CALL_EXTERNAL` avoids the overhead of process creation and parameter passing.
- The shareable object library containing the called routine is only loaded the first time it is referenced, whereas a `SPAWN`d process must be created for each use of the external code.

Compilation and Linking of External Code

Each operating system requires different compilation and link statements for producing a shareable object suitable for usage with `CALL_EXTERNAL`. This is even true between different implementations of a common operating system family. For example, most UNIX systems require unique options despite their shared heritage. You must consult your system and compiler documentation to find the appropriate options for your system.

The IDL `MAKE_DLL` procedure, documented in the *IDL Reference Guide*, provides a portable high level mechanism for building sharable libraries from code written in the C programming language. In many situations, this procedure can completely handle the task of building sharable libraries to be used with `CALL_EXTERNAL`. `MAKE_DLL` requires that you have a C compiler installed on your system that is compatible with the compiler described by the IDL `!MAKE_DLL` system variable.

The IDL `!MAKE_DLL` system variable is used by the `MAKE_DLL` procedure to construct C compiler and linker commands appropriate for the target platform. If you do not use `MAKE_DLL` to compile and link your code, you may find the contents of `!MAKE_DLL.CC` and `!MAKE_DLL.LD` helpful in determining which options to specify to your compiler and linker, in conjunction with your system and compiler documentation. For the C language, the options in `!MAKE_DLL` should be very close to what you need. For other languages, the `!MAKE_DLL` options should be helpful in determining which options to use, as on most systems, all the language compilers accept similar options.

AUTO_GLUE

As described in “[Passing Parameters](#)” on page 54, `CALL_EXTERNAL` uses the *IDL Portable Calling Convention* to call external code. This convention uses an `(argc, argv)` style interface to allow `CALL_EXTERNAL` to call routines with arbitrary numbers and types of arguments. Such an interface is necessary, because IDL, like any compiled program, cannot generate arbitrary function calls at runtime.

Of course, most C functions are not written to the IDL portable convention. Rather, they are written using the natural form of argument passing used in compiled programs. It is therefore common for IDL programmers to write so-called *glue functions* to match the IDL calling interface to that of the target function. On systems that have a C compiler installed that is compatible with the one described by the IDL `!MAKE_DLL` system variable, the `AUTO_GLUE` keyword to `CALL_EXTERNAL` can be used to instruct IDL to automatically write, compile, and load this glue code on demand, and using a cache to preserve this glue code for future invocations of functions with the same interface.

`AUTO_GLUE` thus allows `CALL_EXTERNAL` to call functions with a natural interface, without requiring the user to write or compile additional code. `AUTO_GLUE` is described in the documentation for “[CALL_EXTERNAL](#)” (*IDL Reference Guide*), as well as in “[Using Auto Glue](#)” on page 56. The examples given in “[Basic C Examples](#)” on page 58 show `CALL_EXTERNAL` used with and without `AUTO_GLUE`.

Input and Output

Input and output actions should be performed within IDL code, using IDL's built-in input/output facilities, or by using `IDL_Message()`. Performing input/output from code external to IDL, especially to the user console or tty (e.g. `stdin` or `stdout`), may generate unexpected results.

Memory Cleanup

IDL has a strict internal policy that it never performs memory cleanup on memory that it did not allocate. This policy is necessary so that external code which allocates memory can use any memory allocation package it desires, and so that there is no confusion about which code is responsible for releasing allocated memory.

Note

The code that allocates memory is always responsible for freeing it. IDL allocates and frees memory for its internal needs, and external code is not allowed to release such memory except through a proper IDL function documented for that purpose. Similarly, IDL will never intentionally free memory that it did not allocate.

As such, IDL does not perform any memory cleanup calls on the values returned from external code called via the `CALL_EXTERNAL` routine. Because of this, any dynamic memory returned to IDL will not be returned to the system, which will result in a memory leak. Users should be aware of this behavior and design their `CALL_EXTERNAL` routines in such a manner as not to return dynamically allocated memory to IDL. The discussion in “[Passing String Data](#)” on page 64 contains an example of doing this with strings.

Memory Access

IDL and your external code share the same address space within the same running program. This means that mistakes common in compiled languages, such as a wild pointer altering memory that it does not own, can cause problems elsewhere. In particular, external code can easily corrupt IDL's data structures and otherwise cause IDL to fail. Authors of such code must be especially careful to guard against such errors.

Argument Data Types

When using `CALL_EXTERNAL` to call external code, IDL passes its arguments to the called code using the data types that were passed to it. It has no way to verify

independently that these types are the actual types expected by the external routine. If the data types passed are not of the types expected by the external code, the results are undefined, and can easily include memory corruption or even crashing of the IDL program.

Warning

You must ensure that the arguments passed to external code are of the exact type expected by that routine. Failure to do so will result in undefined behavior.

Mapping IDL Data Types to External Language Types

When writing external code for use with CALL_EXTERNAL, your code must use data types that are compatible with the C data types used internally by IDL to represent the IDL data types. This mapping is the topic of [Chapter 5, “IDL Internals: Types”](#).

By-Value and By-Reference Arguments

There are two basic forms in which arguments can be passed between functions in compiled languages such as C/C++ and Fortran. To use CALL_EXTERNAL successfully, you should be comfortable with these terms and their meanings. In particular, Fortran programmers are often unaware that Fortran code passes everything by reference, and that C code defaults to passing everything by value. By default, CALL_EXTERNAL passes arguments by reference (unless this behavior is explicitly altered by the use of the ALL_VALUE or VALUE keywords), so no special action is typically required to call Fortran code via CALL_EXTERNAL.

Warning

You must ensure that the arguments passed to external code are passed using the correct method — by value, or by reference. Failure to do so will result in undefined behavior.

Arguments Passed by Value

A copy of the value of the argument is passed to the called routine. Any changes made to such a value by the called routine are local to that routine, and do not change the original value of the variable in the calling routine. C/C++ pass everything by value, but have an explicit *address-of* operator (&) that is used to pass addresses of variables and get by-reference behavior.

Arguments Passed by Reference

The machine address of the argument is passed to the called routine. Any changes made to such a value by the called routine are immediately visible to the caller, because both routines are actually modifying the same memory addresses. Fortran passes everything by reference, but most Fortran implementations support intrinsic operators that allow the programmer control over this (sometimes called %LOC and %VAL, or just LOC and VAL). Consult your compiler documentation for details.

Microsoft Windows Calling Conventions

All operating system/hardware combinations define an inter-routine calling convention. A *calling convention* defines the rules used for passing arguments between routines, and specifies such details as how arguments of different types are passed (*i.e.* in registers or on the system stack) and how and when such arguments are cleaned up.

A stable and efficient calling convention is critical to the stability of an operating system, and can affect most aspects of the system:

- The efficiency of the entire system depends on the efficiency of the core calling convention.
- Backwards compatibility, and thus the longevity of binary software written for the platform depends on the stability of the calling convention.
- Calling routines from different languages within a single program depends on all the language compilers adhering to the same calling convention. Even within the same language, the ability to mix code compiled by different compilers requires those compilers to adhere to the same conventions. For example, at the time of this writing, the C++ language standard lacks an Application Binary Interface (ABI) that can be targeted by all C++ compilers. This can lead to situations in which the same compiler must be used to build all of the code within a given program.

Microsoft Windows is unique among the platforms supported by IDL in that it has two distinct calling conventions in common use, whereas other systems define a single convention. On single-convention systems, the calling convention is unimportant to application programmers, and of concern only to hardware designers and the authors of compilers, and operating systems. On a multiple convention system, application programmers sometimes need to be aware of the issue, and ensure that their code is compiled to use the proper convention and that calls to that code use the same convention. The Microsoft Calling Conventions are:

STDCALL

STDCALL is the calling convention used by the majority of the Windows operating system API. In a STDCALL call, the calling routine places the arguments in the proper registers and/or stack locations, and the called routine is responsible for cleaning them up and unwinding the stack.

CDECL

CDECL is the calling convention used by C/C++ code by default. This default can be changed via compiler switches, declspec declarations, or #pragmas. With CDECL, the caller is responsible for both setup and cleanup of the arguments. CDECL is able to call functions with variable numbers of arguments (*varargs* functions) because the caller knows the actual number of arguments passed at runtime, whereas STDCALL cannot call such functions. This is because the STDARGS routine cannot know efficiently at compile time how many arguments it will be passed at runtime in these situations.

The inconvenience of having two distinct and incompatible calling conventions is usually minor, because the header files that define functions for C/C++ programs include the necessary definitions such that the compiler knows to generate the proper code to call them and the programmer is not required to be aware of the issue. However, CALL_EXTERNAL does have a problem: Unlike a C/C++ program, IDL determines how to call a function solely by the arguments passed to CALL_EXTERNAL, and not from a header file.

IDL therefore has no way to know how your external code was compiled. It uses the STDARG convention by default, and the CDECL keyword can be used to change the default. CALL_EXTERNAL therefore relies on the IDL user to tell it which convention to use. If IDL calls your code using the correct convention, it will work correctly. If it calls using the wrong convention, the results are undefined, including memory corruption and possible crashing of the IDL program.

Warning

The default calling convention for CALL_EXTERNAL is STDCALL, whereas the default convention for the Microsoft C compiler is CDECL. Hence, Windows users must usually specify the CDECL keyword when calling such code from IDL. Non-Windows versions of IDL ignore the CDECL keyword, so it is safe to always include it in cross platform code.

Here is what happens when external code is called via the wrong calling convention:

- If a `STDARG` call is made to a `CDECL` function, the caller places the arguments in the proper registers/stack locations, and relies on the called routine to cleanup and unwind the stack. The called routine, however, does not do these things because it is a `CDECL` routine. Hence, cleanup does not happen.
- If a `CDECL` call is made to a `STDARG` function, the caller places the arguments in the proper register/stack locations. The called routine cleans up on exit, and then the caller cleans up again.

Either combination is bad, and can corrupt or kill the program. Sometimes this happens, and sometimes it doesn't, so the results can be random and mysterious to programmers who are not aware of the issue.

Note

When the wrong calling convention is used, it is common for the process stack to become confused. A “smashed stack” visible from the C debugger following a `CALL_EXTERNAL` is usually indicative of having used the wrong calling convention.

Common `CALL_EXTERNAL` Pitfalls

Following are a list of common errors and mistakes commonly seen when using `CALL_EXTERNAL`.

- The number of arguments and their types, as passed to `CALL_EXTERNAL`, must be the exact types expected by the external routine. In particular, it is common for programmers to forget that the default IDL integer is a 16-bit value and that most C compilers define the `int` type as being a 32-bit value. You should be careful to use IDL `LONG` integers, which are 32-bit, in such cases. See [“Argument Data Types”](#) on page 47 for additional details.
- Passing data using the wrong form: Using by-value to pass an argument to a function expecting it by-reference, or the reverse. See [“By-Value and By-Reference Arguments”](#) on page 48 for additional details.
- Under Microsoft Windows, using the incorrect calling convention for a given external function. See [“Microsoft Windows Calling Conventions”](#) on page 49 for additional details.

- Failure to understand that IDL uses IDL_STRING descriptors to represent strings, and not just a C style NULL terminated string. Passing a string value by reference passes the address of the IDL_STRING descriptor to the external code. See [Chapter 8, “IDL Internals: String Processing”](#) for additional details.
- Attempting to make IDL data structures use memory allocated by external code rather than using the proper IDL API for creating such data structures. For instance, attempting to give an IDL_STRING descriptor a different value by using C `malloc()` to allocate memory for the string and then storing the address of that memory in the IDL_STRING descriptor is not supported, and can easily crash or corrupt IDL. Although IDL uses `malloc()/free()` internally on most platforms, you should be aware that this is not part of IDL’s public interface, and that we can change this at any time and without notice. Even on platforms where IDL does use these functions, its use of them is not directly compatible with similar calls made by external code because IDL allocates additional memory for bookkeeping that is generally not present in memory allocations from other sources. See [Chapter 8, “IDL Internals: String Processing”](#) for information on changing the value of an IDL_STRING descriptor using supported IDL interfaces. See [Chapter 3, “Memory Cleanup”](#) for more on memory allocation and cleanup.
- IDL is written in the C language, and when IDL starts, any necessary runtime initialization code required by C programs is automatically executed by the system before the IDL `main()` function is called. Hence, calling C code from IDL usually does not require additional runtime initialization. However, when calling external code written in languages other than C, you may find that your code does not run properly unless you arrange for the necessary runtime support for that language to run first. Such details are highly system specific, and you must refer to your system and compiler documentation for details. Code that is largely computational rarely encounters this issue. It is more common for code that performs Input/Output directly.
- Programming errors in the external code. It is easy to make mistakes in compiled languages that have bad global consequences for unrelated code within the same program. For example, a wild memory pointer can lead to the corruption of unrelated data. If you are lucky, such an error will immediately kill your program, making it easy to locate and fix. Less fortunate is the situation in which the program dies much later in a seemingly unrelated part of the program. Finding such problems can be difficult and time consuming. When IDL crashes following a call to external code, an error in the external code or in the call to CALL_EXTERNAL is the cause in the vast majority of cases.

- Some compilers and operating systems have a convention of adding leading or trailing underscore characters to the names of functions they compile. These conventions are platform specific, and as they are of interest only to system linker and compiler authors, not generally well documented. This is usually transparent to the user, but can sometimes be an issue with inter language function calls. If you find that a function you expect to call from a library is not being found by CALL_EXTERNAL, and the obvious checks do not uncover the error (usually a simple misspelling), this might be the cause. Under UNIX, the **nm** command can be helpful in diagnosing such problems.
- C++ compilers use a technique commonly called *name munging* to encode the types of method arguments and return values into the name of the routine as written to their binary object files. Such names often have only a passing resemblance to the name seen by the C++ programmer in their source code. IDL can only call C++ code that has C linkage, as discussed in “C++” on page 25. C linkage code does not use name munging.
- When calling external code written in other languages, there are sometimes platform and language specific hidden arguments that must be explicitly supplied. Such arguments are usually provided by the compiler when you work strictly within the target language, but become visible in inter-language calls. An example of this can be found in “[Hidden Arguments](#)” on page 73. In this example, the Fortran compiler provides an extra hidden length argument when a NULL terminated string is passed to a function.

Passing Parameters

IDL calls routines within a shareable library using the *IDL portable calling convention*, in which the routine is passed two arguments:

argc

A count of the number of arguments being passed to the routine

argv

An array of **argc** memory pointers, which are the addresses of the arguments (by reference) or the actual value of the argument (by value) depending on the types of arguments passed to CALL_EXTERNAL and the setting of the VALUE keyword to that function. You should note that while all types of data can be passed by reference, there are limitations on data types that can be passed by value, as described in the documentation for “CALL_EXTERNAL” (*IDL Reference Guide*).

The CALL_EXTERNAL portable convention is necessary because IDL, like any program written in a compiled language, cannot generate arbitrary function calls at runtime. Only calls to interfaces that were known to it when it was compiled are possible. Naturally, most existing C functions are not written to use this interface. Calling such functions typically requires IDL users to write *glue functions*, the sole purpose of which is to be called by CALL_EXTERNAL with the portable convention, and then to take the arguments and pass them to the real target function using the natural interface for that function. The AUTO_GLUE keyword to CALL_EXTERNAL can be used to generate, compile, and load such glue routines automatically and on demand, without requiring user intervention. Auto Glue is described in “Using Auto Glue” on page 56. AUTO_GLUE does not eliminate the need for, or use of, the portable convention, but it can relieve the IDL user of the requirement to handle it explicitly. The end result is that calling existing function interfaces is easier to do, and less error prone.

Routines called by CALL_EXTERNAL with the portable convention are defined with a prototype similar to the following:

```
return_type example(int argc; void *argv[])
```

where *return_type* is one of the data types which CALL_EXTERNAL can return. If this *return_type* is not IDL_LONG, a keyword must be used in the CALL_EXTERNAL call to indicate the actual type of the result.

The parameter `argc` gives the number of arguments passed to the external routine by `CALL_EXTERNAL` in the `argv` array, while `argv` is an array containing the arguments. Arguments are passed either by value or by reference. Those passed by value are copied directly into the `argv` array, with the exception of scalar strings, which place a pointer to a null-terminated string in `argv[i]`. All arrays are passed by reference. Scalar items passed by reference (the default) place a pointer to the datum in `argv[i]`. Strings and string arrays passed by reference place a pointer to an `IDL_STRING` structure in `argv[i]`. This structure is defined as follows:

```
typedef struct {
    IDL_STRING_SLEN_T slen; /* Length of string */
    short stype; /* type of string: (0) static, (!0) dynamic */
    char *s; /* Addr of string, invalid if slen == 0. */
} IDL_STRING;
```

See “[CALL_EXTERNAL](#)” (*IDL Reference Guide*) for additional details about passing parameters by value.

It is important to note that IDL integer variables correspond to a 16-bit integer (a C *signed* short integer). For example, an integer variable could be defined in an IDL routine as follows:

```
IDL> A = 5 ;default type of integer, not LONG
```

The variable could then be passed by reference in a `CALL_EXTERNAL` call. The declaration and cast statement in the called C routine should be:

```
short *a;
a = (short *) argv[0];
```

or

```
IDL_INT *a;
a = (IDL_INT *) argv[0];
```

`IDL_INT` corresponds to a C short (16-bit integer), so either form is correct. The corresponding type in Fortran would be `INTEGER*2`.

Using Auto Glue

Users of CALL_EXTERNAL frequently write small functions with the sole purpose of matching the CALL_EXTERNAL portable calling convention with its (argc, argv) interface to the actual interface presented by some existing function that they wish to call. Such functions are often called *glue functions*.

It quickly becomes obvious to anyone who has written a few glue functions that there isn't much to them, and that producing such functions is a purely mechanical operation. As you read the examples in this chapter, you will see many such functions, and will notice that they are all essentially the same. Further examination should serve to convince you that IDL already has all of the information, in the form of the arguments and keywords specified to the CALL_EXTERNAL function, to generate such functions without requiring human intervention. Examining the CALL_EXTERNAL routine's interface, we see that:

- the number and types of arguments to the CALL_EXTERNAL function provide the same information about the arguments for the target external function;
- the VALUE keyword, and CALL_EXTERNAL's built in rules for deciding whether or not to pass arguments by value or by reference determine how the arguments should be passed;
- in the case of Microsoft Windows, the CDECL keyword tells it which system calling convention to employ;
- keywords to CALL_EXTERNAL determine the result type.

Furthermore, other than the actual name of the user function being called, these glue functions are generic in the sense that they could be used to call any function that accepted arguments of the same types and produce a result of the same type.

The AUTO_GLUE keyword to CALL_EXTERNAL exploits these facts to allow you to call functions with natural interfaces, without the need to write, compile, and load a glue function to do the job. The sole requirement is that your system must have a C compiler installed that is compatible with the compiler described by the IDL !MAKE_DLL system variable. This is almost always the case if you are interested in calling external code, since a compiler is necessary to compile such code.

AUTO_GLUE automatically writes the C code for the glue function, uses the MAKE_DLL procedure to build a sharable library containing it, loads that library, and then calls the glue function, passing it a pointer to the target function and all of its arguments. It maintains a cache of glue functions that have been built previously, and never builds the same glue function more than once. From the user perspective, there is a slight pause the first time a given glue function is used. In that brief moment, AUTO_GLUE performs the steps described above, and then makes the call to the user function. All of this happens transparently to the IDL user — no user interaction is required, and no output is produced by the process. Subsequent calls to the same glue function happen instantaneously, as IDL loads the existing glue function from the MAKE_DLL cache without rebuilding it. In principle, it is similar to the way IDL automatically compiles IDL language programs on demand, only with C code instead of IDL code.

See “[CALL_EXTERNAL](#)” (*IDL Reference Guide*) for additional details about how AUTO_GLUE works, and the options for controlling its use.

Generating Glue Without Executing It

AUTO_GLUE is the preferred option for most calls to functions with natural interfaces, due to its simplicity and ease of use. However, you might find yourself in a situation where you would like your glue functions to be automatically generated, but wish to simply get the resulting C code so that you can modify it or incorporate it into a larger library. For example, you might have a large library of IDL specific code, and wish to give it all IDL callable interfaces without requiring the overhead of AUTO_GLUE for all of them.

The WRITE_WRAPPER keyword to CALL_EXTERNAL can be used to produce such code without compiling or using the results. See “[CALL_EXTERNAL](#)” (*IDL Reference Guide*) for additional information on this keyword.

Basic C Examples

All of the code for the examples in this section can be found in the `/external/call_external/C` subdirectory of the IDL distribution. Please read the README file in that directory for details on how to run the examples. In many cases, the files in that directory go into more detail, and are more fully commented than the versions shown here. Also, the examples provide IDL wrapper routines that perform the necessary CALL_EXTERNAL calls, while the examples shown here use CALL_EXTERNAL directly in order to explain how it is used. It is worth reading the contents of the `.c` and IDL `.pro` files in that directory in addition to reading the code shown here.

Example: Passing Parameters by Reference to IDL

The following routine, found in `simple_vars.c`, accepts several of IDL's basic data types as arguments. The parameters are passed in by reference and the new squared values of the numbers are passed back to IDL. This is implemented as a function with a natural C interface, and a second glue routine that implements the

IDL portable convention, using the one with the natural interface to do the actual work.

C

```

1  #include <stdio.h>
2  #include "idl_export.h"          /* IDL external definitions */
3
4  int simple_vars_natural(char *byte_var, short *short_var,
5                          IDL_LONG *long_var, float *float_var,
6                          double *double_var)
7  {
8      /* Square each variable. */
9      *byte_var    *= *byte_var;
10     *short_var   *= *short_var;
11     *long_var    *= *long_var;
12     *float_var   *= *float_var;
13     *double_var  *= *double_var;
14
15     return 1;
16 }
17
18 int simple_vars(int argc, void* argv[])
19 {
20     /* Insure that the correct number of arguments were passed in */
21     if(argc != 5) return 0;
22
23     return simple_vars_natural((char *) argv[0], (short *) argv[1],
24                               (IDL_LONG *) argv[2], (float *) argv[3],
25                               (double *) argv[4]);
26 }

```

Table 3-1: Passing Parameters by Reference to IDL — simple_vars.c

The IDL statements necessary to call the `simple_vars()` function from IDL can be written:

```

B=2B & I=3 & L=3L & F=0.0 & D=0.0D
R = CALL_EXTERNAL(GET_CALLEXT_EXLIB(), 'simple_vars', $
                 b,i,l,f,d, /CDECL)

```

Note

`GET_CALLEXT_EXLIB()` is a function provided with the `CALL_EXTERNAL` examples; it builds the necessary sharable library of external C code and returns the path to the library as its result.

Using the AUTO_GLUE keyword to CALL_EXTERNAL, you can call the function with the natural C interface directly:

```
B=2B & I=3 & L=3L & F=0.0 & D=0.0D
R = CALL_EXTERNAL(GET_CALLEXT_EXLIB(), 'simple_vars_natural', $
                  b,i,l,f,d, /CDECL, /AUTO_GLUE)
```

Example: Calling a C Routine to Perform Computation

The following example demonstrates an external function that returns the sum of a floating point array. It is similar in function to the TOTAL function in IDL. The code for this example is found in the file `sum_array.c` in the IDL distribution. As with the previous example, this function is implemented by a function that has a natural C interface, and a second glue function is provided that matches the IDL portable calling convention to the natural interface:

```

1  #include <stdio.h>
2  #include "idl_export.h"
3
4  float sum_array_natural(float *fp, IDL_LONG n)
5  {
6      float s = 0.0;
7
8      while (n--) s += *fp++;
9      return(s);
10 }
11
12 float sum_array(int argc, void *argv[])
13 {
14     return sum_array_natural((float *) argv[0], (IDL_LONG) argv[1]);
15 }
```

Table 3-2: Calling a C routine — example.c

The IDL statements necessary to call the `sum_array()` function from IDL can be written:

```
X = FINDGEN(10)
S = CALL_EXTERNAL(GET_CALLEXT_EXLIB(), 'sum_array'$
                  X, N_ELEMENTS(X), VALUE=[0,1], /F_VALUE, /CDECL)
```

Note

`GET_CALLEXT_EXLIB()` is a function provided with the `CALL_EXTERNAL` examples; it builds the necessary sharable library of external C code and returns the path to the library as its result.

Using the `AUTO_GLUE` keyword, you can call the function with the natural C interface directly:

```
X = FINDGEN(10)
S = CALL_EXTERNAL(GET_CALLEXT_EXLIB(), 'sum_array_natural'$
                  X, N_ELEMENTS(X), VALUE=[0,1], /F_VALUE,/CDECL,$
                  /AUTO_GLUE)
```

In this example, `sum_array` and `sum_array_natural` are the names of the entry points for the external functions, and `X` and `N_ELEMENTS(X)` are passed to the called routine as parameters. The `F_VALUE` keyword specifies that the returned value is a floating-point number rather than an `IDL_LONG`.

Wrapper Routines

CALL_EXTERNAL routines are very sensitive to the number and type of the arguments they receive. Calling a CALL_EXTERNAL routine with the wrong number of arguments or with arguments of the wrong type can cause IDL to crash. For this reason, it is a good practice to provide an IDL *wrapper routine* that is used to make the actual CALL_EXTERNAL call. The job of this wrapper, which is written in the IDL language, is to ensure that the arguments that are passed to the external code are always of the correct number and type. The following IDL procedure is the wrapper used in the **simple_vars()** example of the previous section (“[Example: Passing Parameters by Reference to IDL](#)” on page 58).

Example Code

This file, `simple_vars.pro`, is located in the `external/call_external/C` subdirectory of the IDL installation directory.

IDL	<pre> 1 PRO SIMPLE_VARS, b, i, l, f, d, AUTO_GLUE=auto_glue, DEBUG=debug, \$ 2 VERBOSE=verbose 3 if ~ (KEYWORD_SET(debug)) THEN ON_ERROR,2 4 5 ; Type checking: Any missing (undefined) arguments will be set 6 ; to a default value. All arguments will be forced to a scalar 7 ; of the appropriate type, which may cause errors to be thrown 8 ; if structures are passed in. Local variables are used so that 9 ; the values and types of the user supplied arguments don't change. 10 b_l = (SIZE(b,/TYPE) EQ 0) ? 2b : byte(b[0]) 11 i_l = (SIZE(i,/TYPE) EQ 0) ? 3 : fix(i[0]) 12 l_l = (SIZE(l,/TYPE) EQ 0) ? 4L : long(l[0]) 13 f_l = (SIZE(f,/TYPE) EQ 0) ? 5.0 : float(f[0]) 14 d_l = (SIZE(d,/TYPE) EQ 0) ? 6.0D : double(d[0]) 15 16 PRINT, 'Calling simple_vars with the following arguments:' 17 HELP, b_l, i_l, l_l, f_l, d_l 18 func = keyword_set(auto_glue) ? 'simple_vars_natural' : 'simple_vars' 19 IF (CALL_EXTERNAL(GET_CALLEXT_EXLIB(VERBOSE=verbose), func, \$ 20 b_l, i_l, l_l, f_l, d_l, /CDECL, \$ 21 AUTO_GLUE=auto_glue, VERBOSE=verbose, \$ 22 SHOW_ALL_OUTPUT=verbose) EQ 1) then BEGIN 23 PRINT, 'After calling simple_vars:' 24 HELP, b_l, i_l, l_l, f_l, d_l 25 ENDIF ELSE MESSAGE, 'External call to simple_vars failed' 26 END </pre>
-----	--

Table 3-3: Wrapper Routine — `simple_vars.pro`

The routine `simple_vars.pro` uses the system routine `SIZE()` to examine the arguments that are passed in by the user to the `simple_vars` routine. If one of the arguments is undefined, a default value will be used in the call to the external routine. Otherwise, the argument will be converted to a scalar of the appropriate type.

Note

`GET_CALLEXT_EXLIB()` is a function provided with the `CALL_EXTERNAL` examples; it builds the necessary sharable library of external C code and returns the path to the library as its result.

Passing String Data

IDL represents strings internally as IDL_STRING descriptors. For more information about IDL_STRING, see [Chapter 7, “IDL Internals: Variables”](#) and [Chapter 8, “IDL Internals: String Processing”](#). These descriptors are defined in the C language as:

```
typedef struct {
    IDL_STRING_SLEN_T slen;
    unsigned short stype;
    char *s;
} IDL_STRING;
```

To pass a string by reference, IDL passes the address of its IDL_STRING descriptor. To pass a string by value the string pointer (the `s` field of the descriptor) is passed. Programmers should be aware of the following when manipulating IDL strings:

- Called code should treat the information in the passed IDL_STRING descriptor and the string itself as read-only, and should not modify these values.
- The `slen` field contains the length of the string without including the NULL termination that is required at the end of all C strings.
- The `stype` field is used internally by IDL to keep track of how the memory for the string was obtained, and should be ignored by CALL_EXTERNAL users.
- `s` is the pointer to the actual C string represented by the descriptor. If the string is NULL, IDL represents it as a NULL (0) pointer, not as a pointer to an empty null terminated string. Hence, called code that expects a string pointer should check for a NULL pointer before dereferencing it.
- You must use the functions discussed in [Chapter 8, “IDL Internals: String Processing”](#) to allocate the memory for an IDL_STRING. Attempting to do this directly by allocating dynamic memory and assigning it to the IDL_STRING descriptor is a common pitfall, as discussed in [“Common CALL_EXTERNAL Pitfalls”](#) on page 51.

Returning a String Value

When returning a string value, a function must allocate the memory used to hold it. On return, IDL will copy this string. You can use a static buffer or dynamic memory, but do not return the address of an automatic (stack-based) variable.

Note

IDL will not free dynamically-allocated memory for this use.

Example

The following routine, found in `string_array.c`, demonstrates how to handle string variables in external code. This routine takes a string or array of strings as input and returns a copy of the longest string that it received. It is important to note that this routine uses a static `char` array as its return value, which avoids the possibility of a memory leak, but which must be long enough to handle the longest string required by the application. This is implemented as a function with a natural C interface, and a second glue routine that implements the IDL portable convention, using the one with the natural interface to do the actual work:

```

1  #include <stdio.h>
2  #include <string.h>
3  #include "idl_export.h"
4  /*
5   * IDL_STRING is declared in idl_export.h like this:
6   *   typedef struct {
7   *     IDL_STRING_SLEN_T slen;           Length of string, 0 for null
8   *     short stype;                     Type of string, static or dynamic
9   *     char *s;                          Address of string
10  *   } IDL_STRING;
11  * However, you should rely on the definition in idl_export.h instead
12  * of declaring your own string structure.
13  */
14
15  char* string_array_natural(IDL_STRING *str_descr, IDL_LONG n)
16  {
17      /*
18       * IDL will make a copy of the string that is returned (if it is
19       * not NULL). One way to avoid a memory leak is therefore to return
20       * a pointer to a static buffer containing a null terminated string.
21       * IDL will copy the contents of the buffer and drop the reference
22       * to our buffer immediately on return.
23       */
24      #define MAX_OUT_LEN 511                /* truncate any string
25                                             longer than this */
26      static char result[MAX_OUT_LEN+1];    /* leave a space for a '\0'
27                                             on the longest string */
28      int max_index;                        /* index of longest string */
29      int max_sofar;                        /* length of longest string*/
30      int i;
31
32      /* Check the size of the array passed in. n should be > 0.*/
33      if (n < 1) return (char *) 0;
34      max_index = 0;
35      max_sofar = 0;
36      for(i=0; i < n; i++) {
37          if (str_descr[i].slen > max_sofar) {
38              max_index = i;
39              max_sofar = str_descr[i].slen;
40          }
41      }

```

C

Figure 3-1: Handling String Variables in External Code — `string_array.c`

```
42  /*
43  * If all strings in the array are empty, the longest
44  * will still be a NULL string.
45  */
46  if (str_descr[max_index].s == NULL) return (char *) 0;
47
48  /*
49  * Copy the longest string into the buffer, up to MAX_OUT_LEN
50  * characters.
51  * Explicitly store a NULL byte in the last byte of the buffer,
52  * because strncpy() does not NULL terminate if the string copied
53  * is truncated.
54  */
55  strncpy(result, str_descr[max_index].s, MAX_OUT_LEN);
56  result[sizeof(result)-1] = '\0';
57  return(result);
58 #undef MAX_OUT_LEN
59 }
60
61 char* string_array(int argc, void* argv[])
62 {
63     /*
64     * Make sure there are the correct # of arguments.
65     * IDL will convert the NULL into an empty string ('').
66     */
67     if (argc != 2) return (char *) NULL;
68     return string_array_natural((IDL_STRING *) argv[0], (IDL_LONG) argv[1]);
69 }
```

Figure 3-1: Handling String Variables in External Code — *string_array.c* (Continued)

Passing Array Data

When you pass an IDL array into a CALL_EXTERNAL routine, that routine gets a pointer to the first memory location in the array. In order to perform any processing on the array, an external routine needs more information—such as the array's size and number of dimensions. With CALL_EXTERNAL, you will need to pass this information explicitly as additional arguments to the routine.

In order to handle multi-dimensional arrays, C needs to know the size of the array at compile time. In most cases, this means that you will need to treat multi-dimensional arrays passed in from IDL as one dimensional arrays. However, you can still build your own indices to access an array as if it had more than one dimension in C. For example, the IDL array index:

```
array[x,y]
```

could be represented in a CALL_EXTERNAL routine as:

```
array_ptr[x + x_size*y];
```

The following routine, found in `sum_2d_array.c`, calculates the sum of a subsection of a two dimensional array. This is implemented as a function with a natural C interface, and a second glue routine that implements the IDL portable convention, using the one with the natural interface to do the actual work:

```

1  #include <stdio.h>
2  #include "idl_export.h"
3  double sum_2d_array_natural(double *arr, IDL_LONG x_start, IDL_LONG x_end,
4                             IDL_LONG x_size, IDL_LONG y_start,
5                             IDL_LONG y_end, IDL_LONG y_size)
6      /* Since we didn't know the dimensions of the array at compile time, we
7       * must treat the input array as if it were a one dimensional vector. */
8      IDL_LONG x,y;
9      double result = 0.0;
10
11     /* Make sure that we don't go outside the array. strictly speaking, this
12      * is redundant since identical checks are performed in the IDL wrapper
13      * routine. IDL_MIN() and IDL_MAX() are macros from idl_export.h */
14     x_start = IDL_MAX(x_start,0);
15     y_start = IDL_MAX(y_start,0);
16     x_end = IDL_MIN(x_end,x_size-1);
17     y_end = IDL_MIN(y_end,y_size-1);
18
19     /* loop through the subsection */
20     for (y = y_start;y <= y_end;y++)
21         for (x = x_start;x <= x_end;x++)
22             result += arr[x + y*x_size]; /* build the 2d index: arr[x,y] */
23     return result;
24 }
25
26 double sum_2d_array(int argc,void* argv[])
27 {
28     if (argc != 7) return 0.0;
29     return sum_2d_array_natural((double *) argv[0], (IDL_LONG) argv[1],
30                                (IDL_LONG) argv[2], (IDL_LONG) argv[3],
31                                (IDL_LONG) argv[4], (IDL_LONG) argv[5],
32                                (IDL_LONG) argv[6]);
33 }

```

Table 3-4: Adding the Elements of a 2D IDL Array — sum_2d_array.c

The IDL system routine interface provides much more support for the manipulation of IDL array variables. See [Chapter 15, “Adding System Routines”](#) for more information.

Passing Structures

IDL structure variables are stored in memory in the same layout that C uses. This makes it possible to pass IDL structure variables into CALL_EXTERNAL routines, as long as the layout of the IDL structure is known. To access an IDL structure from an external routine, you must create a C structure definition that has the exact same layout as the IDL structure you want to process.

For example, for an IDL structure defined as follows:

```
s = {ASTRUCTURE, zero:0B, one:0L, two:0.0, three:0D, four: intarr(2)}
```

the corresponding C structure would look like the following:

```
typedef struct {
    unsigned char zero;
    IDL_LONG one;
    float two;
    double three;
    short four[2];
} ASTRUCTURE;
```

Then, cast the pointer from *argv* to the structure type, as follows:

```
ASTRUCTURE* mystructure;
mystructure = (ASTRUCTURE*) argv[0];
```

The following routine, found in `incr_struct.c`, increments each field of an IDL structure of type ASTRUCTURE. This is implemented as a function with a natural C interface, and a second glue routine that implements the IDL portable convention, using the one with the natural interface to do the actual work:

```

1  #include <stdio.h>
2  #include "idl_export.h"
3
4  /*
5   * C definition for the structure that this routine accepts. The
6   * corresponding IDL structure definition would look like this:
7   *     s = {zero:0B,one:0L,two:0.,three:0D,four: intarr(2)}
8   */
9  typedef struct {
10     unsigned char zero;
11     IDL_LONG one;
12     float two;
13     double three;
14     short four[2];
15 } ASTRUCTURE;
16
17 int incr_struct_natural(ASTRUCTURE *mystructure, IDL_LONG n)
18 {
19     /* for each structure in the array, increment every field */
20     for (; n--; mystructure++) {
21         mystructure->zero++;
22         mystructure->one++;
23         mystructure->two++;
24         mystructure->three++;
25         mystructure->four[0]++;
26         mystructure->four[1]++;
27     }
28
29     return 1;
30 }
31 int incr_struct(int argc, void *argv[])
32 {
33     if (argc != 2) return 0;
34     return incr_struct_natural((ASTRUCTURE*) argv[0], (IDL_LONG)
35     argv[1]);
36 }

```

Table 3-5: Accessing an IDL Structure from a C Routine — incr_struct.c

It is not possible to access structures with arbitrary definitions using the CALL_EXTERNAL interface. The system routine interface, discussed in [Chapter 15, “Adding System Routines”](#), does provide support for determining the layout of a structure at runtime.

Fortran Examples

Example: Calling a Fortran Routine Using a C Interface Routine

Calling Fortran is similar to calling C, with the significant difference that Fortran code expects all arguments to be passed by reference and not by value (the C default). This means that the *address* of the argument is passed rather than the argument itself. This issue is discussed in “[By-Value and By-Reference Arguments](#)” on page 48.

A C interface routine can easily extract the addresses of the arguments from the `argv` array and pass them to the actual routine which will compute the sum. The arguments `f`, `n`, and `s` are pointers that are being passed by value. Fortran expects all arguments to be passed by reference — that is, it expects all arguments to be addresses. If C passes a pointer (an address) by value, Fortran will interpret it correctly as the address of an argument. The following code segments illustrate this. The `example_c2f.c` file contains the C interface routine, which would be compiled as illustrated above. The `example.f` file contains the Fortran routine that actually sums the array.

In these examples, we assume that the routines are being compiled under Sun Solaris. The object name of the Fortran subroutine will be `sum_array1_` to match the output of the Solaris Fortran compiler. The following are the contents of `example_c2f.c` and `example.f`:

C	1	<code>#include <stdio.h></code>
	2	
	3	<code>void sum_array(int argc, void *argv[])</code>
	4	<code>{</code>
	5	<code>extern void sum_array1_(); /* Fortran routine */</code>
	6	<code>int *n;</code>
	7	<code>float *s, *f;</code>
	8	
	9	<code>f = (float *) argv[0]; /* Array ptr */</code>
	10	<code>n = (int *) argv[1]; /* Get # of elements */</code>
	11	<code>s = (float *) argv[2]; /* Pass back result a parameter */</code>
	12	
	13	<code>sum_array1_(f, n, s); /* Compute sum */</code>
	14	<code>}</code>

Table 3-6: C Wrapper Used to Call Fortran Code (`example_c2f.c`)

f77

```

1  c This subroutine is called by SUM_ARRAY and has no IDL-specific code.
2  c
3  SUBROUTINE sumarray1(array, n, sum)
4  INTEGER*4 n
5  REAL*4 array(n), sum
6
7  sum=0.0
8  DO i=1,n
9  sum = sum + array(i)
10 PRINT *, sum, array(i)
11 ENDDO
12
13 RETURN
14 END

```

Table 3-7: Fortran Code Called from IDL via C Wrapper (example.f)

This example is compiled and linked in a manner similar to that used in the C example above. For more information on compiling and linking on your platform, see the README file contained in the `external/call_external/Fortran` subdirectory of the IDL distribution. This directory also contains a makefile, which builds this example on UNIX platforms. To call the example program from within IDL:

```

;Make an array.
X = FINDGEN(10)
;A floating result
SUM = 0.0
S = CALL_EXTERNAL('example.so', $
    'sum_array', X, N_ELEMENTS(X), sum)

```

In this example, `example.so` is the name of the sharable image file, `sum_array` is the name of the entry point, and `X` and `N_ELEMENTS(X)` are passed to the called routine as parameters. The returned value is contained in the variable `sum`.

Hidden Arguments

When passing C null-terminated character strings into a Fortran routine, the C function should also pass in the string length. This extra parameter is added to the end of the Fortran routine call in the C function, but does not explicitly appear in the Fortran routine.

For example, in C:

```

char * str1= 'IDL';
char * str2= 'ITT';

```

```

int len1=3;
int len2=3;
double data, info;
/* Call a Fortran sub-routine named example1 */
example1_(str1, data, str2, info, len1, len2)

```

In Fortran:

```

SUBROUTINE EXAMPLE1(STR1, DATA, STR2, INFO)
CHARACTER*(*)STR1, STR2
DOUBLE PRECISIONDATA, INFO

```

Example: Calling a Fortran Routine Using a Fortran Interface Routine

Calling Fortran is similar to calling C, with the significant difference that Fortran expects all arguments to be passed by reference. This means that the *address* of the argument is passed rather than the argument itself. See “[By-Value and By-Reference Arguments](#)” on page 48 for more on this subject.

A Fortran interface routine can be written to extract the addresses of the arguments from the `argv` array and pass them to the actual routine which will compute the sum. Passing the contents of each `argv` element by value has the same effect as converting the parameter to a normal Fortran parameter.

This method uses the OpenVMS Extensions to Fortran, `%LOC` and `%VAL`.

Some Fortran compilers may not support these extensions. If your compiler does not, use the method discussed in the previous section for calling Fortran with a C interface routine.

The contents of the file `example1.f` are shown in the following figure. This example is compiled, linked, and called in a manner similar to that used in the C example above. For more information on compiling and linking on your platform, see the README file contained in the `external/fortran` subdirectory of the IDL distribution. This directory also contains a makefile, which builds this example on UNIX platforms.

Note

This example is written to run under a 32-bit operating system. To run the example under a 64-bit operating system would require modifications; most notably, to declare `argv` as `INTEGER*8` rather than `INTEGER*4`.

```

1  SUBROUTINE SUM_ARRAY(argc, argv) !Called by IDL
2  INTEGER*4 argc, argv(*)         !Argc and Argv are integers
3
4  j = LOC(argc)                   !Obtains the number of arguments (argc)
5                                  !Because argc is passed by VALUE.
6
7  c Call subroutine SUM_ARRAY1, converting the IDL parameters
8  c to standard Fortran, passed by reference arguments:
9
10 CALL SUM_ARRAY1(%VAL(argv(1)), %VAL(argv(2)), %VAL(argv(3)))
11 RETURN
12 END
f77
13
14 c This subroutine is called by SUM_ARRAY and has no
15 c IDL specific code.
16 c
17 SUBROUTINE SUM_ARRAY1(array, n, sum)
18 INTEGER*4 n
19 REAL*4 array(n), sum
20
21 sum=0.0
22 DO i=1,n
23 sum = sum + array(i)
24 ENDDO
25 RETURN
26 END

```

Table 3-8: Fortran Code Called Directly From IDL

To call the example program from within IDL:

```

X = FINDGEN(10) ; Make an array.
sum = 0.0
S = CALL_EXTERNAL('example1.so', $
    'sum_array_', X, N_ELEMENTS(X), sum)

```

In this example, `example1.so` is the name of the sharable image file, `sum_array_` is the name of the entry point, and `X` and `N_ELEMENTS(X)` are passed to the called routine as parameters. The returned value is contained in the variable `sum`.

Note

The entry point name generated by the Fortran compiler may be different than that produced by the C compiler. One of the best ways to find out what name was generated is to use the UNIX `nm` utility on the object file. See your system's man page for `nm` for details.



Chapter 4

Remote Procedure Calls

This chapter discusses the following topics:

IDL and Remote Procedure Calls	78	Compatibility with Older IDL Code	83
Using IDL as an RPC Server	79	The IDL RPC Library	85
Client Variables	80	RPC Examples	110
Linking to the Client Library	81		

IDL and Remote Procedure Calls

Remote Procedure Calls (RPCs) allow one process (the *client* process) to have another process (the *server* process) execute a procedure call just as if the caller process had executed the procedure call in its own address space. Since the client and server are separate processes, they can reside on the same machine or on different machines. RPC libraries allow the creation of network applications without having to worry about underlying networking mechanisms.

IDL supports RPCs so that other applications can communicate with IDL. A library of C language routines is included to handle communication between client programs and the IDL server.

A startup file is executed only when a command line is present. Running an application using an IDL Remote Procedure Call server does not execute the startup file. See [“Understanding When Startup Files are Not Executed”](#) (Chapter 1, *Using IDL*) for details.

Note

Remote procedure calls are supported only on UNIX platforms.

The current implementation allows IDL to be run as an RPC server and your own program to be run as a client. IDL commands can be sent from your application to the IDL server, where they are executed. Variable structures can be defined in the client program and then sent to the IDL server for creation as IDL variables. Similarly, the values of variables in the IDL server session can be retrieved into the client process.

With the release of IDL version 5.0, IDL’s RPC functionality has been completely revised and a new API created. The new RPC interface mirrors the API used by callable IDL. See [“Compatibility with Older IDL Code”](#) on page 83 for details.

Using IDL as an RPC Server

The IDL RPC Directory

All of the files related to using IDL's RPC capabilities are found in the `rpc` subdirectory of the `external` subdirectory of the main IDL directory. The main IDL directory is referred to here as *idldir*.

Running IDL in Server Mode

To use IDL as an RPC server, run IDL in server mode by using the `idlrpc` command. The RPC server can be invoked one of two ways:

```
idlrpc
```

or

```
idlrpc -server=server_number
```

where *server_number* is the hexadecimal server ID number (between 0x20000000 and 0x3FFFFFFF) for IDL to use. For example, to run IDL with the server ID number 0x20500000, use the command:

```
idlrpc -server=20500000
```

If a server ID number is not supplied, IDL uses the default, `IDL_RPC_DEFAULT_ID`, defined in the file `idldir/external/rpc/idl_rpc.h`. This value is originally set to 0x2010CAFE.

Client Variables

The IDL RPC client API uses the same data structure as IDL to represent a variable, namely an **IDL_VARIABLE** structure. By not using a unique data structure to represent a variable, the IDL RPC client API can follow a format that is similar to the API of Callable IDL.

When a variable is created by the IDL RPC client API (when a variable is returned from the **IDL_RPCGetMainVariable** function, for example) dynamic memory is allocated for the variable and for its value. These dynamic variables are similar to temporary variables which are used in IDL.

The IDL RPC client API provides routines to create, manipulate and delete dynamic or IDL RPC client temporary variables. These API routines follow the same format as the Callable IDL API and most have the same calling sequence.

When a client dynamic or temporary variable is no longer needed by the IDL RPC client program, use the **IDL_RPCDeltmp()** function to delete or free up the memory associated with the variable. Failure to delete a client temporary variable could result a memory “leak” in the client program.

Linking to the Client Library

To make use of the IDL RPC functionality, you will need to do the following:

- Include the file `idl_rpc.h` in your application.
- Have a copy of `idl_export.h` in the include path when you compile the client application.
- Link your client application to the IDL client shared object library (`libidl_rpc`).
- If the client library is linked as a shared object, you must set the `LD_LIBRARY_PATH` environment variable so that it includes the directory that contains the IDL client library. If this variable is not set correctly, an error message will be issued by the system loader when the client program is started.

The command used to compile and link a client program to the IDL RPC client library follows the following format:

```
% cc -o example $(PRE_FLAGS) example.o -lidl_rpc  
$(POST_FLAGS)
```

where `PRE_FLAGS` and `POST_FLAGS` are platform dependent. The proper flags for each UNIX operating system supported by IDL are contained in the file `Makefile`, located in the `rpc` subdirectory of the `external` subdirectory of the main IDL directory.

Example of IDL RPC Client API

To use the IDL client side API, execute the following sequence of steps:

1. Call `IDL_RPCInit()` to connect to the server
2. Perform actions on the server—get and set variables, run IDL commands, etc.
3. Call `IDL_RPCCleanup()` to disconnect from the server.

The code shown in the following figure is an example that can be used to set up a remote session of IDL using the RPC features. Note that this C program will need to be linked against the supplied shared library `libidl_rpc`. This code is included in the `idldir/external/rpc` directory as `example.c`.

```

1  #include "idl_rpc.h"
2  int main()
3  {
4      CLIENT *pClient;
5      char    cmdBuffer[512];
6      int     result;
7
8      /* Connect to the server */
9      if( (pClient = IDL_RPCInit(0, (char*)NULL)) == (CLIENT*)NULL){
10         fprintf(stderr, "Can't register with IDL server\n");
11         exit(1);
12     }
13
14  /* Start a loop that will read commands and then send them to idl */
15     for(;;){
16         printf("RMTIDL> ");
17         cmdBuffer[0]='\0';
18         gets(cmdBuffer);
19         if( cmdBuffer[0] == '\n' || cmdBuffer[0] == '\0')
20             break;
21         result = IDL_RPCExecuteStr(pClient, cmdBuffer);
22     }
23
24     /* Now disconnect from the server and kill it. */
25     if(!IDL_RPCCleanup(pClient, 1))
26         fprintf(stderr, "IDL_RPCCleanup: failed\n");
27     exit(0);
28 }

```

Table 4-1: Remote Execution of IDL via RPC

Compile `example.c` with the appropriate flags for your platform, as described in [“Linking to the Client Library”](#) on page 81. Once this example is compiled, execute it using the following commands:

```
% idlrpc
```

Then, in another process:

```
% example
```

Compatibility with Older IDL Code

With the release of IDL 5.0, IDL's Remote Procedure Call functionality has been completely reworked. While RPC code built for older versions of IDL can still be used with IDL 5.0 and later, the new RPC functionality has the following advantages:

- The new API mirrors the Callable IDL API.
- The RPC client-side library is provided as a pre-built sharable library, eliminating the need to build the library on your system.
- The RPC server-side executable, `idlrpc`, is built using Callable IDL, providing an example of how Callable IDL can be used.
- Source code is provided for both the Server and Client side programs, allowing you to enhance IDL's RPC functionality.

RPC code built for versions of IDL prior to version 5.0 can be linked with IDL version 5 and later using a compatibility layer. This layer is contained in the files `idl_rpc_obsolete.c` and `idl_rpc_obsolete.h`.

To use the compatibility routines, include the file `lib_rpc_obsolete.h` in your application and use the following link statement as a template:

```
% cc -o old_example $(PRE_FLAGS) old_example.o \  
idl_rpc_obsolete.o -lidlrpc $(POST_FLAGS)
```

where the macros `PRE_FLAGS` and `POST_FLAGS` are the same as those described in [“Linking to the Client Library”](#) on page 81.

While the compatibility layer covers most of the old IDL RPC functionality, some of the more obscure operations have either been modified or are no longer supported. The features which have changed are as follows:

- **idl_server_interactive**: This function is no longer supported.
- **get_idl_variable**: The following return values are no longer supported:

Value	Description
-2	Illegal variable name (for example, “213xyz”, “#a”, “!DEVICE”)
-3	Variable not transportable (for example, the variable is a structure or associated variable)

Table 4-2: get_idl_variable Unsupported Values

- **set_idl_timeout**: the **tv_usec** field of the **timeval** struct is ignored.
- **idl_set_verbosity()**: This function is no longer supported.

All other functionality is supported.

The IDL RPC Library

The IDL RPC library contains several C language interface functions that facilitate communication between your application and IDL. There are functions to register and unregister clients, set timeouts, get and set the value of IDL variables, send commands to the IDL server, and cause the server to exit. These functions are:

- `IDL_RPCCleanup`
- `IDL_RPCDeltmp`
- `IDL_RPCExecuteStr`
- `IDL_RPCGetMainVariable`
- `IDL_RPCGettmp`
- `IDL_RPCGetVariable`
- `IDL_RPCImportArray`
- `IDL_RPCInit`
- `IDL_RPCMakeArray`
- `IDL_RPCOutputCapture`
- `IDL_RPCOutputGetStr`
- `IDL_RPCSetMainVariable`
- `IDL_RPCSetVariable`
- `IDL_RPCStoreScalar`
- `IDL_RPCStrDelete`
- `IDL_RPCStrDup`
- `IDL_RPCStrEnsureLength`
- `IDL_RPCStrStore`
- `IDL_RPCTimeout`
- `IDL_RPCVarCopy`
- `IDL_RPCVarGetData`
- Variable Accessor Macros

IDL_RPCCleanup

Calling Sequence

```
int IDL_RPCCleanup( CLIENT *pClient, int iKill)
```

Description

Use this function to release the resources associated with the given CLIENT structure or to kill the IDL RPC server.

Parameters

pClient

A pointer to the CLIENT structure for the client/server connection to be disconnected.

iKill

Set **iKill** to a non-zero value to kill the server when the connection is broken.

Return Value

This function returns 1 on success, or 0 on failure.

IDL_RPCDeltmp

Calling Sequence

```
void IDL_RPCDeltmp( IDL_VPTR vTmp)
```

Description

Use this function to de-allocate all dynamic memory associated with the **IDL_VPTR** that is passed into the function. Once this function returns, any dynamic portion of **vTmp** is deallocated and should not be referenced.

Parameters

vTmp

The variable that will be de-allocated.

Return Value

None.

IDL_RPCExecuteStr

Calling Sequence

```
int IDL_RPCExecuteStr(CLIENT *pClient, char * pCommand)
```

Description

Use this function to send IDL commands to the IDL RPC server. The command is executed just as if it had been entered from the IDL command line.

This function cannot be used to send multiple line commands and will return an error if a “\$” is detected at the end of the command string. It will also return an error if “\$” is the first character, since this would spawn an interactive process and hang the IDL RPC server.

Parameters

pClient

A pointer to the CLIENT structure that corresponds to the desired IDL session.

pCommand

A null-terminated IDL command string.

Return Value

This function returns the following values:

- 1** — Success.
- 0** — Invalid command string.

For all other errors, the value of !ERROR_STATE.CODE is returned. This number could be passed as an argument to the IDL function **STRMESSAGE()** to determine the exact cause of the error.

IDL_RPCGetMainVariable

Calling Sequence

```
IDL_VPTR IDL_RPCGetMainVariable(CLIENT *pClient, char *Name)
```

Description

Call this function to get the value of an IDL RPC server main level variable referenced by the name contained in **Name**. **IDL_RPCGetMainVariable** will then return a pointer to an **IDL_VARIABLE** structure that contains the value of the variable.

Parameters

pClient

A pointer to the CLIENT structure that corresponds to the desired IDL session.

Name

The name of the variable to find.

Return Value

On success, this function returns a pointer to an **IDL_VARIABLE** structure that contains the value of the desired IDL RPC main level variable. On failure this function returns NULL.

Note that the returned variable is marked as **temporary** and should be deleted when the variable is no longer needed. For more information on IDL RPC variables, see [“Client Variables”](#) on page 80.

IDL_RPCGettmp

Calling Sequence

```
IDL_VPTR IDL_RPCGettmp(void)
```

Description

Use this function to create an **IDL_VPTR** to a dynamically allocated **IDL_VARIABLE** structure. When you are finished with this variable, pass it to **IDL_RPCDeltmp()** to free any memory allocated by the variable.

Parameters

None.

Return Value

On success, this function returns an **IDL_VPTR**. On failure, it returns **NULL**.

IDL_RPCGetVariable

Calling Sequence

```
IDL_VPTR IDL_RPCGetVariable(CLIENT *pClient, char *Name)
```

Description

Use this function to get a pointer to an **IDL_VARIABLE** structure that contains the value of an IDL RPC server variable referenced by **Name**. The current scope of the IDL program is used to get the value of the variable.

Parameters

pClient

A pointer to the CLIENT structure that corresponds to the desired IDL session.

Name

The name of the variable to find.

Return Value

On success, this function returns a pointer to an **IDL_VARIABLE** structure that contains the value of the desired IDL RPC variable. On failure this function returns NULL.

Note that the returned variable is marked as **temporary** and should be deleted when the variable is no longer needed. For more information on IDL RPC variables, see [“Client Variables”](#) on page 80.

IDL_RPCImportArray

Calling Sequence

```
IDL_VPTR IDL_RPCImportArray(int n_dim, IDL_MEMINT dim[],
                             int type, UCHAR *data, IDL_ARRAY_FREE_CB free_cb)
```

Description

Use this function to create an IDL array variable whose data the server supplies, rather than having the client API allocate the data space.

Parameters

n_dim

The number of dimensions in the array.

dim

An array of **IDL_MAX_ARRAY_DIM** elements, containing the size of each dimension.

type

The IDL type code describing the data. IDL type codes are discussed in “[Type Codes](#)” on page 114.

data

A pointer to your array data.

free_cb

If non-NULL, **free_cb** is a pointer to a function that will be called when the IDL RPC client routines frees the array. This feature gives the caller a sure way to know when the data is no longer referenced. Use the called function to perform any required cleanup, such as freeing dynamic memory or releasing shared or mapped memory.

Return Value

An **IDL_VPTR** that points to an **IDL_VARIABLE** structure containing a reference to the imported array. This function returns NULL if the operation was unsuccessful.

IDL_RPCInit

Calling Sequence

```
Client *IDL_RPCInit(long ServerId, char* pHostname)
```

Description

Use this function to initialize an IDL RPC client session.

The client program is registered as a client of the IDL RPC server. The server that the client is registered with depends on the values of the parameters passed to the function.

Parameters

ServerId

The ID number of the IDL server that the program is to be registered with. If this value is 0, the default server ID (0x2010CAFE) is used.

pHostname

This is the name of the machine where the IDL server is running. If this value is NULL or "", the default, "localhost", is used.

Return Value

A pointer to the new CLIENT structure is returned upon successful completion. This opaque data structure is then later used by the client program to perform operations with the server. This function returns NULL if the operation was unsuccessful.

IDL_RPCMakeArray

Calling Sequence

```
char * IDL_RPCMakeArray( int type, int n_dim, IDL_MEMINT dim[],
                        int init, IDL_VPTR *var)
```

Description

This function creates an IDL RPC client temporary array variable with a data area of the specified size.

Parameters

type

The IDL type code for the resulting array. IDL type codes are discussed in “[Type Codes](#)” on page 114.

n_dim

The number of array dimensions. The constant **IDL_MAX_ARRAY_DIM** defines the upper limit of this value.

dim

A C array of **IDL_MAX_ARRAY_DIM** elements containing the array dimensions. The number of dimensions in the array is given by the **n_dim** argument.

init

This parameter specifies the sort of initialization that should be applied to the resulting array. **init** must be one of the following:

- **IDL_ARR_INI_NOP** — No initialization is done. The data area of the array will contain whatever garbage was left behind from its previous use.
- **IDL_ARR_INI_ZERO** — The data area of the array is zeroed.

var

The address of an **IDL_VPTR** containing the address of the resulting IDL RPC client temporary variable.

Return Value

On success, this function returns a pointer to the data area of the allocated array. The value returned is the same as is contained in the **var->value.arr->data** field of the variable. On failure, it returns NULL.

As with variables returned from **IDL_RPCGettmp()**, the variable allocated via this function must be de-allocated using **IDL_RPCDeltmp()** when the variable is no longer needed.

IDL_RPCOutputCapture

Calling Sequence

```
int IDL_RPCOutputCapture( CLIENT *pClient, int n_lines)
```

Description

Use this routine to enable and disable capture of lines output from the IDL RPC server. Normally, IDL will write any output to the terminal on which the server was started. This function can be used to save this information so that the client program can request the lines sent to the output buffer.

Parameters

pClient

A pointer to the CLIENT structure that corresponds to the desired IDL session.

n_lines

If this value is less than or equal to zero, no output lines will be buffered in the IDL RPC server and output will be sent to the normal output device on the IDL RPC server. If the value of this parameter is greater than zero, the specified number of lines will be stored by the IDL RPC server.

Return Value

This function returns 1 on success, or 0 on failure.

IDL_RPCOutputGetStr

Calling Sequence

```
int IDL_RPCOutputGetStr(CLIENT *pClient, IDL_RPC_LINE_S *pLine,
                        int first)
```

Description

Use this function to get an output line from the line queue being maintained on the RPC server. The routine **IDL_RPCOutputCapture()** *must* have been called to initialize the output queue on the RPC server before this routine is called.

Parameters

pClient

A pointer to the CLIENT structure that corresponds to the desired IDL session.

pLine

A pointer to a valid **IDL_RPC_LINE_S** structure. The **buf** field of this structure will contain the output string returned from the IDL RPC server and the **flags** field will be set to one of the following (from `idl_export.h`):

- **IDL_TOUT_F_STDERR** — Send the text to **stderr** rather than **stdout**, if that distinction means anything to your output device.
- **IDL_TOUT_F_NLPOST** — After outputting the text, start a new output line. On a tty, this is equivalent to sending a new line (`'\n'`) character.

first

If **first** is set equal to a non-zero value, the first line is popped from the output buffer on the IDL RPC server (the output buffer is treated like a stack). If **first** is set equal to zero, the last line is de-queued from the output buffer (the output buffer is treated like a queue).

Return value

A true value (1) is returned upon success. A false value (0) is returned when there are no more lines available in the output buffer or when an RPC error is detected.

IDL_RPCSetMainVariable

Calling Sequence

```
int IDL_RPCSetMainVariable( CLIENT *pClient, char *Name,  
                           IDL_VPTR pVar)
```

Description

Use this routine to assign a value to a main level IDL variable in the IDL RPC server session referred to by **pClient**. If the variable does not already exist, a new variable will be created.

Parameters

pClient

A pointer to the CLIENT structure that corresponds to the desired IDL session.

Name

A pointer to the null-terminated name of the variable, which must be in upper-case.

pVar

A pointer to an **IDL_VARIABLE** structure that contains the value that the IDL RPC main level variable referenced by **Name** should be set to. For more information on creating this variable, see “[Client Variables](#)” on page 80.

Return Value

This function returns 1 on success, or 0 on failure.

IDL_RPCSetVariable

Calling Sequence

```
int IDL_RPCSetVariable( CLIENT *pClient,  char *Name,
                      IDL_VPTR pVar)
```

Description

Use this routine to assign a value to an IDL variable in the IDL RPC server session referred to by **pClient**. If the variable does not already exist, a new variable will be created. Unlike **IDL_RPCSetMainVariable()**, this routine sets the variable in the current IDL program scope.

Parameters

pClient

A pointer to the CLIENT structure that corresponds to the desired IDL session.

Name

A pointer to the null-terminated name of the variable, which must be in upper-case.

pVar

A pointer to an **IDL_VARIABLE** structure that contains the value that the IDL RPC variable referenced by **Name** should be set to. For more information on creating this variable, see “[Client Variables](#)” on page 80.

Return Value

This function returns 1 on success, or 0 on failure.

IDL_RPCStoreScalar

Calling Sequence

```
void IDL_RPCStoreScalar(IDL_VPTR dest, int type,  
                        IDL_ALLTYPES *value)
```

Description

Use this function to store a scalar value into an **IDL_VARIABLE** structure. Before the scalar is stored, any dynamic part of the existing **IDL_VARIABLE** is deallocated.

Parameters

dest

An **IDL_VPTR** to the **IDL_VARIABLE** in which the scalar should be stored.

type

The type code for the scalar value. IDL type codes are discussed in “[Type Codes](#)” on page 114.

value

The address of an **IDL_ALLTYPES** union that contains the value to store.

Return Value

None.

IDL_RPCStrDelete

Calling Sequence

```
void IDL_RPCStrDelete(IDL_STRING *str, IDL_MEMINT n)
```

Description

Use this function to delete a string. See the description of **IDL_StrDelete()** in [“Deleting Strings”](#) on page 187.

IDL_RPCStrDup

Calling Sequence

```
void IDL_RPCStrDup(IDL_STRING *str, IDL_MEMINT n)
```

Description

Use this function to duplicate a string. See the description of **IDL_StrDup()** in [“Copying Strings”](#) on page 186.

IDL_RPCStrEnsureLength

Calling Sequence

```
void IDL_RPCStrEnsureLength(IDL_STRING *s, int n)
```

Description

Use this function to check the length of a string. See the description of **IDL_StrEnsureLength()** in [“Obtaining a String of a Given Length”](#) on page 189.

IDL_RPCStrStore

Calling Sequence

```
void IDL_RPCStrStore( IDL_STRING *s, char *fs)
```

Description

Use this function to store a string. See description of **IDL_StrStore** in [“Setting an IDL_STRING Value”](#) on page 188.

IDL_RPCTimeout

Calling Sequence

```
int IDL_RPCTimeout(long lTimeOut)
```

Description

Use this function to set the timeout value used when the RPC client makes requests of the server.

Parameters

lTimeOut

A integer value, in seconds, specifying the timeout value that will be used in RPC operations.

Return Value

This function returns 1 on success, or 0 on failure.

IDL_RPCVarCopy

Calling Sequence

```
void IDL_RPCVarCopy(IDL_VPTR src, IDL_VPTR dst)
```

Description

Use this function to copy the contents of the **src** variable to the **dst** variable. Any dynamic memory associated with **dst** is de-allocated before the source data is copied. This function emulates the callable IDL function **IDL_VarCopy()**.

Parameters

src

The source variable to be copied. If this variable is marked as temporary (returned from **IDL_RPCGettmp()**, for example) the dynamic data will be moved rather than copied to the destination variable.

dst

The destination variable that **src** is copied to.

Return Value

None.

IDL_RPCVarGetData

Calling Sequence

```
void IDL_RPCVarGetData(IDL_VPTR v, IDL_MEMINT *n, char **pd,  
    int ensure_simple)
```

Description

Use this function to obtain a pointer to a variable's data, and to determine how many data elements the variable contains.

Parameters

v

The variable for which data is desired.

n

The address of a variable that will contain the number of elements in **v**.

pd

The address of a variable that will contain a pointer to **v**'s data, cast to be a pointer to pointer to char (e.g. (char **) &myptr).

ensure_simple

If TRUE, this routine calls the **ENSURE_SIMPLE** macro on the argument **v** to screen out variables of the types it prevents. Otherwise, **EXCLUDE_FILE** is called, because file variables have no data area to return.

Return Value

On exit, **IDL_RPCVarGetData()** stores the data count and pointer into the variables pointed at by **n** and **pd**, respectively.

Variable Accessor Macros

The following macros can be used to get information on IDL RPC variables. These macros are defined in `idl_rpc.h`.

All of these macros accept a single argument, *v*, of type **IDL_VPTR**.

IDL_RPCGetArrayData(*v*)

This macro returns a pointer (*char**) to the data area of an array block.

IDL_RPCGetArrayDimensions(*v*)

This macro returns a C array which contains the array dimensions.

IDL_RPCGetArrayNumDims(*v*)

This macro returns the number of dimensions of the array.

IDL_RPCGetVarByte(*v*)

This macro returns the value of a 1-byte, unsigned *char* variable.

IDL_RPCGetVarComplex(*v*)

This macro returns the value (as a *struct*, not a pointer) of a complex variable.

IDL_RPCGetVarComplexR(*v*)

This macro returns the real field of a complex variable.

IDL_RPCGetVarComplexI(*v*)

This macro returns the imaginary field of a complex variable.

IDL_RPCGetVarDComplex(*v*)

This macro returns the value (as a *struct*, not a pointer) of a double precision, complex variable.

IDL_RPCGetVarDComplexR(*v*)

This macro returns the real field of a double-precision complex variable.

IDL_RPCGetVarDComplexI(*v*)

This macro returns the imaginary field of a double-precision complex variable.

IDL_RPCGetVarDouble(*v*)

This macro returns the value of a double-precision, floating-point variable.

IDL_RPCGetVarFloat(*v*)

This macro returns the value of a single-precision, floating-point variable.

IDL_RPCGetVarInt(*v*)

This macro returns the value of a 2-byte integer variable.

IDL_RPCGetVarLong(*v*)

This macro returns the value of a 4-byte integer variable.

IDL_RPCGetVarLong64(*v*)

This macro returns the value of a 8-byte integer variable.

IDL_RPCVarIsArray(*v*)

This macro returns non-zero if *v* is an array variable.

IDL_RPCGetVarString(*v*)

This macro returns the value of a string variable (as a *char**).

IDL_RPCGetVarType(*v*)

This macro returns the type code of the variable. IDL type codes are discussed in [“Type Codes”](#) on page 114.

IDL_RPCGetVarUInt(*v*)

This macro returns the value of an unsigned 2-byte integer variable.

IDLRPCGetVarULong(*v*)

This macro returns the value of an unsigned 4-byte integer variable.


IDL_RPCGetVarULong64(*v*)

This macro returns the value of an unsigned 8-byte integer value.

RPC Examples

A number of example files are included in the `IDL_DIR/external/rpc` directory. A `Makefile` for these examples is also included. These short C programs demonstrate the use of the IDL RPC library.

Source files for the `idlrpc` server program are located in the `IDL_DIR/external/rpc` directory. Note that you do not need to build the `idlrpc` server; it is pre-built and included in the IDL distribution. The `idlrpc` server source files are provided as examples only.



Part II: IDL's Internal API



Chapter 5

IDL Internals: Types

This chapter describes the following topics:

Type Codes	114	IDL_MEMINT and IDL_FILEINT Types	119
Mapping of Basic Types	116		

Type Codes

Every IDL variable has a data type. The possible type codes and their mapping to C language types are listed in the following table. The undefined type code (**IDL_TYP_UNDEF**) will always have the value zero.

Although it is rare, the number of types could change someday. Therefore, you should always use the symbolic names when referring to any type except **IDL_TYP_UNDEF**. Even in the case of **IDL_TYP_UNDEF**, using the symbolic name will add clarity to your code. Note that all IDL structures are considered to be of a single type (**IDL_TYP_STRUCT**).

Clearly, distinctions must be made between various structures, but such distinctions are made at a different level. There are a few constants that can be used to make your code easier to read and less likely to break if/when the `idl_export.h` file changes. These are:

- **IDL_MAX_TYPE**—The value of the largest type.
- **IDL_NUM_TYPES**—The number of types. Since the types are numbered starting at zero, **IDL_NUM_TYPES** is one greater than **IDL_MAX_TYPE**.

Name	Type	C Type
IDL_TYP_UNDEF	Undefined	<None>
IDL_TYP_BYTE	Unsigned byte	UCHAR
IDL_TYP_INT	16-bit integer	IDL_INT
IDL_TYP_LONG	32-bit integer	IDL_LONG
IDL_TYP_FLOAT	Single precision floating	float
IDL_TYP_DOUBLE	Double precision floating	double
IDL_TYP_COMPLEX	Single precision complex	IDL_COMPLEX
IDL_TYP_STRING	String	IDL_STRING
IDL_TYP_STRUCT	Structure	See “ Structure Variables ” on page 159
IDL_TYP_DCOMPLEX	Double precision complex	IDL_DCOMPLEX

Table 5-1: IDL Types and Mapping to C

Name	Type	C Type
IDL_TYP_PTR	32-bit integer	IDL_ULONG
IDL_TYP_OBJREF	32-bit integer	IDL_ULONG
IDL_TYP_UINT	Unsigned 16-bit integer	IDL_UINT
IDL_TYP_ULONG	Unsigned 32-bit integer	IDL_ULONG
IDL_TYP_LONG64	64-bit integer	IDL_LONG64
IDL_TYP_ULONG64	Unsigned 64-bit integer	IDL_ULONG64

Table 5-1: IDL Types and Mapping to C (Continued)

Type Masks

There are some situations in which it is necessary to specify types in the form of a bit mask rather than the usual type codes, for example when a single argument to a function can represent more than a single type. For any given type, the bit mask value can be computed as: $\text{Mask} = 2^{\text{TypeCode}}$

The **IDL_TYP_MASK** preprocessor macro is provided to calculate these masks. Given a type code, it returns the bit mask. For example, to specify a bit mask for all the integer types:

```
IDL_TYP_MASK (IDL_TYP_BYTE) | IDL_TYP_MASK (IDL_TYP_INT) |
                               IDL_TYP_MASK (IDL_TYP_LONG)
```

Specifying all the possible types would require a long statement similar to the one above. To avoid having to type so much for this common case, the **IDL_TYP_B_ALL** constant is provided.

Mapping of Basic Types

Within IDL, the IDL data types are mapped into data types supported by the C language. Most of the types map directly into C primitives, while **IDL_TYP_COMPLEX**, **IDL_TYP_DCOMPLEX**, and **IDL_TYP_STRING** are defined as C structures. The mappings are given in the following table. Structures are built out of the basic types by laying them out in memory in the specified order using the same alignment rules used by the C compiler for the target machine.

Unsigned Byte Data

UCHAR is defined to be unsigned char in `idl_export.h`.

Integer Data

IDL_INT represents the signed 16-bit data type and is defined in `idl_export.h`.

Unsigned Integer Data

IDL_UINT represents the unsigned 16-bit data type and is defined in `idl_export.h`.

Long Integer Data

IDL long integers are defined to be 32-bits in size. The C long data type is not correct on all systems because C compilers for 64-bit architectures usually define long as 64-bits. Hence, the **IDL_LONG** typedef, declared in `idl_export.h` is used instead.

Unsigned Long Integer Data

IDL_ULONG represents the unsigned 32-bit data type and is defined in `idl_export.h`.

64-bit Integer Data

IDL_LONG64 represents the 64-bit data type and is defined in `idl_export.h`.

Unsigned 64-bit Integer Data

`IDL_ULONGLONG64` represents the unsigned 64-bit data type and is defined in `idl_export.h`.

Complex Data

The `IDL_TYP_COMPLEX` and `IDL_TYP_DCOMPLEX` data types are defined by the following C declarations:

```
typedef struct { float r, i; } IDL_COMPLEX;
typedef struct { double r, i; } IDL_DCOMPLEX;
```

This is the same mapping used by Fortran compilers to implement their complex data types, which allows sharing binary data with such programs.

String Data

The `IDL_TYP_STRING` data type is implemented by a string descriptor:

```
typedef struct {
    IDL_STRING_SLEN_T slen; /* Length of string */
    short stype;           /* Type of string */
    char *s;              /* Pointer to string */
} IDL_STRING;
```

The fields of the `IDL_STRING` struct are defined as follows:

slen

The length of the string, not counting the null termination. For example, the string “Hello” has 5 characters.

stype

If **stype** is zero, the string pointed at by **s** (if any) was not allocated from dynamic memory, and should not be freed. If non-zero, **s** points at a string allocated from dynamic memory, and should be freed before being replaced. For information on dynamic memory, see “[Dynamic Memory](#)” on page 252 and “[Getting Dynamic Memory](#)” on page 174.

S

If **slen** is non-zero, **s** is a pointer to a null-terminated string of **slen** characters. If **slen** is zero, **s** should not be used. The use of a string pointer to memory located outside the **IDL_STRING** structure itself allows IDL strings to have dynamically-variable lengths.

Note

Strings are the most complicated basic data type, and as such, are at the root of more coding errors than the other types. See “[IDL Internals: String Processing](#)” on page 183.

IDL_MEMINT and IDL_FILEINT Types

Some of the IDL-supported operating systems limit memory and file lengths to a signed 32-bit integer (approximately 2.3 GB). Some systems have 64-bit memory capabilities and others allow files longer than $2^{31}-1$ bytes despite being 32-bit memory limited. To gracefully handle these differences without using conditional code, IDL internals use two special types, IDL_TYP_MEMINT (data type IDL_MEMINT) and IDL_TYP_FILEINT (data type IDL_FILEINT) to represent memory and file length limits.

IDL_MEMINT and IDL_FILEINT are not separate and distinct types; they are actually mappings to the IDL types discussed in [“Mapping of Basic Types”](#) on page 116. Specifically, they will be IDL_LONG for 32-bit quantities, and IDL_LONG64 for 64-bit quantities.

As an IDL internals programmer, you should not write code that depends on the actual machine type represented by these abstract types. To ensure that your code runs properly on all systems, use IDL_MEMINT and IDL_FILEINT in place of more specific types. These types can be used anywhere that a normal IDL type can be used, such as in keyword processing. Their systematic use for these purposes will ensure that your code is correct on any IDL platform.

Programmers should be aware of the IDL_MEMINTScalar() and IDL_FILEINTScalar() functions, described in [“Converting Arguments to C Scalars”](#) on page 206.



Chapter 6

IDL Internals: Keyword Processing

This chapter discusses the following topics:

IDL and Keyword Processing	122	Keyword Processing Options	130
Creating Routines that Accept Keywords	123	The KW_RESULT Structure	132
Overview Of IDL Keyword Processing	124	Cleaning Up	136
The IDL_KW_ARR_DESC_R Structure	129	Keyword Examples	137

IDL and Keyword Processing

Keyword arguments are an important IDL language feature. They allow a multitude of options to be specified to a routine in a straightforward, easily understood way. The price of this added power is that it is somewhat more complicated to write a routine that accepts keywords than one that doesn't. However, the additional effort is well worth it.

Creating Routines that Accept Keywords

As described in “[Adding System Routines](#)” on page 269, you must register your system routine before IDL will recognize it. When registering the routine, you indicate that it accepts keyword arguments in one of the following ways:

- Specifying the **KEYWORDS** option for the routine in the module definition file of a Dynamically Loadable Module (DLM)
- Setting the **KEYWORDS** keyword in a call to **LINKIMAGE**.
- OR-ing the constant **IDL_SYSFUN_DEF_F_KEYWORDS** into the **flags** field of the **IDL_SYSFUN_DEF2** struct passed to **IDL_SysRtnAdd()**

Routines that accept keywords must perform keyword processing. A routine that does not allow keyword processing knows that its **argc** argument gives the number of positional arguments, and **argv** contains only those positional arguments. In contrast, a routine that accepts keywords receives an **argc** that gives the total number of positional and keyword arguments, and these arguments are delivered in **argv** mixed together in an undefined order.

The function **IDL_KWProcessByOffset()** is used to process keywords and separate the positional and keyword arguments. It is passed an array of **IDL_KW_PAR** structures that give information about the allowed keywords and their attributes. The keyword data resulting from this process is stored in a user defined **KW_RESULT** structure. Finally, the **IDL_KW_FREE** macro is used to clean up.

More information about these routines and structures can be found in the following sections.

Overview Of IDL Keyword Processing

IDL keyword processing can seem confusing at first glance, due to the interrelated data structures involved. However, as the examples that follow in this chapter will show, the concepts involved are relatively straightforward once you have seen and understood a concrete example such as “[Keyword Examples](#)” on page 137.

Following is a skeleton of a system routine that accepts keyword arguments. These elements must be present in any such system routine:

```
void keyword_sysrtn_skeleton(int argc, IDL_VPTR *argv, char *argk)
{
    typedef struct {
        IDL_KW_RESULT_FIRST_FIELD; /* Must be first entry in struct */
        ... /* Variables specific to your keywords go here */
    } KW_RESULT;
    static IDL_KW_PAR kw_pars[] = {
        /*
         * Keyword definitions for the keywords you accept go here,
         * one definition per keyword. The keyword definitions refer
         * to fields within the KW_RESULT type defined above.
         */
        ...
        { NULL } /* List must be NULL terminated */
    };
    KW_RESULT kw; /* Variable which will hold the keyword values */

    (void) IDL_KWProcessByOffset(argc, argv, argk, kw_pars,
                                (IDL_VPTR *) 0, 1, &kw);

    /* The body of your routine */

    IDL_KW_FREE;
}
```

IDL keyword processing is made up of the following data structures and steps:

- A NULL terminated array of **IDL_KW_PAR** structures must be present. Each entry in this array describes the keyword processing required for a single keyword.
- If a keyword represents an input-only, by-value array, the **IDL_KW_PAR** structure that describes it points at an auxiliary **IDL_KW_ARR_DESC_R** structure that supplies the additional array specific information.
- The system routine must declare a local type definition named **KW_RESULT**, and a variable of this type named **kw**. The **KW_RESULT** type contains all of

the data fields that will be set as a result of processing the keywords described by the **IDL_KW_PAR** and **IDL_KW_ARR_DESC_R** structures described above. The **IDL_KW_PAR** and **IDL_KW_ARR_DESC_R** structures refer to the fields of the **KW_RESULT** structure by their offset from the beginning of the structure. The **IDL_KW_OFFSETOF()** macro is used to compute this offset.

- The system routine calls the **IDL_KWProcessByOffset()** function, passing it the address of the **IDL_KW_PAR** array, and the **KW_RESULT** variable (**kw**).
- After **IDL_KWProcessByOffset()** is called, the **KW_RESULT** structure (**kw**) contains the results, which can be accessed freely by the system routine.
- Before returning, the system routine must invoke the **IDL_KW_FREE** macro. This macro ensures that any dynamic memory used by **IDL_KWProcessByOffset()** is properly released.
- System routines are not required to, and generally do not, call **IDL_KW_FREE** before throwing errors using **IDL_Message()** with the **IDL_MSG_LONGJMP** or **IDL_MSG_IO_LONGJMP** action codes. In these cases, the IDL interpreter automatically knows to release the resources used by keyword processing on your behalf.

All of these data structures and routines are discussed in detail in the sections that follow.

The IDL_KW_PAR Structure

The **IDL_KW_PAR** struct provides the basic specification for keyword processing. The **IDL_KWProcessByOffset()** function is passed a null-terminated array of these structures. **IDL_KW_PAR** structures specify which keywords a routine accepts, the attributes required of them, and the kinds of processing that should be done to them. **IDL_KW_PAR** structures must be defined in lexical order according to the value of the **keyword** field.

The definition of **IDL_KW_PAR** is:

```
typedef struct {
    char *keyword;
    UCHAR type;
    unsigned short mask;
    unsigned short flags;
    int *specified;
    char *value;
} IDL_KW_PAR;
```

where:

keyword

A pointer to a null-terminated string. This is the name of the keyword, and must be entirely upper case. The array of **IDL_KW_PAR** structures passed to **IDL_KWProcessByOffset()** must be lexically sorted by the strings pointed to by this field. The final element in the array is signified by setting the keyword field to NULL ((**char ***) 0).

type

IDL_KWProcessByOffset() automatically converts the keywords to the IDL type specified by the **type** field. Specify 0 (**IDL_TYPE_UNDEF**) in cases where **IDL_KW_VIN** or **IDL_KW_OUT** are specified in the **flags** field.

mask

The enable mask. This field is ANDed with the mask argument to **IDL_KWProcessByOffset()** and if the result is non-zero, the keyword is accepted. If the result is 0, the keyword is ignored. This ability allows you to share an array of **IDL_KW_PAR** structures between several routines, and enable or disable the keywords used by each one.

As an example of this, the IDL graphics and plotting routines have a large number of keywords in common. In addition, each routine has a few keywords that are unique to it. Keywords are implemented using a single shared array of **IDL_KW_PAR** with appropriate values of the mask field. This technique dramatically reduces the amount of data that would otherwise be required by graphics keyword processing, and makes IDL easier to maintain.

flags

This field specifies special processing instructions. It is a bit mask made by ORing the following values:

- **IDL_KW_ARRAY** — Set this bit to specify that the keyword must be an array. Otherwise, a scalar is required. If **IDL_KW_ARRAY** is specified, the value field must point at an associated **IDL_KW_ARR_DESC_R** structure.
- **IDL_KW_OUT** — Set this bit to indicate that the keyword specifies an output parameter, passed by reference. Expressions and constants are excluded. In other words, the routine is going to change the value of the keyword argument, as opposed to the more usual case of simply reading it. The address of the **IDL_VARIABLE** will be placed in a user supplied field of type **IDL_VPTR** in the **KW_RESULT** structure (**kw**). The offset of this field in the **KW_RESULT** structure is specified by the **value** field (discussed below). **IDL_KW_OUT** implies that no type checking or processing will be performed on the keyword—it is up to the routine to perform the same sort of type checking normally carried out for plain positional arguments.

A standard approach to find out if an **IDL_KW_OUT** parameter is present in a call to a system routine is to specify **IDL_TYP_UNDEF** (0) for the type field and **IDL_KW_OUT | IDL_KW_ZERO** for flags. The **IDL_VPTR** referenced by the value field will either contain NULL, or a pointer to the **IDL_VARIABLE**.

- **IDL_KW_VIN** — Set this bit to indicate that the keyword parameter is an input parameter (expressions and/or constants are valid) passed by reference. The address of the **IDL_VARIABLE** or expression is stored in a user-supplied field of the **KW_RESULT** structure (**kw**) referenced by the value field, as with **IDL_KW_OUT**. **IDL_KW_VIN** implies that no type checking or processing will be performed on the keyword—it is up to the routine to perform the same sort of type checking normally carried out for plain positional arguments.
- **IDL_KW_ZERO** — Set this bit in order to *zero* the C variable pointed to by the value field before parsing the keywords. This means that the object pointed

to by value will always be zero unless it was specified by the user. Use this technique to create keywords that have Boolean (on or off) meanings.

- **IDL_KW_VALUE** — If this bit is set and the specified keyword is present and non-zero, the low 12 bits of this field (**flags**) will be bitwise ORed with the **IDL_LONG** field of the **KW_RESULT** structure referenced by the **value** field. Note that this implies the **IDL_TYP_LONG** type code, and is incompatible with the **IDL_KW_ARRAY**, **IDL_KW_VIN**, and **IDL_KW_OUT** flags.

specified

NULL, or the offset of the user supplied field within the **KW_RESULT** structure (**kw**) of a C int variable that will be set to TRUE (non-zero) or FALSE (0) based on whether the routine was called with the keyword present. The **IDL_KW_OFFSETOF()** macro should be used to calculate the offset. Setting this field to NULL (0) indicates that this information is not needed.

value

If the keyword is a read-only scalar, this field is the offset of the user supplied field in the **KW_RESULT** structure (**kw**) into which the keyword value will be copied. The **IDL_KW_OFFSETOF()** macro should be used to calculate the offset.

In the case of a read-only array, value is the memory address of an **IDL_KW_ARR_DESC_R**, structure, which is discussed in [“The IDL_KW_ARR_DESC_R Structure”](#) on page 129.

In the case of an input (**IDL_KW_VIN**) or output (**IDL_KW_OUT**) variable, this field should contain the offset to the **IDL_VPTR** field within the user supplied **KW_RESULT** that will be filled by **IDL_KWProcessByOffset()** with the address of the keyword argument. The **IDL_KW_OFFSETOF()** macro should be used to calculate the offset.

The IDL_KW_ARR_DESC_R Structure

When a keyword is specified to be a read-only array (i.e., the **IDL_KW_ARRAY** flag is set), the value field of the **IDL_KW_PAR** struct should be set to point to an **IDL_KW_ARR_DESC_R** structure. This structure is defined as:

```
typedef struct {
    char *data;
    IDL_MEMINT nmin;
    IDL_MEMINT nmax;
    IDL_MEMINT* n_offset;
} IDL_KW_ARR_DESC_R;
```

where:

data

The offset of the field within the user supplied **KW_RESULT** structure, of the C array to receive the data. This offset is computed using the **IDL_KW_OFFSETOF()** macro. This array must be of the C type specified by the **type** field of the **IDL_KW_PAR** struct. For example, **IDL_TYP_LONG** maps into a C **IDL_LONG**. There must be **nmax** elements in the array.

nmin

The minimum number of elements allowed.

nmax

The maximum number of elements allowed.

n_offset

The offset of the field within the user defined **KW_RESULT** structure into which **IDL_KWProcessByOffset()** will store the number of elements actually stored into the array field. This offset is computed using the **IDL_KW_OFFSETOF()** macro.

Keyword Processing Options

The following cases occur in keyword processing:

Scalar Input-Only

For scalar, input-only keywords, the user never sees the **IDL_VARIABLE** passed as the keyword argument. Instead, the value of the **IDL_VARIABLE** is converted to the type specified by the **type** field of the **IDL_KW_PAR** struct and is placed into the field of the user specified **KW_RESULT** structure, the offset of which is given by the **value** field. This offset is calculated using the **IDL_KW_OFFSETOF()** macro.

Array Input-Only

Array input-only keywords work similarly to the scalar case, except that the **value** field contains the address of an **IDL_KW_ARR_DESC_R** struct that supplies the added information required to convert the passed array elements to the specified type and place them into a C array for easy access. The array data is copied into a array within the user supplied **KW_RESULT** structure. The **data** field of the **IDL_KW_ARR_DESC_R** struct supplies the offset of the array field within the **KW_RESULT** structure. This offset is calculated using the **IDL_KW_OFFSETOF()** macro.

As part of this process, the number of array elements passed is checked to be within the range specified in the **IDL_KW_ARR_DESC_R** struct, and if no error results, the number is stored into a field of the user supplied **KW_RESULT** struct. The **n_offset** field of the **IDL_KW_ARR_DESC_R** struct supplies the offset of this “number of elements” field within the **KW_RESULT** structure. This offset is calculated using the **IDL_KW_OFFSETOF()** macro.

It is worth noting that input-only array keywords don’t pass information about the number of dimensions or their sizes, only the total number of elements. Therefore, they are treated as 1-dimensional vectors. For more flexibility, use an Input/Output keyword instead.

Input/Output

This case occurs if the **IDL_KW_VIN** or **IDL_KW_OUT** flag is set in the **IDL_KW_PAR** struct. In this case, the value field contains the offset of the **IDL_VPTR** field (computed with the **IDL_KW_OFFSETOF()** macro) in the user defined **KW_RESULT** struct into which the actual keyword argument is copied. In this case, you must do all error checking and type conversion yourself, just like with

positional arguments. This is certainly the most flexible method. However, the other two cases are much easier to use, and are suitable for the vast majority of keywords.

The KW_RESULT Structure

Each system routine that processes keywords is required to define a structure variable into which **IDL_KWProcessByOffset()** will store all the results of keyword processing. This variable must follow the following rules:

- The name of the structure type must be defined as **KW_RESULT**. This requirement exists so that the **IDL_KW_OFFSETOF()** macro can properly do its work.
- The first field within any **KW_RESULT** structure must be defined using the **IDL_KW_RESULT_FIRST_FIELD** macro. The contents of this first field are private, and should not be examined. It contains the information required by IDL to properly track its resource use.
- The name of the **KW_RESULT** variable must be **kw**. This requirement exists so that the **IDL_KW_FREE** macro can properly do its work.

Hence, all system routines that process keywords will contain statements similar to the following:

```
typedef struct {
    IDL_KW_RESULT_FIRST_FIELD; /* Must be first entry in struct */
    ...                       /* Additional user specified fields */
} KW_RESULT;

KW_RESULT kw;
```

All fields within the **KW_RESULT** structure after the required first field can have arbitrary user selected names. The types of these fields are dictated by the **IDL_KW_PAR** and **IDL_KW_ARR_DESC_R** structures that refer to them.

Processing Keywords

The `IDL_KWProcessByOffset()` function is used to process keywords.

`IDL_KWProcessByOffset()` performs the following actions on behalf of the calling system routine:

- Verify that the keywords passed to the routine are all allowed by the routine.
- Carry out the type checking and conversions required for each keyword.
- Find the positional (non-keyword) arguments that are scattered among the keyword arguments in `argv` and copy them in order into the `plain_args` array.
- Return the number of plain arguments copied into `plain_args`.

`IDL_KWProcessByOffset()` has the form:

```
int IDL_KWProcessByOffset(int argc, IDL_VPTR *argv, char *argk,
                        IDL_KW_PAR *kw_list,
                        IDL_VPTR plain_args[], int mask,
                        void * base)
```

where:

`argc`

The number of arguments passed to the caller. This is the first parameter to all system routines.

`argv`

The array of `IDL_VPTR` to arguments that was passed to the caller. This is the second parameter to all system routines.

`argk`

The pointer to the keyword list that was passed to the caller. This is the third parameter to all system routines that accept keyword arguments.

`kw_list`

An array of `IDL_KW_PAR` structures (see [“Overview Of IDL Keyword Processing”](#) on page 124) that specifies the acceptable keywords for this routine. This array is terminated by setting the keyword field of the final struct to `NULL ((char *) 0)`.

plain_args

NULL, or an array of **IDL_VPTR** into which the **IDL_VPTRs** of the positional arguments will be copied. This array must have enough elements to hold the maximum possible number of positional arguments, as defined in **IDL_SYSFUN_DEF2**. See “[Registering Routines](#)” on page 296.

Note

IDL_KWProcessByOffset() sorts the plain arguments into the front of the input **argv** argument. Hence, **plain_args** is often not necessary, and can be set to NULL.

mask

Mask enable. This variable is ANDed with the mask field of each **IDL_KW_PAR** struct in the array given by **kw_list**. If the result is non-zero, the keyword is accepted as a valid keyword for the called system routine. If the result is zero, the keyword is ignored.

base

Address of the user supplied **KW_RESULT** structure, which must be named **kw**.

Speeding Keyword Processing

As mentioned above, the **kw_list** argument to **IDL_KWProcessByOffset()** is a null terminated list of **IDL_KW_PAR** structures. The time required to scan each item of the keyword array and zero the required fields (those fields specified, and value fields with **IDL_KW_ZERO** set), can become significant, especially when more than a few keyword array elements (e.g., 5 to 10 elements) are present.

To speed things up, specify **IDL_KW_FAST_SCAN** as the first keyword array element. If **IDL_KW_FAST_SCAN** is the first keyword array element, the keyword array is compiled by **IDL_KWProcessByOffset()** into a more efficient form the first time it is used. Subsequent calls use this efficient version, greatly speeding keyword processing. Usage of **IDL_KW_FAST_SCAN** is optional, and is not worthwhile for small lists. For longer lists, however, the improvement in speed is noticeable. For example, the following list does not use fast scanning:

```
static IDL_KW_PAR kw_pars[] = {
    { "DOUBLE", IDL_TYP_DOUBLE, 1, 0,
      IDL_KW_OFFSETOF(d_there), IDL_KW_OFFSET_OF(d) },
    { "FLOAT", IDL_TYP_FLOAT, 1, IDL_KW_ZERO, 0, IDL_KW_OFFSET_OF(f) },
    { NULL }
};
```

To use fast scanning, it would be written as:

```
static IDL_KW_PAR kw_pars[] = {
    IDL_KW_FAST_SCAN,
    { "DOUBLE", IDL_TYP_DOUBLE, 1, 0,
      IDL_KW_OFFSET_OF(d_there), IDL_KW_OFFSET_OF(d) },
    { "FLOAT", IDL_TYP_FLOAT, 1, IDL_KW_ZERO, 0, IDL_KW_OFFSET_OF(f) },
    { NULL }
};
```

Cleaning Up

All normal exit paths from your system routine are required to call the **IDL_KW_FREE** macro prior to returning. This macro must be called exactly once for every call to **IDL_KWProcessByOffset()**. You must therefore structure your code so that **IDL_KW_FREE** executes before any return statement. Many functions do not use an explicit return statement, relying on the implicit return that occurs when execution comes to the end of the function. In such a case, **IDL_KW_FREE** must be the last statement in the function.

Keyword Examples

The following C function implements `KEYWORD_DEMO`, a system procedure intended to demonstrate how to write the keyword processing code for a routine. It prints the values of its keywords, changes the value of `READWRITE` to 42 if it is present, and returns. Each line is numbered to make discussion easier. These numbers are not part of the actual program.

Note

The following code is designed to demonstrate keyword processing in a simple, uncluttered example. In actual code, you would not use the **printf** mechanism used on lines 42-53.

```

1 void keyword_demo(int argc, IDL_VPTR *argv, char *argk)
2 {
3     typedef struct {
4         IDL_KW_RESULT_FIRST_FIELD; /* Must be first entry in structure */
5         IDL_LONG l;
6         float f;
7         double d;
8         int d_there;
9         IDL_STRING s;
10        int s_there;
11        IDL_LONG arr_data[10];
12        int arr_there;
13        IDL_MEMINT arr_n;
14        IDL_VPTR var;
15    } KW_RESULT;
16    static IDL_KW_ARR_DESC_R arr_d = { IDL_KW_OFFSETOF(arr_data), 3, 10,
17                                       IDL_KW_OFFSETOF(arr_n) };
18
19    static IDL_KW_PAR kw_pars[] = {
20        IDL_KW_FAST_SCAN,
21        { "ARRAY", IDL_TYP_LONG, 1, IDL_KW_ARRAY,
22          IDL_KW_OFFSETOF(arr_there), CHARA(arr_d) },
23        { "DOUBLE", IDL_TYP_DOUBLE, 1, 0,
24          IDL_KW_OFFSETOF(d_there), IDL_KW_OFFSETOF(d) },
25        { "FLOAT", IDL_TYP_FLOAT, 1, IDL_KW_ZERO, 0, IDL_KW_OFFSETOF(f) },
26        { "LONG", IDL_TYP_LONG, 1, IDL_KW_ZERO|IDL_KW_VALUE|15, 0,
27          IDL_KW_OFFSETOF(l) },
28        { "READWRITE", IDL_TYP_UNDEF, 1, IDL_KW_OUT|IDL_KW_ZERO,
29          0, IDL_KW_OFFSETOF(var) },
30        { "STRING", TYP_STRING, 1, 0,
31          IDL_KW_OFFSETOF(s_there), IDL_KW_OFFSETOF(s) },
32        { NULL }
33    };

```

C

Figure 6-1: Keyword processing example.

```

34
35 KW_RESULT kw;
36     int i;
37     IDL_ALLTYPES newval;
38
39     (void) IDL_KWProcessByOffset(argc, argv, argk, kw_pars,
40                               (IDL_VPTR *) 0, 1, &kw);
41
42     printf("LONG: <%spresent>\n", kw.l ? "": "not ");
43     printf("FLOAT: %f\n", kw.f);
44     printf("DOUBLE: <%spresent>\n", kw.d_there ? "": "not ");
45     printf("STRING: %s\n",
46           kw.s_there ? IDL_STRING_STR(&kw.s) : "<not present>");
47     printf("ARRAY: ");
48     if (kw.arr_there)
49         for (i = 0; i < kw.arr_n; i++)
50             printf(" %d", kw.arr_data[i]);
51     else
52         printf("<not present>");
53     printf("\n");
54
55     printf("READWRITE: ");
56     if (kw.var) {
57         IDL_Print(1, &kw.var, (char *) 0);
58         newval.l = 42;
59         IDL_StoreScalar(kw.var, TYP_LONG, &newval);
60     } else {
61         printf("<not present>");
62     }
63     printf("\n");
64
65     IDL_KW_FREE;
    }

```

Figure 6-1: Keyword processing example. (Continued)

Executing this routine from the IDL command line, by entering:

```
KEYWORD_DEMO
```

gives the output:

```

LONG: <not present>
FLOAT: 0.000000
DOUBLE: <not present>
STRING: <not present>
ARRAY: <not present>
READWRITE: <not present>

```

Executing it again with keywords specified:

```
A = 56
KEYWORD_DEMO, /LONG, FLOAT=2, DOUBLE=34,$
  STRING="hello", ARRAY=FOUNDGEN(10), READWRITE=A
PRINT, 'Final Value of A: ', A
```

gives the output:

```
LONG: <present>
FLOAT: 2.000000
DOUBLE: <present>
STRING: hello
ARRAY: 0 1 2 3 4 5 6 7 8 9
READWRITE:      56
Final Value of A:      42
```

Those features of this procedure that are interesting in terms of keyword processing are, by line number:

3-15

Every system routine that processes keywords must define a **KW_RESULT** structure type. All output from keyword processing is stored in the fields of this structure. The first field in the **KW_RESULT** structure must always be **IDL_KW_RESULT_FIRST_FIELD**. The remaining fields are dictated by the keywords defined in **kw_pars** below, starting on line 19. The fields with names ending in **_there** are used for the specified field of the **IDL_KW_PAR** structs, and must be type int. The types of the other fields must match their definitions in the relevant **IDL_KW_PAR** and **IDL_KW_ARR_DESC_R** structs.

16-17

The **ARRAY** keyword, defined on line 21, is a read-only array, and requires this array description. Note that the data field specifies the location of the **arr_data** array within **KW_RESULT** where the array contents should be copied, and the **n_offset** field specifies the location of the **arr_n** field where the number of elements actually seen is to be written. Both of these are specified as offsets into **KW_RESULT**, using the **IDL_KW_OFFSET()** macro to compute this. The minimum number of elements allowed is 3, the maximum is 10.

19

The start of the keyword definition array. Notice that all of the keywords are ordered lexicographically (ASCII) and that there is a NULL entry at the end (line 32). Also, all of the mask fields are set to 1, as is the mask argument to **IDL_KWProcessByOffset()** on

line 39. This means that all of the keywords in the list are to be considered valid in this routine.

20

This routine is requesting fast keyword processing. You can learn more about this option in [“Speeding Keyword Processing”](#) on page 134.

21-22

ARRAY is a read-only array. Its value field is therefore the actual address (and not an offset into **KW_RESULT**) of the **IDL_KW_ARR_DESC_R** struct that completes the array definition. This program wants to know explicitly if ARRAY was specified, so the **specified** field is set to the offset within **KW_RESULT** of the **arr_there** field.

23-24

DOUBLE is a scalar keyword of **IDL_TYP_DOUBLE**. It uses the variable **kw.d_there** to know if the keyword is present. Both the specified and value fields are specified as offsets into **KW_RESULT**.

25

FLOAT is an **IDL_TYP_FLOAT** scalar keyword. It does not use the **specified** field of the **IDL_KW_PAR** struct to get notification of whether the keyword is present, so that field is set to 0. Instead, it uses the **IDL_KW_ZERO** flag to make sure that the variable **kw.f** is always zeroed. If the keyword is present, the value will be written into **kw.f**, otherwise it will remain 0. The important point is that the routine can’t tell the difference between the keyword being absent, or being present with a user-supplied value of zero. If this distinction doesn’t matter, such as when the keyword is to serve as an on/off toggle, use this method. If it does matter, use the **specified** field as demonstrated with the DOUBLE keyword, above.

26-27

LONG is a scalar keyword of **IDL_TYP_LONG**. It sets the **IDL_KW_ZERO** flag to get the variable **kw.l** zeroed prior to keyword parsing. The use of the **IDL_KW_VALUE** flag indicates that if the keyword is present, the value 15 (the lower 12 bits of the flags field) will be ORed into the variable **kw.l**.

28-29

The **IDL_KW_OUT** flag indicates that the routine wants the **IDL_VPTR** for READWRITE if it is present. Since **IDL_KW_ZERO** is also set, the variable

kw.var will be zero unless the keyword is present. The specification of **IDL_TYP_UNDEF** here indicates that there is no type conversion or processing applied to **IDL_KW_OUT** keywords.

30-31

The **STRING** keyword demonstrates scalar string keywords.

32

All **IDL_KW_PAR** arrays must be terminated with a **NULL** entry.

35

Every system routine that processes keywords must declare a variable named **kw**, of type **KW_RESULT**. This variable should be a C stack based local variable (C auto class).

37

The **IDL_StoreScalar()** function used on line 59 requires the scalar value to be provided in an **IDL_ALLTYPES** struct.

39-40

Do the keyword processing. The first three arguments are simply the arguments the interpreter passed to the routine. The **plain_args** argument is set to **NULL** because this routine is registered as not accepting any plain arguments. Since no plain arguments will be present, the return value from **IDL_KWProcessByOffset()** is discarded. The final argument is the address of the **KW_RESULT** variable (**kw**) into which the results will be written.

42

The **kw.l** variable will be 0 if **LONG** is not present, and 1 if it is.

43

The **kw.f** variable will always have some usable value, but if it is zero there is no way to know if the keyword was actually specified or not.

44-46

These keywords use the variables from the **specified** field of their **IDL_KW_PAR** struct to determine if they were specified or not. Use of the **IDL_STRING_STR** macro is described in [“Accessing IDL_STRING Values”](#) on page 185.

47-53

Accessing the ARRAY keyword is simple. The **kw.arr_there** variable indicates if the keyword is present, and **kw.arr_n** gives the number of elements.

55-63

Since the READWRITE keyword is accessed via the argument's **IDL_VPTR**, we use the **IDL_Print()** function to print its value. This has the same effect as using the user-level PRINT procedure when running IDL. See “[Output of IDL Variables](#)” on page 248. Then, we change its value to 42 using **IDL_StoreScalar()**.

Again, please note that we use this mechanism in order to create a simple example. You will probably want to avoid the use of this type of output (**printf** and **IDL_Print()**) in your own code.

65

Normal exit from any routine that calls **IDL_KWProcessByOffset()** must be preceded by a call to **IDL_KW_FREE**. This macro releases any dynamic resources that were allocated by keyword processing.

The Pre-IDL 5.5 Keyword API

Versions of IDL prior to IDL 5.5 used a different, but similar, keyword processing API to that found in the current versions. The remainder of this chapter provides information of interest to programmers maintaining older system routines that were written to that API.

Note

ITT Visual Information Solutions recommends that all new code be written using the new keyword processing API. The older API continues to be supported for backwards compatibility, and there is no urgent reason to convert code that uses it. However, the effort of converting old code to the new API is minimal, and can be beneficial.

Background

If you have system routines that were written for use with versions of IDL older than IDL 5.5, your code uses an older keyword processing API, described in “[Processing Keywords With IDL_KWGetParams\(\)](#)” on page 390, that including the following obsolete elements:

- **IDL_KWGetParams()**
- **IDL_KW_ARR_DESC**
- **IDL_KWCleanup()**, **IDL_KW_MARK**, **IDL_KW_CLEAN**

This old API served for many years, but it had some unfortunate features that made it hard to use correctly:

- The rules for when and how to use **IDL_KWCleanup()** were difficult to remember. The programmer had to decide whether or not to call it based on the types of the keywords being processed. If you didn’t call it when you should, memory would be leaked.
- In order to ensure correctness, many programmers would resort to always calling **IDL_KWCleanup()** whether it was is needed or not. This is inefficient, but otherwise fine.
- The use of **IDL_KWCleanup()** is based on a worst case assumption that the keywords that require cleaning will actually be invoked by the IDL user. This is often not the case, and is therefore inefficient.

- Imagine an existing system routine that does not need to use **IDL_KWCleanup()**, and therefore is coded not to use it. If a new keyword should later be added to that routine, and that new keyword should require the use of **IDL_KWCleanup()**, it is very likely that the programmer adding this new keyword will fail to recognize that issue. This leads to memory leaking from a formerly correct routine.
- If a future version of IDL should get a new data type that requires cleaning, that will change the rules for when **IDL_KWCleanup()** needs to be called. Existing code may need to be examined to fix this situation.
- The old keyword API is not reentrant, because it requires static variable addresses to be embedded in the keyword list. This has always been a problem for recursively callable routines (*e.g.* **WIDGET_CONTROL**, which can cause an IDL procedure callback to execute, which can in turn call **WIDGET_CONTROL** again). In the past, we have carefully coded these complex routines to avoid problems, but the large amount of code required is difficult to write and verify. The proper solution is a reentrant keyword API that uses offsets to variables within a structure, instead of actual statically scoped variable addresses. This is what the current API provides.

Advantages Of The IDL 5.5 API

In contrast, keyword processing, in IDL 5.5 and later is built around the **IDL_KWProcessByOffset()** function, has the following advantages:

- The old API remains in place with full functionality. Hence, you are not required to update your old code (There are benefits to such updating, however).
- A transitional API, build around the **IDL_KWProcessByAddr()** function, exists to help ease updating your code. See “[The Transitional API](#)” on page 147 for details. The transitional API is a halfway measure designed to solve the worst problems of the old API while requiring the minimum amount of change.
- The new API, and the transitional API both eliminate the confusing **IDL_KWCleanup()** routine and its requirement to **KW_MARK** before, and **KW_CLEAN** after keyword processing based on the types of the keywords. Instead, the keyword processing API keeps track of the need to cleanup itself, and handles this efficiently. The user is freed from guesswork about how and when to do the required cleanup.

- Keyword cleanup will only happen if the keyword module determines that it is necessary as it processes the actual keywords used. This is more efficient, and it can be easily extended within IDL if a new data type is added to the IDL system, without requiring any change to your code.
- The internal data structures used to maintaining keyword state are now dynamically allocated, and do not have the static limits that the old implementation did.
- The new API is able to process keywords using a re-entrant keyword description. Results are written to stack based (C auto) variables rather than statically defined variables. This can be used to greatly simplify the implementation of recursive system routines, and has been used extensively for that purpose within IDL.

Differences And Similarities Between APIs

The current IDL keyword processing API was designed to minimize the changes necessary to convert existing older code. The differences and similarities between these APIs are summarized below:

- The basic **IDL_KW_PAR** data structure is unchanged between the two. However, in the old API, the **specified**, and **value** fields are addresses to statically allocated C variables or **IDL_KW_ARR_DESC** structures. In the new API, **specified** is always an offset into a user defined **KW_RESULT** structure. The **value** field is an offset into **KW_RESULT** when it is used to refer to data. It is an address when used to refer to an **IDL_KW_ARR_DESC_R** structure.
- The old API uses the **IDL_KW_ARR_DESC** structure to define **IDL_KW_ARRAY** read-only arrays. The new API uses the very similar **IDL_KW_ARR_DESC_R** structure. This is necessary because **IDL_KW_ARR_DESC** is not reentrant (the number of array elements used is written into the struct), while **IDL_KW_ARR_DESC_R** causes them to be written into a field in the **KW_RESULT** variable instead.
- The new API requires all keyword variables to be contained in a single **KW_RESULT** structure, while the old API allowed them to be independent variables. This is important to the offset-based scheme used in the new API, as well as having the nice side effect of improving the organization and readability of most code.
- The old API uses **IDL_KWGetParams()** to process keywords. The new API uses **IDL_KWProcessByOffset()**.

- The old API uses **IDL_KWCleanup()** to free resources. The rules for using it are complicated and lead to latent coding errors. The new API uses the **IDL_KW_FREE** macro, and has a simple consistent rule for use.

Converting Existing Code To The New API

To convert code that uses the old API to the new version:

- Define a typedef for a struct named **KW_RESULT**, containing the keyword variables. Make the first field be the predefined **IDL_KW_RESULT_FIRST_FIELD**.
- Modify your keyword definition list so that the specified and value fields of each **IDL_KW_PAR** struct contain offsets instead of addresses in all cases except when the value field references an **IDL_KW_ARR_DESC** struct. To do this, use the **IDL_KW_OFFSETOF()** macro.
- Any reference to an **IDL_KW_ARR_DESC** structure for an **IDL_KW_ARRAY** keyword must be converted to an **IDL_KW_ARR_DESC_R** struct.
- Replace the call to **IDL_KWGetParams()** with a call to **IDL_KWProcessByOffset()**.
- Remove any **IDL_KWCleanup(IDL_KW_MARK)** calls.
- Replace any **IDL_KWCleanup(IDL_KW_CLEAN)** calls with the **IDL_KW_FREE** macro. Check to ensure that all exit paths from your function other than via **IDL_Message()** include a call to this macro.

The Transitional API

We recommend that you convert your code to the reentrant keyword API based around the **IDL_KWProcessByOffset()** and **IDL_KWFree()** functions. This is almost always a straightforward operation, and the resulting code has all of the advantages discussed in [“Advantages Of The IDL 5.5 API”](#) on page 145. However, there is another alternative that may be useful in some situations. A third keyword API, built around the **IDL_KWProcessByAddr()** function exists that provides the benefits of eliminating the confusing **IDL_KWCleanup()** function, while not requiring the use of static non-reentrant separate variables to change.

The transitional API is a halfway measure designed to solve the worst problems of the old API while requiring the minimum amount of change to your code:

```
int IDL_KWProcessByAddr(int argc, IDL_VPTR *argv, char *argk,
                       IDL_KW_PAR *kw_list, IDL_VPTR *plain_args,
                       int mask, int *free_required)

void IDL_KWFree(void)
```

where:

argc, argv, argk, plain_args, mask

These arguments are the same as those required by **IDL_KWProcessByOffset()**

kw_list

An array of **IDL_KW_PAR** structures, in the absolute address form required by the old **IDL_KWGetParams()** keyword API (the **specified** and **value** fields use address to static C variables).

free_required

The address of an integer to be filled in by **IDL_KWProcessByAddr()**. If set to **TRUE**, the caller must call **IDL_KWFree()** prior to exit from the routine.

Example: Converting From The Old Keyword API

To illustrate the use of the old keyword API, the transitional API, and the new reentrant API, this section provides an extremely simple example, written three times, once with each API.

Another useful comparison is to compare the example [“Keyword Examples”](#) on page 137 with its original version written with the old API which can be found in [“Keyword Examples”](#) on page 394.

Old API

```
IDL_VPTR IDL_someroutine(int argc, IDL_VPTR *argv, char *argk)
{
    static IDL_VPTR count_var;
    static IDL_LONG debug;
    static IDL_STRING name;
    static IDL_KW_PAR kw_pars[] = {
        { "COUNT", 0, 1, IDL_KW_OUT | IDL_KW_ZERO, 0, IDL_CHARA(count_var) },
        { "DEBUG", IDL_TYP_LONG, 1, IDL_KW_ZERO, 0, IDL_CHARA(debug) },
        { "NAME", IDL_TYP_STRING, 1, IDL_KW_ZERO, 0, IDL_CHARA(name) },
    }
```

```

        { NULL }
    };
    IDL_VPTR result;

    IDL_KWcleanup(IDL_KW_MARK);
    argc = IDL_KWGetParams(argc,argv,argk,kw_pars,(IDL_VPTR *)0,1);

    /* Your code goes here. Keyword values are available in the
     * static variables.*/

    /* Cleanup keywords before leaving */
    IDL_KWcleanup(IDL_KW_CLEAN);
    return(result);
}

```

Transitional API

The transitional API provides the benefits of simplified and straightforward cleanup, but does not require you to alter your `IDL_KW_PAR` array or gather the keyword variables into a common structure. The resulting code is very similar to the old API.

```

IDL_VPTR IDL_someroutine(int argc, IDL_VPTR *argv, char *argk)
{
    static IDL_VPTR count_var;
    static IDL_LONG debug;
    static IDL_STRING name;
    static IDL_KW_PAR kw_pars[] = {
        {"COUNT", 0, 1, IDL_KW_OUT|IDL_KW_ZERO,
         0,IDL_KW_ADDROF(count_var) },
        { "DEBUG", IDL_TYP_LONG,1,IDL_KW_ZERO,0,IDL_KW_ADDROF(debug) },
        { "NAME", IDL_TYP_STRING,1,IDL_KW_ZERO,0,IDL_KW_ADDROF(name) },
        { NULL }
    };

    int kw_free;
    IDL_VPTR result;

    argc = IDL_KWProcessByAddr(argc, argv, argk, kw_pars,
                               (IDL_VPTR *) 0, 1, &kw_free);

    /* Your code goes here. Keyword values are available in the
     * static variables.*/

    /* Cleanup keywords before leaving */
    if (kw_free) IDL_KWFree();

    return(result);
}

```

New Reentrant API

```

IDL_VPTR IDL_someroutine(int argc, IDL_VPTR *argv, char *argk)
{
    typedef struct {
        IDL_KW_RESULT_FIRST_FIELD; /* Must be first entry in struct */
        IDL_VPTR count_var;
        IDL_LONG debug;
        IDL_STRING name;
    } KW_RESULT;
    static IDL_KW_PAR kw_pars[] = {
        { "COUNT", 0, 1, IDL_KW_OUT | IDL_KW_ZERO,
          0, IDL_KW_OFFSETOF(count_var) },
        { "DEBUG", IDL_TYP_LONG, 1, IDL_KW_ZERO,
          0, IDL_KW_OFFSETOF(debug) },
        { "NAME", IDL_TYP_STRING, 1, IDL_KW_ZERO,
          0, IDL_KW_OFFSETOF(name) },
        { NULL }
    };

    KW_RESULT kw;
    IDL_VPTR result;

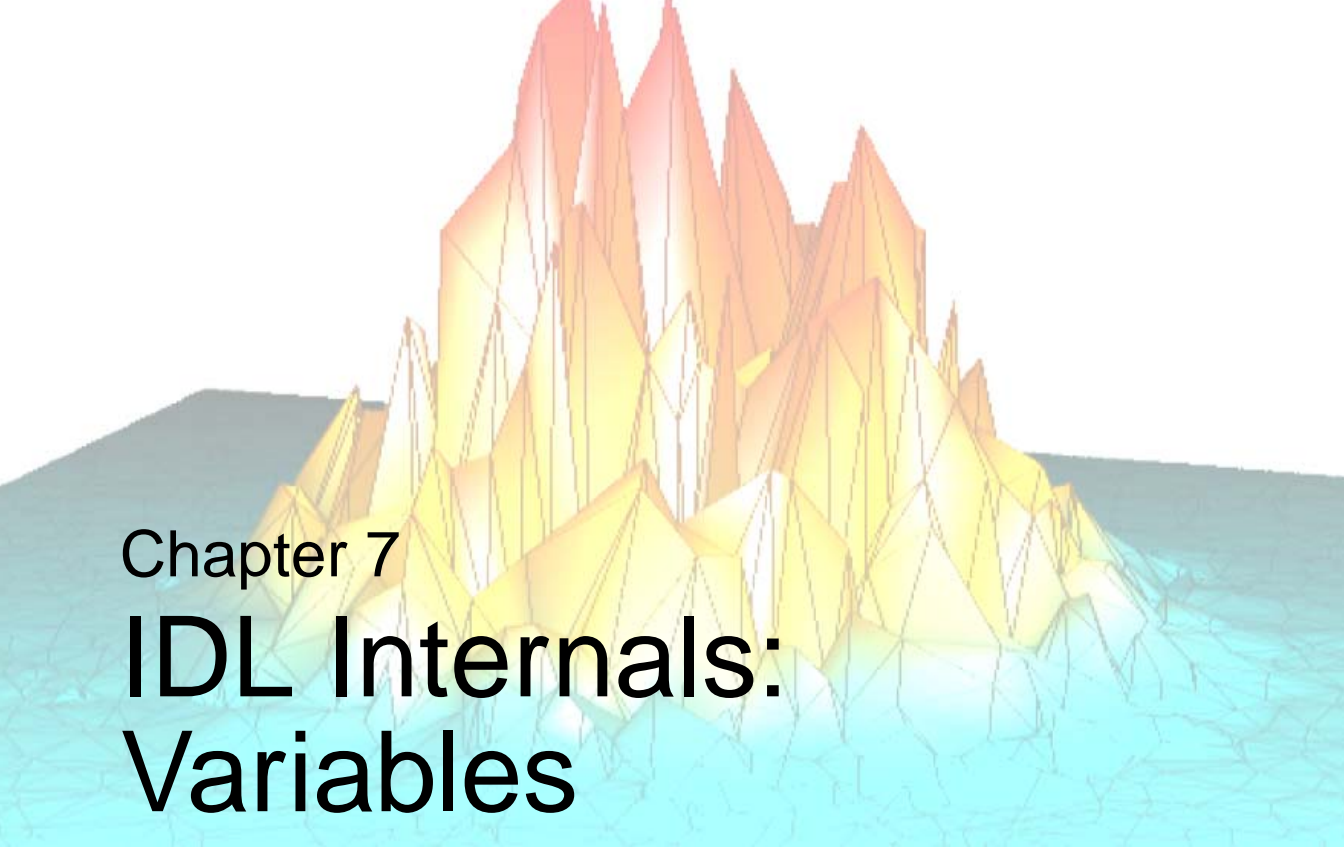
    argc = IDL_KWProcessByOffset(argc, argv, argk, kw_pars,
                                (IDL_VPTR *) 0, 1, &kw);

    /* Your code goes here. Keyword values are available in the
     * kw struct.*/

    /* Cleanup keywords before leaving if necessary */
    IDL_KW_FREE;

    return(result);
}

```



Chapter 7

IDL Internals: Variables

This chapter discusses the following topics:

IDL and Internal Variables	152	Getting Dynamic Memory	174
The IDL_VARIABLE Structure	153	Accessing Variable Data	176
Scalar Variables	156	Copying Variables	177
Array Variables	157	Storing Scalar Values	178
Structure Variables	159	Obtaining the Name of a Variable	180
Heap Variables	164	Looking Up Main Program Variables	181
Temporary Variables	165	Looking Up Variables in Current Scope	182
Creating an Array from Existing Data	172		

IDL and Internal Variables

This chapter describes how variables are created and managed within IDL. While reading this chapter, you should refer to the following figure to see how each part fits into the overall structure of an IDL variable.

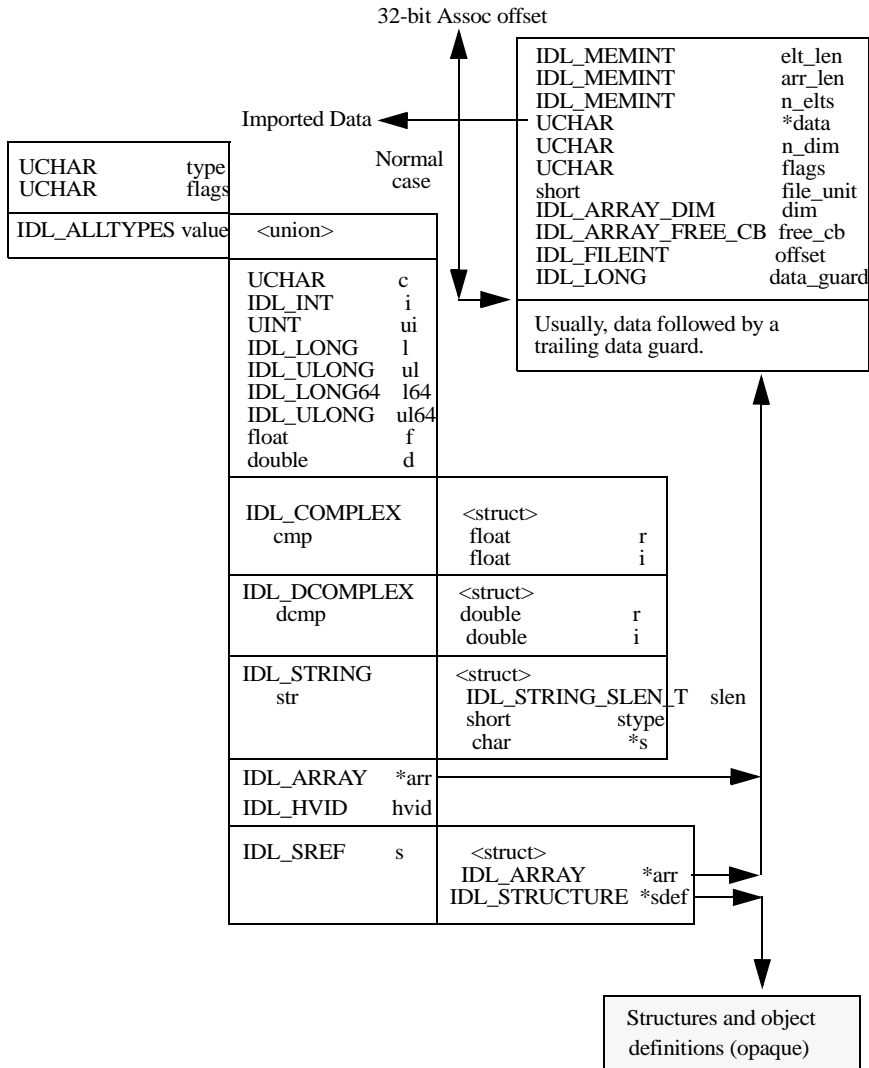


Figure 7-1: Structure of an IDL variable

The IDL_VARIABLE Structure

IDL variables are represented by **IDL_VARIABLE** structures. The definition of **IDL_VARIABLE** is as follows:

```
typedef struct {
    UCHAR type;
    UCHAR flags;
    IDL_ALLTYPES value;
} IDL_VARIABLE;
```

An **IDL_VPTR** is a pointer to an **IDL_VARIABLE** structure:

```
typedef IDL_VARIABLE *IDL_VPTR;
```

The **IDL_ALLTYPES** union is defined as:

```
typedef union {
    UCHAR c; /* Scalar IDL_TYP_BYTE */
    IDL_INT i; /* Scalar IDL_TYP_INT */
    IDL_UINT ui; /* Unsigned short integer value */
    IDL_LONG l; /* Scalar IDL_TYP_LONG */
    IDL_ULONG ul; /* Unsigned long value */
    IDL_LONG64 l64; /* 64-bit integer value */
    IDL_ULONG64 ul64; /* Unsigned 64-bit integer value */
    float f; /* Scalar IDL_TYP_FLOAT */
    double d; /* Scalar IDL_TYP_DOUBLE */
    IDL_COMPLEX cmp; /* Scalar IDL_TYP_COMPLEX */
    IDL_DCOMPLEX dcmp; /* Scalar IDL_TYP_DCOMPLEX */
    IDL_STRING str; /* Scalar IDL_TYP_STRING */
    IDL_ARRAY *arr; /* Pointer to array descriptor */
    IDL_SREF s; /* Structure descriptor */
    IDL_HVID hvid; /* Heap variable identifier */
} IDL_ALLTYPES;
```

The basic scalar types are contained directly in this union, while arrays and structures are represented by the **IDL_ARRAY** and **IDL_SREF** structures that are discussed later in this chapter. The type field of the **IDL_VARIABLE** structure contains one of the type codes discussed in “[Type Codes](#)” on page 114. When a variable is initially created, it is given the type code **IDL_TYP_UNDEF**, indicating that the variable contains no value.

The **flags** field is a bit mask that specifies information about the variable. As a programmer adding code to the IDL system, you will rarely need to set bits in this mask. These bits are set by whatever portion of IDL created the variable. You can check them to make sure the characteristics of the variable fit the requirements of your routine (see “[Checking Arguments](#)” on page 202).

The defined bits in the mask are:

IDL_V_CONST

If this flag is set, the variable is actually a constant. This means that storage for the **IDL_VARIABLE** resides inside the code section of the user procedure or function that used it. The IDL compiler generates such **IDL_VARIABLES** when an expression involving a constant occurs. For example, the IDL statement:

```
PRINT, 23 * A
```

causes the compiler to generate a constant for the “23”. You must not change the value of this type of “variable”.

IDL_V_TEMP

If this flag is set, the variable is a temporary variable. IDL maintains a pool of nameless **IDL_VARIABLES** that can be checked out and returned as needed. Such variables are used by the interpreter to temporarily store the results of expressions on the stack. For example, the statement:

```
PRINT, 2 * 3
```

will cause the interpreter to go through a sequence of events similar to:

1. Push a constant variable for the 2 on the stack.
2. Push a constant variable for the 3 on the stack.
3. Allocate a temporary variable, pop the two constants from the stack, perform the multiplication with the result going into the temporary variable.
4. Push the temporary variable onto the stack.
5. Call the **PRINT** system procedure specifying one argument.
6. Remove the argument to **PRINT** from the stack, and return the temporary variable.

Temporary variables are also used inside user procedures and functions. See [“Temporary Variables”](#) on page 165.

IDL_V_ARR

If this flag is set, the variable is an array, and the value field of the **IDL_VARIABLE** points to an array descriptor.

IDL_V_FILE

If this flag is set, the variable is a file variable, as created by IDL’s **ASSOC** function.

IDL_V_DYNAMIC

If this flag is set, the memory used by this **IDL_VARIABLE** is dynamically allocated. This bit is set for arrays, structures, and for variables of **IDL_TYP_STRING** (because the memory referenced via the string pointer is dynamic).

IDL_V_STRUCT

If this flag is set, the variable is a structure, and the value field of the **IDL_VARIABLE** points to the structure descriptor. For implementation reasons, all structure variables are also arrays, so **IDL_V_STRUCT** also implies **IDL_V_ARR**. Therefore, it is impossible to have a scalar structure. However, single-element structure arrays are quite common.

Because structure variables have their type field set to **IDL_TYP_STRUCT**, the **IDL_V_STRUCT** bit is redundant. It exists for efficiency reasons.

Scalar Variables

A scalar **IDL_VARIABLE** is distinguished by not having the **IDL_V_ARR** bit set in its **flags** field. A scalar variable must have one of the basic data types (IDL structures are never scalar) shown in [Table 7-1](#). The data for a scalar variable is stored in the **IDL_VARIABLE** itself, using the **IDL_ALLTYPES** union. The following table gives the relationship between the data type and the field used.

Scalar Data Type	Field that Stores Data
IDL_TYP_UNDEF	None.
IDL_TYP_BYTE	value.c
IDL_TYP_INT	value.i
IDL_TYP_UINT	value.ui
IDL_TYP_LONG	value.l
IDL_TYP_ULONG	value.ul
IDL_TYP_LONG64	value.l64
IDL_TYP_ULONG64	value.ul64
IDL_TYP_FLOAT	value.f
IDL_TYP_DOUBLE	value.d
IDL_TYP_COMPLEX	value.cmp
IDL_TYP_DCOMPLEX	value.dcmp
IDL_TYP_STRING	value.str
IDL_TYP_PTR	value.hvid
IDL_TYP_OBJ	value.hvid

Table 7-1: Scalar Variable Data Locations

Array Variables

Array variables have the `IDL_V_ARR` bit of their `flags` field set, and the `value.arr` field points to an array descriptor defined by the `IDL_ARRAY` structure:

```
typedef IDL_MEMINT IDL_ARRAY_DIM[IDL_MAX_ARRAY_DIM];

typedef struct {
    IDL_MEMINT elt_len;
    IDL_MEMINT arr_len;
    IDL_MEMINT n_elts;
    UCHAR *data;
    UCHAR n_dim;
    UCHAR flags;
    short file_unit;
    IDL_ARRAY_DIM dim;
} IDL_ARRAY;
```

The meaning of the fields of an array descriptor are:

`elt_len`

The length of each array element in bytes (chars). The array descriptor does not keep track of the types of the array elements, only their lengths. Single elements can get quite long in the case of structures.

For IDL structures, this value includes any padding necessary to properly align the data along required boundaries. On a given platform, IDL creates structures the same way a C compiler does on that platform. As a result, you should not assume that the size of a structure is the sum of the sizes of the structure fields, or that the field offsets are in specific locations.

`arr_len`

The length of the entire array in bytes. This value could be calculated as (`elt_len * n_elts`), but is used so frequently that it is maintained as a separate field in the `IDL_ARRAY` struct.

`n_elts`

The number of elements in the array.

data

A pointer to the data area for the array. This is a region of dynamically allocated memory **arr_len** bytes long. This pointer should be cast to be a pointer of the correct type for the data being manipulated. For example, if the array variable being processed is pointed at by an **IDL_VPTR** named **v** and contains **IDL_TYP_INT** data:

```
IDL_INT *data;          /* Declare a pointer variable */
data = (IDL_INT *) v->value.arr->data;
```

n_dim

The number of array dimensions. The constant **IDL_MAX_ARRAY_DIM** defines the upper limit of this value.

flags

A bit mask that specifies characteristics of the array. Allowed values are:

IDL_A_FILE — This flag indicates that the array is a file variable, as created by the **ASSOC** function. The variable has an array block to specify the structure of the variable, but it has no data area. The data field of the **IDL_ARRAY** structure does not contain useful information, and should not be used.

IDL_A_PACKED — If array is an **IDL_A_FILE** variable and the data type is **IDL_TYP_STRUCT**, then Input/Output to this struct should use a packed data layout compatible with **WRITEU** instead of being a direct mapping onto the struct (which reflects the C compiler layout of the structure including its alignment holes).

file_unit

When the **IDL_A_FILE** bit is set in the **flags** field, **file_unit** contains the IDL Logical Unit Number associated with the variable.

dim

An array that contains the dimensions of the IDL variable. There can be up to **IDL_MAX_ARRAY_DIM** dimensions. The *number* of dimensions in the current array is given by the **n_dim** field.

Structure Variables

Structure variables have the type code **IDL_TYP_STRUCT**. They also have the **IDL_V_STRUCT** bit set in their **flags** field. The **value.s** field of such a variable contains a structure descriptor defined by the **IDL_SREF** structure:

```
typedef struct {
    IDL_ARRAY *arr;      /* ^ to IDL_ARRAY containing data */
    void *sdef;         /* ^ to structure definition */
} IDL_SREF;
```

The **arr** field points at an array block, as described in “[Array Variables](#)” on page 157. It is worth noting that in the definition of the **IDL_ALLTYPES** union (described in “[The IDL_VARIABLE Structure](#)” on page 153), the **arr** field is a pointer to **IDL_ARRAY**, while the **s** field is an **IDL_SREF**, a structure that contains a pointer to **IDL_ARRAY** as its first member.

The resulting definition looks like:

```
union {
    IDL_ARRAY arr;
    struct {
        IDL_ARRAY arr;
        void *sdef;
    } s;
} value;
```

Due to the way C lays out fields in structs and unions, **value.arr** will have the same offset and size within the value union as **value.s.arr**. Therefore, it is possible to access the array block of a structure variable as **var->value.arr** rather than the more correct **var->value.s.arr**. You should avoid use of this shorthand, however, because it is not strictly correct usage and because we reserve the right to change the **IDL_SREF** definition in a way that could cause the memory layout of the **ALLTYPES** union to change.

Creating Structures

The actual structure definition is accessed through the **sdef** field, which is a pointer to an opaque IDL structure definition. Although the implementation of structure definitions is not public information, they can be created using the **IDL_MakeStruct()** function from a structure name and a list of tags:

```
void *IDL_MakeStruct(char *name, IDL_STRUCT_TAG_DEF *tags)
```

name

The name of the structure definition, or NULL for anonymous structures.

tags

An array of **IDL_STRUCT_TAG_DEF** elements, one for each tag.

The result from this function can be passed to **IDL_ImportArray()** or **IDL_ImportNamedArray()**, as described in [“Creating an Array from Existing Data”](#) on page 172.

IDL_STRUCT_TAG_DEF is defined as:

```
typedef struct {
    char *name;
    IDL_MEMINT *dims;
    void *type;
    UCHAR flags;
} IDL_STRUCT_TAG_DEF;
```

name

Null-terminated uppercase name of the tag.

dims

An array that contains information about the dimensions of the structure. The first element of this array is the number of dimensions. Following elements contain the size of each dimension.

type

Either a pointer to another structure definition, or a simple IDL type code cast to void (e.g., **(void *) IDL_TYP_BYTE**).

flags

A bit mask that specifies additional characteristics of the tag. Allowed values are:

IDL_STD_INHERIT — Type must be **IDL_TYP_STRUCT**. This flag indicates that the structure is inherited (inlined) instead of making it a sub-structure as usual.

The following example shows how to define an anonymous structure. Suppose that you want to create a structure whose definition in the IDL language is:

```
{TAG1: 0L, TAG2: FLTARR(2,3,4), TAG3: STRARR(10)}
```


It can be created with **IDL_MakeStruct()** as follows:

```
static IDL_MEMINT one = 1;
static IDL_MEMINT tag2_dims[] = { 3, 2, 3, 4};
static IDL_MEMINT tag3_dims[] = { 1, 10 };
static IDL_STRUCT_TAG_DEF s_tags[] = {
    { "TAG1", 0, (void *) IDL_TYP_LONG},
    { "TAG2", tag2_dims, (void *) IDL_TYP_FLOAT},
    { "TAG3", tag3_dims, (void *) IDL_TYP_STRING},
    { 0 }
};
typedef struct data_struct {
    IDL_LONG tag1_data;
    float tag2_data [4] [3] [2];
    IDL_STRING tag_3_data [10];
} DATA_STRUCT;
static DATA_STRUCT s_data;
void *s;
IDL_VPTR v;

/* Create the structure definition */
s = IDL_MakeStruct(0, s_tags);
/* Import the data area s_data into an IDL structure,
   note that no data are moved. */
v = IDL_ImportArray(1, &one, IDL_TYP_STRUCT,
    (UCHAR *) &s_data, 0, s);
```

Accessing Structure Tags

Given an opaque IDL structure definition, you can determine the offset of the data and a description of its size and form (scalar, array, etc) for a given tag.

IDL_StructTagInfoByName() returns this information given the name of the tag.

IDL_StructTagInfoByIndex() does the same thing, given the numeric index of the tag. They are essentially the same routine, although **IDL_StructTagInfoByIndex()** is slightly more efficient:

```
IDL_MEMINT IDL_StructTagInfoByName(IDL_StructDefPtr sdef,
                                   char *name, int msg_action,
                                   IDL_VPTR *var)
IDL_MEMINT IDL_StructTagInfoByIndex(IDL_StructDefPtr sdef,
                                     int index, int msg_action,
                                     IDL_VPTR *var)
```

where:

sdef

Structure definition for which offset is needed.

name (IDL_StructTagInfoByName)

Name of tag for which information is required.

index (IDL_StructTagInfoByIndex)

Zero based index of tag for which information is required.

msg_action

The parameter that will be passed directly to **IDL_Message()** if the specified tag cannot be found in the supplied structure definition.

var

NULL, or the address of an **IDL_VPTR** to be filled in with a pointer to the variable description for the specified field.

On success, these functions return the data offset of the tag. On error, if the resulting call to **IDL_Message()** returns to the caller, a -1 is returned. The data offset can be added to the data pointer of an IDL variable of this structure type to obtain a pointer to the actual data for that tag.

If the tag is successfully located and the var argument is non-NULL, the **IDL_VPTR** it points at is filled in with a pointer to an **IDL_VARIABLE** structure that describes the type and organization of the tag. It is important to understand that this **IDL_VARIABLE** does not contain any actual data (or in the case of an array tag, a valid data pointer). Hence, the data part of the **IDL_VARIABLE** description should be ignored.

Determining the Number Of Structure Tags

One often needs to know how many tags a structure definition has in order to make use of the information supplied by the routines described above. The **IDL_StructNumTags()** function returns this information:

```
int IDL_StructNumTags(IDL_StructDefPtr sdef)
```

where:

sdef

Structure definition for which offset is needed.

Determining the Names Of Structures and their Tags

The **IDL_StructTagNameByIndex()** function returns the name of a specified tag from a structure definition, and optionally the name of the structure:

```
char *IDL_StructTagNameByIndex(IDL_StructDefPtr sdef, int index,  
                               int msg_action, char **struct_name)
```

where:

sdef

Structure definition for which name information is needed.

index

Zero based index of tag within the structure.

msg_action

The parameter that will be passed directly to **IDL_Message()** if the specified tag cannot be found in the supplied structure definition.

struct_name

NULL, or the address of a character pointer (char *) to be filled in with a pointer to the name of the structure. If the structure is anonymous, the string "<Anonymous>" is returned.

On success, a pointer to the tag name is returned. On error, if the resulting call to **IDL_Message()** returns to the caller, a NULL pointer is returned.

All strings returned by this function must be considered read-only, and must not be modified by the caller.

Heap Variables

Direct access to pointer and object reference heap variables (types **IDL_TYP_PTR** and **IDL_TYP_OBJREF**, respectively) is not allowed. Rather than accessing the heap variable directly, store the value of the heap variable (an IDL pointer or object reference) in a regular IDL variable at the IDL user level. Access the data in the regular variable, then store the results back in the heap variable (via the pointer or object reference) when done.

Note

You can use IDL's **TEMPORARY** function to avoid making copies of the data.

Temporary Variables

As discussed previously, IDL maintains a pool of nameless variables known as temporary variables. These variables are used by the interpreter to hold temporary results from evaluating expressions, and are also used within system procedures and functions that need temporary workspace. In addition, system functions often obtain a temporary variable to return the result of their operation to the interpreter.

Temporary variables have the following characteristics:

- All temporaries, when initially allocated, are of type **IDL_TYP_UNDEF**.
- Temporary variables do not have a name associated with them.
- Routines that check out temporaries must either check them back in or return them as the result of the function. Once you return a temporary variable, you cannot access it again.
- Temporary variables are reclaimed by the interpreter when it is about to exit after executing a program, so it is not possible to lose them and leak dynamic memory by allocating them and failing to return them. If the interpreter is exiting normally and it detects temporaries that have not been returned, it issues an error message. Such an error message indicates an error in the implementation of your system routine. If your routine exits by issuing an **IDL_MSG_LONGJMP** or **IDL_MSG_IO_LONGJMP** error via **IDL_Message()** however, allocated temporaries are expected, and are reclaimed quietly. Hence, your routines need only return temporaries on normal return, and not before issuing errors. See [“IDL Internals: Error Handling”](#) on page 191.

The interpreter uses temporary variables to hold values that are the result of evaluating expressions. Such temporaries are pushed on the interpreter stack where they are often passed as arguments to other routines. For example, the IDL statement:

```
PRINT, MAX(FINDGEN(100))
```

causes the interpreter to perform the following steps:

1. Push a constant variable with the value 100 onto the stack.
2. Call the system function **FINDGEN**, passing it one argument.
3. **FINDGEN** returns a temporary variable which is a 100-element vector with each element set to the value of its index.
4. The interpreter removes the arguments to **FINDGEN** from the stack (the constant 100) and pushes the resulting temporary variable onto the stack.

5. The MAX system function is called with a single argument—the temporary result from FINDGEN.
6. MAX finds the largest element in its argument (99), places that value into a temporary scalar variable, and returns that temporary variable as its result.
7. The interpreter removes the argument to MAX from the stack. This was the temporary array from FINDGEN, so it is returned to the pool of temporary variables. The resulting temporary variable from MAX is then pushed onto the stack.
8. The PRINT system procedure is called with a single argument, which is the temporary scalar variable from MAX. It prints the value of the variable and returns.
9. The interpreter removes the argument to PRINT from the stack, and returns it to the pool of temporary variables.

Getting a Temporary Variable

Temporary variables are obtained via the **IDL_Gettmp()** function:

```
IDL_VPTR IDL_Gettmp(void);
```

IDL_Gettmp() requires no arguments, and returns an IDL_VPTR to a temporary variable. This variable must be returned to the pool of temporary variables (with a call to **IDL_Deltmp()**) or be returned as the value of a system function before control returns to the interpreter, or an error will occur.

A number of variants on **IDL_Gettmp()** exist, as convenience routines for creating temporary scalar variables of a given type and value. In all cases, the value is supplied as the sole argument, and the resulting type is indicated by the name of the routine:

```
IDL_VPTR IDL_GettmpInt(IDL_INT value);  
IDL_VPTR IDL_GettmpUInt(IDL_UINT value);  
IDL_VPTR IDL_GettmpLong(IDL_LONG value);  
IDL_VPTR IDL_GettmpULong(IDL_ULONG value);  
IDL_VPTR IDL_GettmpFILEINT(IDL_FILEINT value);  
IDL_VPTR IDL_GettmpMEMINT(IDL_MEMINT value);
```

Creating a Temporary Array

Temporary array variables can be obtained via the **IDL_MakeTempArray()** function:

```
char *IDL_MakeTempArray(int type, int n_dim, IDL_MEMINT dim[],
                       int init, IDL_VPTR *var)
```

where:

type

The type code for the resulting array. See “[Type Codes](#)” on page 114.

n_dim

The number of array dimensions. The constant **IDL_MAX_ARRAY_DIM** defines the upper limit of this value.

dim

An array of **IDL_MAX_ARRAY_DIM** elements containing the array dimensions. The number of dimensions in the array is given by the **n_dim** argument.

init

Specifies the sort of initialization that should be applied to the resulting array. The **init** argument must be one of the following:

- **IDL_ARR_INI_INDEX** — Each element of the array is set to the value of its index. The INDGEN family of built-in system functions is implemented using this feature.
- **IDL_ARR_INI_NOP** — No initialization is done. The data area of the array will contain whatever garbage was left behind from its previous use. Experience has shown that **IDL_TYP_STRING** data should never be left uninitialized due to the risk of dereferencing an invalid string pointer and crashing IDL. Therefore, **IDL_TYP_STRING** data is zeroed when **IDL_ARR_INI_NOP** is specified.
- **IDL_ARR_INI_ZERO** — The data area of the array is zeroed.

var

The address of an **IDL_VPTR** where the address of the resulting temporary variable will be put.

The data area of an array **IDL_VARIABLE** is accessible from its **IDL_VPTR** as **var->value.arr->data**. However, since most routines that create an array need to access the data area, **IDL_MakeTempArray()** returns the data area pointer as its value. As with **IDL_Gettmp()**, the variable allocated via **IDL_MakeTempArray()** must be returned to the pool of temporary variables or be returned as the value of a system function before control returns to the interpreter, or an error will occur.

Creating a Temporary Vector

IDL_MakeTempArray() can be used to create arrays with any number of dimensions, but the common case of creating a 1-dimensional vector can be carried out more conveniently using the **IDL_MakeTempVector()** function:

```
char *IDL_MakeTempVector(int type, IDL_MEMINT dim, int init,
                        IDL_VPTR *var)where:
```

type, init, var

These arguments are the same as for **IDL_MakeTempArray()**.

dim

The number of elements in the resulting vector.

Creating a Temporary Structure

The **IDL_MakeTempStruct()** allows you to create an IDL structure variable using memory allocated by IDL, in much the same way that **IDL_MakeStruct()** and **IDL_ImportArray()** allow you to create an IDL structure variable using memory you provide. Temporary structure variables can be obtained via the **IDL_MakeTempStruct()** function:

```
char *IDL_MakeTempStruct(IDL_StructDefPtr sdef, int n_dim,
                        IDL_MEMINT dim[], IDL_VPTR *var, int zero)
```

where:

sdef

A pointer to the structure definition.

n_dim

The number of structure dimensions. The constant **IDL_MAX_ARRAY_DIM** defines the upper limit of this value.

dim

A C array of **IDL_MAX_ARRAY_DIM** elements containing the structure dimensions. The number of dimensions in the array is given by the **n_dim** argument.

var

The address of an **IDL_VPTR** where the address of the resulting temporary variable will be put.

The data area of an array **IDL_VARIABLE** is accessible from its **IDL_VPTR** as **var->value.arr->data**. However, since most routines that create an array need to access the data area, **IDL_MakeTempStruct()** returns the data area pointer as its value. As with **IDL_Gettmp()**, the variable allocated via **IDL_MakeTempStruct()** must be returned to the pool of temporary variables (with a call to **IDL_Deltmp()**) or be returned as the value of a system function before control returns to the interpreter, or an error will occur.

zero

Set to TRUE if the data area of the resulting variable should be zeroed, or to FALSE otherwise. Unless the caller intends to immediately copy a valid result into the variable, this argument should be set to TRUE to prevent memory corruption.

Creating a Temporary Vector

IDL_MakeTempStruct() can be used to create arrays with any number of dimensions, but the common case of creating a 1-dimensional vector can be carried out more conveniently using the **IDL_MakeTempStructVector()** function:

```
char *IDL_MakeTempStructVector(IDL_StructDefPtr sdef, IDL_MEMINT
dim,
                                IDL_VPTR *var, int zero)
```

where:

sdef, var, zero

These arguments are the same as for **IDL_MakeTempStruct()**.

dim

The number of elements in the resulting vector.

Creating A Temporary Variable Using Another Variable As A Template

It is common to want to create a temporary variable with a form that mimics that of a variable you already have access to. Often, such a temporary variable has the same number of elements and dimensions, but may vary in type. It is possible to do this by using the basic temporary variable creation routines discussed earlier in this chapter, but the resulting code will be complex, and this sort of code occurs frequently. The best way to create such a variable is using the **IDL_VarMakeTempFromTemplate()** function.

IDL_VarMakeTempFromTemplate() creates a temporary variable of the desired type, using the **template_var** argument to specify its dimensionality. The address of this temporary variable is stored at the address specified by the **result_addr** argument. The address of the start of this variable's data area is returned as the value of the function.

```
char *IDL_VarMakeTempFromTemplate(IDL_VPTR template_var,int type,
                                  IDL_StructDefPtr sdef,
                                  IDL_VPTR *result_addr,int zero);
```

where:

template_var

Source variable to take dimensionality from. This can be a scalar or array of any type.

type

The IDL type code for the desired temporary variable.

sdef

NULL, or a pointer to a structure definition. This argument is ignored if **type** is not **IDL_TYP_STRUCT**. If type is **IDL_TYP_STRUCT**, **sdef** supplies the structure definition for the result. It is an error to specify a result type of **IDL_TYP_STRUCT** without providing a value for **sdef**, with one exception: If type is **IDL_TYP_STRUCT** and **template_var** is a variable of **IDL_TYP_STRUCT**, and **sdef** is NULL, then **IDL_VarMakeTempFromTemplate()** will use structure definition of **template_var**.

result_addr

Address of **IDL_VPTR** to receive a pointer to the newly allocated temporary variable.

zero

TRUE if the resulting variable should be zeroed, and FALSE to not do this. Variables of **IDL_TYP_STRING**, and structure types that contain strings, are always zeroed.

Freeing A Temporary Variable

Use **IDL_Deltmp()** to free a temporary variable:

```
void IDL_Deltmp(IDL_VPTR p)
```

where **p** is an **IDL_VPTR** to the temporary variable to be returned. **IDL_Deltmp()** frees the dynamic parts of the temporary variable (if any) and then returns the variable to the pool of available temporaries. Once you have deallocated a temporary variable, you may not access it again. There is also a macro named **IDL_DELTMP** which checks its argument to make sure it's a temporary, and if so, calls **IDL_Deltmp()** to return it.

Creating an Array from Existing Data

There are two functions that allow you to create an IDL array variable whose data points at data you supply rather than having IDL allocate the data space. The routine **IDL_ImportArray()** returns a temporary variable, while **IDL_ImportNamedArray()** returns a named variable in the current execution scope, creating the new variable if necessary. Your data must already exist in memory. The data area, which can be quite large, is not copied. These functions simply create variable and array descriptors that point to the data you supply and return the pointer to the resulting variable. Their definitions are:

```
IDL_VPTR IDL_ImportArray(int n_dim, IDL_MEMINT dim[], int type,
    UCHAR *data, IDL_ARRAY_FREE_CB free_cb, void *s)

IDL_VPTR IDL_ImportNamedArray(char *name, int n_dim,
    IDL_MEMINT dim[], int type, UCHAR *data,
    IDL_ARRAY_FREE_CB free_cb, void *s)

typedef void (* IDL_ARRAY_FREE_CB) (UCHAR *);
```

where:

name

The name of the variable to be created or modified.

n_dim

The number of dimensions in the array.

dim

An array of **IDL_MAX_ARRAY_DIM** elements, containing the size of each dimension.

type

The IDL type code describing the data. See [“Type Codes”](#) on page 114.

data

A pointer to your array data. Your data will not be modified unless the user explicitly modifies elements of the array using subscripts.

The temporary variable returned by **IDL_ImportArray()** can be used immediately in an expression, in which case the descriptors are freed immediately. It can also be assigned to a longer-lived variable using **IDL_VarCopy()**.

Note

IDL frees only the memory that it allocates for the descriptors, *not* the memory that you supply containing your data. You can arrange to be notified when IDL is finished with your data by using the **free_cb** argument, described below.

free_cb

If non-NULL, **free_cb** is a pointer to a function that will be called when IDL frees the array. This feature gives the caller a sure way to know when IDL is no longer referencing data. Use the called function to perform any required cleanup such as freeing dynamic memory or releasing shared or mapped memory. The called function should have no return value and should accept as its argument a (**uchar ***), which is a pointer to the memory to be freed.

S

If the type of the variable is **IDL_TYP_STRUCT**, **s** points to the opaque structure definition, as returned by **IDL_MakeStruct()**.

Getting Dynamic Memory

Many programs need to get dynamic memory for some temporary calculation. In the C language, the functions **malloc()** and **free()** are used for this purpose, while other languages have their own facilities. IDL provides its own memory allocation routines (see “[Dynamic Memory](#)” on page 252). Use of such facilities within the IDL interpreter and the system routines can lead to the loss of usable dynamic memory. The following code fragment demonstrates how this can happen.

For example, assume that there is a need for 100 IDL_LONG integers:

```
char *c;

c = (char *) IDL_MemAlloc((unsigned) (sizeof(IDL_LONG) * 100)
                        (char *) 0, IDL_MSG_RET);
.
.
.
if (some_error_condition) IDL_Message(..., IDL_MSG_LONGJMP,...);
.
.
.
IDL_MemFree((void *) c, (char *) 0, IDL_MSG_RET);
```

In the normal case, the allocated memory is released exactly as it should be. However, if an error causes the **IDL_Message()** function to be called, execution will return directly to the interpreter and this code will never have a chance to clean up. The dynamic memory allocated will therefore leak, and although it will continue to occupy space in the IDL processes, will not be used again.

The IDL_GetScratch Function

To solve this problem, use a temporary variable to obtain dynamic memory. Then, if an error should cause execution to return to the interpreter, the interpreter will reclaim the temporary variable and no dynamic memory will be lost. This frequently-needed operation is provided by the **IDL_GetScratch()** function:

```
char *IDL_GetScratch(IDL_VPTR *p, IDL_MEMINT n_elts,
                    IDL_MEMINT elt_size)
```

where:

p

The address of an **IDL_VPTR** that should be set to the address of the temporary variable allocated.

n_elts

The number of elements for which memory should be allocated.

elt_size

The length of each element, in bytes.

Once the need for the temporary memory has passed, it should be returned using the **IDL_Deltmp()** function. Using these functions, the above example becomes:

```
char *c;
IDL_VPTR v;

c = IDL_GetScratch(&v, 100L, (IDL_LONG) sizeof(IDL_LONG));
.
.
.
if (some error condition) IDL_Message(...,MSG_LONGJMP,...);
.
.
.
IDL_Deltmp(v);
```

Using the **IDL_GetScratch()** and **IDL_Deltmp()** functions is similar to using **IDLMemAlloc()** directly. In fact, IDL uses **IDL_MemAlloc()** and **IDL_MemFree()** internally to implement array variables. The important difference is that dynamic memory doesn't leak when error conditions occur.

To avoid losing dynamic memory, always use the **IDL_GetScratch()** function in preference to other ways of allocating dynamic memory, and use **IDL_Deltmp()** to return it.

Accessing Variable Data

Often, we are not concerned with the distinction between a scalar and array variable—all that is desired is a pointer to the data and to know how many elements there are. **IDL_VarGetData()** can be used to obtain this information:

```
void IDL_VarGetData(IDL_VPTR v, IDL_MEMINT *n, char **pd,
                   int ensure_simple)
```

where:

v

The variable for which data is desired.

n

The address of a variable that will hold the number of elements.

pd

The address of variable that will hold a pointer to data, cast to be a pointer to a pointer to a character (for example (**char ****) **&myptr**).

ensure_simple

If TRUE, this routine calls the **IDL_ENSURE_SIMPLE** macro on the argument **v** to screen out variables of the types it prevents. Otherwise, **IDL_EXCLUDE_FILE** is called, because file variables have no data area to return.

On exit, **IDL_VarGetData()** stores the data count and pointer into the variables pointed at by **n** and **pd**, respectively.

Copying Variables

To copy the contents of one variable to another, use the **IDL_VarCopy()** function:

```
void IDL_VarCopy(IDL_VPTR src, IDL_VPTR dst)
```

Arguments `src` and `dst` are the source and destination, respectively.

IDL_VarCopy() uses the following rules when copying variables:

- If the destination variable already has a dynamic part, this dynamic part is released.
- The destination becomes a copy of the source.
- If the source is a temporary variable, **IDL_VarCopy()** does not make a duplicate of the dynamic part for the destination. Instead, the dynamic part of the source is given to the destination, and the source variable itself is returned to the pool of free temporary variables. This is the equivalent of freeing the temporary variable. Therefore, the variable must not be used any further and the caller should not explicitly free the variable. This optimization significantly improves resource utilization and performance because this special case occurs frequently.

Storing Scalar Values

The `IDL_StoreScalar()` function sets an `IDL_VARIABLE` to a scalar value:

```
void IDL_StoreScalar(IDL_VPTR dest, int type,  
                    IDL_ALLTYPES *value)
```

where:

dest

An `IDL_VPTR` to the `IDL_VARIABLE` in which the scalar should be stored.

type

The type code for the scalar value. See [“Type Codes”](#) on page 114.

value

The address of the `IDL_ALLTYPES` union that contains the value to store.

If `dest` is a location that cannot be stored into (for example, a temporary variable, constant, and so on), an error is issued and control returns to the interpreter.

Otherwise, any dynamic part of `dest` is freed and `value` is stored into it.

The `IDL_StoreScalarZero()` function is a specialized variation of `IDL_StoreScalar()`. It stores a zero scalar value of any specified type into the specified variable:

```
void IDL_StoreScalarZero(IDL_VPTR dest, int type)
```

where:

dest

An `IDL_VPTR` to the `IDL_VARIABLE` in which the scalar zero should be stored.

type

The type code for the scalar zero value. See [“Type Codes”](#) on page 114.

Using `IDL_StoreScalar()` to Free Dynamic Resources

In addition to performing its primary function, `IDL_StoreScalar()` and `IDL_StoreScalarZero()` have two very useful side effects:

1. Storing a scalar value in a variable causes IDL to free any dynamic memory currently used by that variable.
2. These routines do the required error checking to make sure the variable allows a new value to be stored into it before performing the actual storage operation.

Often, a system routine accepts an input argument that will have a new value assigned to it before the routine returns to its caller, and the initial value of that argument is of no interest to the routine. Storing a scalar value into such an argument at the start of the routine will automatically check it for storability and free unnecessary dynamic memory. In one easy operation, the required error checking is done, and you've improved the dynamic memory behavior of the IDL system by minimizing dynamic memory fragmentation. For example:

```
IDL_StoreScalarZero(&v, IDL_TYP_LONG);
```

Error handling is discussed further in [“IDL Internals: Error Handling”](#) on page 191.

Obtaining the Name of a Variable

The `IDL_VarName()` function returns the name of a variable, constant, or expression given its address. If the item is a named variable, it must be in the currently active program unit:

```
char *IDL_VarName (IDL_VPTR v)
```

Looking Up Main Program Variables

The `IDL_GetVarAddr()` function returns the address of a *main program* variable, given its name:

```
IDL_VPTR IDL_GetVarAddr(char *name)
```

name

Points to the null terminated name of the variable, which must be in upper case.

The return value is NULL if the variable does not exist, otherwise the pointer to the variable is returned.

Alternatively, `IDL_GetVarAddr1()` will enter a new variable into the symbol table of the main program if called with the parameter **ienter** set to TRUE, and the specified variable name does not already exist. Otherwise, its operation is the same as `IDL_GetVarAddr()`. Note that new variables cannot be created if a user procedure or function is active. `IDL_GetVarAddr1()` is called as shown following:

```
IDL_VPTR IDL_GetVarAddr1(char *name, int enter)
```

name

Points to the null-terminated name of the variable, which must be in upper case.

ienter

Set this parameter to TRUE to create the variable if it does not already exist.

If **ienter** is TRUE and the specified variable name does not already exist, it will be added to the symbol table of the main program. If **ienter** is FALSE, `IDL_GetVarAddr1()` is equivalent to `IDL_GetVarAddr()`.

Looking Up Variables in Current Scope

The `IDL_FindNamedVariable()` function returns the address of a variable in the *current execution scope* given its name:

```
IDL_VPTR IDL_FindNamedVariable(char *name, int ienter)
```

name

Name of the variable to find.

ienter

Set this parameter to `TRUE` to create the variable if it does not already exist.

If the variable is found (or created if **ienter** is `TRUE`), its `IDL_VPTR` is returned. Otherwise, `NULL` is returned.

Note

Even if **ienter** is `TRUE`, this routine can return `NULL` if creating the variable is not possible due to memory constraints.



Chapter 8

IDL Internals: String Processing

This chapter discusses the following topics:

String Processing and IDL	184	Deleting Strings	187
Accessing IDL_STRING Values	185	Setting an IDL_STRING Value	188
Copying Strings	186	Obtaining a String of a Given Length ...	189

String Processing and IDL

A number of functions exist to simplify the processing of **IDL_STRING** descriptors. By using these functions instead of doing your own string management, you can eliminate a common source of errors.

Accessing IDL_STRING Values

It is important to realize that the `s` field of an **IDL_STRING** struct does not contain a valid string pointer in the case of a null string (i.e., when `slen` is zero). A common error that can cause IDL to crash is illustrated by the following code fragment:

```
void print_str(IDL_STRING *s)
{
    printf("%s", s->s);
}
```

The problem with this code is that it fails to consider the case where the argument `s` describes a null string. The proper way to write this code is as follows:

```
void print_str(IDL_STRING *s)
{
    printf("%s", IDL_STRING_STR(s));
}
```

The macro **IDL_STRING_STR** takes as its argument a pointer to an **IDL_STRING** struct. If the string is null, it returns a pointer to a zero length null-terminated string, otherwise it returns the string pointer from the struct. Consistent use of this macro will avoid the most common sort of error involving strings.

It is common for IDL system routines to accept arguments that provide names. Such arguments must be scalar strings, or string arrays that contain a single element. To properly process such an argument, it is necessary to screen out non-string types or multi-element arrays, locate the string descriptor, and use the **IDL_STRING_STR()** macro to extract a usable NULL terminated C string from it. The **IDL_VarGetString()** is used for this purpose. It encapsulates all of the error checking, and always returns a pointer to a NULL terminated C string, throwing the proper **IDL_MSG_LONGJMP** error via the **IDL_Message()** function when this is not possible:

```
char *IDL_VarGetString(IDL_VPTR v)
```

where

v

Variable from which string value is desired.

Copying Strings

It is often necessary to copy one string to another. Assume, for example, that there are two string descriptors `s_src` and `s_dst`, and that `s_dst` contains garbage. It would almost suffice to simply copy the contents of `s_src` into `s_dst`. The reason this is not quite correct is that both descriptors would then contain a pointer to the same string. This aliasing can cause some strange effects, or even cause IDL to crash if one of the two descriptors is freed and the string from the other is accessed.

`IDL_StrDup()` takes care of this problem by allocating memory for a second copy of the string, and replacing the string pointer in the descriptor with a pointer to the fresh copy. Naturally, if the string descriptor is for a null string, nothing is done.

```
void IDL_StrDup(IDL_STRING *str, IDL_MEMINT n)
```

where:

str

Pointer to one or more `IDL_STRING` descriptors which need their strings duplicated.

n

The number of descriptors.

The proper way to copy a string is:

```
s_dst = s_src;          /* Copy the descriptor */
IDL_StrDup(&s_dst, 1L); /* Duplicate the string */
```

Deleting Strings

Before an **IDL_STRING** can be discarded or re-used, it is important to release any dynamic memory it might be using. The **IDL_StrDelete()** function should be used to delete strings:

```
void IDL_StrDelete(IDL_STRING *str, IDL_MEMINT n)
```

where:

str

Pointer to one or more **IDL_STRING** descriptors which need their contents freed.

n

The number of descriptors.

IDL_StrDelete() deletes all dynamic memory used by the **IDL_STRING**s. The descriptors contain garbage once this has been done, and their contents should not be used.

The **IDL_Deltmp()** function automatically calls **IDL_StrDelete()** when returning temporary variables of type **IDL_TYP_STRING**, so it is not necessary or desirable to call **IDL_StrDelete()** explicitly in this case.

Setting an IDL_STRING Value

The **IDL_StrStore()** function should be used to store a null-terminated C string into an **IDL_STRING** descriptor:

```
void IDL_StrStore(IDL_STRING *s, char *fs)
```

where:

s

Pointer to an **IDL_STRING** descriptor. This descriptor is assumed to contain garbage, so call **IDL_StrDelete()** on it first if this is not the case.

fs

Pointer to the null-terminated string to be copied into s.

IDL_StrStore() is useful for placing a string value into an **IDL_STRING**. This **IDL_STRING** does not need to be a component of a **VARIABLE**, which makes this function very flexible.

One often needs a temporary, scalar **VARIABLE** of type **IDL_TYP_STRING** with a given value. The function **IDL_StrToSTRING()** fills this need:

```
IDL_VPTR IDL_StrToSTRING(char *s)
```

where:

s

Pointer to the null-terminated string to be copied into the resulting temporary variable.

Obtaining a String of a Given Length

Sometimes you need to make sure that the string in an **IDL_STRING** descriptor has a specific length. The **IDL_StrEnsureLength()** function can be used in this case:

```
void IDL_StrEnsureLength(IDL_STRING *s, int n)
```

where:

s

A pointer to the **IDL_STRING** that will have its length checked.

n

The number of characters the string must be able to contain, not including the terminating null character.

If the **IDL_STRING** passed already has enough room for the specified number of characters, it is not re-allocated. Otherwise, the existing string is freed and a new string of sufficient length is allocated. In either case, the **slen** field of the **IDL_STRING** will be set to the requested length.

If a new dynamic string is allocated, it will contain garbage values because **IDL_StrEnsureLength()** only allocates memory of the specified size, it does not copy a value into it. Therefore, the calling routine must copy a null-terminated string into the new dynamic string.



Chapter 9

IDL Internals: Error Handling

This chapter discusses the following topics:

Message Blocks	192	Looking Up A Message Code by Name ..	201
Issuing Error Messages	195	Checking Arguments	202

Message Blocks

IDL maintains messages in opaque data structures known as *Message Blocks*. A message block contains all the messages for a logically related area.

When IDL starts, there is only one defined block named **IDL_MBLK_CORE**, containing all messages defined by the core IDL product. Typically, dynamically loadable modules (DLMs) each define a message block for their error messages when they are loaded (See “[Dynamically Loadable Modules](#)” on page 309 for a description of DLMs).

There are often two versions of IDL message module functions. Those with names that end in **FromBlock** require an explicit message block. The versions that do not end in **FromBlock** use the **IDL_MBLK_CORE** message block.

To define a message block, you must supply an array of **IDL_MSG_DEF** structures:

```
typedef struct {
    char *name;
    char *format;
} IDL_MSG_DEF;
```

where:

name

A string giving the name of the message. We suggest that you adopt a consistent unique prefix for all your error codes. All message codes defined by ITT Visual Information Solutions start with the prefix **IDL_M_**. You should not use this prefix when naming your blocks in order to avoid unnecessary name collisions.

format

A format string, in printf(3) format. There is one extension to the printf formatting codes: If the first two letters of the format are “%N”, then IDL will substitute the name of the currently executing IDL procedure or function (if any) followed by a colon and a space when this message is issued. For example:

```
IDL> print, undefined_var
% PRINT: Variable is undefined: UNDEFINED_VAR.
```

The **IDL_MessageDefineBlock()** function is used to define a new message block:

```
IDL_MSG_BLOCK IDL_MessageDefineBlock
(char *block_name, int n, IDL_MSG_DEF *defs)
```


The arguments to **IDL_MessageDefineBlock()** are as follows:

block_name

Name of the message block. This can be any string, but it will be case folded to upper case. We suggest a single word be used. It is important to pick names that are unlikely to be used by any other application. All blocks defined by ITT Visual Information Solutions start with the prefix **IDL_MBLK_**. You should not use this prefix when naming your blocks in order to avoid unnecessary confusion.

n

of message definitions pointed at by defs.

defs

An array of message definition structs, each one supplying the name and format string for a message in printf(3) format. The memory used for this array, including the strings it points at, must be in permanently allocated read-only storage. IDL does not copy this memory, but simply uses it in place.

If possible, the new message block is defined and an opaque pointer to it is returned. This pointer must be supplied to subsequent calls to the “FromBlock” message module functions to identify the message block a given error is being issued from. If it is not possible to define the message block, this function returns NULL.

The message functions require a message block pointer and the negative index of the specific message to be issued. Hence, message codes start and zero and grow negatively. For mnemonic convenience, it is standard practice to define preprocessor macros to represent the error codes.

Example: Defining A Message Block

The following code defines a message block named TESTMODULE that contains two messages:

```
static IDL_MSG_DEF msg_arr[] =
{
#define M_TM_INPRO 0
  { "M_TM_INPRO", "%NThis is from a loadable module procedure."
},
#define M_TM_INFUN -1
  { "M_TM_INFUN", "%NThis is from a loadable module function."
},
};
```

```
msg_block = IDL_MessageDefineBlock("Testmodule",  
                                   sizeof(msg_arr)/sizeof(msg_arr[0]),  
                                   msg_arr);
```

Issuing Error Messages

Errors are reported using one of the following functions:

- **IDL_Message()**
- **IDL_MessageFromBlock()**
- **IDL_MessageSyscode()**
- **IDL_MessageSyscodeFromBlock()**

These functions are patterned after the standard C library **printf()** function. They are really the same function, differing in which message block the error is issued from (the **FromBlock** versions allow you to specify the block) and their reporting of system errors that might accompany IDL errors (the **Syscode** versions allow you to specify a system error). IDL documentation often refers to **IDL_Message()**. This should be understood to be a generic reference to any of these four functions.

```
void IDL_Message(int code, int action, ...)
void IDL_MessageFromBlock(IDL_MSG_BLOCK block, int code,
                          int action, ...)
void IDL_MessageSyscode(int code, IDL_MSG_SYSCODE_T syscode_type,
                        int syscode, int action, ...)
void IDL_MessageSyscodeFromBlock(IDL_MSG_BLOCK block, int code,
                                  IDL_MSG_SYSCODE_T syscode_type,
                                  int syscode, int action, ...)
```

The arguments to are as follows:

block

Pointer to the IDL message block from which the error should be issued. If block is a NULL pointer, the default IDL core block (**IDL_MBLK_CORE**) is used.

code

This is the error code associated with the error message to be issued. There are two error codes in the default IDL core block (**IDL_MBLK_CORE**) that are available to programmers adding system routines to IDL. The use of these codes is described below. See “[IDL_M_GENERIC](#)” on page 199 and “[IDL_M_NAMED_GENERIC](#)” on page 199.

Note

For any significant development involving an IDL system routine, we recommend your code be packaged as a Dynamically Loadable Module (DLM), and that your DLM define a message block to contain its errors instead of using the GENERIC core block messages.

syscode_type

IDL_Message() always issues a single-line error message that describes the problem from IDL's point of view. Often, however, there is an underlying system reason for the error that should also be displayed to give the user a complete picture of what went wrong. For example, the IDL view of the problem might be "Unable to open file," while the underlying system reason for the error is "no such directory." The **IDL_MessageSyscode()** functions allow you to include the relevant system error code, and have it incorporated into the IDL message on a second line of output. There are several different types of system error code that can be specified. The **syscode_type** argument is used to tell **IDL_MessageSyscode()** which type of system error is present:

IDL_MSG_SYSCODE_NONE — Indicates that there is no system error. In this case, the **syscode** argument is ignored, and **IDL_MessageSyscode()** is functionally equivalent to **IDL_Message()**.

IDL_MSG_SYSCODE_ERRNO — The UNIX operating system uses a system provided global variable named **errno** for communicating system level errors. Whenever a call to a system function fails, it returns a value of -1, and puts an error code into **errno** that specifies the reason for the failure. Other functions, such as those provided by the standard C library, do not set **errno**. The system documentation (man pages) describes which functions do and do not set **errno**, and the rules for interpreting its value.

The C programming language and UNIX operating system share a common heritage, as C was originally created by its authors as an implementation language for UNIX. Since then, C has found broad acceptance on non-UNIX platforms, bringing along with standard POSIX libraries that provide functionality commonly expected by C programs. Hence, although **errno** is a UNIX concept, non-UNIX C implementations generally provide it as a convenience. Hence, IDL supports **IDL_MSG_SYSCODE_ERRNO** on all platforms.

You should specify **IDL_MSG_SYSCODE_ERRNO** only if you are calling **IDL_MessageSyscode()** as the result of a failed function that is documented to set **errno** on your target platform. Otherwise, **errno** might contain an unrelated

garbage value resulting in an incorrect error message. When specifying **IDL_MSG_SYSCODE_ERRNO**, you should supply the current value of **errno** as the **syscode** argument to **IDL_MessageSyscode()**.

The Microsoft Windows operating system has **errno** for compatibility with the expectations of C programmers, but typically does not set it. On this operating system, specifying **IDL_MSG_SYSCODE_ERRNO** may have no effect.

IDL_MSG_SYSCODE_WIN (Microsoft Windows Only) — Microsoft Windows system error codes. The value supplied to the **syscode** argument to **IDL_MessageSyscode()** should be a system error code, as returned by the Windows **GetLastError()** system function.

IDL_MSG_SYSCODE_WINSOCK (Microsoft Windows Only) — Microsoft Windows winsock error codes. The value supplied to the **syscode** argument to **IDL_MessageSyscode()** should be a system error code, as returned by the Windows **WSAGetLastError()** system function.

syscode

Value of the system error code that should be reported. This argument is ignored if its value is zero (0), or if **syscode_type** is **IDL_MSG_SYSCODE_NONE**. Otherwise, it is interpreted as an error code of the type given by **syscode_type**, and the text of the specified system error will be output along with the IDL message on a separate second line.

action

IDL_Message() can take a number of different actions after issuing the error message. The action to take is specified by the **action** argument:

IDL_MSG_RET

Use this argument to make **IDL_Message()** return to the caller after issuing the error message. In this case, the calling routine can either continue or return to the interpreter as it sees fit.

IDL_MSG_INFO

Use this argument to issue a message that is not an error, but is simply informational in nature. The message is output and **IDL_Message()** returns to the caller. Normally, **IDL_Message()** sets the values of IDL's **!ERROR_STATE** system variables, but not in this case.

IDL_MSG_EXIT

Use this argument to cause the IDL process to exit after the message is issued. This code should never be used in a system function or procedure—it is intended for use in other sections of the system.

IDL_MSG_LONGJMP

Use this argument to cause **IDL_Message()** to exit directly back to the interpreter after issuing the message. In this case, **IDL_Message()** does not return to its caller. It is an error to use this action code in code not called by the IDL interpreter since the resulting call to **longjmp()** will be invalid.

IDL_MSG_IO_LONGJMP

This action code is exactly like **IDL_MSG_LONGJMP**, except that it is issued in response to an input/output error. This code is only used by the I/O module. User written system routines should use the existing I/O routines, so they do not need to use this action.

In addition, the following modifier codes can be ORed into the action code. They modify the normal behavior of **IDL_Message()**:

IDL_MSG_ATTR_NOPRINT

Suppress the printing of the error message, but do everything else in the normal way.

IDL_MSG_ATTR_MORE

Use paging in the style of the UNIX **more** command to display the output. This option exists primarily for use by the IDL compiler, and is unlikely to be of interest to authors of system routines.

IDL_MSG_ATTR_NOPREFIX

Normally, **IDL_Message()** prefixes the output message with the string contained in IDL's **!MSG_PREFIX** system variable.

IDL_MSG_ATTR_NOPREFIX suppresses this prefix string.

IDL_MSG_ATTR_QUIET

If the **IDL_MSG_INFO** action has been specified and this bit mask has been included, and the IDL user has IDL's **!QUIET** system variable, **IDL_Message()** returns without issuing a message.

IDL_MSG_ATTR_NOTRACE

Set this code to inhibit the traceback portion of the error message.

IDL_MSG_ATTR_BELL

Set this code to ring the bell when the message is output.

...

The message format string (specified by the **code** argument) specifies a format string to be used for the error message. This format string is exactly like those used by the standard C library **printf()** function. Any arguments following action are taken to be arguments for this format string.

Error Codes

As mentioned above, ITT Visual Information Solutions has reserved two error codes for use by writers of system routines. They are:

IDL_M_GENERIC

This message code simply specifies a format string of “%s”. The first argument after **action** is taken to be a null-terminated string that is substituted into the format string. For example, the C statement:

```
IDL_Message(IDL_M_GENERIC, IDL_MSG_LONGJMP, "Error! Help!")
```

causes IDL to abort the current routine and issue the message:

```
% Error! Help!
```

IDL_M_NAMED_GENERIC

This message code is exactly like the one above, except that it prints the name of the system routine in front of the error string. For example, assuming that the name of the routine is **MY_PROC**, the C statement:

```
IDL_Message(IDL_M_NAMED_GENERIC, IDL_MSG_LONGJMP,
            "Error! Help!")
```

causes IDL to interrupt the current routine and issue the message:

```
% MY PROC: Error! Help!
```

Choosing an Error Code

Note

For any significant development involving an IDL system routine, we recommend your code be packaged as a Dynamically Loadable Module (DLM), and that your DLM define a message block to contain its errors instead of using the `GENERIC` messages described here.

The choice of which code to use depends on the context in which the message is issued, but `IDL_M_NAMED_GENERIC` is usually preferred.

If you wish to include arguments into your message string, you should use the `sprintf()` function from the C standard library to format a string into a temporary buffer, and then supply the buffer as the argument to `IDL_Message()`. For example, executing the code:

```
char buf[128];
int unit = 23;

sprintf(buf, "Help! Error number %d.", unit);
IDL_Message(IDL_M_GENERIC, IDL_MSG_LONGJMP, buf);
```

interrupts the current routine and issues the message:

```
% Help! Error number 23.
```


Looking Up A Message Code by Name

Given a message block pointer and the name of a message from that block, the `IDL_MessageNameToCode()` function returns the message code that corresponds to it. This is especially useful for dynamically loadable modules that need to throw errors from the IDL core block. The actual error codes are subject to change between IDL releases, so looking them up this way at run-time allows a given DLM to work with different IDL versions.

```
int IDL_MessageNameToCode(IDL_MSG_BLOCK block, char *name)
```

where:

block

Message block name should be translated against, or NULL to use the default core IDL block.

name

The message name for which the code is desired. Name is case sensitive, and should usually be specified as uppercase.

`IDL_MessageNameToCode ()` returns the message code, or 0 if it is not found.

Checking Arguments

IDL allows a user to provide any number of arguments, of any type, to system functions and procedures. IDL checks for a valid number of arguments, but the routine itself must check the validity of types. This task consists of examining the **argv** argument to the routine checking the type and flags field of each argument for suitability. The **IDL_StoreScalar()** function (see “[Storing Scalar Values](#)” on page 178) can be very useful in checking write-only arguments.

A number of macros exist in order to simplify testing of variable attributes. All of these macros accept a single argument—the VPTR to the argument in question. The macros check for a desired condition and use the **IDL_Message()** function with the **IDL_MSG_LONGJMP** action to return to the interpreter if an argument type doesn’t agree. Some of these macros overlap, and some are contradictory. You should select the smallest set that covers your requirements for each argument. For an example that uses one of these macros, see “[Example: A Complete Numerical Routine Example \(FZ_ROOTS2\)](#)” on page 277.

IDL_EXCLUDE_UNDEF

The argument must not be of type **IDL_TYP_UNDEF**. This condition is usually imposed if the argument is intended to provide some input information to the routine.

IDL_EXCLUDE_CONST

The argument must not be a constant. This condition should be specified if your routine intends to change the value of the argument.

IDL_EXCLUDE_EXPR

The argument must not be a constant or a temporary variable (i.e., the argument must be a named variable). Specify this condition if you intend to return a value in the argument. Returning a value in a temporary variable is pointless because the interpreter will remove it from the stack as soon as the routine completes, causing it to be freed for re-use.

The **IDL_VarCopy()** and **IDL_StoreScalar()** functions automatically check their destination and issue an error if it is an expression. Therefore, if you are using one of these functions to write the new value into the argument variable, you do not need to perform this check first.

IDL_EXCLUDE_FILE

The argument cannot be a file variable (as returned by the IDL ASSOC) function. Most system routines exclude file variables—they are handled by a small set of existing routines. This check is also handled by the **IDL_ENSURE_SIMPLE** macro, which also excludes structure variables.

IDL_EXCLUDE_STRUCT

The argument cannot be a structure.

IDL_EXCLUDE_COMPLEX

The argument cannot be **IDL_TYP_COMPLEX**.

IDL_EXCLUDE_STRING

The argument cannot be **IDL_TYP_STRING**.

IDL_EXCLUDE_SCALAR

The argument cannot be a scalar.

IDL_ENSURE_ARRAY

The argument must be an array.

IDL_ENSURE_OBJREF

The argument must be an object reference heap variable.

IDL_ENSURE_PTR

The argument must be a pointer heap variable.

IDL_ENSURE_SCALAR

The argument must be a scalar.

IDL_ENSURE_STRING

The argument must be **IDL_TYP_STRING**.

IDL_ENSURE_SIMPLE

The argument cannot be a file variable, a structure variable, a pointer heap variable, or an object reference heap variable.

IDL_ENSURE_STRUCTURE

The argument must be **IDL_TYP_STRUCT**.



Chapter 10

IDL Internals: Type Conversion

This chapter discusses the following topics:

Converting Arguments to C Scalars	206	Converting to Specific Types	208
General Type Conversion	207		

Converting Arguments to C Scalars

The routines described in this section convert the value of their `IDL_VARIABLE` argument to the C scalar type indicated by their name. `IDL_MEMINTScalar()` and `IDL_FILEINTScalar()` exist for processing memory and file sizes without the need to know their actual types, as discussed in “[IDL_MEMINT and IDL_FILEINT Types](#)” on page 119. The converted value is returned as the function value. The functions are defined as:

```
IDL_LONG IDL_LongScalar(IDL_VPTR p)
IDL_ULONG IDL_ULongScalar(IDL_VPTR v)
IDL_LONG64 IDL_Long64Scalar(IDL_VPTR v)
IDL_ULONG64 IDL_ULong64Scalar(IDL_VPTR v)
double IDL_DoubleScalar(IDL_VPTR p)
IDL_MEMINT IDL_MEMINTScalar(IDL_VPTR p)
IDL_FILEINT IDL_FILEINTScalar(IDL_VPTR p)
```

If these functions are unable to perform the conversion (e.g., the argument is a file variable, an array, etc.), they issue a descriptive error and jump back to the interpreter. By using these functions, you avoid having to do any of the type checking described in “[Checking Arguments](#)” on page 202.

For example, the following IDL system function (named `PRINT_LONG`) prints the value of its first argument, converted to an `IDL_LONG` 32-bit integer:

```
IDL_VPTR print_long(int argc, IDL_VPTR argv[], char *argk)
{
    printf("%d\n", IDL_LongScalar(argv[0]));
}
```

Executing it as:

```
PRINT_LONG, 23D
```

gives the output:

```
23
```

as expected, while the statement:

```
PRINT_LONG, FINDGEN(10)
```

causes the error:

```
% PRINT_LONG: Expression must be a scalar in this context:
                <FLOAT Array(10)>
% Execution halted at $MAIN$ .
```

because it is not possible to convert an array (the result of `FINDGEN`) to a scalar.

General Type Conversion

The `IDL_BasicTypeConversion()` function provides general purpose type conversion:

```
IDL_VPTR IDL_BasicTypeConversion(int argc, IDL_VPTR argv[]
                                int type)
```

where:

argc

The number of `IDL_VPTR`s contained in **argv**.

argv

An array of pointers to `VARIABLE` arguments.

type

The desired type code of the result. See “[Type Codes](#)” on page 114.

If **argc** is 1, this function returns a pointer to a temporary `VARIABLE` containing the value of **argv[0]** converted to the type specified by the **type** argument. If the variable is already of the correct type, the variable itself is returned.

If **argc** is greater than 1, **argv[1]** is taken to be an offset into the variable specified by **argv[0]**, and following arguments are taken as the dimensions to be used for the result. In this case, enough bytes are copied (starting from the offset) to satisfy the requirements of the dimensions given. This second form does not work for variables of type string, and if **argc** is greater than 1 an error is generated if **argv[0]** is of string type. You should ensure that variables of appropriate type are used with this function.

The IDL `BYTE` and `STRING` system routines (implemented by the `IDL_CvtByte()` and `IDL_CvtString()` functions, described below) treat conversions between variables of type byte and string in a special way. `IDL_BasicTypeConversion()` does not handle this special case. Instead, it simply performs a straightforward type conversion between those types.

Converting to Specific Types

A series of functions exist to convert **VARIABLES** to specific types:

```
IDL_VPTR IDL_CvtByte(int argc, IDL_VPTR argv[])
IDL_VPTR IDL_CvtBytscl(int argc, IDL_VPTR argv[], char *argk)
IDL_VPTR IDL_CvtFix(int argc, IDL_VPTR argv[])
IDL_VPTR IDL_CvtUInt(int argc, IDL_VPTR argv[])
IDL_VPTR IDL_CvtLng(int argc, IDL_VPTR argv[])
IDL_VPTR IDL_CvtULng(int argc, IDL_VPTR argv[])
IDL_VPTR IDL_CvtLng64(int argc, IDL_VPTR argv[])
IDL_VPTR IDL_CvtULng64(int argc, IDL_VPTR argv[])
IDL_VPTR IDL_CvtFlt(int argc, IDL_VPTR argv[])
IDL_VPTR IDL_CvtDbl(int argc, IDL_VPTR argv[])
IDL_VPTR IDL_CvtComplex(int argc, IDL_VPTR argv[])
IDL_VPTR IDL_CvtDComplex(int argc, IDL_VPTR argv[])
IDL_VPTR IDL_CvtString(int argc, IDL_VPTR argv[], char *argk)
```

When calling these functions, you should set the **argk** argument to **NULL**.

These functions are the direct implementations of the IDL commands **BYTE**, **BYTSCL**, **FIX**, **UINT**, **LONG**, **ULONG**, **LONG64**, **ULONG64**, **FLOAT**, **DOUBLE**, **COMPLEX**, **DCOMPLEX**, and **STRING**. See the description of these functions in the *IDL Reference Guide* for details on their arguments and calling sequences.

The behavior of these functions is the same as **IDL_BasicTypeConversion()** except when converting between bytes and strings. Calling **IDL_CvtByte()** with a single argument of string type causes each string to be converted to a byte vector of the same length as the string. Each array element is the character code of the corresponding character in the string. Calling **IDL_CvtString()** with a single argument of **IDL_TYP_BYTE** has the opposite effect.



Chapter 11

IDL Internals: UNIX Signals

This chapter discusses the following topics:

IDL and Signals	210	Removing a Signal Handler	215
Signal Handlers	213	UNIX Signal Masks	216
Establishing a Signal Handler	214		

IDL and Signals

Signals pose one of the more difficult challenges faced by the UNIX programmer. Although seemingly simple, they are a tough portability problem because there are several variants, and their semantics are subtle, inconvenient, and easily confused. These issues are only magnified when signals are used in a program like IDL that employs multiple threads. IDL has always done whatever is necessary with signals in order to get its job done, but its signal assumptions can also affect user written code linked to it.

Note

This discussion refers primarily to UNIX IDL. Microsoft Windows uses different mechanisms to solve the problems solved by signals under UNIX.

The following is a brief list of problems and contradictions inherent in UNIX signals. For a more complete description, see Chapter 10 of *External Programming in the UNIX Environment* by W. Richard Stevens.

- POSIX signals (sigaction) promise to unify and simplify signals, but not all platforms support them fully.
- You can only have one signal handler function registered for each signal. This means that if one part of a program uses a signal, the rest of the program must leave that signal alone.
- In order to meet the needs of programs originally developed under different UNIX systems (AT&T System V, BSD, Posix), most UNIX implementations provide more than one package of signal functions. Typically, a given program is restricted to one of these libraries. If a programmer links code into IDL that chooses a library or signal options different from that used by IDL itself, unexpected results may occur.
- The number and exact semantics of some signals differ in different versions.
- Details of signal blocking differ.
- Some System V implementations of signals are unreliable, meaning that signals can occur in a process and be missed.
- Some older System V systems reset the handling action to **SIG_DFL** before calling the handler. This opens a window in time where two signals in a row can cause the process to be killed. Also, the signal handler must re-establish itself every time it is called.

- On most platforms, if a signal is generated more than once while it is blocked, the second and subsequent occurrences are lost. In other words, most UNIX implementations do not queue signals.

These are among the reasons that most libraries avoid signals, and leave their use to the end programmer. IDL, however, must use signals to function properly. In order to allow users to link their code into IDL while using signals, IDL provides a signal API built on top of the signal mechanism supported by the target platform. The IDL signal API has the following attributes:

- It disallows use of **SIGTRAP** and **SIGFPE**. These signals are reserved to IDL.
- It disallows use of **SIGALRM**. Most uses for **SIGALRM** are provided by the IDL timer API.
- It works with all other signals, including those IDL doesn't currently use, so the interface won't change over time.
- It allows multiple signal handlers for each signal, so IDL and other code can use the same signal simultaneously.
- It unifies the signal interface by supplying a stable consistent interface with known behavior to the underlying system signal mechanism.
- It keeps IDL in charge of which signal package is used and how.

This is not a perfect solution, it is a compromise between the needs of IDL and programmers wishing to link code with it. Usually, the IDL signal abstraction is sufficient, but it does have the following limitations:

- The calling program must not attempt to catch **SIGTRAP** or **SIGFPE**, either directly or through library routines that use these signals to achieve their ends. Furthermore, the IDL signal abstraction does not allow the caller to catch these signals, so your program must leave exception handling to IDL.
- The caller loses control over signal package choice and some minor signal abilities.
- Having multiple signal handler routines for a given signal opens the possibility that one handler might do something that causes problems for the others (like change the signal mask, or `longjmp()`). To minimize such problems, user code linked into IDL must not call the actual system signal routines, and must not `longjmp()` out of signal handlers—a tactic that is usually allowed, but which would seriously damage IDL's signal state.

- Since there may be more than one signal handler registered for a given signal, the signal dispositions of **SIG_IGN** and **SIG_DFL** are not directly available to the caller as they would be if you were allowed to use the system signal facilities directly.

If you find that these restrictions are too limiting for your application, chances are your code is not compatible with IDL and should be executed in a separate process. We then encourage you to consider running IDL in a separate process and to use an interprocess communication mechanism such as RPC.

Signal Handlers

IDL signal handler functions are defined as:

```
typedef void (* IDL_SignalHandler_t)(int signo);
```

When a signal is delivered to the process, all registered signal handlers are called. `signo` is the integer number of the signal delivered, as defined by the C language header file `signal.h` (found in `/usr/include/signal.h` on most UNIX systems). `signo` can be used by a signal handler registered for more than one signal to tell which signal called it.

Establishing a Signal Handler

To register a signal handler, use the **IDL_SignalRegister()** function:

```
int IDL_SignalRegister(int signo, IDL_SignalHandler_t func,
                     int msg_action)
```

where:

signo

The numeric value of the signal to register for, as defined in `signal.h`.

func

The signal handler to be called when the signal specified by `signo` is raised.

msg_action

One of the **IDL_MSG_*** action codes for **IDL_Message()**. If there is an error in registering the signal handler, this action code is passed to **IDL_Message()** to direct its recovery action. Note that it is incorrect to use any of the message codes that cause **IDL_Message()** to **longjmp()** back to the IDL interpreter if your code is running in a context where the IDL interpreter is not active—specifically as part of using Callable IDL.

If `func` is successfully registered for `signo`, this routine returns **TRUE**. Otherwise, **FALSE** is returned and **IDL_Message()** is called with `msg_action` to control its behavior. Note that there are values of `msg_action` that result in this routine not returning on error. Multiple registration of the same function is allowed, but has no additional effect—the handler will only be called once.

Removing a Signal Handler

To remove a signal handler, use the **IDL_SignalUnregister()** function:

```
export int IDL_SignalUnregister(int signo,  
                               IDL_SignalHandler_t func, int msg_action)
```

where:

signo

The signal to unregister.

func

The handler to be unregistered.

msg_action

One of the **IDL_MSG_*** action codes for **IDL_Message()**. If there is an error in removing the signal handler, this action code is passed to **IDL_Message()** to direct its recovery action.

Once **IDL_SignalUnregister()** has been called, **func** is unregistered and will no longer be called if the signal is raised. **IDL_SignalUnregister()** returns **TRUE** for success, **FALSE** for failure. Requests to unregister a function that has not been previously registered are ignored.

UNIX Signal Masks

UNIX processes contain a signal mask that defines which signals can be delivered and which are blocked from delivery at any given time. When a signal arrives, the UNIX kernel checks the signal mask: If the signal is in the process mask, it is delivered, otherwise it is noted as undeliverable and nothing further is done until the signal mask changes. Sets of signals are represented within IDL with the opaque type **IDL_SignalSet_t**. UNIX IDL provides several functions that manipulate signal sets to change the process mask and allow/disallow delivery of signals.

IDL_SignalSetInit()

IDL_SignalSetInit() initializes a signal set to be empty, and optionally sets it to contain one signal.

```
void IDL_SignalSetInit(IDL_SignalSet_t *set, int signo)
```

where:

set

The signal set to be emptied/initialized.

signo

If non-zero, a signal to be added to the new set. This is provided as a convenience for the large number of cases where a set contains only one signal. Use **IDL_SignalSetAdd()** to add additional signals to a set.

IDL_SignalSetAdd()

IDL_SignalSetAdd() adds the specified signal to the specified signal set:

```
void IDL_SignalSetAdd(IDL_SignalSet_t *set, int signo)
```

where:

set

The signal set to be added to. The signal set must have been initialized by **IDL_SignalSetInit()**.

signo

The signal to be added to the signal set.

IDL_SignalSetDel()

IDL_SignalSetDel() deletes the specified signal from a signal set:

```
void IDL_SignalSetDel(IDL_SignalSet_t *set, int signo)
```

where:

set

The signal set to delete from. The signal set must have been initialized by **IDL_SignalSetInit()**.

signo

The signal to be removed from the signal set.

IDL_SignalSetIsMember()

IDL_SignalSetIsMember() tests a signal set for the presence of a specified signal, returning TRUE if the signal is present and FALSE otherwise:

```
int IDL_SignalSetIsMember(IDL_SignalSet_t *set, int signo)
```

where:

set

The signal set to test. The signal set must have been initialized by **IDL_SignalSetInit()**.

signo

The signal to be removed from the signal set.

IDL_SignalMaskGet()

IDL_SignalMaskGet() sets a signal set to contain the signals from the current process signal mask:

```
void IDL_SignalMaskGet(IDL_SignalSet_t *set)
```

where:

set

The signal set in which the current process signal mask will be stored.

IDL_SignalMaskSet()

IDL_SignalMaskSet() sets the current process signal mask to contain the signals specified in a signal mask:

```
void IDL_SignalMaskSet(IDL_SignalSet_t *set,  
                      IDL_SignalSet_t *omask)
```

where:

set

The signal set from which the current process signal mask will be set.

omask

If **omask** is non-NULL, the unmodified process signal mask is stored in it. This is useful for restoring the mask later using **IDL_SignalMaskSet()**.

There are some signals that cannot be blocked. This limitation is silently enforced by the operating system.

IDL_SignalMaskBlock()

IDL_SignalMaskBlock() adds signals to the current process signal mask:

```
void IDL_SignalMaskBlock(IDL_SignalSet_t *set,  
                        IDL_SignalSet_t *oset)
```

where:

set

The signal set containing the signals that will be added to the current process signal mask.

oset

If **oset** is non-NULL, the unmodified process signal mask is stored in it. This is useful for restoring the mask later using **IDL_SignalMaskSet()**.

There are some signals that cannot be blocked. This limitation is silently enforced by the operating system.

IDL_SignalBlock()

IDL_SignalBlock() does the same thing as **IDL_SignalMaskBlock()** except it accepts a single signal number instead of requiring a mask to be built:

```
void IDL_SignalBlock(int signo, IDL_SignalSet_t *oset)
```

where:

signo

The signal to be blocked.

There are some signals that cannot be blocked. This limitation is silently enforced by the operating system.

IDL_SignalSuspend()

IDL_SignalSuspend() replaces the process signal mask with the ones in set and then suspends the process until a signal is delivered. On return, the original process signal mask is restored:

```
void IDL_SignalSuspend(IDL_SignalSet_t *set)
```

where:

set

The signal set containing the signals that will be added to the current process signal mask.



Chapter 12

IDL Internals: Timers

This chapter discusses the following topics:

IDL and Timers	222	Canceling Asynchronous Timer Requests	225
Making Timer Requests	223	Blocking UNIX Timers	226

IDL and Timers

The details of how timers work varies widely between operating systems and between variants of the same operating system (different “flavors” of UNIX, for example). IDL’s timer module is intended to provide a stable interface to the rest of IDL, and to isolate the non-portable code in one place.

Under UNIX, IDL’s timer module performs a more important function. UNIX processes contain a single timer that must be shared by the code in the process. When the timer fires, it raises the **SIGALRM** signal which must be caught and handled by the process. The IDL timer routines transparently multiplex this single timer to provide multiple virtual timers.

Under UNIX, IDL provides both blocking and non-blocking timers. Blocking timers put the calling process to sleep until they go off. Non-blocking timers are delivered asynchronously when they fire.

Under Microsoft Windows, only the blocking form of timer requests are supported.

Making Timer Requests

The `IDL_TimerSet()` function registers a timer request. IDL timer requests are one-shot timers. If you wish to have a timer go off repeatedly, your callback function must make a new request each time it is called to set up the next timer.

```
void IDL_TimerSet(length, callback, from_callback, context)
```

where:

length

The length of time to delay before issuing the alarm, in microseconds. You should be aware that other activity on the system, overhead incurred in managing the timers, and non-realtime attributes of the operating system can cause the actual duration of the timer to be longer than requested.

callback

Under UNIX, if **callback** is non-NULL, the timer request is queued and `IDL_TimerSet()` returns immediately. When the alarm is due, the function pointed at by **callback** is called. If **callback** is NULL (and not **from_callback**), the request is queued and `IDL_TimerSet()` blocks until the requested time expires.

Warning

When called, the callback function will be running in signal scope, meaning that it has been called from a signal handler running asynchronously from the rest of the program. There are significant restrictions on what code running in signal scope is allowed to do. Most common C library functions (such as `printf()`) are disallowed. Consult a book on UNIX programming or your system documentation for details.

Under Windows, **callback** should always be NULL. `IDL_TimerSet()` does not support non-blocking timers on these platforms.

from_callback

Set this argument to TRUE if this invocation is from a callback function previously set up via a call to `IDL_TimerSet()`. Set this argument to FALSE if this is the first invocation. In other words, this argument should only be TRUE if you call `IDL_TimerSet()` from within a timer callback.

context

This argument is a pointer to a variable of type `IDL_TIMER_CONTEXT`, an opaque IDL data type that uniquely identifies a timer request. If this is a top level request (if `from_callback` is `FALSE`), the context pointed at will be assigned a unique value that identifies the request.

If this request is coming from within a timer callback in order to make another request on the same timer, the context pointed at should contain the value from the previous request.

If `context` is `NULL`, no context value is returned.

Note

It is an error to queue more than one request using the same callback. The results are undefined.

For the timer module to perform adequately, the time request must be large compared to the run-time of the called function. Re-queuing an extremely short request repeatedly will cause any other requests to starve.

Canceling Asynchronous Timer Requests

Under UNIX, **IDL_TimerCancel()** can be used to cancel a timer request that has not yet been delivered:

```
void IDL_TimerCancel(context)
```

where:

context

A timer request context returned by a previous call to **IDL_TimerSet()**.

Blocking UNIX Timers

Under UNIX operating systems, the delivery of signals such as **SIGALRM** (used to manage timers) can cause system calls to be interrupted. In such cases, the system call returns a status of **-1** and the global **errno** variable is set to the value **EINTR**. It is the caller's responsibility to check for this result and restart the system call when it occurs.

It is easy enough to handle this case when you make system calls directly, but sometimes the problem surfaces in libraries (even those provided by the system, such as `libc`) that are not properly coded against this possibility because the author assumed that no interrupts would occur. There is very little that the end user can do about such libraries except take steps that prevent signals from being raised during these critical sections.

If the IDL timer module is being used to deliver asynchronous events, it is inevitable that the delivery of **SIGALRM** will interfere with this sort of library code. The **IDL_TimerBlock()** function is available under UNIX to suspend the delivery of the timer signal. This can be used to provide a window in which no timer will fire. This routine should always be called in pairs, so the timer doesn't get turned off permanently. It is important to be sure a `longjmp()` (such as caused by calling **IDL_Message()** with the **IDL_MSG_LONGJMP** action code) doesn't happen in the critical region. In addition, this function is not re-entrant.

The effect of blocking timer delivery is that the UNIX **SIGALRM** signal is masked to prevent delivery. If the timer fires during this window of time, the signal will not be delivered until timers are unblocked. At that time, the timer module resumes managing the single real UNIX timer. In the meantime, timer requests are arbitrarily delayed from being queued and processed. Clearly, excessive blocking of the timer can lead to poor timer performance and should only be performed when necessary and on the smallest possible critical section of code. Of course, the act of blocking and unblocking signals requires a context switch into the UNIX kernel and back, making them relatively computationally expensive operations. It is therefore better to block a longer section of code rather than block and unblock around every critical library call.

It has been our experience that some UNIX platforms have more problem with this issue than others. You should let experience guide you in deciding when to block signals and when to let them go.

```
void IDL_TimerBlock(stop)
```

where:

stop

TRUE if the timer should be suspended, FALSE to restart it.

Chapter 13

IDL Internals: Files and Input/Output

This chapter discusses the following topics:

IDL and Input/Output Files	230	Allocating and Freeing File Units	243
File Information	232	Detecting End of File	245
Opening Files	236	Flushing Buffered Data	246
Closing Files	239	Reading a Single Character	247
Preventing File Closing	240	Output of IDL Variables	248
Checking File Status	241	Adding to the Journal File	249

IDL and Input/Output Files

IDL provides extensive Input/Output facilities at the user level. Internally, it uses native Input/Output facilities (UNIX system calls or Windows system API) in addition to the standard C library stream package (stdio). The choice of which facilities to use are made based on the target platform and the requested features for the file.

Most external code linked with IDL (CALL_EXTERNAL, system routines, etc.) should not do Input/Output directly, for the following reasons:

- Part of the IDL philosophy is that Input/Output is handled by dedicated I/O facilities provided by IDL, and that computational code should accept data from IDL variables and return results in the same way. This gives the user of your code the freedom and flexibility to save their data in any of the many forms supported by IDL's core I/O facilities, and frees you from writing complex and error prone input/output code.
- Using IDL's Input/Output facilities frees you from having to code around platform specific differences in I/O behavior.
- Input/Output from languages other than C often require runtime library support code to run at program startup before your code and successfully perform I/O. For example, Fortran Input/Output may depend on a Fortran runtime subsystem having been initialized. IDL, as a C program, does not perform initialization of such libraries for other languages. If you know enough about your Fortran system, you can often supply the missing initialization call, but such workarounds are usually not well documented, and are inherently platform specific.

For the reasons above, only minimal I/O abilities are available from IDL's internals, and only for files that explicitly use the standard C stdio library. Therefore, if your application must directly perform I/O to a file managed by IDL, it is necessary to use the standard C library streampackage (stdio) by specifying the IDL_F_STDIO flag to **IDL_FileOpen()**. Most of the routines associated with the standard C library I/O package can be used in the normal manner.

Note, however, that the C library routines listed in the following table should not be used; use the IDL-specific functions instead:

C Library Function	IDL Function
fclose()	IDL_FileClose()
fdopen()	IDL_FileOpen()
feof()	IDL_FileEOF()
fflush()	IDL_FileFlushUnit()
fopen()	IDL_FileOpen()
freopen()	IDL_FileOpen()

Table 13-1: Disallowed C Library Routines and Their IDL Counterparts

Note

In order to access a file opened using `IDL_FileOpen()` in this manner, you must ensure that it is stdio compatible by specifying `IDL_F_STDIO` as part of the `extra_flags` argument to `IDL_FileOpen()`. Failure to do this will cause your code to fail to execute as expected.

File Information

IDL maintains a file table in which it keeps a file descriptor for each file opened with `IDL_FileOpen()`. This table is indexed by the file Logical Unit Number, or LUN. These are the same LUNs IDL users use.

The `IDL_FileStat()` function is used to get information about a file.

IDL_FileStat()

```
void IDL_FileStat(int unit, IDL_FILE_STAT *stat_blk)
```

unit

The logical unit number (LUN) of the file unit to be checked. This function should only be called on file units that are known to be open.

stat_blk

A pointer to an `IDL_FILE_STAT` struct to be filled in with information about the file. The information returned is valid only as long as the file remains open. You must not access the fields of an `IDL_FILE_STAT` once the file it refers to has been closed. This struct has the definition:

```
typedef struct {
    char *name;
    short access;
    IDL_SFILE_FLAGS_T flags;
    FILE *fptr;
} IDL_FILE_STAT;
```

The fields of this struct are listed below:

name

A pointer to a null-terminated string containing the name the file was opened with.

access

A bit mask describing the access allowed to the file. The allowed bit values are listed in the following table:

Bit Value	Description
IDL_OPEN_R	The file is open for input.
IDL_OPEN_W	The file is open for output.
IDL_OPEN_TRUNC	The file was truncated when it was opened. This implies that IDL_OPEN_W is also set.
IDL_OPEN_APND	The file was opened with the file pointer set just past the last byte of data in the file (the file is open for appending).

Table 13-2: Bit values for the access field

flags

A bit mask that gives special information about the file. The defined bits are listed in the following table:

Bit Value	Description
IDL_F_ISATTY	The file is a terminal.
IDL_F_ISAGUI	The file is a Graphical User Interface.
IDL_F_NOCLOSE	The CLOSE command will refuse to close the file.
IDL_F_MORE	If the file is a terminal, output is sent through a pager similar to the UNIX <code>more</code> command. Details on this pager are not included in this document, and it is therefore not available for general use.
IDL_F_XDR	The file is read/written using XDR (eXternal Data Representation).
IDL_F_DEL_ON_CLOSE	The file will be deleted when it is closed.

Table 13-3: Bit values for the flags field

Bit Value	Description
IDL_F_SR	The file is a SAVE/RESTORE file.
IDL_F_SWAP_ENDIAN	The file has opposite byte order than that of the current system.
IDL_F_VAX_FLOAT	Binary float and double are in VAX F and D format.
IDL_F_COMPRESS	The file is in compressed gzip format. If IDL_F_SR is set (the file is a SAVE/RESTORE file), the file contains zlib compressed data.
IDL_F_UNIX_F77	The file is read/written in the format used by the UNIX Fortran (f77) compiler for unformatted binary data.
IDL_F_PIPE	The file is a bi-directional data path connecting IDL to a child process created by the SPAWN procedure.
IDL_F_UNIX_RAWIO (formerly called IDL_F_UNIX_NOSTDIO)	No application level buffering will be performed for the file and all data transfers will go directly to the operating system for processing (e.g. read() and write() system calls under UNIX, Windows system-level API for MS Windows). Note that setting this bit does not guarantee that data will be written to the file immediately, because the operating system may buffer the data. This bit value was formerly called IDL_F_UNIX_NOSTDIO. IDL_F_UNIX_RAWIO is the preferred form, but both names are supported.
IDL_F_UNIX_SPECIAL	The file is a UNIX device special file, most likely a pipe. This differs from IDL_F_PIPE because it applies to any file, not only those opened with the SPAWN procedure.

Table 13-3: Bit values for the flags field (Continued)

Bit Value	Description
IDL_F_STDIO	Use the C standard I/O library (stdio) to perform I/O on this file instead of any other native OS API that might be otherwise used. People intending to access IDL files via their own code should specify this flag if they intend to access the file from their external code as a stdio stream.
IDL_F_SOCKET	File is an internet TCP/IP socket.

Table 13-3: Bit values for the flags field (Continued)

fptr

The stdio stream file pointer to the file. This field can be used with standard library functions to perform I/O. This field is always valid, although you shouldn't use it if the file is an XDR file. You can check for this by looking for the **IDL_F_XDR** bit in the flags field.

If the file is not opened with the **IDL_F_STDIO** flag, **fptr** may be returned as an unusable NULL pointer, reflecting the fact that IDL is not using stdio to perform I/O on the file. If access to a valid **fptr** is important to your application, you should be sure to specify **IDL_F_STDIO** to the **extra_flags** argument to **IDL_FileOpen**, or use the **STDIO** keyword to **OPEN** if opening the file from the IDL user level.

In addition to the requirement to set the **IDL_F_STDIO** flag, you should be aware that IDL buffers I/O at a layer above the stdio package. If your code does I/O directly to a file that is also being written to from the IDL user level, the IDL buffer may cause data to be written to the file in a different order than you expect. There are several approaches you can take to prevent this:

- Tell IDL not to buffer, by opening the file from the IDL user level and specifying a value of -1 to the **BUFSIZE** keyword.
- Disable stdio buffering by calling the stdio **setbuf()** function.
- Ensure that you flush IDL's buffer before you do any Input/Output, as discussed in [“Flushing Buffered Data”](#) on page 246.

Opening Files

Files are opened using the `IDL_FileOpen()` function.

IDL_FileOpen()

```
int IDL_FileOpen(int argc, IDL_VPTR *argv, char *argk,
                int access_mode, IDL_SFILE_FLAGS_T extra_flags,
                int longjmp_safe, int msg_attr)
```

IDL_FileOpen() returns TRUE if the file has been successfully opened and FALSE otherwise.

Note

If `longjmp_safe` is TRUE, the usual course is to jump back to the IDL interpreter, in which case the routine won't return.

argc

The number of arguments in *argv*. This value should always be 2.

argv

The arguments to `IDL_File_Open()`. `argv[0]` should be a scalar integer value giving the file unit number (LUN) to be opened. `argv[1]` is a scalar string giving the file name.

argk

Keywords. Set this argument to NULL.

access_mode

A bit mask that specifies the access to be allowed to the file being opened. The allowed bit values are listed in the following table:

Bit Value	Description
IDL_OPEN_R	The file is open for input.
IDL_OPEN_W	The file is open for output.

Table 13-4: Bit Values for the `access_mode` Argument

Bit Value	Description
IDL_OPEN_TRUNC	The file was truncated when it was opened. This implies that IDL_OPEN_W is also set.
IDL_OPEN_APND	The file was opened with the file pointer set just past the last byte of data in the file (the file is open for appending).

Table 13-4: Bit Values for the access_mode Argument (Continued)

It is important that conflicting bits not be set together (for example, do not specify IDL_OPEN_TRUNC | IDL_OPEN_APND). Also, one or both of IDL_OPEN_R and IDL_OPEN_W must always be specified.

extra_flags

Used to specify additional file attributes using the flags defined in the description of the flags field of the IDL_FILE_STAT struct (see “[File Information](#)” on page 232). Note that some flags are set by IDL based on the actual attributes of the opened file (e.g. IDL_F_ISTTY) and that it makes no sense to set such flags in this mask.

If you intend to use the opened file as a C standard I/O (stdio) stream file, you must specify the **IDL_F_STDIO** flag when calling **IDL_FileOpen()**. Otherwise, IDL may choose not to use stdio.

longjmp_safe

If set to TRUE, **IDL_FileOpen()** is being called in a context where an IDL_MSG_LONGJMP **IDL_Message** action code is okay. If set to FALSE, the routine won't `longjmp()`.

IDL_FileOpen() returns TRUE if the file has been successfully opened and FALSE otherwise. Of course, if `longjmp_safe` is TRUE, the usual course is to jump back to the IDL interpreter, in which case the routine won't return.

msg_attr

A zero (0), or any combination of the IDL_MSG_ATTR_ flags, used to fine tune the error handling specified by the `longjmp_safe` argument. Note that you must not specify any of the base IDL_MSG_ codes, but only the attributes. The base code (e.g. IDL_MSG_LONGJMP) is determined by the value of `longjmp_safe`. For a discussion of the IDL_MSG_ATTR_ flags, see “[Issuing Error Messages](#)” on page 195.

Special File Units

There are three files that are always open. The three units are:

- **IDL_STDIN_UNIT** — Unit 0 (zero) is the standard input for the IDL process.
- **IDL_STDOUT_UNIT** — Unit -1 is the standard output.
- **IDL_STDERR_UNIT** — Unit -2 is the standard error.

Note

The constant `IDL_NON_UNIT` always has a value that is *not* a valid file unit.

Closing Files

Files are closed using the `IDL_FileClose()` function.

IDL_FileClose()

```
void IDL_FileClose(int argc, IDL_VPTR *argv, char *argk)
```

argc

The number of arguments in *argv*.

argv

The arguments to the close function. These should be scalar integer values giving the Logical Unit Numbers of the file units to close.

argk

Keywords. Set this argument to `NULL`.

Preventing File Closing

Use the `IDL_FileSetClose()` function to prevent files from closing. It does this by setting or clearing the `IDL_F_NOCLOSE` bit in the flags field of the internal file descriptor maintained by IDL for the file (see “[File Information](#)” on page 232). This feature is used primarily in graphics drivers for printers. For example, the PostScript driver uses this feature to prevent the user from closing the plot data file prematurely.

When IDL exits, it only closes open files that do not have the `IDL_F_NOCLOSE` bit set. Files with close inhibited are simply left alone. Often, you will want to declare an exit handler which takes care of closing such files.

IDL_FileSetClose()

```
void IDL_FileSetClose(int unit, int allow)
```

unit

The Logical Unit Number (LUN) of the file in question. The file must be open for this function to have effect.

allow

Set this field to `TRUE` if users are allowed to close the file. Set to `FALSE` if users should be prevented from closing the file.

Checking File Status

System routines that deal with files must verify that the files have the proper attributes for the intended operation. Use the function `IDL_FileEnsureStatus()` for this.

IDL_FileEnsureStatus()

```
int IDL_FileEnsureStatus(int action, int unit, int flags)
```

action

If the file unit does not satisfy the requirements of the flags argument, `IDL_FileEnsureStatus()` will issue an error using the `IDL_Message()` function (see [“Issuing Error Messages”](#) on page 195). This action is the action argument to `IDL_Message()` and should be `IDL_MSG_RET`, `IDL_MSG_LONGJMP`, or `IDL_MSG_IO_LONGJMP`.

unit

The Logical Unit Number of the file to be checked.

flags

`IDL_FileEnsureStatus()` always checks to make sure unit is a valid logical file unit. In addition, flags is a bit mask specifying the file attributes that should be checked. The possible bit values are listed in the following table:

Bit Value	Description
<code>IDL_EFS_USER</code>	The file must be a user unit. This means that the file is not one of the three special files, <code>stdin</code> , <code>stdout</code> , or <code>stderr</code> .
<code>IDL_EFS_IDL_OPEN</code>	The file unit must be open.
<code>IDL_EFS_CLOSED</code>	The file unit must be closed.
<code>IDL_EFS_READ</code>	The file unit must be open for input.
<code>IDL_EFS_WRITE</code>	The file unit must be open for output.
<code>IDL_EFS_NOTTY</code>	The file unit cannot be a tty.

Table 13-5: Bit Values for the flags Argument

Bit Value	Description
IDL_EFS_NOGUI	The file unit cannot be a Graphical User Interface.
IDL_EFS_NOPIPE	The file unit cannot be a pipe.
IDL_EFS_NOXDR	The file unit cannot be a XDR file.
IDL_EFS_ASSOC	The file unit can be ASSOC'ed. This implies IDL_EFS_USER, IDL_EFS_OPEN, IDL_EFS_NOTTY, IDL_EFS_NOPIPE, IDL_EFS_NOXDR, IDL_EFS_NOCOMPRESS, and IDL_EFS_NOSOCKET.
IDL_EFS_NOT_RAWIO (formerly called IDL_EFS_NOT_NOSTDIO)	The file was not opened with the IDL_F_UNIX_RAWIO attribute. This bit was formerly called IDL_F_NOTSTDIO. IDL_EFS_NOT_RAWIO is the preferred form, but both names are accepted.
IDL_EFS_NOCOMPRESS	The file unit cannot have been opened for compressed input/output (IDL_F_COMPRESS).
IDL_EFS_STDIO	The file must be using the C stdio package (IDL_F_STDIO).
IDL_EFS_NOSOCKET	The file unit cannot be a socket (IDL_F_SOCKET).

Table 13-5: Bit Values for the flags Argument (Continued)

Note

Some of these values are contradictory. The caller must select a consistent set.

If the file unit meets the desired conditions, IDL_FileEnsureStatus() returns TRUE. If it does not meet the conditions, and action was IDL_MSG_RET, then it returns FALSE.

Allocating and Freeing File Units

System routines must allocate and deallocate file units in order to avoid conflicts. When writing IDL procedures, the `GET_LUN` and `FREE_LUN` procedures are used. When writing system-level routines, you can access the same routines by calling `IDL_FileGetUnit()` and `IDL_FileFreeUnit()`.

Use `IDL_FileGetUnit()` to allocate file units:

`IDL_FileGetUnit()`

```
void IDL_FileGetUnit(int argc, IDL_VPTR *argv)
```

argc

`argc` should always be 1.

argv

`argv[0]` contains an `IDL_VPTR` to the `IDL_VARIABLE` that will be filled in with the resulting unit number.

Use `IDL_FileFreeUnit()` to free file units:

`IDL_FileFreeUnit()`

```
void IDL_FileFreeUnit(int argc, IDL_VPTR *argv)
```

argc

argc gives the number of elements in **argv**.

argv

argv should contain scalar integer values giving the Logical Unit Numbers of the file units to be returned.

Read the description of `GET_LUN` and `FREE_LUN` in the *IDL Reference Guide* for additional details about these functions. The following code fragment demonstrates how these functions might be used to open and close a file named `junk.dat`:

```
IDL_VPTR argv[2];
IDL_VARIABLE unit;
IDL_VARIABLE name;
.
```

```
.
/* Allocate a file unit. */
argv[0] = &unit;
unit.type = IDL_TYP_LONG;
unit.flags = 0;
IDL_FileGetUnit(1, argv);

/* Set up the file name */
name.type = IDL_TYP_STRING;
name.flags = IDL_V_CONST;
name.value.str.s = "junk.dat";
name.value.str.slen = sizeof("junk.dat") - 1;
name.value.str.stype = 0;
argv[1] = &name;
.
.
.
IDL_FileOpen(2, argv, (char *) 0, IDL_OPEN_R, 0, 1, 0);

/* Perform any required actions. */
.
.
.
/* Free the file unit. This will also close the file. */
IDL_FileFreeUnit(1, argv);
```

Detecting End of File

IDL_FileEOF()

The IDL_FileEOF() function returns TRUE if the file specified by the Logical Unit Number unit is at EOF, and FALSE otherwise:

```
int IDL_FileEOF(int unit)
```

unit

The Logical Unit Number (LUN) of the file in question.

Flushing Buffered Data

IDL_FileFlushUnit()

File data might be buffered in system memory in order to boost input/output performance. The `IDL_FileFlushUnit()` function forces any buffered data for the file specified by the Logical Unit Number unit to be written out:

```
int IDL_FileFlushUnit(int unit)
```

unit

The Logical Unit Number (LUN) of the file in question.

Reading a Single Character

IDL_GetKbrd()

The IDL_GetKbrd() function returns the character code of the next available character from IDL_STDIN_UNIT:

```
int IDL_GetKbrd(int should_wait)
```

should_wait

Set this argument to TRUE if IDL_GetKbrd() should wait for a key to be struck, FALSE otherwise.

If should_wait is FALSE and no input characters are waiting in the input stream, IDL_GetKbrd() returns NULL. Otherwise, it waits until a key is struck (if necessary) and then returns its ASCII value. This function will generate an error and return to the interpreter if IDL_STDIN_UNIT is not a terminal.

Output of IDL Variables

IDL_Print() and IDL_Printf()

The `IDL_Print()` and `IDL_Printf()` functions output the value of `IDL_VARIABLES`. `IDL_Print()` simply outputs to `IDL_STDOUT_UNIT`, while `IDL_Printf()` outputs to a specified unit:

```
void IDL_Print(int argc, IDL_VPTR *argv, char *argk)
void IDL_Printf(int argc, IDL_VPTR *argv, char *argk)
```

argc

The number of arguments to `argv`.

argv

`IDL_VPTR`s of the `IDL_VARIABLES` to be output.

argk

Keywords. Set this argument to `NULL ((char *) 0)`.

These functions are the implementation of the built-in IDL system procedures `PRINT` and `PRINTF`. See “[PRINT/PRINTF](#)” (*IDL Reference Guide*) for information on the available arguments and the order in which you must specify them.

Adding to the Journal File

IDL_Logit()

The IDL_Logit() function can be used to add lines of output to the journal log file:

```
void IDL_Logit(char *s)
```

s

A pointer to a NULL terminated string to be added to the journal log file.

If a journal log file is currently open, this function writes the specified string to it on a new line. If no journal file is open, IDL_Logit() returns quietly. The only way to open or close the journal file is via the user-system-level JOURNAL procedure.



Chapter 14

IDL Internals: Miscellaneous

This chapter discusses the following topics:

Dynamic Memory	252	Ensuring UNIX TTY State	260
Exit Handlers	255	Type Information	261
User Interrupts	256	User Information	263
Functions for Returning System Variables	257	Constants	264
Terminal Information	258	Macros	265

Dynamic Memory

IDL provides access to the dynamic memory allocation routines it uses internally. Use these routines rather than system-provided routines such as **malloc()/free()** when possible.

Warning

The memory pointers returned by the IDL memory allocation routines discussed in this chapter do not necessarily correspond directly to **malloc()/free()** calls, or to any other system memory allocation package. You must be careful not to mix memory allocation packages. Memory allocated via a given API can only be freed by the corresponding free call provided by that API. For example, memory allocated by an IDL memory allocation routine can only be freed by the IDL **IDL_MemFree()** function. Memory allocated by **malloc()** can only be freed by **free()**.

Failure to follow this rule can lead to memory corruption, including possible crashing of the IDL program.

Please note that code called via **CALL_EXTERNAL**, or as a system routine (**LINKIMAGE**, Dynamically Loadable Modules) should not use the IDL dynamic memory routines. Instead, use **IDL_GetScratch()** (see “[Getting Dynamic Memory](#)” on page 174) which prevents memory from being lost under error conditions.

Warning

Our experience shows that in situations where **IDL_GetScratch()** is appropriate, use of any other memory allocation mechanism should raise a warning flag to the programmer that something is wrong in their code. Rarely if ever is a direct call to **malloc()/free()** reasonable in such a situation — even if it appears to work correctly, you will have to work harder to provide the error handling functionality that **IDL_GetScratch()** provides automatically, or your code will leak memory in such situations.

IDL_MemAlloc()

IDL_MemAlloc() is used to allocate dynamic memory.

```
void *IDL_MemAlloc(IDL_MEMINT n, char *err_str, int action)
```

where:

n

The number of bytes to allocate.

err_str

NULL, or a null terminated text string describing the memory being allocated.

action

An action parameter to be passed to **IDL_Message()** if **IDL_MemAlloc()** is unable to allocate the desired memory and **err_str** is non-NULL.

IDL_MemAlloc() attempts to allocate the desired amount of memory. If the requested amount is allocated, a pointer to the memory is returned. The memory is aligned strictly enough to be suitable for any object.

If the attempt to allocate memory fails and **err_str** is non-NULL, **IDL_Message()** is called as:

```
IDL_Message(IDL_M_CNTGETMEM, action, err_str)
```

If **IDL_Message()** returns, or if **err_str** is NULL and **IDL_Message()** is not called, **IDL_MemAlloc()** returns a NULL pointer indicating its failure.

IDL_MemFree()

Memory allocated via **IDL_MemAlloc()** should only be returned via **IDL_MemFree()**:

```
void IDL_MemFree(REGISTER void *m, char *err_str, int action)
```

m

A pointer to memory previously allocated via **IDL_MemAlloc()**.

err_str

NULL, or a null terminated text string describing the memory being freed.

action

An action parameter to be passed to **IDL_Message()** if unable to free memory and **err_str** is non-NULL.

IDL_MemFree() attempts to free the specified memory. If the attempt to free memory fails and **err_str** is non-NULL, **IDL_Message()** is called as:

```
IDL_Message(IDL_M_CNTRFEMEM, action, err_str)
```

The following actions have undefined consequences, and should not be done:

- Returning memory allocated from a source other than **IDL_MemAlloc()**.
- Freeing the same allocation more than once.
- Dereferencing memory once it has been freed.

IDL_MemAllocPerm()

Another memory allocation routine, **IDL_MemAllocPerm()**, exists to allocate dynamic memory that will not be returned for reuse. **IDL_MemAllocPerm()** allocates memory in moderately large units and carves out pieces of these blocks to satisfy its requests. Use of this routine can help minimize the effects of memory fragmentation.

```
void *IDL_MemAllocPerm(IDL_MEMINT n, char *err_str, int action)
```

IDL_MemAllocPerm() takes the same arguments as **IDL_MemAlloc()**, differing only in that the memory allocated will not be freed until the process exits. Do not attempt to free memory allocated by **IDL_MemAllocPerm()**. The results of such an action are undefined.

Exit Handlers

IDL maintains a list of exit handler functions that it calls as part of its shutdown operations. These handlers perform actions such as closing files, wrapping up graphics output, and restoring the user environment to its initial state. Exit handlers accept no arguments and return no value.

A typical declaration would be:

```
void my_exit_handler(void)
{
    /* Cleanup Code Here */
}
```

IDL_ExitRegister()

To register an exit handler, use the **IDL_ExitRegister()** function:

```
void IDL_ExitRegister(IDL_EXIT_HANDLER_FUNC)
```

where `IDL_EXIT_HANDLER_FUNC` is defined as:

```
typedef void(* IDL_EXIT_HANDLER_FUNC)(void);
```

proc

IDL will call **proc** just before it exits.

The order in which exit handlers are called is undefined, and you should not depend on any particular ordering. If you have several exit handlers and the order in which they are called is important, you should register a single handler that calls all the others in the required order.

Note

Under some operating systems, it is possible that the IDL process will die in an abnormal way that prevents the calling of the exit handlers. For example, under UNIX, receiving some signals (possibly via the **kill(1)** command) will cause the process to die immediately. IDL always calls exit handlers when possible, so this is rarely a significant problem.

User Interrupts

IDL catches certain operating system signals including **SIGINT**, which occurs when the user types the interrupt character (usually Control-C). When the interpreter detects the interrupt character, it sets an internal flag which causes execution of the program to stop at the next sequence statement. The interpreter clears this variable every time it is invoked, and checks to see if it has been set before it executes each statement. This means that when the user presses the interrupt character, the current statement must complete before the interpreter checks the value of the variable and halts execution.

Typical statements do not take long to complete, so this delay is not noticeable. However, some system routines take a long time to complete, and the user can be fooled by the long delay into thinking that IDL is ignoring the interrupt. While the occasional long delay can be annoying, this method of handling interrupts is the only way to maintain acceptable performance in the usual case where no interrupt is pending. Therefore, it is the responsibility of system routines that take a long time to complete to check the value of this internal variable and to clean up and return if **SIGINT** is seen. IDL's Input/Output and FFT routines, among others, do this.

IDL_BailOut()

The **IDL_BailOut()** function is used to sense or set the state of IDL's internal interrupt flag. It returns TRUE if the keyboard interrupt character has been typed, otherwise FALSE.

```
int IDL_BailOut(int stop)
```

where:

stop

Set to FALSE to sense the state of the keyboard interrupt flag without changing its value. Set to TRUE to set the keyboard interrupt flag.

Functions for Returning System Variables

The following functions return the values of certain system variables. Note that these values should be considered READ-ONLY.

IDL_STRING *IDL_SysvVersionArch(void)

This function returns a pointer to the !VERSION.ARCH system variable.

IDL_STRING *IDL_SysvVersionOS(void)

This function returns a pointer to the !VERSION.OS system variable.

IDL_STRING *IDL_SysvVersionOSFamily(void)

This function returns a pointer to the !VERSION.OS_FAMILY system variable.

IDL_STRING *IDL_SysvVersionRelease(void)

This function returns a pointer to the !VERSION.RELEASE system variable.

IDL_STRING *IDL_SysvDirFunc(void)

This function returns a pointer to the !DIR system variable.

IDL_STRING *IDL_SysvErrStringFunc(void)

This function returns a pointer to the !ERROR_STATE.MSG system variable.

IDL_STRING *IDL_SysvSyserrStringFunc(void)

This function returns a pointer to !ERROR_STATE.SYS_MSG system variable.

IDL_LONG IDL_SysvErrorCodeValue(void)

This function returns the value of the !ERROR_STATE system variable.

IDL_LONG IDL_SysvOrderValue(void)

This function returns the value of the !ORDER system variable.

For more information on IDL system variables, see [Appendix D, “System Variables”](#) (*IDL Reference Guide*).

Terminal Information

The global variable **IDL_FileTerm** is a structure of type **IDL_TERMINFO**:

```
typedef struct {
    char *name;        /* Name Of Terminal Type */
    char is_tty;      /* True if stdin is a terminal */
    int lines;        /* Lines on screen */
    int columns;     /* Width of output */
} IDL_TERMINFO;
```

Note

Under operating systems that do not support the concept of a terminal (Microsoft Windows) the **name** and **is_tty** fields are not present.

IDL_FileTerm is initialized when IDL is started. Few, if any, user routines will need this information, because user routines should not do their own I/O. User routines that must do their own I/O should use this variable instead of making assumptions about the output device.

Note

Under Microsoft Windows, the **IDL_FileTerm** is not accessible outside of the IDL sharable library, and cannot be directly accessed by user code. Instead, use the functions described in the following section.

Functions for Returning IDL_FileTerm Variable Values

The following functions can be used to return values from the **IDL_FileTerm** variable. They return the same information contained in the global variable, but in a functional form. This is the preferred way to access the **IDL_FileTerm** information, as it will work on any platform.

char *IDL_FileTermName(void)

This function returns the value of **IDL_FileTerm.name**. This function is only available under UNIX.

int IDL_FileTermIsTty(void)

This function returns the value of **IDL_FileTerm.is_tty**. This function is only available under UNIX.

int IDL_FileTermLines(void)

This function returns the value of **IDL_FileTerm.lines**.

int IDL_FileTermColumns(void)

This function returns the value of **IDL_FileTerm.columns**.

Ensuring UNIX TTY State

Under some UNIX operating systems, IDL keeps the users terminal in a *raw mode*, required to implement command line editing. On these platforms, externally linked code that performs output to the terminal will find that the output does not appear as expected. A usual symptom of this is that newline characters (`'\n'`) do not move the cursor to the left margin of the screen, and commands such as `more(1)` (perhaps started via the C runtime library `system()` function) do not control the screen properly.

This is not an issue for IDL routines such as `SPAWN` that start sub-programs, because they are written to be aware of this issue and to ensure the TTY is in the correct state before they do their work. Externally linked code can call the `IDL_TTYReset()` function to do the same thing:

```
void IDL_TTYReset(void)
```

This function is available under all operating systems. On systems where such an operation is not needed, it is a stub. On platforms that require the TTY to be managed in this way, this operation ensures that the terminal is returned to its standard configuration.

Type Information

The following read-only global variables provide information about IDL data.

Note

Under Microsoft Windows, these global variables are not available; use the functions listed below to retrieve the values contained in the global variables.

IDL_OutputFormat

An array of pointers to character strings. **IDL_OutputFormat** is indexed by type code, and specifies the default output formats for the different data types (see “[Type Codes](#)” on page 114). The default formats are used by the PRINT and STRING built-in routines as well as for type conversions.

IDL_OutputFormatLen

An array of integers. **IDL_OutputFormatLen** gives the length in characters of the corresponding elements of **IDL_OutputFormat**.

IDL_TypeSize

An array of long integers. **IDL_TypeSize** is indexed by type code, and gives the size of the data object used to represent each type.

IDL_TypeName

An array of pointers to character strings. **IDL_TypeName** is indexed by type code, and gives a descriptive string for each type.

Functions for Returning Data Type Variable Values

The following functions can be used to return the values contained in the global variables described above, but in a functional form.

char *IDL_OutputFormatFunc(int type)

Given an IDL type code, this function returns the default output format for that type. This is equivalent to accessing the **IDL_OutputFormat** array.

int IDL_OutputFormatLenFunc(int type)

Given an IDL type code, this function returns the default output format length for that type. This is equivalent to accessing the **IDL_OutputFormatLen** array.

int IDL_TypeSizeFunc(int type)

Given an IDL type code, this function returns the size of the data object used to represent that type. This is equivalent to accessing the **IDL_TypeSize** array.

char *IDL_TypeNameFunc(int type)

Given an IDL type code, this function returns the name of the type as a null terminated character string. This is equivalent to accessing the **IDL_TypeName** array.

User Information

Use the `IDL_GetUserInfo()` function to get information about the current session. This is the sort of information that can be used in the header of files produced by graphics drivers. It is used, for example, by the PostScript driver:

```
void IDL_GetUserInfo(IDL_USER_INFO *user_info)
```

where the `IDL_USER_INFO` struct is defined as:

```
typedef struct {
    char *logname;           /* User's login name */
    char *homedir;          /* User's home directory */
    char *pid;               /* The process ID */
    char host[64];           /* Machine name */
    char wd[IDL_MAXPATH+1]; /* Working Directory */
    char date[25];          /* Current System Time */
} IDL_USER_INFO;
```

Constants

Preprocessor constants defined in the `idl_export.h` file should be used in preference to hardwired values. To accommodate the needs of various operating systems, some of these constants have different values in different versions of IDL. Those constants that are not discussed elsewhere in this book are listed below.

IDL_TRUE

A more readable alternative to the constant 1.

IDL_FALSE

A more readable alternative to the constant 0.

IDL_REGISTER

Some C compilers are good at allocating registers, and using the C register declaration can cause efficiency to suffer. On the other hand, some C compilers won't put any variables into registers unless register definitions are used. Our solution is to use **IDL_REGISTER** to declare variables we feel should be placed into registers. For machines that we feel have a good register allocation scheme, we define **IDL_REGISTER** to be a null macro. For lesser compilers, it is defined to be the C `register` keyword.

IDL_MAX_ARRAY_DIM

The maximum number of dimensions an array can have.

IDL_MAXIDLEN

The maximum number of characters IDL allows in an identifier (variable names, program names, and so on).

IDL_MAXPATH

The maximum number of characters allowed in a filepath.

Macros

The macros defined in `idl_export.h` handle recurring small jobs. Those macros not discussed elsewhere in this book are covered here.

IDL_ABS(x)

IDL_ABS() accepts a single argument of any numeric C type, and returns its absolute value. **IDL_ABS()** evaluates its argument more than once, so be careful to avoid unwanted side effects, and for efficiency do not call it with a complex expression.

IDL_CARRAYELTS(arr)

This macro encapsulates a common C language idiom for determining the number of elements in a statically defined array without requiring the programmer to provide a constant or otherwise hardwire the length. Its use improves the robustness of code that uses it by automatically adapting to any change in the definition of the array without requiring additional programmer effort. This macro corresponds directly to the C expression:

```
sizeof(arr)/sizeof(arr[0])
```

The C compiler evaluates this expression at compile time, so there is no additional runtime cost for using this macro instead of a hardwired constant.

IDL_CHAR(ptr)

IDL_CHAR() casts its argument to be a pointer to **char**. It is used to convert an existing pointer into a generic pointer to a memory location.

IDL_CHARA(addr)

IDL_CHARA() takes the address of its argument and casts it to be a pointer to **char**. It is used to get a generic pointer to a memory location.

IDL_MIN(x,y) and IDL_MAX(x,y)

The arguments can be of any numeric C type as long as they are compatible with each other. **IDL_MIN()** and **IDL_MAX()** return the smaller and larger of their two arguments, respectively. These macros evaluate their arguments more than once, so be careful to avoid unwanted side effects, and for efficiency do not call them with a complex expression.


IDL_ROUND_UP(x, m)

IDL_ROUND_UP() returns the value of **x** rounded up modulo **m**. **m** must be a power of 2. This macro is useful for extending data regions out to a specified alignment.

IDL_TRUE and IDL_FALSE

When performing logical expression evaluation the C programming language, in which IDL is written, treats zero (0) as False, and non-zero as True, and when returning the result of such an expression, uses 1 for True and 0 for False.

IDL_TRUE is defined as the constant 1, and **IDL_FALSE** is defined as the constant 0. These constants are used internally by IDL.



***Part III: Techniques
That Use IDL's Internal
API***



Chapter 15

Adding System Routines

This chapter discusses the following topics:

IDL and System Routines	270	Example: An Example Using Routine Design	
The System Routine Interface	272	Iteration (RSUM)	286
Example: Hello World	273	Registering Routines	296
Example: Doing a Little More (MULT2)	274	Enabling and Disabling System Routines	299
Example: A Complete Numerical Routine		LINKIMAGE	307
Example (FZ_ROOTS2)	277	Dynamically Loadable Modules	309

IDL and System Routines

An IDL system routine is an IDL procedure or function that is written in a compiled language with an IDL specific interface, and linked into IDL, instead of being written in the IDL language itself. The best way to create an IDL system routine is to compile and link the routine into a sharable library and then to add the routine to IDL at runtime using either the LINKIMAGE procedure or by making your routines part of a Dynamically Loadable Module (DLM).

Note

We recommend the use of Dynamically Loadable Modules rather than LINKIMAGE whenever possible. The small additional effort is more than compensated for by the superior integration into IDL.

This chapter explains how to write a system routine, including several examples, and discusses the various options for adding such routines to IDL.

Debugging System Routines

To debug system routines and DLMs, you must attach your debugger to the IDL process that is running your library code. In IDL 7.0 and later, the process to which you should attach the debugger varies, depending on your platform and which version of IDL you are using.

If the IDL Workbench is Running

When the IDL Workbench is running, IDL comprises two processes:

IDL Workbench — This process runs the IDL Workbench interface. It will show up in a process list as `idlde.exe` on Windows systems, and as a Java command containing the string `idlde` on Unix-based systems.

IDL Out-of-Process server — This process runs the IDL interpreter and loads shared libraries (DLM's). It will show up in a process list as `idl_opserver.exe` on Windows systems, and as `idl_opserver` on Unix-based systems.

To debug your shared library code when the IDL Workbench is running, attach your debugger (`gdb`, Visual Studio, *etc.*) to one of the following:

- On Windows, using Visual Studio, you can select **Attach to Process** and choose the `idl_opserver.exe` process.

- On Unix-based systems, using `gdb` or `dbx`, you can use the `ps` command to retrieve the ID of the `idl_opserver` process, and then attach your debugger to that process ID.

If the IDL Workbench is Not Running

When the IDL is running in command-line mode, IDL comprises either one or two processes:

IDL — On Unix-based systems, the `idl` process is the only IDL process. On Windows systems, the `idl.exe` process is a front-end to the `idl_opserver.exe` process.

IDL Out-of-Process server — On Windows systems, the `idl_opserver.exe` process runs the IDL interpreter and loads shared libraries (DLM's). Unix-based systems do not use an `idl_opserver` process when running in command-line mode.

To debug your shared library code when IDL is running in command-line mode, attach your debugger (`gdb`, Visual Studio, *etc.*) to one of the following:

- On Windows, using Visual Studio, you can select **Attach to Process** and choose the `idl_opserver.exe` process.
- On Unix-based systems, using `gdb` or `dbx`, you can use the `ps` command to retrieve the ID of the `idl` process, and then attach your debugger to that process ID.

The System Routine Interface

All IDL system routines must supply the same calling interface to the system, differing only in that system functions must return an **IDL_VPTR** to the **IDL_VARIABLE** that contains the result while system procedures do not return anything. Typical system routine definitions are:

```
IDL_VPTR my_function(int argc, IDL_VPTR argv[], char *argk)
void my_procedure(int argc, IDL_VPTR argv[], char *argk)
```

System routines that do not accept keywords are called with two arguments:

argc

The number of elements in **argv**.

argv

An array of **IDL_VPTR**s. These point to the **IDL_VARIABLE**s which comprise the arguments to the function.

System routines that accept keywords are called with an additional third argument:

argk

The keywords which were present when the routine was called. **argk** is an opaque object—the called routine is not intended to understand its contents. **argk** is provided to the function **IDL_KWProcessByOffset()**, which processes the keywords in a standard way. For more information on keywords, see [“IDL Internals: Keyword Processing”](#) on page 121.

Example: Hello World

Thanks to the definitive text on the C language (Kernighan and Ritchie, *The C Programming Language*, Prentice Hall, NJ, Second Edition, 1988), the “Hello World” program is often used as an example of a trivial program. Our version of this program is a system function that returns a scalar string containing the text “Hello World!”:

```
#include <stdio.h>
#include "idl_export.h"

IDL_VPTR hello_world(int argc, IDL_VPTR argv[])
{
    return(IDL_StrToSTRING("Hello World!"));
}
```

This is about as simple as an IDL system routine can be. The function **IDL_StrToSTRING()**, returns a temporary variable which contains a scalar string. Since this is exactly what is wanted, **hello_world()** simply returns the variable.

After compiling this function into a sharable object (named, for example, **hello_exe**), we can link it into IDL with the following LINKIMAGE call:

```
LINKIMAGE, 'HELLO_WORLD', 'hello_exe', 1, 'hello_world', $
    MAX_ARGS=0, MIN_ARGS=0
```

We can now issue the IDL command:

```
PRINT, HELLO_WORLD()
```

In response, IDL writes to the screen:

```
Hello World!
```

Example: Doing a Little More (MULT2)

The system function shown in the following figure does a little more than the previous one, though not by much. It expects a single argument, which must be an array. It returns a single-precision, floating-point array that contains the values from the argument multiplied by two.

```

1  #include <stdio.h>
2  #include "idl_export.h"
3
4  IDL_VPTR mult2(int argc, IDL_VPTR argv[])
5  {
6      IDL_VPTR dst, src;
7      float *src_d, *dst_d;
8      int n;
9      src = dst = argv[0];
10
11     IDL_ENSURE_SIMPLE(src);
12     IDL_ENSURE_ARRAY(src);
13
14     if (src->type != IDL_TYP_FLOAT)
15         src = dst = IDL_CvtFlt(1, argv);
16
17     src_d = dst_d = (float *) src->value.arr->data;
18
19     if (!(src->flags & IDL_V_TEMP))
20         dst_d = (float *)
21             IDL_MakeTempArray(IDL_TYP_FLOAT, src->value.arr->n_dim,
22                               src->value.arr->dim,
23                               IDL_ARR_INI_NOP, &dst);
24
25     for (n = src->value.arr->n_elts; n--; )
26         *dst_d++ = 2.0 * *src_d++;
27
28     return(dst);
29 }

```

Table 15-1: mult2.c

Each line is numbered to make discussion easier. These numbers are not part of the actual program. Each line of this routine is discussed below:

1-2

Include the required header files.

4

Every system routine takes the same two or three arguments. **argc** is the number of arguments, **argv** is an array of arguments. This routine does not accept keywords, so **argk** is not present.

6

dst will become a pointer to the resulting variable's descriptor. **src** points at the input variable which is found in **argv[0]**.

7

src_d and **dst_d** will point to the source and destination data areas.

8

n will contain the number of elements in **src**.

10

Assume, for now, that the input variable will serve as both the source and destination. This will only be true if the parameter is a temporary floating-point array.

11-12

Screen out any input that is not of a basic type, and only allow arrays. A better version of this routine would handle scalar input also, but we want to keep the example simple.

14

If the input is not of **IDL_TYP_FLOAT**, we call the **IDL_CvtFlt()** function to create a floating-point copy of the argument (see [“Converting to Specific Types”](#) on page 208 for information about converting types).

Note that the routine could also be written, more efficiently, with a C switch statement which would handle each of the eight possible data types, eliminating conversion of the input parameter. This would be more in the spirit of the IDL language, where system routines work with all possible data types and sizes, but is outside the scope of this example.

17

Here, we initialize the pointers to the source and destination data areas from the array block structure pointed to by the input variable descriptor.

19-23

If the input variable is not a temporary variable, we cannot change its value and return it as the function result. Instead, we allocate a new temporary floating point array into which the result will be placed. Notice how the number of dimensions and their sizes are taken from the source variable array block. See “[Array Variables](#)” on page 157 and “[Temporary Variables](#)” on page 165.

25

Loop over each element of the arrays.

26

Do the multiplication for each element.

28

Return the temporary variable containing the result.

Testing the Example

Once we have compiled the function and linked it into IDL (possibly using LINKIMAGE), we can use the built-in IDL function INDGEN to test the new function, which we name MULT2. INDGEN returns an array of values with each element set to the value of its array index. Therefore, the statement:

```
PRINT, INDGEN(5)
```

prints the following on the screen:

```
0 1 2 3 4
```

To test our new function we use INDGEN to provide an input argument:

```
PRINT, MULT2(INDGEN(5))
```

The result, as expected, is:

```
0.00000 2.00000 4.00000 6.00000 8.00000
```

Example: A Complete Numerical Routine Example (FZ_ROOTS2)

The following is a complete implementation of the IDL system function `FZ_ROOTS`, used to find the roots of an m -degree complex polynomial, using Laguerre's method. The result is an m -element complex vector. We call this version `FZ_ROOTS2` to avoid a name clash with the real routine. `FZ_ROOTS2` has an additional keyword, `TC_INPUT`, that is not present in the real routine.

`FZ_ROOTS2` uses the routine `zroots()`, described in section 9.5 of *Numerical Recipes in C: The Art of Scientific Computing* (Second Edition), published by Cambridge University Press:

```
void zroots(fcomplex a[], int m, fcomplex roots[], int polish)
```

Quoting from the referenced book:

Given the degree m and the $m+1$ complex coefficients `a[0..m]` of the polynomial,

$$\sum_{i=0}^m a(i)x^i$$

this routine successively calls `laguer` and finds all m complex roots in `roots[1..m]`. The boolean variable `polish` should be input as true (1) if polishing (also by Laguerre's method) is desired, false (0) if the roots will be subsequently polished by other means.

`FZ_ROOTS2` will support both single and double precision complex values as well as give the caller control over the error tolerance, which is hard wired into the Numerical Recipes code as a C preprocessor constant named `EPS`. In order to support these requirements, we have copied the `zroots()` function given in the book and altered it to support both data types and make `EPS` a user specified parameter, giving two functions:

```
void zroots_f(fcomplex a[], int m, fcomplex roots[], int polish,
             float eps);
```

```
void zroots_d(dcomplex a[], int m, dcomplex roots[], int polish,
             double eps);
```

Note that **fcomplex** and **dcomplex** are Numerical Recipes defined types that happen to have the same definition as the IDL types **IDL_COMPLEX** and **IDL_DCOMPLEX**, a convenient fact that eliminates some type conversion issues.

The definition of `FZ_ROOTS2` from the IDL user perspective is:

Calling Sequence

Result = FZ_ROOTS2(C)

Arguments

C

A vector of length $m+1$ containing the coefficients of the polynomial, in ascending order.

Keywords

DOUBLE

FZ_ROOTS2 normally uses the type of C to determine the type of the computation. If DOUBLE is specified, it overrides this default. Setting DOUBLE to a non-zero value causes the computation type and the result to be double precision complex. Setting it to zero forces single precision complex.

EPS

The desired fractional accuracy. The default value is 2.0×10^{-6} .

NO_POLISH

Set this keyword to suppress the usual polishing of the roots by Laguerre's method.

TC_INPUT

If present, TC_INPUT specifies a named variable that will be assigned the input value C, with its type converted to the type of the result.

Example

The following figure gives the code for fzroots2.c,. This is ANSI C code that implements FZ_ROOTS2. The line numbers are not part of the code and are present to make the discussion easier to follow. Each line of this routine is discussed below.

C

```

1  #include <stdio.h>
2  #include <stdarg.h>
3  #include "idl_export.h"
4  #include <nr/nr.h>
5
6  IDL_VPTR fzroots2(int argc, IDL_VPTR *argv, char *argk)
7  {
8      typedef struct {
9          IDL_KW_RESULT_FIRST_FIELD; /* Must be first entry in this
10 structure */
11      int force_type;
12      IDL_LONG do_double;
13      double eps;
14      IDL_LONG no_polish;
15      IDL_VPTR tc_input;
16  } KW_RESULT;
17  static IDL_KW_PAR kw_pars[] = {
18      {"DOUBLE", IDL_TYP_LONG, 1, 0,
19      IDL_KW_OFFSETOF(force_type), IDL_KW_OFFSETOF(do_double) },
20      { "EPS", IDL_TYP_DOUBLE, 1, 0, 0, IDL_KW_OFFSETOF(eps) },
21      { "NO_POLISH", IDL_TYP_LONG, 1, IDL_KW_ZERO,
22      0, IDL_KW_OFFSETOF(no_polish) },
23      { "TC_INPUT", 0, 1, IDL_KW_OUT|IDL_KW_ZERO,
24      0, IDL_KW_OFFSETOF(tc_input) },
25      { NULL }
26  };
27
28  KW_RESULT kw;
29  IDL_VPTR result;
30  IDL_VPTR c_raw;
31  IDL_VPTR c_tc;
32  IDL_MEMINT m;
33  void *outdata;
34  IDL_ARRAY_DIM dim;
35  int rtype;
36  static IDL_ALLTYPES zero;
37
38  kw.eps = 2.0e-6;
39  (void) IDL_KWProcessByOffset(argc, argv, argk,
40  kw_pars, &c_raw, 1, &kw);
41
42  IDL_ENSURE_ARRAY(c_raw);
43  IDL_ENSURE_SIMPLE(c_raw);
44  if (c_raw->value.arr->n_dim != 1)
45  IDL_Message(IDL_M_NAMED_GENERIC, IDL_MSG_LONGJMP,
46      "Input argument must be a column vector.");

```

Table 15-2: fzroots2.c

```

47  m = c_raw->value.arr->dim[0];
48  if (--m <= 0)
49      IDL_Message(IDL_M_NAMED_GENERIC, IDL_MSG_LONGJMP,
50                  "Input array does not have enough elements");
51  if (kw.tc_input)
52      IDL_StoreScalar(kw.tc_input, IDL_TYP_LONG, &zero);
53
54  if (kw.force_type) {
55      rtype = kw.do_double ? IDL_TYP_DCOMPLEX : IDL_TYP_COMPLEX;
56  } else {
57      rtype = ((c_raw->type == IDL_TYP_DOUBLE)
58              || (c_raw->type == IDL_TYP_DCOMPLEX))
59              ? IDL_TYP_DCOMPLEX : IDL_TYP_COMPLEX;
60  }
61  dim[0] = m;
62  outdata = (void *)
63      IDL_MakeTempArray(rtype, 1, dim, IDL_ARR_INI_NOP, &result);
64
65  if (c_raw->type == rtype) {
66      c_tc = c_raw;
67  } else {
68      c_tc = IDL_BasicTypeConversion(1, &c_raw, rtype);
69  }
70
71  if (rtype == IDL_TYP_COMPLEX) {
72      zroots_f((fcomplex *) c_tc->value.arr->data, m,
73              ((fcomplex *) outdata) - 1, !kw.no_polish, (float) kw.eps);
74  } else {
75      zroots_d((dcomplex *) c_tc->value.arr->data, m,
76              ((dcomplex *) outdata) - 1, !kw.no_polish, kw.eps);
77  }
78
79  if (kw.tc_input) IDL_VarCopy(c_tc, kw.tc_input);
80  else if (c_raw != c_tc) IDL_Deltmp(c_tc);
81
82  IDL_KW_FREE;
83  return result;
}

```

Table 15-2: fzroots2.c (Continued)

4

nr.h is the header file provided with Numerical Recipes in C code.

6

FZROOTS2 has the usual three standard arguments.

10

kw.force_type will be TRUE if the user specifies the DOUBLE keyword. In this case, the value of the DOUBLE keyword will determine the result type without regard for the type of the input argument.

If the user specifies DOUBLE, a zero value forces a single precision complex result and non-zero forces double precision complex.

12

The value of the EPS keyword.

13

The value of the NO_POLISH keyword.

15

The value of the TC_INPUT keyword.

16

This array defines the keywords accepted by FZ_ROOTS2.

17

Since setting DOUBLE to 0 has a different meaning than not specifying the keyword at all, **kw.force_type** is used to detect the fact that the keyword is set independent of its value.

20

The EPS keyword allows the user to specify the **kw.eps** tolerance parameter. **kw.eps** is specified as double precision to avoid losing accuracy for double precision computations—it will be converted to single precision if necessary. The default value for this keyword is non-zero, so no zeroing is specified here. If the user includes the EPS keyword, the value will be placed in **kw.eps**, otherwise **kw.eps** will not be changed.

21

This keyword lets the user suppress the usual polishing performed by **zroots()**. Since specifying a value of 0 is equivalent to not specifying the keyword at all, **IDL_KW_ZERO** is used to initialize the variable.

23

If present, `TC_INPUT` is an output keyword that will have the type converted value of the input argument stored in it. By specifying **`IDL_KW_OUT`** and **`IDL_KW_ZERO`**, we ensure that `TC_INPUT` is either zero or a pointer to a valid IDL variable.

28

The results of keyword processing will all be written to this variable by **`IDL_KWProcessByOffset()`**.

29

This variable will receive the function result.

30

The input argument prior to any type conversion.

31

The type converted input variable. If the input variable is already of the correct type, this will be the same as **`c_raw`**, otherwise it will be different.

32

The value of m to be passed to **`zroots()`**.

33

Pointer to the data area of the result variable. We declare it as `(void *)` so that it can point to data of any type.

34

Used to specify dimensions of the result. This will always be a vector of m elements.

35

IDL type code for result variable.

36

Used by **`IDL_StoreScalar()`** to type check the `TC_INPUT` keyword. It is declared as static to ensure it is initialized to zero.

38

Set the default EPS value before doing keyword processing. If the user specifies EPS, the supplied value will override this. Otherwise, this value will still be in **kw.eps** and will be passed to **zroots()** unaltered.

39

Perform keyword processing.

42-43

Ensure that the input argument is an array, and is one of the basic types (not a file variable or structure).

44-46

The input variable must be a vector, and therefore should have only a single dimension.

47-50

Ensure that the input variable is long enough for m to be non-zero. m is one less than the number of elements in the input vector, so this is equivalent to saying that the input must have at least 2 elements.

51

If the TC_INPUT keyword was present, use **IDL_StoreScalar()** to make sure the named variable specified can receive the converted input value. A nice side effect of this operation is that any dynamic memory currently being used by this variable will be freed now instead of later after we have allocated other dynamic memory. This freed memory might be immediately reusable if it is large enough, which would reduce memory fragmentation and lower overall memory requirements.

54

If the user specified the DOUBLE keyword, it is used to control the resulting type, otherwise the input argument type is used to decide.

55

The DOUBLE keyword was specified. If it is non-zero, use **IDL_TYP_DCOMPLEX**, otherwise **IDL_TYP_COMPLEX**.

56-60

Use the input type to decide the result type. If the input is **IDL_TYP_DOUBLE** or **IDL_TYP_DCOMPLEX**, use **IDL_TYP_DCOMPLEX**, otherwise **IDL_TYP_COMPLEX**.

61-63

Create the output variable that will be passed back as the result of **FZ_ROOTS2**.

65-69

If necessary, convert the input argument to the result type. This is done *after* creation of the output variable because it is likely to have a short lifetime. If it does get freed at the end of this routine, it won't cause memory fragmentation by leaving a hole in the process virtual memory.

71

The version of **zroots()** to call depends on the data type of the result.

72-73

Single precision complex. Note that the outdata pointer is decremented by one element. This compensates for the fact that the Numerical Recipe routine will index it from [1..m] rather than [0..m-1] as is the usual C convention. Also, **kw.eps** is cast to single precision.

74-76

Double precision complex case.

79

If the user specified the **TC_INPUT** keyword, copy the type converted input into the keyword variable. Since **VarCopy()** frees its source variable if it is a temporary variable, we are relieved of the usual responsibility to call **IDL_Deltmp()** if **c_tc** contains a temporary variable created on line 66.

80

The user didn't specify the **TC_INPUT** keyword. In this case, if we allocated **c_tc** on line 66, we must free it before returning.

82

Free any resources allocated by keyword processing.

83

Return the result.

Example: An Example Using Routine Design Iteration (RSUM)

We now show how a simple routine can be developed in stages. RSUM is a function that returns the running sum of the values in its single input argument. We will present three versions of this routine, each one of which represents an improvement in functionality and flexibility.

All three versions use the function **IDL_MakeTempFromTemplate()**, described in [“Creating A Temporary Variable Using Another Variable As A Template”](#) on page 170. The result of RSUM always has the same general shape and dimensions as the input argument. **IDL_MakeTempFromTemplate()** encapsulates the task of creating a temporary variable of the desired type and shape using the input argument as a template.

Running Sum (Example 1)

The first example of RSUM is very simple. Here is a simple “Reference Manual” style description of it:

RSUM1

Compute a running sum on the array input. The result is a floating point array of the same dimensions.

Calling Sequence

Result = RSUM1(Array)

Arguments

Array

Array for which a running sum will be computed.

This is a minimal design that lacks some important characteristics that IDL system routines usually embody:

- It does not handle scalar input.
- The type of the input is inflexible. IDL routines usually try to handle any input type and do whatever type conversions are necessary.
- The result type is always single precision floating point. IDL routines usually perform computations in the type of the input arguments and return a value of the same type.

We will improve on this design in our subsequent attempts. The code to implement RSUM1 is shown in the following figure. The line numbers are not part of the code and are present to make the discussion easier to follow. Each line of this routine is discussed below:

C	<pre> 1 IDL_VPTR IDL_rsum1(int argc, IDL_VPTR argv[]) 2 { 3 IDL_VPTR v; 4 IDL_VPTR r; 5 float *f_src; 6 float *f_dst; 7 IDL_MEMINT n; 8 9 10 v = argv[0]; 11 if (v->type != IDL_TYP_FLOAT) 12 IDL_Message(IDL_M_NAMED_GENERIC, IDL_MSG_LONGJMP, 13 "argument must be float"); 14 IDL_ENSURE_ARRAY(v); 15 IDL_EXCLUDE_FILE(v); 16 17 f_dst = (float *) 18 IDL_VarMakeTempFromTemplate(v, IDL_TYP_FLOAT, 19 (IDL_StructDefPtr) 0, &r, FALSE); 20 f_src = (float *) v->value.arr->data; 21 n = v->value.arr->n_elts - 1; 22 *f_dst++ = *f_src++; /* First element */ 23 for (; n--; f_dst++) *f_dst = *(f_dst - 1) + *f_src++; 24 25 return r; 26 } </pre>
----------	--

Table 15-3: Code for IDL_rsum1()

1

The standard signature for an IDL system function that does not accept keywords.

3

This variable is used to access the input argument in a convenient way.

4

This **IDL_VPTR** will be used to return the result.

5–6

As the running sum is computed, **f_src** will point at the input data and **f_dst** will point at the output data.

7

The number of elements in the input.

10

Obtain the input variable pointer from `argv[0]`.

11

If the input is not single precision floating point, throw an error and quit. This is overly rigid. Real IDL routines would attempt to either type convert the input or do the computation in the input type.

14

This version can only handle arrays. If the user passes a scalar, it throws an error.

15

This routine cannot handle ASSOC file variables. Most IDL routines exclude such variables as they do not contain any data to work with. ASSOC variable data usually comes into a routine as the result of an expression that yields a temporary variable (e.g. `TMP = RSUM(MY_ASSOC_VAR(2))`).

17

Create a single precision floating point temporary variable of the same size as the input variable and get a pointer to its data area.

20

Get a pointer to the data area of the input variable. At this point we know this variable is always a floating point array.

21

The number of data elements in the input.

22-23

The running sum computation.

25

Return the result.

Running Sum (Example 2)

In our second example of RSUM, we improve on version 1 in several ways:

- RSUM2 accepts scalar input.
- If the input is not of floating type, we type convert it instead of throwing an error.
- If the input is a temporary variable of the correct type, we will do the running sum computation in place and return the input as our result variable rather than creating an extra temporary. This optimization reduces memory use, and can have positive effects on dynamic memory fragmentation.

As always, the first step in writing a system routine is to write a simple description of its interface and intended behavior:

RSUM2

Compute a running sum on the input. The result is a floating point variable with the same structure.

Calling Sequence

Result = RSUM2(Input)

Arguments

Input

Scalar or array data of any numeric type for which a running sum will be computed.

The following is the code for RSUM2:

```

1  IDL_VPTR IDL_rsum2(int argc, IDL_VPTR argv[])
2  {
3      IDL_VPTR v;
4      IDL_VPTR r;
5      float *f_src;
6      float *f_dst;
7      IDL_MEMINT n;
8
9
10     v = IDL_BasicTypeConversion(1, argv, IDL_TYP_FLOAT);
11     /* IDL_BasicTypeConversion calls IDL_ENSURE_SIMPLE, so
12        skip it here. */
13     IDL_VarGetData(v, &n, (char **) &f_src, FALSE);
14
15     /* Get a result var, reusing the input if possible */
16     if (v->flags & V_TEMP) {
17         r = v;
18         f_dst = f_src;
19     } else {
20         f_dst = (float *)
21             IDL_VarMakeTempFromTemplate(v, IDL_TYP_FLOAT,
22                                         (IDL_StructDefPtr) 0, &r, FALSE);
23     }
24
25     *f_dst++ = *f_src++; /* First element */
26     n--;
27     for (; n--; f_dst++) *f_dst = *(f_dst - 1) + *f_src++;
28
29     return r;
30 }

```

Table 15-4: Code for IDL_rsum2().

Discussion of the code for the improvements introduced in this version follow:

10

If the input has the wrong type, obtain one of the right type. If it was already of the correct type, **IDL_BasicTypeConversion()** will simply return the input value without allocating a temporary variable. Hence, no explicit check for that is required. Also, if the input argument cannot be converted to the desired type (e.g. it is a structure or file variable) **IDL_BasicTypeConversion()** will throw an error. Hence, we know that the result from this function will be what we want without requiring any further checking.

13

IDL_VarGetData() is a more elegant way to obtain a pointer to variable data along with a count of elements. A further benefit is that it automatically handles scalar variables which removes the restriction from **RSUM1**.

15–23

If the input variable is a temporary, we will do the computation in place and return the input. Otherwise, we create a temporary variable of the desired type to be the result.

Note that if **IDL_BasicTypeConversion()** returned a pointer to anything other than the passed in value of **argv[0]**, that value will be a temporary variable which will then be turned into the function result by this code. Hence, we never free the value from **IDL_BasicTypeConversion()**.

Running Sum (Example 3)

RSUM2 is a big improvement over **RSUM1**, but it still suffers from the fact that all computation is done in a single data type. A real IDL system routine always tries to perform computations in the most significant type presented by its arguments. In a single argument case like **RSUM**, that would mean doing computations in the input data type whatever that might be. Our final version, **RSUM3**, resolves this shortcoming.

RSUM3

Compute a running sum on the input. The result is a variable with the same type and structure as the input.

Calling Sequence

Result = **RSUM3**(Input)

Arguments

Input

Scalar or array data of any numeric type for which a running sum will be computed.

The code for **RSUM3** is given in the following figure. Discussion of the code for the improvements introduced in this version follow:

```

1  cx_public IDL_VPTR IDL_rsum3(int argc, IDL_VPTR argv[])
2  {
3      IDL_VPTR v, r;
4      union {
5          char *sc;           /* Standard char */
6          UCHAR *c;          /* IDL_TYP_BYTE */
7          IDL_INT *i;        /* IDL_TYP_INT */
8          IDL_UINT *ui;     /* IDL_TYP_UINT */
9          IDL_LONG *l;      /* IDL_TYP_LONG */
10         IDL_ULONG *ul;    /* IDL_TYP_ULONG */
11         IDL_LONG64 *ul64; /* IDL_TYP_LONG64 */
12         IDL_ULONG64 *ul64; /* IDL_TYP_ULONG64 */
13         float *f;         /* IDL_TYP_FLOAT */
14         double *d;        /* IDL_TYP_DOUBLE */
15         IDL_COMPLEX *cmp; /* IDL_TYP_COMPLEX */
16         IDL_DCOMPLEX *dcmp; /* IDL_TYP_DCOMPLEX */
17     } src, dst;
18     IDL_LONG n;
19
20
21     v = argv[0];
22     if (v->type == IDL_TYP_STRING)
23         v = IDL_BasicTypeConversion(1, argv, IDL_TYP_FLOAT);
24     IDL_VarGetData(v, &n, &(src.sc), TRUE);
25     n--; /* First is a special case */
26
27     /* Get a result var, reusing the input if possible */
28     if (v->flags & IDL_V_TEMP) {
29         r = v;
30         dst = src;
31     } else {
32         dst.sc = IDL_VarMakeTempFromTemplate(v, v->type,
33             (IDL_StructDefPre) 0, &r, FALSE);
34     }
35
36     #define DOCASE(type, field) \
37     case type: for (*dst.field++ = *src.field++; n--; dst.field++) \
38         *dst.field = *(dst.field - 1) + *src.field++; break

```

Table 15-5: Code for IDL_rsum3

```

39 #define DOCASE_CMP(type, field) case type: \
40 for (*dst.field++ = *src.field++; n--; \
41     dst.field++, src.field++) { \
42     dst.field->r = (dst.field - 1)->r + src.field->r; \
43     dst.field->i = (dst.field - 1)->i + src.field->i; } \
44 break
45
46     switch (v->type) {
47     DOCASE(IDL_TYP_BYTE, c);
48     DOCASE(IDL_TYP_INT, i);
49     DOCASE(IDL_TYP_LONG, l);
50     DOCASE(IDL_TYP_FLOAT, f);
51     DOCASE(IDL_TYP_DOUBLE, d);
52     DOCASE_CMP(IDL_TYP_COMPLEX, cmp);
53     DOCASE_CMP(IDL_TYP_DCOMPLEX, dcmp);
54     DOCASE(IDL_TYP_UINT, ui);
55     DOCASE(IDL_TYP_ULONG, ul);
56     DOCASE(IDL_TYP_LONG64, l64);
57     DOCASE(IDL_TYP_ULONG64, ul64);
58     default: IDL_Message(IDL_M_NAMED_GENERIC, IDL_MSG_LONGJMP,
59                          "unexpected type");
60     }
61 #undef DOCASE
62 #undef DOCASE_CMP
63
64 return r;
65 }

```

Table 15-5: Code for `IDL_rsum3` (Continued)

17

`f_src` and `f_dst` are no longer pointers to float. They are now the `IDL_ALLPTR` type, which can point to data of any IDL type. To reflect this change in scope, the leading `f_` prefix has been dropped.

22-23

Strings are the only input type that now require conversion. The other types can either support the computation, or are not convertible to a type that can.

36-38

The code for the running sum computation is logically the same for all non-complex data types, differing only in the **IDL_ALLPTR** field that is used for each type. Using a macro for this means that the expression is only typed in once, and the C compiler automatically fills in the different parts for each data type. This is less error prone than entering the expression manually for each type, and leads to more readable code. This is one of the rare cases where a macro makes things *more* reliable and readable.

39-44

A macro for the 2 complex types.

46-60

A switch statement that uses the macros defined above to perform the running sum on all possible types. Note the default case, which traps attempts to compute a running sum on structures.

61-62

Don't allow the macros used in the above switch statement to remain defined beyond the scope of this function.

Registering Routines

The `IDL_SysRtnAdd()` function adds system routines to IDL's internal tables of system functions and procedures. As a programmer, you will need to call this function directly if you are linking a version of IDL to which you are adding routines, although this is very rare and not considered to be a good practice for maintainability reasons. More commonly, you use `IDL_SysRtnAdd()` in the `IDL_Load()` function of a Dynamically Loadable Module (DLM). DLMs are discussed in [“Dynamically Loadable Modules”](#) on page 309.

Note

LINKIMAGE or DLMs are the preferred way to add system routines to IDL because they do not require building a separate IDL program. Of the two, we recommend the use of DLMs whenever possible. These mechanisms are discussed in the following sections of this chapter.

Syntax

```
int IDL_SysRtnAdd(IDL_SYSFUN_DEF2 *defs, int is_function, int cnt)
```

It returns True if it succeeds in adding the routine or False in the event of an error.

Arguments

defs

An array of `IDL_SYSFUN_DEF2` structures, one per routine to be declared. This array must be defined with the C language `static` storage class because IDL keeps pointers to it. `defs` must be sorted by routine name in ascending lexical order.

is_function

Set this parameter to `IDL_TRUE` if the routines in `defs` are functions, and `IDL_FALSE` if they are procedures.

cnt

The number of `IDL_SYSFUN_DEF2` structures contained in the `defs` array.

The definition of `IDL_SYSFUN_DEF2` is:

```
typedef IDL_VARIABLE *(* IDL_SYSRTN_GENERIC)();
```



```
typedef struct {
    IDL_SYSRTN_GENERIC funct_addr;
    char *name;
    unsigned short arg_min;
    unsigned short arg_max;
    int flags
    void *extra;
} IDL_SYSFUN_DEF2;
```

IDL_VARIABLE structures are described in “[The IDL_VARIABLE Structure](#)” on page 153.

funct_addr

Address of the function implementing the system routine.

name

The name by which the routine is to be invoked from within IDL. This should be a pointer to a null terminated string. The name should be capitalized. If the routine is an object method, the name should be fully qualified, which means that it should include the class name at the beginning followed by two consecutive colons, followed by the method name (e.g. CLASS : METHOD).

arg_min

The minimum number of arguments allowed for the routine.

arg_max

The maximum number of arguments allowed for the routine. If the routine does not place an upper value on the number of arguments, use the value **IDL_MAXPARAMS**.

flags

A bitmask that provides additional information about the routine. Its value can be any combination of the following values (bitwise OR-ed together to specify more than one at a time) or zero if no options are necessary:

IDL_SYSFUN_DEF_F_OBSOLETE

IDL should issue a warning message if this routine is called and **!WARN.OBS_ROUTINE** is set.

IDL_SYSFUN_DEF_F_KEYWORDS

This routine accepts keywords as well as plain arguments.

IDL_SYSFUN_DEF_F_METHOD

This routine is an object method.

extra

Reserved to ITT Visual Information Solutions. The caller should set this to 0.

Example

The following example shows how to register a system routine linked directly with IDL. For simplicity, everything is placed in a single file. Normally, you would modularize things to allow easier code maintenance.

```
#include <stdio.h>
#include "idl_export.h"

void prox1(int argc, IDL_VPTR argv[])
{
    printf("prox1 %d\n", IDL_LongScalar(argv[0]));
}

main(int argc, char *argv[])
{
    static IDL_SYSFUN_DEF2 new_pros[] = {
        {(IDL_SYSRtn_GENERIC) prox1, "PROX1", 1, 1, 0, 0}
    };

    if (!IDL_SysRtnAdd(new_pros, IDL_FALSE, 1))
        IDL_Message(IDL_M_GENERIC, IDL_MSG_RET,
            "Error adding system routine");
    return IDL_Main(0, argc, argv);
}
```

This adds a system procedure named **PROX1** which accepts a single argument. It converts this argument to a scalar longword integer and prints it.

Enabling and Disabling System Routines

The following IDL internal functions allow the enabling and/or disabling of IDL system routines. Disabled routines throw an error when called from IDL code instead of performing their usual functions.

These routines are primarily of interest to authors of Runtime or Callable IDL applications.

Enabling Routines

The `IDL_SysRtnEnable()` function is used to enable and/or disable system routines.

Syntax

```
void IDL_SysRtnEnable(int is_function, IDL_STRING *names,
                    IDL_MEMINT n, int option,
                    IDL_SYSRTN_GENERIC disfcn)
```

Arguments

is_function

Set to `TRUE` if functions are being manipulated, `FALSE` for procedures.

names

`NULL`, or an array of names of routines.

n

The number of names in **names**.

option

One of the values from the following table which specify what this routine should do.

Bit	Description
<code>IDL_SRE_ENABLE</code>	Enable specified routines.
<code>IDL_SRE_ENABLE_EXCLUSIVE</code>	Enable specified routines and disable all others.
<code>IDL_SRE_DISABLE</code>	Disable specified routines.
<code>IDL_SRE_DISABLE_EXCLUSIVE</code>	Disable specified routines and enable all others.

Table 15-6: Values for **option** Argument

disfcn

NULL, or address of an IDL system routine to be called by the IDL interpreter for these disabled routines. If this argument is not provided, a default routine is used.

Result

All routines are enabled/disabled as specified. If a non-existent routine is specified, it is quietly ignored. Attempts to enable routines disabled for licensing reasons are also quietly ignored.

Note

The routines CALL_FUNCTION, CALL_METHOD (function and procedure), CALL_PROCEDURE, and EXECUTE are not real system routines, but are actually special cases that result in different IDL pcode. For this reason, they cannot be disabled. However, anything they *can* call can be disabled, so this is not a serious drawback.

Obtaining Enabled/Disabled Routine Names

The `IDL_SysRtnGetEnabledNames()` function can be used to obtain the names of all system routines which are currently enabled or disabled, either due to licensing reasons (i.e., some routines are disabled in IDL demo mode) or due to a call to `IDL_SysRtnEnable()`.

Syntax

```
void IDL_SysRtnGetEnabledNames(int is_function,  
                               IDL_STRING *str, int enabled)
```

Arguments

is_function

Set to TRUE if a list of functions is desired, FALSE for a list of procedures.

str

Points to a buffer of IDL_STRING descriptors to fill in. The caller must call `IDL_SysRtnNumEnabled()` to determine how many such routines exist, and this buffer must be large enough to hold that number.

enabled

Set to TRUE to receive names of enabled routines, FALSE to receive names of disabled ones.

Result

The memory supplied via `str` is filled in with the desired names.

Obtaining the Number of Enabled/Disabled Routines

The `IDL_SysRtnGetEnabledNames()` function requires you to supply a buffer large enough to hold all of the names to be returned. `IDL_SysRtnNumEnabled()` can be called to obtain the number of such routines, allowing you to properly size the buffer.

Syntax

```
IDL_MEMINT IDL_SysRtnNumEnabled(int is_function, int enabled)
```

Arguments

is_function

Set to TRUE if the number of functions is desired, FALSE for procedures.

enabled

Set to TRUE to receive number of enabled routines, FALSE to receive number of disabled ones.

Result

Returns the requested count.

Obtaining the Real Function Pointer

The `IDL_SysRtnGetRealPtr()` routine returns the pointer to the actual internal IDL function that implements the system function or procedure of the specified name.

This routine can be used to interpose your own code in between IDL and the actual routine. This process is sometimes called *hooking* in other systems. To implement such a hook function, you must use the `IDL_SysRtnEnable()` function to register the interposed routine, which in turn uses `IDL_SysRtnGetRealPtr()` to obtain the actual IDL function pointer for the routine.

Syntax

```
IDL_SYSRTN_GENERIC IDL_SysRtnGetRealPtr(int is_function,  
                                         char *name)
```

Arguments

`is_function`

Set to TRUE if functions are being manipulated, FALSE for procedures.

`name`

The name of function or procedure for which the real function pointer is required.

Result

If the specified routine...

- exists and is not disabled, its function pointer is returned.
- does not exist, a NULL pointer is returned.
- has been disabled by the user, its actual function pointer is returned.
- has been disabled for licensing reasons, the real function pointer does not exist, and the pointer to its stub is returned.

Note

This routine can cause an `IDL_MSG_LONGJMP` message to be issued if the function comes from a DLM and the DLM load fails due to memory allocation errors. Therefore, it must not be called unless the IDL interpreter is active. The prime intent for this routine is to call it from the stub routine of a disabled function when the interpreter invokes the associated system routine.

Obtaining the IDL Name of the Current System Routine

To get the IDL name for the currently executing system routine, use the `IDL_SysRtnGetCurrentName()`.

Syntax

```
char *IDL_SysRtnGetCurrentName(void)
```

This function returns a pointer to the name of the currently executing system routine. If there is no currently executing system routine, a NULL (0) pointer is returned.

This routine will never return NULL if called from within a system routine.

LINKIMAGE

The IDL user level `LINKIMAGE` procedure makes the functionality of the `IDL_SysRtnAdd()` function available to IDL programs. It allows IDL programs to merge routines written in other languages with IDL at run-time. Each call to `LINKIMAGE` defines a new system procedure or function by specifying the routine's name, the name of the file containing the code, and the entry point name. The name of your routine is added to IDL's internal system routine table, making it available in the same manner as any other IDL built-in routine.

`LINKIMAGE` is the easiest way to add your system routines to IDL. It does not require linking a separate version of the IDL program with your code the way a direct call to `IDL_SysRtnAdd()` does, and it does not require writing the extra code required for a Dynamically Loadable Module (DLM). It is therefore commonly used for simple applications, and for testing during the development of a system routine.

If you are developing a larger application, or if you intend to redistribute your work, you should package your routines as Dynamically Loadable Modules, which are much easier for end-users to install and use than `LINKIMAGE` calls. You will find that the small additional programming effort is more than repaid from the time saved providing support for your code to your users.

If your IDL application relies on code written in languages other than IDL and linked into IDL using the `LINKIMAGE` procedure, you must make sure that the routines declared with `LINKIMAGE` are linked into IDL before any code that calls them is restored. In practice, the best way to do this is to make the calls to `LINKIMAGE` in your `MAIN` procedure, and include the code that uses the linked routines in a secondary `.SAV` file. In this case your `MAIN` procedure may look something like this:

```
PRO main

;Link the external code.
LINKIMAGE, 'link_function', 'new.dll'

;Restore code that uses linked code.
RESTORE, 'secondary.sav'

;Run your application.
myapp

END
```

In this scenario, the IDL code that calls the `LINK_FUNCTION` routine (the routine linked into IDL in the `LINKIMAGE` call) is contained in the secondary `.SAV` file `'secondary.sav'`.

Note

When creating your secondary `.SAV` file, you will need to issue the `LINKIMAGE` command before calling the `SAVE` procedure to link your routine into IDL after you have exited and restarted. The `RESOLVE_ALL` routine does not resolve routines linked to IDL with the `LINKIMAGE` procedure.

Dynamically Loadable Modules do not have this issue, and are the best way to avoid the problem.

Dynamically Loadable Modules

LINKIMAGE can be used to make IDL load your system routines in a simple and efficient manner. However, it quickly becomes inconvenient if you are adding more than a few routines. Furthermore, the limitation that the LINKIMAGE call must happen before any code that calls it is compiled makes it difficult to use and complicates the process of redistributing your routines to others. IDL offers an alternative method of packaging your system routines, called Dynamically Loadable Modules (DLMs), that address these and other problems.

This section covers the following topics:

- [DLM Concepts](#) (page 309)
- [How DLMs Work](#) (page 310)
- [The Module Description \(.dlm\) File](#) (page 312)
- [The IDL_Load\(\) function](#) (page 314)
- [DLM Example](#) (page 315)
- [Packaging and Installing DLMs](#) (page 317)

DLM Concepts

The IDL_SYSFUN_DEF2 structure, which is described in “[Registering Routines](#)” on page 296, contains all the information required by IDL for it to be able to compile calls to a given system routine and call it:

- A routine signature (Name, minimum and maximum number of arguments, if the routine accepts keywords).
- A pointer to a compiled language function (usually C) that supplies the standard IDL system routine interface (argc, argv, argk) and which implements the desired operation.

IDL does not require the actual code that implements the function until the routine is called: It is able to compile other routines and statements that reference it based only on its signature.

DLMs exploit this fact to load system routines on an “as needed” basis. The routines in a DLM are not loaded by IDL unless the user calls one of them. A DLM consists of two files:

1. A module description file (human readable text) that IDL reads when it starts running. This file tells IDL the signature for all system routines contained in the loadable module.
2. A sharable library that implements the actual system routines. This library must be coded to present a specific IDL mandated interface (described below) that allows IDL to automatically load it when necessary without user intervention.

DLMs are a powerful way to extend IDL's built-in system routines. This form of packaging offers many advantages:

- Unlike LINKIMAGE, IDL automatically discovers DLMs when it starts up without any user intervention. This makes them easy to install — you simply copy the two files into a directory on your system where IDL will look for them.
- DLM routines work exactly like standard built-in routines, and are indistinguishable from them. There is no need for the user to load them (for example, using LINKIMAGE) before compiling code that references them.
- As the amount of code added to IDL grows, using sharable libraries in this way prevents name collisions in unrelated compiled code from fooling the linker into linking the wrong code together. DLMs thus act as a firewall between unrelated code. For example, there are instances where unrelated routines both use a common third party library, but they require different versions of this library. A specific example is that the HDF support in IDL requires its own version of the NetCDF library. The NetCDF support uses a different incompatible version of this library with the same names. Use of DLMs allows each module to link with its own private copy of such code.
- Since DLMs are separate from the IDL program, they can be built and distributed on their own schedule independent of IDL releases.
- System routines packaged as DLMs are effectively indistinguishable from routines built-into IDL by ITT Visual Information Solutions.

Use of sharable libraries in this manner has ample precedent in the computer industry. Most modern operating systems use loadable kernel modules to keep the kernel small while the functionality grows. The same technique is used in user programs in the form of sharable libraries, which allows unrelated programs to share code and memory space (e.g. a single copy of the C runtime library is used by all running programs on a given system).

How DLMs Work

IDL manages DLMs in the following manner:

1. When IDL starts, it looks in the current working directory for module definition (`.dlm`) files. It reads any file found and adds the routines and structure definitions thus defined to its internal routine and structure lookup tables as “stubs”. In the system routine dispatch table, stubs are entries that inform IDL of the routine’s existence, but which lack an actual compiled function to call. They contain sufficient information for IDL to properly compile calls to the routines, but not to actually call them. Similarly, stub entries in the structure definition table allow IDL to know that the DLM supplies the structure definition, but the actual definition is not present.

After it looks in the current working directory, IDL searches `!DLM_PATH` for `.dlm` files and adds them to the table in the same manner. The default value of `!DLM_PATH` is the directory in the IDL distribution where the binary executables are kept (`bin/bin.platform`), followed by the `idlde/plugins` directory, both in the IDL installation. This default can be changed by defining the `IDL_DLM_PATH` preference (this is similar to the way the `IDL_PATH` preference works with `!PATH`). This process happens once at startup, and never again. This means that IDL’s knowledge of loadable modules is static and unchangeable once the session is underway. This is very different from the way `!PATH` works, and reflects the static nature of built-in routines. The format of `.dlm` files is discussed in “[The Module Description \(.dlm\) File](#)” on page 312.

Warning

If you redefine the `IDL_DLM_PATH` preference, be sure to include the token `<IDL_DEFAULT>`. IDL will not run correctly if the default DLM directories are not included in `!DLM_PATH`.

See “[Packaging and Installing DLMs](#)” on page 317 for additional information about how IDL selects sharable libraries on different platforms.

2. The IDL session then continues in the usual fashion until a call to a routine from a loadable module occurs. At that time, the IDL interpreter notices the fact that the routine is a stub, and loads the sharable library for the loadable module that supplies the routine. It then looks up and calls a function named **`IDL_Load()`**, which is required to exist, from the library. It’s job is to replace the stubs from that module with real entries (by using **`IDL_SysRtnAdd()`**) and otherwise prepare the module for use.
3. Once the module is loaded, the interpreter looks up the routine that caused the load one more time. If it is still a stub then the module has failed to load properly and an error is issued. Normally, a full routine entry is found and the interpreter successfully calls the routine.

4. At this point the module is fully loaded, and cannot be distinguished from a compiled part of IDL. A module is only loaded once, and additional calls to any routine, or access to any structure definition, from the module are made immediately and without requiring any additional loading.

The Module Description (.d1m) File

The module description file is a simple text file that is read by IDL when it starts. Module description files have the file suffix `.d1m`. The information in the `.d1m` file tells IDL everything it needs to know about the routines supplied by a loadable module. With this information, IDL can compile calls to these routines and otherwise behave as if it contains the actual routine. The loadable module itself remains unloaded until a call to one of its routines is made, or until the user forces the module to load by calling the IDL `DLM_LOAD` procedure.

Empty lines are allowed in `.d1m` files. Comments are indicated using the `#` character. All text from a `#` to the end of the line is ignored by IDL and is for the user's benefit only.

All other lines start with a keyword indicating the type of information being conveyed, possibly followed by arguments. The syntax of each line depends on the keyword. Possible lines are:

MODULE Name

Gives the name of the DLM. This should always be the first non-comment line in a `.d1m` file. There can only be one `MODULE` line.

```
MODULE JPEG
```

DESCRIPTION DescriptiveText

Supplies a short one line description of the purpose of the module. This information is displayed by **HELP, /DLM**. This line is optional.

```
DESCRIPTION IDL JPEG support
```

VERSION VersionString

Supplies a version string that can be used by the IDL user to determine which version of the module will be used. IDL does not interpret this string, it only displays it as part of the **HELP, /DLM** output. This line is optional.

```
VERSION 6a
```


BUILD_DATE **DateString**

If present, IDL will display this information as part of the output from **HELP**, **/DLM**. IDL does not parse this string to determine the date, it is simply for the users benefit. This line is optional.

```
BUILD_DATE JAN 8 1998
```

SOURCE **SourceString**

A short one line description of the person or organization that is supplying the module. This line is optional.

```
SOURCE ITT Visual Information Solutions
```

CHECKSUM **CheckSumValue**

This directive is used by ITT Visual Information Solutions to sign the authenticity of the DLMs supplied with IDL releases. It is not required for user-written DLMs.

STRUCTURE **StructureName**

There should be one **STRUCTURE** line in the DLM file for every named structure definition supplied by the loadable module. If you refer to such a structure before the DLM is loaded, IDL uses this information to cause the DLM to load. The **IDL_Init()** function for the DLM will define the structure.

GLOBAL_SYMBOLS

This line is optional. Including this line in the DLM file will cause the shared library to load all of its symbols (functions or procedures) as *globally accessible* rather than locally accessible. If a symbol is globally accessible, then libraries that are loaded later will be able to access the symbol. In practice, adding this line to the DLM file will cause IDL to set the **RTL_D_GLOBAL** flag when calling the `dlopen()` operating system function to load the module.

On Microsoft Windows and Macintosh OS X systems, symbols are automatically loaded as global. A **GLOBAL_SYMBOLS** line in the DLM file will be quietly ignored.

Use caution when making a DLM's symbols globally accessible. Judicious naming of the DLM's symbol names will help ensure that symbols exported by the DLM will not cause namespace collisions with symbols from other libraries.

FUNCTION RtnName [MinArgs] [MaxArgs] [Options...]

PROCEDURE RtnName [MinArgs] [MaxArgs] [Options...]

There should be one FUNCTION or PROCEDURE line in the DLM file for every IDL routine supplied by the loadable module. These lines give IDL the information it needs to compile calls to these routines before the module is loaded.

RtnName

The IDL user level name for the routine. The routine name can be a simple procedure or function name (e.g. MY_PROCEDURE or MY_FUNCTION), or the name of an object method (e.g. MY_OBJECT::PROCEDURE_METHOD or MY_OBJECT::FUNCTION_METHOD).

MinArgs

The minimum number of arguments accepted by this routine. If not supplied, 0 is assumed.

MaxArgs

The maximum number of arguments accepted by this routine. If not supplied, 0 is assumed.

Options

Zero or more of the following:

OBSOLETE — IDL should issue a warning message if this routine is called and !WARN.OBS_ROUTINE is set.

KEYWORDS — This routine accepts keywords as well as plain arguments.

For example, a procedure named READ_JPEG that accepts a minimum of one argument, a maximum of three arguments, and also accepts keyword arguments would have the following definition in the .d1m file:

```
PROCEDURE READ_JPEG 1 3 KEYWORDS
```

The IDL_Load() function

Every loadable module sharable library must export a single symbol called **IDL_Load()**. This function is called when IDL loads the module, and is expected to do all the work required to load real definitions for the routines supplied by the function and prepare the module for use. This always requires at least one call to **IDL_SysRtnAdd()**. It usually also requires a call to **IDL_MessageDefineBlock()** if the module defines any messages. Any other initialization needed would also go here:

```
int IDL_Load(void)
```

This function takes no arguments. It is expected to return *True* (non-zero) if it was successful, and *False* (0) if some initialization step failed.

DLM Example

This example creates a loadable module named **TESTMODULE**.

Note

Code for this example is included in the `external/dlm` subdirectory of the IDL installation.

TESTMODULE provides 2 routines:

TESTFUN

A function that issues a message indicating that it was called, and then returns the string “TESTFUN” This function accepts between 0 and **IDL_MAXPARAMS** arguments, but it does not use them for anything.

TESTPRO

A procedure that issues a message indicating that it was called. This procedure accepts between 0 and **IDL_MAX_ARRAY_DIM** arguments, but it does not use them for anything.

The intent of this example is to show the support code required to write a DLM for a completely trivial application. This framework can be easily adapted to real modules by replacing TESTFUN and TESTPRO with other routines.

The first step is to create the module definition file for TESTMODULE, named `testmodule.dlm`:

```
MODULE testmodule
DESCRIPTION Test code for loadable modules
VERSION 1.0
SOURCE ITT Visual Information Solutions
BUILD_DATE JAN 8 1998
FUNCTION TESTFUN 0 IDL_MAXPARAMS
PROCEDURE TESTPRO 0 IDL_MAX_ARRAY_DIM
```

The next step is to write the code for the sharable library. The contents of `testmodule.c` are shown in the following figure. Comments in the code explain what each step is doing.

```

1  #include <stdio.h>
2  #include "idl_export.h"
3
4  /* Define message codes and their corresponding printf(3) format
5   * strings. Note that message codes start at zero and each one is
6   * one less than the previous one. Codes must be monotonic and
7   * contiguous. */
8  static IDL_MSG_DEF msg_arr[] = {
9  #define M_TM_INPRO          0
10 { "M_TM_INPRO", "%NThis is from a loadable module procedure." },
11 #define M_TM_INFUN         -1
12 { "M_TM_INFUN", "%NThis is from a loadable module function." },
13 };
14
15 /* The load function fills in this message block handle with the
16 * opaque handle to the message block used for this module. The other
17 * routines can then use it to throw errors from this block. */
18 static IDL_MSG_BLOCK msg_block;
19
20 /* Implementation of the TESTPRO IDL procedure */
21 static void testpro(int argc, IDL_VPTR *argv)
22 { IDL_MessageFromBlock(msg_block, M_TM_INPRO, IDL_MSG_RET); }
23
24 /* Implementation of the TESTFUN IDL function */
25 static IDL_VPTR testfun(int argc, IDL_VPTR *argv)
26 {
27     IDL_MessageFromBlock(msg_block, M_TM_INFUN, IDL_MSG_RET);
28     return IDL_StrToSTRING("TESTFUN");
29 }
30
31 int IDL_Load(void)
32 {
33     /* These tables contain information on the functions and procedures
34     * that make up the TESTMODULE DLM. The information contained in these
35     * tables must be identical to that contained in testmodule.dlm.
36     */
37     static IDL_SYSFUN_DEF2 function_addr[] = {
38         { testfun, "TESTFUN", 0, IDL_MAXPARAMS, 0, 0},
39     };
40     static IDL_SYSFUN_DEF2 procedure_addr[] = {
41         { (IDL_SYSRTN_GENERIC) testpro, "TESTPRO", 0, IDL_MAX_ARRAY_DIM, 0, 0},
42     };
43
44     /* Create a message block to hold our messages. Save its handle where
45     * the other routines can access it. */
46     if (!(msg_block = IDL_MessageDefineBlock("Testmodule",
47                                             IDL_CARRAYELTS(msg_arr),
48                                             msg_arr))) return IDL_FALSE;
49
50     /* Register our routine. The routines must be specified exactly the same
51     * as in testmodule.dlm. */
52     return IDL_SysRtnAdd(function_addr, TRUE,
53                          IDL_CARRAYELTS(function_addr))
54        && IDL_SysRtnAdd(procedure_addr, FALSE,
55                          IDL_CARRAYELTS(procedure_addr));
56 }

```

C

Table 15-7: testmodule.c

If building a DLM for Microsoft Windows, a linker definition file (testmodule.def) is also needed. All of these files, along with the commands

required to build the module can be found in the `d1m` subdirectory of the `external` directory of the IDL distribution.

Once the loadable module is built, you can cause IDL to find it by doing one of the following:

- Move to the directory containing the `.d1m` and sharable library for the module.
- Define the `IDL_DLM_PATH` preference to include the directory.

Running IDL to demonstrate the resulting module:

```
IDL> HELP,/DLM,'testmodule'
** TESTMODULE - Test code for loadable modules (not loaded)
Version:1.0,Build Date:JAN 8 1998,Source:ITT Visual Information
Solutions.
Path: /home/user/testmodule/external/testmodule.so
IDL> testpro
% Loaded DLM: TESTMODULE.
% TESTPRO: This is from a loadable module procedure.
IDL> HELP,/DLM,'testmodule'
** TESTMODULE - Test code for loadable modules (loaded)
Version:1.0,Build Date:JAN 8 1998,Source:ITT Visual Information
Solutions.
Path: /home/user/testmodule/external/testmodule.so
IDL> print, testfun()
% TESTFUN: This is from a loadable module function.
TESTFUN
```

The initial `HELP` output shows that the module starts out unloaded. The call to `TESTPRO` causes the module to be loaded. As IDL loads the module, it prints an announcement of the fact (similar to the way it announces the `.pro` files it automatically compiles to satisfy calls to user routines). Once the module is loaded, subsequent calls to `HELP` show that it is present. Calls to routines from this module do not cause the module to be reloaded (as evidenced by the fact that calling `TESTFUN` did not cause an announcement message to be issued).

Packaging and Installing DLMs

Once you have created sharable library (`.so` or `.dll`) and module description (`.d1m`) files, you will need to ensure that the files are installed in a location where IDL can find and load the libraries. Your approach may be slightly different depending on whether your dynamically loadable module supports a single platform or multiple platforms.

Single-Platform DLMS

If your module will be installed only on computers of a single architecture (32-bit Windows machines, for example, or 64-bit Linux machines), the process is relatively simple:

1. Create the sharable library file. The file will have the extension `.dll` for Microsoft Windows platforms, or `.so` for UNIX-like platforms (Macintosh, Linux, Solaris).
2. Create the module description file (`.dlm`) as described “[The Module Description \(.dlm\) File](#)” on page 312.
3. Place both the sharable library file and the module description file in a directory included in IDL’s `IDL_DLM_PATH` preference. See “[Installing DLMS Using the IDL Workbench Update Mechanism](#)” on page 322 for additional notes.
4. Restart IDL.

Even if your module supports only one platform, consider following the naming rules described in “[How IDL Selects the Correct Sharable Library File](#)” on page 319.

Using the multi-platform naming rules incurs no performance penalty, and may save effort if you end up supporting other platforms in the future.

Multi-Platform DLMS

If your module will be installed on computers of different architectures, you must create a unique sharable library file for each architecture. To install the DLM on a user’s machine, you have the following options:

Create Platform-Specific Installations

If you create a separate installation package for each architecture, creating a multi-platform DLM is essentially just creating a series of [Single-Platform DLMS](#), one for each platform. Use caution with this approach, since you will have to ensure that if your end-user installs more than one platform’s version of the DLM, the module description and shared library files for the different platforms are installed in the correct directories.

Create a Multi-Platform Installation

You can create a single installation package that supports multiple architectures if you follow a simple set of naming rules when creating your sharable library files. To create a multi-platform installation package:

1. Create a sharable library file for each platform, following the naming rules described in “[How IDL Selects the Correct Sharable Library File](#)” on page 319.
2. Create a single module description file (.d1m) as described in “[The Module Description \(.d1m\) File](#)” on page 312.
3. Place the module description file and all of the sharable libraries in a single directory included in IDL’s `IDL_DLM_PATH` preference. See “[Installing DLMs Using the IDL Workbench Update Mechanism](#)” on page 322 for additional notes.
4. Restart IDL.

How IDL Selects the Correct Sharable Library File

When IDL starts, it searches for DLMs in the directories included in IDL’s `IDL_DLM_PATH` preference as described in “[How DLMs Work](#)” on page 310. When IDL finds a module description file, it adds the routines and structure definitions defined by the DLM to its internal routine and structure lookup tables.

It is not until later, when a user calls a routine defined by the DLM, that IDL actually loads the sharable library. At this point, IDL searches for a sharable library file built for the current platform.

Note

IDL’s ability to search for platform-specific library file names was introduced in IDL 7.1.

IDL uses the following process to search for the sharable library file:

1. IDL constructs the base name of the library file by removing the .d1m suffix from the module definition file’s name.
2. To the library’s base name, IDL appends a platform-specific string. The specific strings are shown in [Table 15-8](#) below. The string is the concatenation of the name of the platform’s platform-specific `bin` subdirectory, along with the suffix `.dll` on Windows systems or `.so` on all UNIX-based systems.

For example, if the name of the DLM file is

```
my_module.d1m
```

then the platform-specific sharable library file name for a 64-bit Linux platform would be

```
my_module.linux.x86_64.so
```

3. IDL searches the directory that contains the module definition file (`.d1m`) for a library file with the platform-specific sharable library file name. If it finds a matching file, it loads the library and executes the routine called by the user.
4. If IDL does not find the platform-specific library file, it searches the directory that contains the module definition file (`.d1m`) for a library file with the same base name as the module definition file, replacing the `.d1m` extension with the suffix `.dll` or `.so`.

For example, if the name of the DLM file is

```
my_module.d1m
```

then the generically-named sharable library file name would be

```
my_module.dll
```

on a Windows system, or

```
my_module.so
```

on a UNIX system.

If IDL finds a generically-named sharable library file (with *either* the `.dll` or the `.so` extension), it attempts to load the library and execute the routine called by the user. Note that IDL will only be able to successfully load the library if the generically-named library file was built for the current platform.

5. If IDL fails to find *either* the platform-specific sharable library file or the generically-named library file, it will issue one of the following error messages:
 - If a platform-specific sharable library file for a different platform exists in the same directory, the error message is


```
Dynamically loadable module is unavailable on this platform:
my_module.
```
 - If no platform-specific sharable library files for any platform are present, the error message is


```
Dynamically loadable module failed to load: my_module.
```

The first message indicates that the DLM exists but is not supported for the current platform, the second indicates that the DLM does not exist, despite the presence of the `.d1m` file.

One benefit of this file naming and search procedure is that you can distribute a DLM package that includes library files for several platforms in a single directory. IDL will load the correct shared library for the end-user's platform, or provide a sensible error message if the platform is not supported.

Platform-Specific Sharable Library File Suffixes

The following table lists the platform-specific file suffixes for IDL's supported platforms:

Platform	Sharable Library File Suffix
Windows 32-bit	.x86.dll
Windows 64-bit	.x86_64.dll
Solaris SPARC 32-bit	.solaris2.sparc.so
Solaris SPARC 64-bit	.solaris2.sparc64.so
Solaris x86 64-bit	.solaris2.x86_64.so
Linux 32-bit	.linux.x86.so
Linux 64-bit	.linux.x86_64.so
Macintosh OS X PPC 32-bit	.darwin.ppc.so
Macintosh OS X Intel 32-bit	.darwin.i386.so
Macintosh OS X Intel 64-bit	.darwin.x86_64.so

Table 15-8: Sharable Library File Suffixes

Example DLM Distribution

For example, suppose you have created a dynamically loadable module named `my_cool_module`, and created sharable libraries for Windows (32- and 64-bit) and Linux (32- and 64-bit) but not for Macintosh OS X or Solaris. Your DLM installation directory would contain the following files:

```
my_cool_module.dlm
my_cool_module.x86.dll
my_cool_module.x86_64.dll
my_cool_module.linux.x86.so
my_cool_module.linux.x86_64.so
```

If a user on a Macintosh OS X or Solaris system attempts to call a routine from the `my_cool_module` DLM, the fact that sharable libraries for other platforms exist informs IDL that the DLM *intentionally* does not provide support for that platform. If a user on one of these unsupported platforms attempts to use the functionality from the DLM, IDL will issue the message

```
Dynamically loadable module is unavailable on this platform:  
my_cool_module.
```

Installing DLMS Using the IDL Workbench Update Mechanism

Beginning in IDL 7.1, the `idlde/plugins` subdirectory of the IDL installation is automatically added to the list of directories in the `IDL_DLM_PATH` preference. This, coupled with IDL's ability to select a sharable library file based on a platform-specific file name, allows you to distribute multi-platform DLMS using the IDL Workbench update mechanism. See [IDL Workbench Questions and Answers](#) for more on the Workbench update mechanism.

Briefly stated, you can distribute your module using the Workbench update mechanism if you create an Eclipse update site that includes (among other things) a plug-in that contains your module description file and sharable libraries. When an end-user installs the plug-in, the module description and library files are placed in the `idlde/plugins` subdirectory of the IDL installation automatically, and are thus found by the DLM search mechanism the next time IDL starts. See [IDL Plug-in Wizard](#) for information on how to generate a plug-in and update site.



Chapter 16

Callable IDL

This chapter discusses the following topics:

Calling IDL as a Subroutine	324	Executing IDL Statements	339
When is Callable IDL Appropriate?	325	Runtime IDL and Embedded IDL	340
Licensing Issues and Callable IDL	328	Cleanup	341
Using Callable IDL	329	Issues and Examples: UNIX	342
Initialization	331	Issues and Examples: Microsoft Windows	358
Diverting IDL Output	337		

Calling IDL as a Subroutine

IDL can be called as a subroutine from other programs. This capability is referred to as Callable IDL to distinguish it from the more common case of calling your code from IDL (as with `CALL_EXTERNAL` or as a system routine (`LINKIMAGE`, Dynamically Loadable Module)).

How Callable IDL is Implemented

IDL is built in a sharable form that allows other programs to call IDL as a subroutine. The specific details of how IDL is packaged depend on the platform:

- IDL for UNIX has a small driver program linked to a sharable object library that contains the actual IDL program.
- IDL for Windows consists of a driver program that implements the user interface (known as the IDE) linked to a dynamic-link library (DLL) that contains the actual IDL program.

In all cases, it is possible to link the sharable portion of IDL into your own programs. Note that Callable IDL is not a separate copy of IDL that implements a library version of IDL. It is in fact the same code, being used in a different context.

When is Callable IDL Appropriate?

Although Callable IDL is very powerful and convenient, it is not always the best method of communication between IDL and other programs. There are usually easier approaches that will solve a given problem. See [“Supported Inter-Language Communication Techniques in IDL”](#) on page 13 for alternatives.

IDL will not integrate with *all* programs. Understanding the issues described in this section will help you decide when Callable IDL is and is not appropriate.

Technical Issues Relating to Callable IDL

IDL makes computing easier by raising the level at which IDL users interface with the computer. It is natural to think that calling IDL from other programs will have the same effect, and under the correct circumstances this is true. However, using Callable IDL is not as easy as using IDL. Programmers who wish to use Callable IDL need to possess the skills described in [“Skills Required to Combine External Code with IDL”](#) on page 23.

Be aware that the same things that make IDL powerful at the user level can make it difficult to include in other programs. As an interactive, interpreted language, IDL is a decidedly non-trivial object to add to a process. Unlike a simple mathematical subroutine, IDL includes a compiler, a language interpreter, and related code that the caller must work around. As an interactive program, IDL must control the process to a high degree, which can conflict with the caller’s wishes. The following (certainly incomplete) list summarizes some of the issues that must be dealt with.

UNIX IDL Signal API

IDL uses UNIX signals to manage many of its features, including exception handling, user interrupts, and child processes. The exact signals used and the manner in which they are used can change from IDL release to release as necessary. Although the IDL signal API (described in [“IDL Internals: UNIX Signals”](#) on page 209) allows you to use signals in an IDL-compatible way, the resulting constraints may require changes to your code.

IDL Timer API

IDL’s use of the process timer requires you to use the IDL timer API instead of the standard system routines. This restriction may require changes to some programs. Under UNIX, the timer module can interrupt system calls. Timers are discussed in [“IDL Internals: Timers”](#) on page 221.

GUI Considerations

Most applications will call IDL and display IDL graphics in an IDL window. However, programmers may want to write applications in which they create the graphical user interface (GUI) and then have IDL draw graphics into windows that IDL did not create. It is not always possible for IDL to draw into windows that it did not create for the reasons described below:

X Windows

The IDL X Windows graphics driver can draw in windows it did not create as long as the window is compatible with the IDL display connection (see [Appendix A, “IDL Direct Graphics Devices”](#) (*IDL Reference Guide*) for details). However, the design of IDL’s X Windows driver requires that it open its own display connection and run its own event loop. If your program cannot support a separate display connection, or if dividing time between two event loops is not acceptable, consider the following options:

- Run IDL in a separate process and use interprocess communication (possibly Remote Procedure Calls, to control it.
- If you choose to use Callable IDL, use the IDL Widget stub interface, described in [“Adding External Widgets to IDL”](#) on page 369, to obtain the IDL display connection, and create your GUI using that connection rather than creating your own. The IDL event loop will dispatch your events along with IDL’s, creating a well-integrated system.

Microsoft Windows

At this time, the IDL for Windows graphics driver does not have the ability to draw into windows that were not created by IDL.

Program Size Considerations

On systems that support preemptive multitasking, a single huge program is a poor use of system capabilities. Such programs inevitably end up implementing primitive task-scheduling mechanisms better left to the operating system.

Troubleshooting

Troubleshooting and debugging applications that call IDL can be very difficult. With standard IDL, malfunctions in the program are clearly the fault of ITT Visual Information Solutions, and given a reproducible bug report, we attempt to fix them promptly. A program that combines IDL with other code makes it difficult to

unambiguously determine where the problem lies. The level of support we can provide in such troubleshooting is minimal. The programmer is responsible for locating the source of the difficulty. If the problem is in IDL, a simple program demonstrating the problem must be provided before we can address the issue.

Threading

IDL uses threads to implement its thread pool functionality, which is used to speed numerical computation on multi-CPU hardware. Despite this, it is essentially a single threaded program, and is not designed to be called from different threads of a threaded application. Attempting to use IDL from any thread other than the main thread is unsupported, and may cause unpredictable results.

Inter-language Calling Conventions

IDL is written in standard ANSI C. Calling it from other languages is possible, but it is the programmer's responsibility to understand the inter-language calling conventions of the target machine and compiler.

Appropriate Applications of Callable IDL

Callable IDL is most appropriate in the following situations:

- Callable IDL is clearly the correct choice when the resulting program is to be a front-end that creates a different interface for IDL. For example, you might wish to turn IDL into an RPC server that uses an RPC protocol not directly supported by IDL, or use IDL as a module in a distributed system.
- Callable IDL is appropriate if either the calling program or IDL handles *all* graphics, including the Graphical User Interface, *without the involvement of the other*. Intermediate situations are possible, but more difficult. In particular, beware of attempts to have two event/message loops.
- Callable IDL is appropriate when the calling program makes little or no use of signals, timers, or exception handling, or is able to operate within the constraints imposed by IDL.

Licensing Issues and Callable IDL

If you intend to distribute an application that calls IDL, note that each copy of your application must have access to a properly licensed copy of the IDL library. For availability of a runtime version of IDL, contact ITT Visual Information Solutions or your IDL distributor.

Using Callable IDL

The process of using Callable IDL has three stages: initialization, IDL use, and cleanup. Between the initialization and the cleanup, your program contains a complete active IDL session, just as if a user were typing commands at an `IDL>` prompt. In addition to the usual IDL abilities, you can import data from your program and cause IDL to see it as an IDL variable. IDL can use such data in computations as if it had created the variable itself. In addition, you can obtain pointers to data currently held by IDL variables and access the results of IDL computations from your program.

Note

The functions documented in this chapter should only be used when calling IDL from other programs—their use in code called by IDL via `CALL_EXTERNAL` or a system routine (`LINKIMAGE`, Dynamically Loadable Module) is not supported and is certain to corrupt and/or crash the IDL process.

Before calling IDL to execute instructions, you must initialize it. This is done by calling `IDL_Initialize()`. This is a one-time operation, and must occur before calling any other IDL function. For complete information on this topic, see [“Initialization”](#) on page 331. Once IDL is initialized, you can:

1. Send IDL commands to IDL for execution. Commands are sent as strings, using the same syntax as interactive IDL. Note that there is not a separate C language function for every IDL command—any valid IDL command can be executed as IDL statements. This approach allows us to keep the callable IDL API small and simple while allowing full access to IDL’s abilities. This is explained in [“Executing IDL Statements”](#) on page 339.
2. Call any of the several routines that interact with IDL through other means to perform operations such as:
 - Importing data into IDL. (See [“Creating an Array from Existing Data”](#) on page 172.)
 - Accessing data within IDL. (See [“Looking Up Variables in Current Scope”](#) on page 182.)
 - Changing items in the process, such as signal handling or timers. (See [“IDL Internals: UNIX Signals”](#) on page 209, or [“IDL Internals: Timers”](#) on page 221.)
 - Redirecting IDL output to your own function for processing. See [“Diverting IDL Output”](#) on page 337.

The above list is not complete, but is representative of the possibilities afforded by Callable IDL.

Cleanup

After all IDL use is complete, but before the program exits, you must call **IDL_Cleanup()** to allow IDL to shutdown gracefully and clean up after itself. Once this has been done, you are not allowed to call IDL again from this process. See [“Cleanup”](#) on page 341.

Initialization

The **IDL_Initialize()** function is used to prepare Callable IDL for use. As a convenience in simpler situations, the **IDL_Init()** function may also be used for this purpose.

Note

IDL can only be initialized once for a given process; calling an IDL initialization function more than once for a process will cause an error. If you need to reinitialize an IDL session that is already running, consider using

```
IDL_ExecuteStr(".reset_session");
```

IDL_Initialize()

The **IDL_Initialize()** function is the primary function used to prepare Callable IDL for use. This must be the first IDL routine called.

```
int IDL_Initialize(IDL_INIT_DATA *init_data)
```

IDL_Initialize() returns TRUE if IDL was successfully initialize, and FALSE otherwise:

init_data

A pointer to an **IDL_INIT_DATA** structure, used to specify initialization options. If no initialization data is required, a NULL pointer may be passed.

The definition of **IDL_INIT_DATA** includes the following fields:

```
typedef struct {
    IDL_INIT_DATA_OPTIONS_T options;

    struct {
        int argc;
        char **argv;
    } clargs;

    void *hwnd;
} IDL_INIT_DATA;
```

The **options** field of **IDL_INIT_DATA** can be set to any combination of the **IDL_INIT_** values described below. Most of these values represent boolean (on/off) options and no other data is required for them. However, a relatively small number of

the options require additional information. This extra information is provided via one of the other fields in the **IDL_INIT_DATA** structure — the appropriate field to use in each case is discussed with each individual option below.

IDL_Initialize() always examines the value of the **options** field of this structure. It will only examine the other fields if a value in **options** requires it to. Otherwise, those other fields are not used and may safely be left uninitialized. This organization allows ITT Visual Information Solutions to add additional initialization options to newer versions of IDL without requiring source code changes to older applications that do not require those new features.

The values allowed in the options field of **IDL_INIT_DATA** are:

IDL_INIT_BACKGROUND

A convenience option, equivalent to setting both the **IDL_INIT_NOCMDLINE** and **IDL_INIT_NOTTYEDIT** options.

IDL_INIT_CLARGS

Execution of C programs starts when the `main()` function is called. Command line arguments are passed to the `main()` function via the standard `argc` and `argv` arguments. Set the **IDL_INIT_DATA** `clargs.argc` and `clargs.argv` fields to these values and set the **IDL_INIT_CFLAGS** bit in options to pass these command line arguments to IDL for processing. On return, IDL will remove any arguments it understands, and will alter the value of `clargs.argc` to reflect the count of remaining items.

The `argc/argv` pair passed must follow the usual convention in which `argv[0]` is the name under which the program was run, and any additional values are the arguments passed to that program.

IDL_INIT_EMBEDDED

IDL is initialized to run applications from a Save/Restore file that contains an embedded license. **IDL_RuntimeExec()** is then used to run the application(s).

IDL_INIT_GUI

Indicates that IDL is being accessed via the IDL Workbench GUI interface rather than using the standard tty based interface.

IDL_INIT_GUI_AUTO (Unix Only)

IDL will try to use the IDL Workbench GUI. If that fails, IDL uses the standard tty interface.

IDL_INIT_HWND (Microsoft Windows only)

Under Microsoft Windows, an application calling IDL will usually want IDL to use its main window as its own. This option is used to pass the application's main window handle to IDL. In addition to setting **IDL_INIT_HWND** in the **options** field, you must set the **hwnd** field to the value of the window handle to use.

IDL_INIT_LMQUEUE

At startup, if no license is immediately available, IDL will wait for an available license before continuing. This is useful for non-interactive IDL based tasks such as batch processing, where waiting is acceptable and processing cannot succeed without a license.

IDL_INIT_NOCMDLINE

Indicates to IDL that it is going to be used in a background mode by some other program, and that IDL will not be in control of the user's input command processing. The main effect of this is that IDL will never prompt for user input from the command line, and execution will never stop in such a situation. Instead, IDL will act as if the desired read returned an end of file (EOF) and IDL will instead return to the caller. Another related effect is that XMANAGER will realize that the active command line functionality for processing widget events is not available, and XMANAGER will block to manage events when it is called rather than return immediately.

IDL_INIT_NOLICALIAS

Our FLEXlm floating license policy is to alias all IDL sessions that share the same user/system/display to the same license. If **IDL_INIT_NOLICALIAS** is set, this IDL session will force a unique license to be checked out. In this case, we allow the user to change the DISPLAY environment variable. This is useful for RPC servers that don't know where their output will need to go before invocation.

IDL_INIT_NOTTYEDIT (Unix Only)

Normally under UNIX, if IDL sees that stdin and stdout are ttys, it puts the tty into raw mode and uses termcap/terminfo to handle command line editing. When using callable IDL in a background process that isn't doing input/output to the tty, the termcap initialization can cause the process to block (because of job control from the shell) with a message like "Stopped (tty output) idl". Setting this option prevents all tty edit functions and disables the calls to termcap. I/O to the tty is then done with a simple fgets()/printf(). This option only has meaning when a Unix tty is in use. It is ignored on non-Unix platforms, or when the **IDL_INIT_GUI** bit is set.

IDL_INIT_QUIET

Setting this bit suppresses the display of the startup announcement and message of the day.

IDL_INIT_RUNTIME

Setting this bit causes IDL to check out a runtime license instead of the normal license. **IDL_RuntimeExec()** is then used to run an IDL application restored from a Save/Restore file.

IDL_Init()

The **IDL_Init()** function offers a simplified interface to **IDL_Initialize()**. When possible, callable IDL programs should call **IDL_Initialize()** to perform the initialization operation. However, **IDL_Initialize()** requires the **IDL_INIT_DATA** structure type to be defined. This definition comes from the `idl_export.h` header file supplied with IDL, which can be used from the C or C++ languages, but which is not directly usable from languages such as Fortran. **IDL_Init()** does not use the **IDL_INIT_DATA** structure, and is therefore more convenient in such cases.

Note

Most Microsoft Windows applications need to pass their main window handle (HWND) to IDL, which is possible using **IDL_Initialize()**, but not **IDL_Init()**. **IDL_Init()** is therefore primarily of interest in Unix environments.

```
int IDL_Init(int options, int *argc, char *argv[]);
```

IDL_Init() returns TRUE if IDL was successfully initialized, and FALSE otherwise.

IDL_Init() is nothing more than a simple convenience wrapper written using **IDL_Initialize()**. As an aid in understanding the relationship between these two routines, the code that implements it is shown in [Table 16-1](#):

options

A bitmask used to specify initialization options. This is equivalent setting the **options** field of the **IDL_INIT_DATA** structure to the desired options value when using the **IDL_Initialize()** function. Allowed values for **options** can be found in [“Initialization”](#) on page 331.

argc, argv

Command line arguments, as passed by the operating system to the program `main()` function. Setting these arguments to non-NULL values is equivalent to the following 3 steps using `IDL_Initialize()`:

1. Set the `clargs.argc` field of the `IDL_INIT_DATA` structure to the number of items in the `argv` array, as given by `argc`.
2. Set the `clargs.argv` field of the `IDL_INIT_DATA` structure to the value of `argv`.
3. Set `IDL_INIT_CLARGS` bit of the `options` field of that structure.

```

1  int IDL_Init(int options, int *argc, char *argv[])
2  {
3      IDL_INIT_DATA init_data;
4      int r;
5
6      init_data.options = options;
7      if (argc) {
8          init_data.options |= IDL_INIT_CLARGS;
9          init_data.clargs.argc = *argc;
10         init_data.clargs.argv = argv;
11     }
12
13     r = IDL_Initialize(&init_data);
14     if (argc) *argc = init_data.clargs.argc;
15
16     return r;
17 }

```

Table 16-1: `IDL_Init()` Implementation Based on `IDL_Initialize()`.

Common Microsoft Windows Initialization Issues

Callable IDL applications intended to run under Microsoft Windows commonly face the following issues:

- Under Windows, it is usually the case that the use of IDL from another program is non-interactive. By default, IDL assumes an interactive environment in which it is communicating with a user directly. It is necessary to set the `IDL_INIT_NOCMDLINE` option to change this.

- Most Microsoft Windows applications have a main window that they wish IDL to use as its main window. The window handle for this window must be specified to **IDL_Initialize()**.

The function **MyAppInitIDL()**, shown below, demonstrates how to specify this information to **IDL_Initialize()**. This function accepts two arguments: **options** allows the caller to supply any other **IDL_INIT_** option values that the program may need, while **hwnd** allows the specification of a window handle to be used as the application main window.

```
1 int MyAppInitIDL(int options, HWND hwnd)
2 {
3     IDL_INIT_DATA init_data;
4
5     /* Combine any other IDL init options with NOCMDLINE */
6     init_data.options = options | IDL_INIT_NOCMDLINE;
7
8     /* If we have a non-NULL HWND, tell IDL to use it */
9     if (hwnd) {
10        init_data.options |= IDL_INIT_HWND;
11        init_data.hwnd = hwnd;
12    }
13
14    return IDL_Initialize(&init_data);
15 }
```

C

Table 16-2: Setting Initialization Information for Microsoft Windows Applications

Diverting IDL Output

When using a tty-based interface (available only on UNIX platforms), IDL sends its output to the screen for the user to see. When using a GUI-based interface (any platform), the output goes to the IDL log window. The default output function is automatically installed by IDL at startup. To divert IDL output to a function of your own design, use **IDL_ToutPush()** and **IDL_ToutPop()** to change the output function called by IDL.

Internally, IDL maintains a stack of output functions, and provides two functions (**IDL_ToutPush()** and **IDL_ToutPop()**) to manage them. The most recently pushed output function is called to output each line of text. Output functions of your own design should have the following type definition:

```
typedef void (* IDL_TOUT_OUTF)(int flags, char *buf, int n);
```

The arguments to an output function are:

flags

A bitmask of flag values that specify how the text should be output. The allowed bit values are:

IDL_TOUT_F_STDERR

Send the text to **stderr** rather than **stdout**, if that distinction means anything to your output device.

IDL_TOUT_F_NLPOST

After outputting the text, start a new output line. On a tty, this is equivalent to sending a newline ('\n') character.

buf

The text to be output. There may or may not be a NULL termination, so the character count provided by **n** must be used to move only the specified number of characters.

n

The number of characters in **buf** to be output.

IDL_ToutPush()

Use **IDL_ToutPush()** to push a new output function onto the stack. The most recently pushed function is the one used by IDL for output.

```
void IDL_ToutPush(IDL_TOUT_OUTF outf);
```

IDL_ToutPop()

IDL_ToutPop() removes the most recently pushed output function. The removed function pointer is returned.

```
IDL_TOUT_OUTF IDL_ToutPop(void);
```

Warning

Do not pop an output function you did not push. It is an error to attempt to remove the last remaining function.

Executing IDL Statements

There are two functions that allow you to execute IDL statements.

IDL_ExecuteStr() executes a single command, while **IDL_Execute()** takes an array of commands and executes them in order. In both cases, the commands are null terminated strings—just as they would be typed by an IDL user at the `IDL>` prompt. It is important to realize that the full abilities of IDL are available at this point.

Typically, the commands you issue will run IDL programs of varying complexity, including support routines written in IDL from the IDL Library (found via the IDL `!PATH` system variable). This ability to “download” complicated programs into IDL and then run them via a simple command can be very powerful.

IDL_Execute()

IDL_Execute() executes the command strings in the order given. It returns the value of `!ERROR_STATE.CODE` after the final command has executed. If the value of `!ERROR_STATE.CODE` is needed for an intermediate command, you should use **IDL_ExecuteStr()** instead of **IDL_Execute()**.

```
int IDL_Execute(int argc, char *argv[]);
```

argc

The number of commands contained in **argv**.

argv

An array of pointers to NULL-terminated strings containing IDL statements to execute.

IDL_ExecuteStr()

IDL_ExecuteStr() returns the value of the `!ERROR_STATE.CODE` system variable after the command has executed.

```
int IDL_ExecuteStr(char *cmd);
```

cmd

A NULL-terminated string containing an IDL statement to execute.

Runtime IDL and Embedded IDL

If you distribute programs that call IDL with a runtime license or an embedded license, use `IDL_RuntimeExec()`. After initialization `IDL_RuntimeExec()` can be used to run self-contained IDL applications from a Save/Restore file.

`IDL_RuntimeExec()` restores the file, then attempts to call an IDL procedure named MAIN. If no MAIN procedure is found, the function attempts to call a procedure with the same name as the restored Save file. (That is, if the Save file is named `myprog.sav`, `IDL_RuntimeExec()` looks for a procedure named `myprog`.)

Note

`IDL_RuntimeExec()` clears the value of the `!ERROR_STATE` system variable before it restores the specified Save file.

`IDL_RuntimeExec()` returns the value of the `!ERROR_STATE.CODE` system variable after IDL attempts to restore the specified file and execute the MAIN or named procedure. Thus, a return value of zero indicates that the specified Save file was restored and the appropriate procedure executed without error.

```
int IDL_RuntimeExec(char *file);
```

where:

file

The complete path specification to the Save file to be restored, in the native syntax of the platform in use.

Checking the Error Status

If the return value from `IDL_RuntimeExec()` is not zero, you may wish to check the values of other fields in the `!ERROR_STATE` structure. The following code fragment populates `buffer` with the values of the `!ERROR_STATE.MSG`, `!ERROR_STATE.SYS_MSG`, and `!ERROR_STATE.CODE` system variable fields:

```
sprintf(buffer, "error_state.msg: %s\nerror_state.sys_msg:
    %s\nerror_state: %d\n",
    IDL_STRING_STR(IDL_SysvErrStringFunc()),
    IDL_STRING_STR(IDL_SysvSyserrStringFunc()),
    IDL_SysvErrorCodeValue()
);
```

See [“Functions for Returning System Variables”](#) on page 257 for additional information.

Cleanup

After your program is finished using IDL (typically just before it exits) it should call **IDL_Cleanup()** to allow IDL to shut down gracefully. **IDL_Cleanup()** returns a status value that can be passed to **Exit()**.

```
int IDL_Cleanup(int just_cleanup);
```

where:

just_cleanup

If **TRUE**, **IDL_Cleanup()** does all the process shutdown tasks, but doesn't actually exit the process. If **FALSE** (the usual), the process exits.

Microsoft Windows applications should place this call in their Main **WndProc** to be called as a result of the **WM_CLOSE** message.

```
switch(msg) {  
    ...  
    case WM_CLOSE:  
        IDL_Cleanup(TRUE);  
        any additional processing  
    ...  
}
```

Issues and Examples: UNIX

Interactive IDL

Under UNIX, **IDL_Main()** implements IDL as seen by the interactive user. In the interactive version of IDL as shipped by ITT Visual Information Solutions, the actual **main()** function simply decodes its arguments to determine which options to specify and then calls **IDL_Main()** to do the rest. **IDL_Main()** calls **exit()** and does not return to its caller.

```
int IDL_Main(int init_options, int argc, char *argv[]);
```

where:

init_options

The **options** value to be passed to **IDL_Initialize()** via the **init_data** argument to that function.

argc, argv

From **main()**. Arguments that correspond to options specified via the **init_options** argument should be removed and converted to **init_options** flags prior to calling this routine.

Compiling Programs That Call IDL

A complete discussion of the issues that arise when compiling and linking C programs is beyond the scope of this manual. The following is a brief list of basic concepts to consider when building programs that call IDL.

- Compilers for some languages add underscores at the beginning or end of user defined names. To check the naming convention employed by your compiler, use the UNIX `nm (1)` command to list the symbols exported from an object file.

If you use only one language, naming details are handled transparently by the compiler, linker, and debugger. If you use more than one language, problems can arise if the different compilers use different naming conventions. For example, the Fortran compiler might add an underscore to the end of each name, while the C compiler does not. To call a Fortran routine from C, you must then include this underscore in your code (to call the function `my_code`, you would refer to it as `my_code_`). Note that you may also need to set a compiler flag to make case significant.

To determine whether your compilers use compatible naming conventions, consult your compiler documentation or experiment with small test programs using the compilers and the `nm` command.

- Every program starts execution at a known routine. In the C language, this routine is explicitly named `main()`. In Fortran, execution begins with the implicit main program. If you are using Callable IDL, you must provide a `main()` function for your program.
- When linking a C program, use the `cc` command instead of the `ld` command. `cc` calls `ld` to perform the link operation, and when necessary adds a directive to `ld` that causes the C runtime library to be used.

If you don't use `cc` to link your program (if you are using `ld` directly or are using a Fortran compiler, for example) and you get “unsatisfied symbol” errors for symbols that are in the standard C library, try including the runtime library explicitly in your link command. Usually, adding the string `-lc` to the end of the command is all that is necessary.

See “[Compilation and Linking Details](#)” on page 31 for advice on how to compile and link programs with the IDL libraries under various operating systems.

Example: Calling IDL From C

The program in the following figure is a simplified Unix-only version of `calltest.c`, found in the `callable` subdirectory of the `external` subdirectory of the IDL distribution. It demonstrates how to import data from a C program into IDL, execute IDL statements, and obtain data from IDL variables. It performs the following actions:

1. Create an array of 10 floating point values with each element set to the value of its index. This is equivalent to the IDL command `FINDGEN(10)`.
2. Initialize Callable IDL.
3. Import the floating point array into IDL as a variable named `TMP`.
4. Have IDL print the value of `TMP`.

5. Execute a short sequence of IDL statements from a string array:

```
tmp2 = total(tmp)
print, 'IDL total is ', tmp2
plot, tmp
```

6. Set TMP to zero, causing IDL to release the pointer to the floating point array.
7. Obtain a pointer to the data contained in TMP2. From examining the IDL statements executed to this point, we know that TMP2 is a scalar floating point value.
8. From our C program, print the value of the IDL TMP2 variable.
9. Execute a small widget program. Pressing the button allows the program to end:

```
a = widget_base()
b = widget_button(a, value='Press When Done', xsize=300,
                 ysize=200)
widget_control, /realize, a
dummy = widget_event(a)
widget_control, /destroy, a
```

See [“Compilation and Linking Statements”](#) on page 357 for details on compiling and linking this program.

Each line is numbered to make discussion easier. The line numbers are not part of the actual program.

C

```

1  #include <stdio.h>
2  #include "idl_export.h"
3
4  static void free_callback(UCHAR *addr)
5  {
6      printf("IDL released(%u)\n", addr);
7  }
8
9  int main(int argc, char **argv)
10 {
11     IDL_INIT_DATA init_data;
12     float f[10];
13     int i;
14     IDL_VPTR v;
15     IDL_MEMINT dim[IDL_MAX_ARRAY_DIM];
16     static char *cmds[] = { "tmp2 = total(tmp)",
17         "print,'IDL total is ',tmp2", "plot,tmp" };
18     static char *cmds2[] = { "a = widget_base()",
19         "b = widget_button(a, value='Press When Done', xsize=300,
20         ysize=200)", "widget_control,/realize, a",
21         "dummy = widget_event(a)",
22         "widget_control,/destroy, a" };
23
24
25     for (i=0; i < 10; i++) f[i] = (float) i;
26     init_data.options = IDL_INIT_CLARGS;
27     init_data.clargs.argc = argc;
28     init_data.clargs.argv = argv;
29     if (IDL_Initialize(&init_data)) {
30         dim[0] = 10;
31         printf("ARRAY ADDRESS(%u)\n", f);
32         if (v=IDL_ImportNamedArray("TMP", 1, dim, IDL_TYP_FLOAT,
33             (UCHAR *) f, free_callback, (void *) 0)) {
34             (void) IDL_ExecuteStr("print, tmp");
35             (void) IDL_Execute(sizeof(cmds)/sizeof(char *), cmds);
36             (void) IDL_ExecuteStr("print, 'Free the user memory'");
37             (void) IDL_ExecuteStr("tmp = 0");
38             if (v = IDL_FindNamedVariable("tmp2", IDL_FALSE))
39                 printf("Program total is %f\n", v->value.f);
40             (void) IDL_Execute(sizeof(cmds2)/sizeof(char *), cmds2);
41             IDL_Cleanup(IDL_FALSE); /* Don't return */
42         }
43     }
44     return 1;
45 }

```

Table 16-3: Calling IDL from C on UNIX

Following is commentary on this program, by line number:

25

C equivalent to IDL command “F = FINDGEN(10)”

26–28

Prepare initialization data.

29

Initialize IDL

30–33

Import C array **F** into IDL as a **FLTARR** vector named **TMP** with 10 elements. Note use of the callback argument **free_callback**. This function will be called when IDL is finished with the array **F**, giving us a chance to properly clean up at that time.

34

Have IDL print the value of **TMP**.

35

Execute the commands contained in the C string array **cmds** defined on lines 16–17. These commands create a new IDL variable named **TMP2** containing the sum of the elements of **TMP**, print its value, and plot the vector.

36–37

Set **TMP** to a new value. This will cause IDL to release the user supplied memory from lines 30–33 and call **free_callback**.

38–39

From C, get a reference to the IDL variable **TMP2** and print its value. This should agree with the value printed by IDL on line 35. It is important to realize that the pointer to the variable or anything it points at can only be used until the next call to execute an IDL statement. After that, the pointer and the contents of the referenced **IDL_VARIABLE** may become invalid as a result of IDL’s execution.

40

Run the simple IDL widget program contained in the array C string array **cmds2** defined on lines 18–22.

41

Shut down IDL. The **IDL_FALSE** argument instructs **IDL_Cleanup()** to exit the process, so this call should not return.

44

This line should never be reached. If it is, return the UNIX failing status.

Example: Calling an IDL Math Function

This example demonstrates how to write a simple C wrapper function that allows calling IDL commands simply from another language. We implement a function named **call_idl_fft()** that calls the IDL FFT function operating on data imported from our C program. It returns **TRUE** on success, **FALSE** for failure:

```
int call_idl_fft(IDL_COMPLEX *data, int n, int direction);
```

data

A pointer to a linear array of complex data to be processed.

n

The number of data points contained in the array data.

dir

The direction of the FFT transform to take. Specify **-1** for a forward transform, **1** for the reverse

The program is shown in the following figure. Each line is numbered to make discussion easier. These numbers are not part of the actual program.

```

1  #include <stdio.h>
2  #include "idl_export.h"
3
4
5  int call_idl_fft(IDL_COMPLEX *data, IDL_MEMINT n, int dir)
6  {
7      int r;
8      IDL_MEMINT dim[IDL_MAX_ARRAY_DIM];
9      char buf[64];
10
11     dim[0] = n;
12     if (IDL_ImportNamedArray("TMP_FFT_DATA", 1, dim,
13         IDL_TYP_COMPLEX, (UCHAR *) data, 0, 0)) {
14         (void) IDL_ExecuteStr("MESSAGE, /RESET");
15         sprintf(buf, "TMP_FFT_DATA=FFT(TMP_FFT_DATA, /OVERWRITE) "
16             ,dir);
17         r = !IDL_ExecuteStr(buf);
18         (void) IDL_ExecuteStr("TMP_FFT_DATA=0");
19     } else {
20         r = FALSE;
21     }
22
23     return r;
24 }
25
26 main(int argc, char **argv)
27 {
28     #define NUM_PNTS 10
29     IDL_COMPLEX data[NUM_PNTS];
30     IDL_INIT_DATA init_data;
31     int i;
32
33     for (i = 0; i < NUM_PNTS; i++) data[i].r = data[i].i = i;
34     init_data.options = IDL_INIT_CLARGS;
35     init_data.clargs argc = argc;
36     init_data.clargs argv = argv;
37     if (IDL_Initialize(&init_data)) {
38         call_idl_fft(data, NUM_PNTS, -1);
39         call_idl_fft(data, NUM_PNTS, 1);
40         for (i = 0; i < NUM_PNTS; i++)
41             printf("(%f, %f)\n", data[i].r, data[i].i);
42         IDL_Cleanup(IDL_FALSE);
43     }
44
45     return 1;
46 }

```

C

Table 16-4: *call_idl_fft()*

Following is commentary on the above program, by line number:

7

The variable **r** holds the result from the function.

8

dim is used to import the data into IDL as an array.

9

A temporary buffer to format the IDL FFT command.

11–13

Import data into IDL as the variable **TMP_FFT_DATA**. We don't set up a **free_callback** because we will explicitly force IDL to release the pointer after the call to FFT.

14

Set the **!ERROR_STATE** system variable back to the “success” state so previous errors don't confuse our results.

15–16

Format an FFT command to IDL into **buf**. Note the use of the **OVERWRITE** keyword. This tells the IDL FFT function to place the results into the input variable rather than creating a separate output variable. Hence, the results end up in our data array without the need to obtain a pointer to the results and copy them out.

17

Have IDL execute the FFT statement. **IDL_ExecuteStr()** returns the value of **!ERROR_STATE.CODE**, which should be zero for success and non-zero in case of error. Hence, negating the result of **IDL_ExecuteStr()** yields the status value we require for the result of this function.

18

Set **TMP_FFT_DATA** to 0 within IDL. This causes IDL to release the data pointer imported previously.

20

If the call to `IDL_ImportNamedArray()` fails, we must report failure.

26

In order to test the `call_idl_fft()` function, this main program calls it twice. Taking numerical error into account the end result should be equal to the original data.

33

Set the real and imaginary part of each element to the index value.

34-36

Prepare initialization data.

37

Initialize Callable IDL.

38

Call `call_idl_fft()` to perform a forward transform.

39

Call `call_idl_fft()` to perform a reverse transform.

40-41

Print the results.

42

Shut down IDL and exit the process.

45

This line should never be reached. If it is, return the UNIX failing status.

Example: Calling IDL from Fortran

The program shown in the following figure (CALLTEST, found in the callable subdirectory of the external subdirectory of the IDL distribution) demonstrates how to import data from a Fortran program into IDL, execute IDL statements, and obtain data from IDL variables. See “[Compilation and Linking Statements](#)” on page 357 for details on compiling and linking this program. The source code for this file can be found in the file `calltest.f`, located in the callable subdirectory of the external subdirectory of the IDL distribution.

Each line is numbered to make discussion easier. The line numbers are not part of the actual program:

1	C-----
2	C Routine to print a floating point value from an IDL variable.
3	
4	SUBROUTINE PRINT_FLOAT(VPTR)
5	
6	C Declare a Fortran Record type that has a compatible form with
7	C the IDL C struct IDL_VARIABLE for a floating point value.
8	C Note this structure contains a union which is the size of
9	C the largest data type. This structure has been padded to
10	C support the union. Fortran records are not part of
11	C F77, but most compilers have this option.
12	
13	STRUCTURE /IDL_VARIABLE/
14	CHARACTER*1 TYPE
15	CHARACTER*1 FLAGS
16	INTEGER*4 PAD !Pad for largest data type
17	REAL*4 VALUE_F
18	END STRUCTURE
19	
20	RECORD /IDL_VARIABLE/ VPTR
21	
22	WRITE(*, 10) VPTR.VALUE_F
23	10 FORMAT('Program total is: ', F6.2)
24	
25	RETURN
26	
27	END
28	

Table 16-5: Calling IDL from Fortran On UNIX

```

29 C-----
30 C   This function will be called when IDL is finished with the
31 C   array F.
32
33         SUBROUTINE FREE_CALLBACK(ADDR)
34
35             INTEGER*4 ADDR
36
37             WRITE(*,20) LOC(ADDR)
38 20      FORMAT ('IDL Released:', I12)
39
40             RETURN
41
42         END
43
44 C-----
45 C   This program demonstrates how to import data from a Fortran
46 C   program into IDL, execute IDL statements and obtain data
47 C   from IDL variables.
48
49
50 PROGRAM CALLTEST
51
52 C   Some Fortran compilers require external defs. for IDL routines:
53     EXTERNAL IDL_Init !$pragma C(IDL_Init)
54     EXTERNAL IDL_Cleanup !$pragma C(IDL_Cleanup)
55     EXTERNAL IDL_Execute !$pragma C(IDL_Execute)
56     EXTERNAL IDL_ExecuteStr !$pragma C(IDL_ExecuteStr)
57     EXTERNAL IDL_ImportNamedArray !$pragma C(IDL_ImportNamedArray)
58     EXTERNAL IDL_FindNamedVariable !$pragma C(IDL_FindNamedVariable)
59
60 C   Define arguments for IDL_Init routine
61     INTEGER*4 ARGV(1)
62     DATA ARGV(1) /2 * 0/
63
64

```

Table 16-5: Calling IDL from Fortran On UNIX (Continued)


```

65 C Define IDL Definitions for IDL_ImportNamedArray
66
67     PARAMETER (IDL_MAX_ARRAY_DIM = 8)
68     PARAMETER (IDL_TYP_FLOAT = 4)
69
70     REAL*4 F(10)
71     INTEGER*4 DIM(IDL_MAX_ARRAY_DIM)
72     DATA DIM /10, 7*0/
73     INTEGER*4 FUNC_PTR      !Address of function
74     INTEGER*4 VAR_PTR      !Address of IDL variable
75     EXTERNAL FREE_CALLBACK !Declare ext routine for use as arg
76
77     PARAMETER (MAXLEN=80)
78     PARAMETER (N=10)
79
80 C Define commands to be executed by IDL
81
82     CHARACTER*(MAXLEN) CMDS(3)
83     DATA CMDS /"tmp2 = total(tmp)",
84     &          "print, 'IDL total is ', tmp2",
f77 85     &          "plot, tmp"/
86     INTEGER*4 CMD_ARGV(10)
87
88 C Define widget commands to be executed by IDL
89
90     CHARACTER*(MAXLEN) WIDGET_CMDS(5)
91     DATA WIDGET_CMDS /"a = widget_base()",
92     & "b = widget_button(a, val='Press When Done', xs=300, ys=200)",
93     & "widget_control, /realize, a",
94     & "dummy = widget_event(a)",
95     & "widget_control, /destroy, a"/
96
97     INTEGER*4 ISTAT
98
99 C Null Terminate command strings and store the address
100 C for each command string in CMD_ARGV
101
102     DO I = 1, 3
103         CMDS(I) (MAXLEN:MAXLEN) = CHAR(0)
104         CMD_ARGV(I) = LOC(CMDS(I))
105     ENDDO
106

```

Table 16-5: Calling IDL from Fortran On UNIX (Continued)

f77

```

107 C Initialize floating point array, equivalent to IDL FINDGEN(10)
108
109     DO I = 1, N
110         F(I) = FLOAT(I-1)
111     ENDDO
112
113 C Print address of F
114
115 WRITE(*,30) LOC(F)
116     30 FORMAT('ARRAY ADDRESS:', I12)
117
118 C Initialize Callable IDL
119
120     ISTAT = IDL_Init(%VAL(0), ARGV, ARGV(1))
121
122     IF (ISTAT .EQ. 1) THEN
123
124 C Import the floating point array into IDL as a variable named TMP
125
126     CALL IDL_ImportNamedArray('TMP'//CHAR(0), %VAL(1), DIM,
127 &         %VAL(IDL_TYP_FLOAT), F, FREE_CALLBACK, %VAL(0))
128
129 C Have IDL print the value of tmp
130
131     CALL IDL_ExecuteStr('print, tmp'//CHAR(0))
132
133 C Execute a short sequence of IDL statements from a string array
134
135     CALL IDL_Execute(%VAL(3), CMD_ARGV)
136
137 C Set tmp to zero, causing IDL to release the pointer to the
138 C floating point array.
139
140     CALL IDL_ExecuteStr('tmp = 0'//CHAR(0))
141
142 C Obtain the address of the IDL variable containing the
143 C the floating point data
144
145     VAR_PTR = IDL_FindNamedVariable('tmp2'//CHAR(0), %VAL(0))
146
147 C Call a Fortran routine to print the value of the IDL tmp2 variable
148     CALL PRINT_FLOAT(%VAL(VAR_PTR))
149

```

Table 16-5: Calling IDL from Fortran On UNIX (Continued)

```

150 C Null Terminate command strings and store the address
151 C for each command string in CMD_ARGV
152
153     DO I = 1, 5
154         WIDGET_CMDS(I)(MAXLEN:MAXLEN) = CHAR(0)
155         CMD_ARGV(I) = LOC(WIDGET_CMDS(I))
156     ENDDO
157
f77 158 C Execute a small widget program. Pressing the button allows
159 C the program to end
160
161     CALL IDL_Execute(%VAL(5), CMD_ARGV)
162
163 C Shut down IDL
164     CALL IDL_Cleanup(%VAL(0))
165
166     ENDDIF
167
168     END

```

Table 16-5: Calling IDL from Fortran On UNIX (Continued)

1-27

In order to print variables returned from IDL, we must define a Fortran structure type for **IDL_VARIABLE**. This subroutine creates the **IDL_VARIABLE** structure and defines a way to print the floating-point value returned in the an IDL variable.

14-17

Define a Fortran structure equivalent to the floating-point portion of the C **IDL_VARIABLE** structure. Since we know our value is a floating-point number, only the floating-point portion of the structure is implemented. The structure is padded for the largest data type contained in the union. With some Fortran compilers, the combination of **UNION** and **MAP** can be used to implement the **ALLTYPES** union portion of the **IDL_VARIABLE** structure.

29-42

This subroutine is called when IDL releases the user-supplied memory.

44-164

This is the main Fortran program.

51-57

External definitions for IDL internal routines. These definitions may not be necessary with some Fortran compilers.

59-62

Define the **argc** and **argv** arguments required by **IDL_Init()**.

66-67

Define constants equivalent to C IDL constants for the maximum array dimensions and type **float**.

69-77

Define parameters necessary for **IDL_ImportNamedArray()**.

79-85

Define an array of IDL commands to be executed.

87-96

Define an array of IDL widget commands to be executed.

98-104

Null-terminate each of the command strings and store the address of each command to pass to IDL.

106-110

Initialize the floating-point array. This is the Fortran equivalent to the IDL command `F=FINDGEN(10)`.

117-121

Initialize IDL.

125-126

Import the Fortran array **F** in the IDL as a 10-element FLTARR vector named **TMP**. Note the use of the callback argument **FREE_CALLBACK()**, which will be called when IDL is finished with the array **F**, giving us a chance to clean up at that time.

134

Execute the commands contained in the character array **CMDS** defined on lines 71-77. The address for each command is stored in the corresponding array element of **CMD_ARGV**.

139

Set the **TMP** variable to a new value. This causes IDL to release the user-supplied memory and call **FREE_CALLBACK()**.

144

Get a reference to the IDL variable **TMP2**.

147

Call the routine **PRINT_FLOAT** to print the value of **TMP2**. This should agree with the value printed by line 130. Note that the address of the IDL variable **TMP2**, and its contents, can only be used until the next call to execute an IDL statement, since IDL may change the value of the referenced **IDL_VARIABLE**.

150-161

Execute the commands contained in the character array **WIDGET_CMDS** defined on lines 79-88.

163-168

Shut down IDL. The 0 argument instructs **IDL_CLEANUP()** to exit the process, so this call should not return.

Compilation and Linking Statements

Compilation and linking procedures used when calling IDL on a UNIX system are described in the file `calltest_unix.txt` in the `callable` subdirectory of the `external` subdirectory of the main IDL directory. Note that different UNIX systems have different compilation and link statements. Note also that the name of the entry point in the object may be different than that shown here, because compilers may add leading or trailing underscores to the name of the source routine.

Note

The `Makefile` in the architecture-specific subdirectory of the `bin` subdirectory of the IDL distribution contains a make rule for building the `calltest` application.

Issues and Examples: Microsoft Windows

Building an Application that Calls IDL

To build your Microsoft Windows application that calls IDL, you must take the following steps:

1. Use a `#include` line to include the declarations from `idl_export.h` into your source code. This include file is found in the `external/include` subdirectory of the IDL distribution.
2. Compile your application.
3. Link your application with `IDL.LIB`.
4. Place `IDL.DLL` in a directory with your application. See the `readme.txt` file located in the `IDL_DIR/external/callable` for more information.

Example: A Simple Application

The following program demonstrates how to display message text sent from IDL, execute IDL statements entered by a user, and how to obtain data from IDL variables. It performs the following actions:

1. Creates a Main window with four client controls; a scrolling edit control to display text messages from IDL, a single line edit control to allow a user to enter an IDL command, a **Send** button to send the user command to IDL, and a **Quit** button to exit the application.
2. Registers a callback function to handle text messages sent by IDL to the application.
3. Initializes Callable IDL.
4. Call `IDL_Cleanup()` when we receive the `WM_CLOSE` message.

Each line is numbered to make discussion easier. These numbers are not part of the actual program. The source code for this program can be found in the file `simple.c`, located in the `callable` subdirectory of the `external` subdirectory of the IDL distribution. See the source code for details of the program not printed here.

```

1  /*-----
2  * simple.c Source code for sample IDL callable application
3  *
4  * Copyright (c) 1992-1995, ITT Visual Information Solutions
9  *-----
*/
10 #include <windows.h>
11 #include <windowsx.h>
12 #include <ctl3d.h>
13 #include <string.h>
14 #include <stdio.h>
15 #include "simple.h"
16 #include "idl_export.h"
17
18 /*-----
19  * WinMain
20  *
21  * This is the required entry point for all windows
22  * applications.
23  * RETURNS:      TRUE if successful
24  *-----*/
25 int WINAPI WinMain(HINSTANCE hInstance, HINSTANCE
hInstancePrev,
26     LPSTR lpszCmdline, int nCmdShow)
27 {
28     IDL_INIT_DATA    init_data;
29     HWND             hwnd;
30     MSG              msg;
31
32     // Register the main window class.
33     if (!RegisterWinClass(hInstance)) {
34         return(0);
35     }
36
37     ...
38
39     // Create and display the main window.
40     if ((hwnd = InitMainWindow(hInstance)) == NULL) {
41         return(0);
42     }
43     MainhWnd = hwnd;
44
45     // Register our output function with IDL.
46     IDL_ToutPush(OutFunc);
47
48     // Initialize IDL
49     init_data.options = IDL_INIT_BACKGROUND;
50     init_data.options |= IDL_INIT_HWND;

```

```

51     init_data.hwnd = hwnd;
52     if (!IDL_Initialize(&init_data))
53         return(FALSE);
54
55     // Main message loop.
56     while (GetMessage(&msg, NULL, 0, 0)) {
57         TranslateMessage(&msg);
58         DispatchMessage(&msg);
59     }
60
61     return(msg.wParam);
62 }
63
64
65 /*-----
66 * RegisterWinClass
67 *
68 * To create a Main window (TLB in IDL speak). You must first
69 * register the class for that window
70 *
71 * RETURNS:      TRUE if successful
72 *-----*/
73 BOOL RegisterWinClass(HINSTANCE hInst)
74 {
75     WNDCLASS          wc;
76
77     wc.style           = CS_HREDRAW | CS_VREDRAW;
78     wc.lpfnWndProc    = MainWndProc;
79     wc.cbClsExtra     = 0;
80     wc.cbWndExtra     = 0;
81     wc.hInstance      = hInst;
82     wc.hIcon          = NULL;
83     wc.hCursor        = LoadCursor(NULL, IDC_ARROW);
84     wc.hbrBackground  = (HBRUSH)(COLOR_BTNFACE + 1);
85     wc.lpszMenuName   = NULL;
86     wc.lpszClassName  = "Simple";
87
88     if (!RegisterClass(&wc)) {
89         return(FALSE);
90     }
91
92     return(TRUE);
93 }
94
95 /*-----
96 * InitMainWindow
97 *
98 * This is where our Main window is created and displayed
99 *

```



```

100 * RETURNS:          Handle to window
101 *-----*/
102 HWND InitMainWindow(HINSTANCE hInst)
103 {
104     HWND          hwnd;
105     CREATESTRUCT  cs;
106
107
108     hwnd = CreateWindow("Simple",
109         "Callable IDL Sample Application",
110         WS_DLGFRAME | WS_SYSMENU | WS_MINIMIZEBOX | WS_VISIBLE,
111         CW_USEDEFAULT,
112         0,
113         600,
114         480,
115         NULL,
116         NULL,
117         hInst,
118         &cs);
119
120     if (hwnd) {
121         ShowWindow(hwnd, SW_SHOWNORMAL);
122         UpdateWindow(hwnd);
123     }
124
125     return(hwnd);
126 }
127
128 /*-----*/
129 * MainWndProc
130 *
131 * The window procedure (event handler) for our main window.
132 * All messages (events) sent to our app are routed through
133 * here
134 * RETURNS:          Depends of message.
135 *-----*/
136 LRESULT WINAPI MainWndProc(HWND hwnd, UINT uMsg, WPARAM wParam, LPARAM lParam)
137 {
138     static int          nDisplayable = 0;
139
140
141     switch (uMsg) {
142         //When our app is first created, we are sent this message.
143         //We take this opportunity to create our child controls and
144         //place them in their desired locations on the window.
145         case WM_CREATE:
146             if (!CreateControls(((LPCREATESTRUCT)lParam)->hInstance, hwnd)) {
147                 return(0);
148             }

```

```

149     if (!LayoutControls(hwnd)) {
150         return(0);
151     }
152     nDisplayable = GetCharacterHeight(GetDlgItem(hwnd, IDE_COMMANDLOG));
153     break;
154
155     ...
156
157     case WM_DESTROY:
158         PostQuitMessage(1);
159         break;
160
161     //Each time a button or menu item is selected, we get this message
162     case WM_COMMAND:
163         OnCommand(hwnd, LOWORD(wParam), wParam, lParam);
164         return(FALSE);
165
166     //This is a message we send ourselves to indicate the need to
167     //display a text message in our log window.
168     case IDL_OUTPUT:
169         OutputMessage(wParam, lParam, nDisplayable);
170         return(FALSE);
171
172     case WM_CLOSE:
173         IDL_Cleanup(TRUE);
174         return(FALSE);
175
176     default:
177         break;
178 }
179
180 return(DefWindowProc(hwnd, uMsg, wParam, lParam));
181 }
182
183 /*-----*/
184 * OnCommand
185 *
186 * This is the message handle for our WM_COMMAND messages
187 *
188 * RETURNS:          FALSE
189 *-----*/
190 BOOL OnCommand(HWND hWnd, UINT uId, WPARAM wParam, LPARAM lParam)
191 {
192
193     switch(uId){
194         case IDB_SENDCOMMAND:{
195             LPSTR    lpCommand;
196             LPSTR    lpOut;
197

```

```

198         lpCommand = GlobalAllocPtr(GHND, 256);
199         lpOut = GlobalAllocPtr(GHND, 256);
200         if(!lpCommand)
201             return(FALSE);
202
203         /* First we get the string that is in the input window */
204         GetDlgItemText(hWnd, IDE_COMMANDLINE, lpCommand,
255);
205
206         /* and then clear the window */
207         SetDlgItemText(hWnd, IDE_COMMANDLINE, "");
208
209         lstrcpy(lpOut, "\r\nSent to IDL: ");
210         lstrcat(lpOut, lpCommand);
211
212         /* Send the string to our "log" window */
213         OutFunc(IDL_TOUT_F_NLPOST, lpOut, strlen(lpOut));
214
215         /* then send the string to IDL */
216         IDL_ExecuteStr(lpCommand);
217
218         /* Now clean up */
219         GlobalFreePtr(lpCommand);
220         GlobalFreePtr(lpOut);
221     }
222     break;
223 }
224 return(FALSE);
225 }
226
227 /*-----
228 * OutFunc
229 *
230 * This is the output function that receives messages from IDL
231 * and displays them for the user
232 *
233 * RETURNS:          NONE
234 *-----*/
235 void OutFunc(long flags, char *buf, long n)
236 {
237     static    fShowMain = FALSE;
238
239     /* If there is a message, post it to our MAIN window */
240     if (n){
241         SendMessage (MainhWnd, IDL_OUTPUT, 0, (LPARAM)buf);
242     }
243
244     /* If we need to post a new line message... */
245     if (flags & IDL_TOUT_F_NLPOST){

```

```

246     SendMessage (MainhWnd, IDL_OUTPUT, 0, (LPARAM) (LPSTR) "\r\n\0");
247 }
248
249 /* This message gets sent to the log window to have it scroll
250    and display the last message at the bottom of the window.
251    With this, the user will always see the last screen full of
252    messages sent
253    */
254     SendMessage (MainhWnd, IDL_OUTPUT, (WPARAM) TRUE,
255                 (LPARAM) (LPSTR) "\0");
256
257     return;
258 }
259
260 /*-----
261  * OutputMessage
262  *
263  * Here we do the actual display of the text to our log window
264  *
265  * RETURNS:          nothing
266  *
267  *-----*/
268 void OutputMessage(WPARAM wParam, LPARAM lParam, int nDisplayable)
269 {
270     LRESULT    lRet;
271     LONG       lBufflen, lNumLines, lFirstView;
272
273     /* Turn off the READONLY bit and postpone redraw */
274     lRet = SendMessage (hwndLog, EM_SETREADONLY, FALSE, 0L);
275     lRet = SendMessage (hwndLog, WM_SETREDRAW, FALSE, 0L);
276
277     /* Get the length of the text in the log window*/
278     lBufflen = SendMessage (hwndLog, WM_GETTEXTLENGTH, 0, 0L);
279     lNumLines = SendMessage (hwndLog, EM_GETLINECOUNT, 0, 0L);
280     lFirstView = SendMessage (hwndLog, EM_GETFIRSTVISIBLELINE, 0, 0L);
281     lRet = SendMessage (hwndLog, EM_SETSEL, lBufflen, lBufflen);
282
283     /* If we are adding text, wParam will be 0 */
284     if (!wParam) {
285         lRet = SendMessage (hwndLog, EM_REPLACESEL, 0, lParam);
286     } else {
287         if (lNumLines > (lFirstView + nDisplayable)) {
288             int    iLineLen = 0;
289             int    iChar;
290             int    iLines = 0;
291             lNumLines--;
292             while (!iLineLen) {
293                 iChar = SendMessage (hwndLog, EM_LINEINDEX,
294                                     (WPARAM) lNumLines, 0L);

```

```

295             iLineLen = SendMessage(hwndLog, EM_LINELENGTH,
296                 iChar, 0L);
297             if(!iLineLen)
298                 lNumLines--;
299         }
300         iLines = lNumLines-(lFirstView + (nDisplayable - 1));
301         iLines = iLines >= 0 ? iLines : 0;
302         SendMessage (hwndLog, EM_LINESCROLL, 0, (LPARAM)iLines);
303     }
304 }
305
306 /* Set the window to redraw and reset the READONLY bit */
307 lRet = SendMessage(hwndLog, WM_SETREDRAW, TRUE, 0L);
308 lRet = SendMessage(hwndLog, EM_SETREADONLY, TRUE, 0L);
309
310 return;
311 }

```

The following is a commentary on the program, by line number:

16

`idl_export.h` contains the **IDL_** function prototypes, IDL specific structures, and IDL constants.

46

Call **IDL_ToutPush()** with the address of the output function (**OutFunc**) as it's only argument. This will register **OutFunc** as a callback for IDL. IDL will call **OutFunc** when it needs to display text.

49–53

Initialize IDL as a non-interactive session and supply it with the handle to the main window.

56

Start the windows message loop.

136-181

This is the Main window procedure. It will handle any messages that are sent to the main window. This includes **WM_COMMAND** messages that occur as a result of user interaction with the client controls. In addition, it handles a user defined message called **IDL_OUTPUT** (the name doesn't matter but this is a clue as to its purpose).

163

When the user presses either the “Send” or “Quit” buttons, route the message to the **OnCommand** function.

169

When we receive an **IDL_OUTPUT** message, call the function that displays text in the scrolling window (**OutputMessage**. See line 263).

173

When we receive the **WM_CLOSE** message, call **IDL_Cleanup()** to unlink IDL from our application.

190-225

OnCommand handles the **WM_COMMAND** messages generated when the user clicks on the application’s buttons.

204

Get the IDL command that the user has entered in the single line edit control and store it in a buffer.

207

Clear the text in the edit control.

213

Display the command sent to IDL in the output window.

216

Call **IDL_ExecuteStr()** with the IDL command retrieved in line 204.

235-258

OutFunc is the callback registered with IDL to handle text messages IDL sends to our application. In addition it will handle text from IDL routines that display information, such as **PRINT**.

268-311

OutputMessage handles displaying the text to the output window. Since this window is a multi-line edit control, we have created it as a read-only window. See the source code for additional information on handling this situation.

285

OutputMessage appends new messages to the existing text in the control.

286-304

When the text has been displayed, **OutputMessage** scrolls the window to display the last line of text in the bottom of the window.



Chapter 17

Adding External Widgets to IDL

This chapter discusses the following topics:

IDL and External Widgets	370	Functions for Use with Stub Widgets	374
WIDGET_STUB	371	Internal Callback Functions	377
WIDGET_CONTROL/WIDGET_STUB .	372	UNIX WIDGET_STUB Example: WIDGET_ARROWB	379

IDL and External Widgets

This chapter describes an IDL widget type not documented in the *IDL Reference Guide*, called the *stub widget*. It also describes a small set of internal functions to manipulate stub widgets. Stub widgets allow `CALL_EXTERNAL`, `LINKIMAGE`, `DLM`, and Callable IDL users to add their own widgets to IDL widget hierarchies.

This feature depends on your system providing the window system libraries used by IDL (particularly the Motif libraries under UNIX) as sharable libraries. It will not work with versions of IDL that statically link against the window system libraries.

The next two sections describe IDL's `WIDGET_STUB` function and changes to `WIDGET_CONTROL` when used with `WIDGET_STUB`. “[Functions for Use with Stub Widgets](#)” on page 374 describes support functions that can be called from your external code to manipulate stub widgets. “[Internal Callback Functions](#)” on page 377 describes how to make stub widgets generate IDL widget events. Finally, “[UNIX WIDGET_STUB Example: WIDGET_ARROWB](#)” on page 379 illustrates the use of stub widgets with an external program.

Note

Although `WIDGET_STUB` can be used under Microsoft Windows, this feature is primarily of interest with UNIX IDL. Under Windows, we recommend the use of the `WIDGET_ACTIVEX` functionality, which allows you to use ActiveX controls with IDL without requiring external programming.

WIDGET_STUB

The `WIDGET_STUB` function creates a widget record that contains no actual underlying widgets. Stub widgets are place holders for integrating external widget types into IDL. Events from those widgets can then be processed in a manner consistent with the rest of the IDL widget system.

First, the programmer calls `WIDGET_STUB` to create the widget, and then uses `CALL_EXTERNAL` to call additional custom code to handle the rest. A number of internal functions are provided to manipulate widgets from this custom code. See [“Functions for Use with Stub Widgets”](#) on page 374.

The returned value of this function is the widget ID of the newly-created stub widget.

Calling Sequence

Result = `WIDGET_STUB(Parent)`

Arguments

Parent

The widget ID of the parent widget. Stub widgets can only have bases or other stub widgets as their parents.

Keywords

The following keywords are accepted by `WIDGET_STUB` and work the same as for other widget creation functions:

<code>EVENT_FUNC</code>	<code>SCR_XSIZE</code>
<code>EVENT_PRO</code>	<code>SCR_YSIZE</code>
<code>FUNC_GET_VALUE</code>	<code>UVALUE</code>
<code>GROUP_LEADER</code>	<code>XOFFSET</code>
<code>KILL_NOTIFY</code>	<code>XSIZE</code>
<code>NO_COPY</code>	<code>YOFFSET</code>
<code>PRO_SET_VALUE</code>	<code>YSIZE</code>

WIDGET_CONTROL/WIDGET_STUB

The WIDGET_CONTROL procedure has some differences and limitations when used with WIDGET_STUB that are not documented in the *IDL Reference Guide*. These differences are described below.

Keywords

Only the most general keywords are allowed with WIDGET_CONTROL when used with stub widgets. All other keywords are ignored. Here is a list of those keywords that behave identically with all widgets including stub widgets:

BAD_ID	PRO_SET_VALUE
CLEAR_EVENTS	RESET
EVENT_FUNC	SET_UVALUE
EVENT_PRO	SHOW
FUNC_GET_VALUE	TIMER
GET_UVALUE	TLB_GET_OFFSET
GROUP_LEADER	TLB_GET_SIZE
HOURGLASS	TLB_SET_TITLE
ICONIFY	TLB_SET_XOFFSET
KILL_NOTIFY	TLB_SET_YOFFSET
MANAGED	XOFFSET
NO_COPY	YOFFSET

The following keywords also work with stub widgets, but require additional commentary:

DESTROY

When a widget hierarchy containing stub widgets is destroyed, the following steps are taken:

- The lower-level code that deals with the system toolkit destroys any real widgets currently used by the stub widgets.
- All IDL widget records are added to the free list for re-use.

- Any requested `KILL_NOTIFY` callbacks are called.

You should register `KILL_NOTIFY` callbacks on the topmost stub widget in each widget subtree. Remember that the actual widgets are gone before the callbacks are issued, so don't attempt to access them. However, the callback provides an opportunity to clean up any related resources used by the widget.

MAP, REALIZE, and SENSITIVE

These keywords cause the toolkit-specific, lower layer of the IDL widgets implementation to be called. In the process of satisfying the specified request, any real widgets used by the stub widgets will be processed, along with the ones created by the non-stub widgets, in the usual way. Any additional processing must be provided via `CALL_EXTERNAL`.

XSIZE, SCR_XSIZE, YSIZE, and SCR_YSIZE

These keywords inform IDL how large the stub widget is expected to be. This information is necessary for IDL to calculate sizes and offsets of the surrounding widgets.

IDL tries to do something reasonable with these requests but, without knowledge of the actual widget being manipulated, it is possible that the results will not be satisfactory. In such cases, the `IDL_WidgetStubSetSizeFunc()` function can be used to specify a routine that IDL can call to perform the necessary sizing for your stub widget.

Functions for Use with Stub Widgets

The following functions present a highly simplified interface to the stub widget class that gives the user enough access to IDL widget internals to make the stub widget work while hiding the details of the actual implementation.

IDL_WidgetStubLock()

Syntax:

```
void IDL_WidgetStubLock(int set);
```

IDL event processing occurs asynchronously, so any code that manipulates widgets *must* execute in a protected region. This function is used to create such a region. Any code that manipulates widgets must be surrounded by two calls to **IDL_WidgetStubLock()** as follows:

```
IDL_WidgetStubLock(TRUE);  
/* Do your widget stuff */  
IDL_WidgetStubLock(FALSE);
```

IDL_WidgetStubLookup()

Syntax:

```
char *IDL_WidgetStubLookup(IDL_ULONG id);
```

When IDL creates a widget, it returns an integer value to the caller of the widget creation function. Internally, however, IDL widgets are represented by a pointer to memory. The **IDL_WidgetStubLookup()** function is used to translate the user-level integer value to this memory pointer. All the other internal routines use the memory pointer to reference the widget.

Id is the integer returned at the user level. Your call to **CALL_EXTERNAL** should pass this integer to your C-level code for use with **IDL_WidgetStubLookup()** which translates the integer to the pointer.

If the specified **id** does not represent a valid IDL widget, this function returns **NULL**. This situation can occur if a widget was killed but its integer handle is still lingering somewhere.

IDL_WidgetIssueStubEvent()

Syntax:

```
void IDL_WidgetIssueStubEvent(char *rec, LONG value);
```

Given a handle to the IDL widget, obtained via **IDL_WidgetStubLookup()**, this function queues an IDL **WIDGET_STUB_EVENT**. Such an event is a structure that contains the three standard fields (ID, TOP, and HANDLER) as well as an additional field named VALUE that contains the specified **value**.

VALUE can provide a way to access additional information about the widget, possibly by providing a memory address to the information.

IDL_WidgetSetStubIds()

Syntax:

```
void IDL_WidgetSetStubIds(char *rec, unsigned long t_id,  
                          unsigned long b_id);
```

IDL widgets are built out of one or more actual widgets. Every IDL widget carries two pointers that are used to locate the top and bottom real widget for a given IDL widget. This function allows you to set these top and bottom pointers in the stub widget for later use.

Since the actual pointer type differs from toolkit to toolkit, this function declares **t_id** (the top real widget) and **b_id** (the bottom real widget) as unsigned long, an integer data type large enough to safely contain any pointer. Use a C cast operator to handle the difference.

After calling **WIDGET_STUB** to create an IDL stub widget, you will need to use **CALL_EXTERNAL** to call additional code that creates the real widgets that represent the stub. Having done that, use **IDL_WidgetSetStubIds()** to save the top and bottom widget pointers.

IDL_WidgetGetStubIds()

Syntax:

```
void IDL_WidgetGetStubIds(char *rec, unsigned long *t_id,  
                          unsigned long *b_id);
```

This function returns the top (**t_id**) and bottom (**b_id**) real widget pointers for any specified widget (not just stub widgets). When using these values for non-stub widgets, it is the caller's responsibility to avoid damaging the IDL-created widgets in any way.

IDL_WidgetStubSetSizeFunc()

Syntax:

```
void IDL_WidgetStubSetSizeFunc(char *rec,  
                               IDL_WIDGET_STUB_SET_SIZE_FUNC func)  
  
typedef void (* IDL_WIDGET_STUB_SET_SIZE_FUNC);  
             (IDL_ULONG id, int width, int height);
```

When IDL needs to set the size of a stub widget, it attempts to set the size of the bottom real widget to the necessary dimensions. Often, this is the desired behavior, but cases can arise where it would be better to handle sizing differently. In such cases, use **IDL_WidgetStubSetSizeFunc()** to register a function that IDL will call to do the actual sizing.

Internal Callback Functions

Real widget toolkits (upon which IDL widgets are built) are event driven. C language programs register interest in specific events by providing callback functions that are called when that event occurs. All but the most basic of widgets are capable of generating events.

In order for IDL stub widgets to generate IDL events, you must use `CALL_EXTERNAL` to invoke code that sets up real widget event callbacks for the events you are interested in. This setup can be done as part of creating the real widgets after the initial call to `WIDGET_STUB`. These callbacks then call `IDL_WidgetIssueStubEvent()` to issue the IDL event.

Your C-language widget toolkit callback functions should be patterned after the following template. Note that the arguments and return type will depend on the widget toolkit used, and so cannot be shown here:

```
stub_widget_call()
{
    char *idl_widget;
    IDL_WidgetStubLock(TRUE);
    /* Get the IDL user-level identifier for this widget */
    if (idl_widget = IDL_WidgetStubLookup(id)) {
        /* Do whatever work is required */
        ...
        /* Optionally, issue an IDL event */
        IDL_WidgetIssueStubEvent(idl_widget, value)
    }
    IDL_WidgetStubLock(FALSE);
}
```

Commentary on the Example Shown Above

Note that `IDL_WidgetStubLock()` is used to protect the critical section where widgets are being manipulated.

Somehow, the callback must be able to find the user-level integer returned by `WIDGET_STUB` when the stub widget was created in IDL. Usually, this is done in one of two ways:

- When registering the callback, it is sometimes possible to specify a value that will be passed to the callback without interpretation. For example, the X windows `XtAddCallback()` function takes an argument named `client_data`. This value is passed to the callback and can be used to supply the user-level identifier.

- Some widget toolkits have a set of attributes that they carry along with each widget. Under the X windows Xt toolkit, these attributes are called resources. Xt widgets usually have a resource capable of holding a single integer or memory address. This resource can be used to supply the user level identifier.

IDL_WidgetStubLookup() is used to translate the user level widget identifier into a memory pointer. If this function returns NULL, no further event processing is done since it would be a fatal error to issue an IDL event for a non-existent widget.

The event is issued via **IDL_WidgetIssueStubEvent()**. This step is not required. Many of the IDL widget types process real widget events via callbacks that do not always result in an IDL widget event being sent.

UNIX WIDGET_STUB Example: WIDGET_ARROWB

The following example adds the Motif ArrowButton widget to UNIX IDL in the form of an IDL program named `widget_arrowb.pro`.

The primary user interface to our arrow button widget is the `WIDGET_ARROWB` function. It presents an interface much like any of the built in `WIDGET_` functions provided by IDL. `WIDGET_ARROWB` uses the `MAKE_DLL` procedure, and the `AUTO_GLUE` keyword to `CALL_EXTERNAL` to automatically build and load the C code required for this widget. This building and loading process is transparent to the IDL user, requiring only that you have a C compiler installed on your system. All the user has to do to use an arrow button widget is to call `WIDGET_ARROWB`

The `WIDGET_ARROWB` widget acts like a normal pushbutton. Events are sent when the button is pressed (`VALUE=1`) and released (`VALUE=0`). If the `USE_OWN_SIZE` keyword is set to zero, IDL performs its default sizing on the stub widget. A non-zero value causes a special routine provided by the `WIDGET_ARROWB` implementation to be registered to handle such sizing.

All of the code used in this example, including all code shown here, is available in the `external/widstub` directory of the UNIX IDL distribution. To run it, execute the following statements from IDL:

```
PUSHD, FILEPATH('', SUBDIRECTORY=['external','widstub'])
WIDGET_ARROWB_TEST
POPD
```

When running `WIDGET_ARROWB_TEST`, you can specify the `VERBOSE` keyword, in which case, it will show you the compilation and linking steps it takes to build the sharable library from the C code. The use of `pushd` and `popd` are due to the fact that your IDL search path (`!PATH`) is unlikely to have the directory containing these examples in it. `PUSHD` changes your working directory to the location where these files are found, and `POPD` restores it to its original location afterwards.

The IDL Program for WIDGET_ARROWB

The following text is the IDL program for `WIDGET_ARROWB`. It is found in the file named `WIDGET_ARROWB.PRO`:

```
function WIDGET_ARROWB, parent, use_own_size, UVALUE=uvalue, $
    VERBOSE=verbose, _EXTRA=extra
; Uses WIDGET_STUB, and a sharable library containing
; the necessary C support code, to provide the IDL user
```

```

; with a Motif Arrow Button widget. The interface is consistent
; with that presented by the built in IDL widgets.
;
; If the sharable library does not exist, it is built using
; MAKE_DLL.

common WIDGET_ARROWB_BLK, shlib

; Build sharable lib if first call or lib doesn't exist
build_lib = n_elements(shlib) eq 0
if (not build_lib) then build_lib = not FILE_TEST(shlib, /READ)
if (build_lib) then begin
    ; Location of the widget_arrowb files from IDL distribution
    arrowb_dir=FILEPATH(' ',SUBDIRECTORY=['external','widstub'])

    ; Use MAKE_DLL to build the widget_arrowb sharable library
    ; in the !MAKE_DLL.COMPILE_DIRECTORY directory.
    ;
    ; Normally, you wouldn't use VERBOSE, or SHOW_ALL_OUTPUT
    ; once your work is debugged, but as a learning exercise it
    ; can be useful to see all the underlying work that gets
    ; done. If the user specified VERBOSE, then use those
    ; keywords to show what MAKE_DLL is doing.
    MAKE_DLL,'widget_arrowb', 'widget_arrowb', $
        DLL_PATH=shlib, INPUT_DIR=arrowb_dir, $
        VERBOSE=verbose,SHOW_ALL_OUTPUT=verbose
endif

; Use a stub widget along with the C code in the library to
; create an arrow button widget. The use of the AUTO_GLUE
; keyword simplifies the call to the sharable library by
; eliminating the need to use the CALL_EXTERNAL portable
; calling convention.
l_parent=LONG(parent)
l_use_own_size = $
    (n_elements(use_own_size) eq 0) ? 0L: LONG(use_own_size)
result = WIDGET_STUB(parent, _extra=extra)
if (n_elements(uvalue) ne 0) then $
    WIDGET_CONTROL, result, set_uvalue=uvalue
JUNK = CALL_EXTERNAL(shlib, 'widget_arrowb',l_parent,result,$
    l_use_own_size, value=[1, 1, 1], /AUTO_GLUE)

RETURN, result
end

```

The C Program for widget_arrowb.c

The C language code invoked by the call to `CALL_EXTERNAL` in the above IDL code is contained in a file named `widget_arrowb.c`. This file can be found in the `widstub` subdirectory of the `external` subdirectory of the IDL distribution. The contents of this file are shown below:

```

/*
 * widget_arrowb.c - This file contains C code to be called from
 * UNIX IDL via CALL_EXTERNAL. It uses the IDL stub widget to add
 * a Motif ArrowButton to an IDL created widget hierarchy. The
 * button issues a WIDGET_STUB_EVENT every time the button is
 * released.
 *
 * While this code is Motif-centric, the principles apply across
 * platforms and could be adapted to Microsoft Windows.
 */
#include <stdio.h>
#include <X11/keysym.h> /* Keysyms for text widget events */
#include <X11/Intrinsic.h>
#include <X11/StringDefs.h>
#include <X11/Shell.h>
#include <Xm/ArrowB.h>
#include "idl_export.h"

/*ARGSUSED*/
static void arrowb_CB(Widget w, caddr_t client_data,
                    caddr_t call_data)
{
    char *rec;
    XmArrowButtonCallbackStruct *abcs;

    IDL_WidgetStubLock(TRUE);
    if (rec = IDL_WidgetStubLookup((unsigned long) client_data)) {
        abcs = (XmArrowButtonCallbackStruct *) call_data;
        IDL_WidgetIssueStubEvent(rec, abcs->reason == XmCR_ARM);
    }
    IDL_WidgetStubLock(FALSE);
}

static void arrowb_size_func(IDL_ULONG stub, int width,
                            int height)
{
    char *stub_rec;
    unsigned long t_id, b_id;
    char buf[128];

```

```

IDL_WidgetStubLock(TRUE);
if (stub_rec = IDL_WidgetStubLookup(stub)) {
    IDL_WidgetGetStubIds(stub_rec, &t_id, &b_id);
    sprintf(buf, "Setting WIDGET %d to width %d and height %d",
            stub, width, height);
    IDL_Message(IDL_M_NAMED_GENERIC, IDL_MSG_INFO, buf);
    XtVaSetValues((Widget) b_id, XmNwidth, width, XmNheight,
                height, NULL);
}
IDL_WidgetStubLock(FALSE);
}
int widget_arrowb(IDL_LONG parent, IDL_LONG stub, IDL_LONG
                use_own_size_func)
{
    Widget parent_w;
    Widget stub_w;
    char *parent_rec;
    char *stub_rec;
    unsigned long t_id, b_id;

    IDL_WidgetStubLock(TRUE);
    if ((parent_rec = IDL_WidgetStubLookup(parent))
        && (stub_rec = IDL_WidgetStubLookup(stub))) {
        /* Bottom widget of parent is parent to arrow button */
        IDL_WidgetGetStubIds(parent_rec, &t_id, &b_id);
        parent_w = (Widget) b_id;
        stub_w = XtVaCreateManagedWidget("arrowb",
                                         xmArrowButtonWidgetClass,
                                         parent_w, NULL);
        IDL_WidgetSetStubIds(stub_rec, (unsigned long) stub_w,
                              (unsigned long) stub_w);
        XtAddCallback(stub_w, XmNarmCallback,
                      (XtCallbackProc) arrowb_CB, (XtPointer) stub);
        XtAddCallback(stub_w, XmNdisarmCallback,
                      (XtCallbackProc) arrowb_CB, (XtPointer) stub);
        if (use_own_size_func)
            IDL_WidgetStubSetSizeFunc(stub_rec, arrowb_size_func);
    }
    IDL_WidgetStubLock(FALSE);
    return stub;
}

```

An IDL Program to Test the External Widget

Shown below is an IDL widget program to test the ARROWB widget. This program is found in the file `widget_arrowb_test.pro` in the IDL distribution:

```

pro widget_arrowb_test_event, ev
  widget_control, get_uvalue=val, ev.id
  if (val eq 0) then begin
    widget_control, /destroy, ev.top
  endif else begin
    HELP, /STRUCT, ev
    if (ev.value eq 1) then begin
      widget_control, val, set_value='New label string'
      tmp = widget_info(ev.id, /GEOMETRY)
      widget_control, xsize=tmp.xsize+25, $
                    ysize=tmp.ysize+25, ev.id
    endif
  endif
endelse
end

pro widget_arrowb_test, VERBOSE=verbose
  a = widget_base(/COLUMN)
  b = widget_button(a, value='Done', uvalue = 0)
  label=widget_label(a, value='A label')
  arrow_w = widget_arrowb(a, 0, xsize=100, ysize=100, $
                        uvalue=label, verbose=verbose)
  arrow_w = widget_arrowb(a, 1, xsize=100, ysize=50, $
                        uvalue=label, verbose=verbose)
  widget_control, /real, a
  xmanager, 'WIDGET_ARROWB_TEST', a, /NO_BLOCK
end

```




Appendix A

Obsolete Internal Interfaces

This chapter discusses the following topics:

Interfaces Obsoleted in IDL 6.3	386	Simplified Routine Invocation	404
Interfaces Obsoleted in IDL 5.5	388	Obsolete Error Handling API	411
Interfaces Obsoleted in IDL 5.2.1	401		

Interfaces Obsoleted in IDL 6.3

Prior to IDL 6.3, the `IDL_Win32Init()` function was used to initialize IDL in callable IDL applications for the Microsoft Windows environment. It was obsoleted in IDL 6.3, replaced by the `IDL_Initialize()` function that offers the same abilities in addition to being usable on a cross platform basis. New code should be written to use `IDL_Initialize()`.

Initialization: Microsoft Windows

Under Microsoft Windows, the `IDL_Win32Init()` function prepares the IDL DLL for use. `IDL_Win32Init()` must be called before any other function except `IDL_ToutPush()`.

Note

Windows applications should not call `IDL_Init()`, described in the previous section. `IDL_Win32Init()` calls `IDL_Init()` on your behalf at the appropriate time.

```
int IDL_Win32Init(int iOpts, void *hinstExe, void *hwndExe,
                 void *hAccel);
```

where:

iOpts

A bitmask used to specify initialization options. The allowed bit values are:

IDL_INIT_RUNTIME

Setting this bit causes IDL to check out a runtime license instead of the normal license. `IDL_RuntimeExec()` is then used to run an IDL application restored from a Save/Restore file.

IDL_INIT_LMQUEUE

Setting this bit causes IDL to wait for an available license before beginning an IDL task such as batch processing.

hinstExe

HINSTANCE from the application that will be calling IDL.

hwndExe

HWND for the application's main window.

hAccel

Reserved. This argument should always be NULL.

IDL_Win32Init() returns TRUE if the initialization is successful, and FALSE for failure.

Interfaces Obsoleted in IDL 5.5

The following areas changed in IDL 5.5, requiring the introduction of new interfaces, and causing some old interfaces to become obsolete. These old interfaces remain in IDL and can be used by user code. However, new code should not use them, and old code might benefit from migration as part of normal maintenance:

- The **IDL_Message()** **IDL_MSG_ATTR_SYS** attribute has been retired, in favor of the more general **IDL_MessageSyscode()** function.
- The **IDL_MessageErrno()** and **IDL_MessageErrnoFromBlock()** functions have been retired in favor of the **IDL_MessageSyscode()** and **IDL_MessageSyscodeFromBlock()** functions, which are more general.
- IDL's keyword API has been redesigned to be easier to use and understand, and to be reentrant.

IDL_MSG_ATTR_SYS

Note

IDL_MSG_ATTR_SYS is one of the possible attribute values that can be included in the **action** argument to the **IDL_Message()** function. Its purpose was to cause **IDL_Message()** to report the system error currently contained in the process **errno** global variable. This functionality is now available in a more general and useful form via the **IDL_MessageSyscode()** and **IDL_MessageSyscodeFromBlock()** functions, documented in [“Issuing Error Messages”](#) on page 195

IDL_MSG_ATTR_SYS

IDL_Message() always issues a single-line error message that describes the problem from IDL's point of view. Often, however, there is an underlying system reason for the error that should also be displayed to give the user a complete picture of what went wrong. For example, the IDL view of the problem might be “Unable to open file”, while the underlying system reason for the error is “no such directory”.

The UNIX system provides a global variable named **errno** for communicating such system level errors. Whenever a call to a system function fails, it returns a 1, and puts an error code into **errno** that specifies the reason for the failure. Other functions, such as those provided by the standard C library, do not set **errno**. These functions do set **errno**.

Specifying **IDL_MSG_ATTR_SYS** tells **IDL_Message()** to check **errno**, and if it is non-null, to issue a second line containing the text of the system error message.

Specify **IDL_MSG_ATTR_SYS** only if you are calling **IDL_Message()** as the result of a failed UNIX system call. Otherwise, **errno** might contain an unrelated garbage value resulting in an incorrect error message.

The Microsoft Windows operating system has **errno** for compatibility with the expectations of C programmers, but typically do not set it. On these operating systems, it is possible to specify **IDL_MSG_ATTR_SYS**, but it has no effect.

Specifying **errno** Explicitly: **IDL_MessageErrno()**

Note

The **IDL_MessageErrno()** and **IDL_MessageErrnoFromBlock()** functions allow you to throw an error message that includes the system error from the UNIX/POSIX **errno** global variable. These functions have been replaced by **IDL_MessageSyscode()** and **IDL_MessageSyscodeFromBlock()** which in addition to being able to throw UNIX/Posix errors, can also throw other types of system error.

There are times when specifying the **IDL_MSG_ATTR_SYS** modifier code in the action argument to **IDL_Message()** is inadequate. This situation usually occurs when your code attempts to perform some cleanup operation when an operating system call fails before calling **IDL_Message()** and this cleanup code might alter the value of **errno**. In such cases, it is preferable to use the **IDL_MessageErrno()** or **IDL_MessageErrnoFromBlock()** functions to issue the message:

```
void IDL_MessageErrno(int code, int errno, int action, ...)
void IDL_MessageErrnoFromBlock(IDL_MSG_BLOCK block, int code, int
errno, int action, ...)
```

These function differs from **IDL_Message()** in two ways:

1. There is an additional argument used to specify the value of **errno**. See the discussion of **errno** in “**IDL_MSG_ATTR_SYS**” on page 388 for additional information about **errno** and its use.
2. The **IDL_MSG_ATTR_SYS** modifier code for the action argument is ignored.-

Processing Keywords With `IDL_KWGetParams()`

Note

Previous versions of IDL used a keyword API based around the `IDL_KWGetParams()` and `IDL_KWCleanup()` functions. This API was confusing to use (It was difficult to know when `IDL_KWCleanup()` was supposed to be called), and was not reentrant (requiring extensive and error prone code in some IDL system routines). The new API, using `IDL_KWProcessByOffset()` and `IDL_KW_FREE`, solve these problems and result in easier to write and maintain code.

To enable rapid conversion from the old API to the new, the new API uses most of the same data structures as the old (with the notable exception of `IDL_KW_ARR_DESC`, which is replaced by `IDL_KW_ARR_DESC_R`).

This section reproduces those parts of the documentation of the original API that differ from the current API, which is described in [Chapter 6, “IDL Internals: Keyword Processing”](#)

The `IDL_KW_PAR` Structure

Note

`IDL_KW_PAR` is used with the old keyword API in largely the same manner as the current API, as described in [“Overview Of IDL Keyword Processing”](#) on page 124. The main difference is that the contents of the **specified** and **value** fields are the addresses of static variables, rather than offsets into a `KW_RESULT` structure as with the new API.

specified

The address of a C int variable that will be set to TRUE (non-zero) or FALSE (0) based on whether the routine was called with the keyword present. This field should be set to NULL (`(int *) 0`) if this information is not needed.

value

If the keyword is a read-only scalar, this field is a pointer to a C variable of the correct type (`IDL_LONG`, `IDL_ULONG`, `IDL_LONG64`, `IDL_ULONG64`, float, double, or `IDL_STRING`).

In the case of a read-only array, value is a pointer to an **IDL_KW_ARR_DESC**, which is discussed in “[The IDL_KW_ARR_DESC Structure](#)” on page 391. In the case of an output variable (i.e., the **IDL_KW_OUT** flag is set), this field should point to an **IDL_VPTR** that will be filled by **IDL_KWGetParams()** with the address of the keyword argument.

The IDL_KW_ARR_DESC Structure

Note

The **IDL_KW_ARR_DESC** structure was superseded by **IDL_KW_ARR_DESC_R** in the current API. The reason for this change is that the **n** field of **IDL_KW_ARR_DESC** is modified by the call to **IDL_KWGetParams()**, requiring the **IDL_KW_ARR_DESC** structure to be defined in static memory, and rendering it non-reentrant.

When a keyword is specified to be a read-only array (i.e., the **IDL_KW_ARRAY** flag is set), the value field of the **IDL_KW_PAR** struct should be set to point to an **IDL_KW_ARR_DESC** structure. This structure is defined as:

```
typedef struct {
    char *data;
    IDL_MEMINT nmin;
    IDL_MEMINT nmax;
    IDL_MEMINT n;
} IDL_KW_ARR_DESC;
```

where:

data

The address of a C array to receive the data. This array must be of the C type mapped into by the **type** field of the **IDL_KW_PAR** struct. For example, **IDL_TYP_LONG** maps into a C **IDL_LONG**. There must be **nmax** elements in the array.

nmin

The minimum number of elements allowed.

nmax

The maximum number of elements allowed.

n

The number of elements actually present. Unlike the other fields, this field is set by **IDL_KWGetParams()**.

Processing Keywords

The **IDL_KWGetParams()** function is used to process keywords.

IDL_KWGetParams() performs the following actions on behalf of the calling system routine:

- Verify that the keywords passed to the routine are all allowed by the routine.
- Carry out the type checking and conversions required for each keyword.
- Find the positional (non-keyword) arguments that are scattered among the keyword arguments in **argv** and copy them in order into the **plain_args** array.
- Return the number of plain arguments copied into **plain_args**.

IDL_KWGetParams() has the form:

```
int IDL_KWGetParams(int argc, IDL_VPTR *argv, char *argk,
                   IDL_KW_PAR *kw_list, IDL_VPTR plain_args[], int mask)
```

where:

argc

The number of arguments passed to the caller. This is the first parameter to all system routines.

argv

The array of **IDL_VPTR** to arguments that was passed to the caller. This is the second parameter to all system routines.

argk

The pointer to the keyword list that was passed to the caller. This is the third parameter to all system routines that accept keyword arguments.

kw_list

An array of **IDL_KW_PAR** structures (see “[Overview Of IDL Keyword Processing](#)” on page 124, and “[The IDL_KW_PAR Structure](#)” on page 390) that specifies the acceptable keywords for this routine. This array is terminated by setting the keyword field of the final struct to NULL ((**char ***) 0).

plain_args

An array of **IDL_VPTR** into which the **IDL_VPTRs** of the positional arguments will be copied. This array must have enough elements to hold the maximum possible number of positional arguments, as defined in **IDL_SYSFUN_DEF2**. See [“Registering Routines”](#) on page 296.

mask

Mask enable. This variable is ANDed with the mask field of each **IDL_KW_PAR** struct in the array given by **kw_list**. If the result is non-zero, the keyword is accepted as a valid keyword for the called system routine. If the result is zero, the keyword is ignored.

Speeding Keyword Processing

As mentioned above, the **kw_list** argument to **IDL_KWGetParams()** is a null terminated list of **IDL_KW_PAR** structures. The time required to scan each item of the keyword array and zero the required fields (those fields specified, and value fields with **IDL_KW_ZERO** set), can become significant, especially when more than a few keyword array elements (e.g., 5 to 10 elements) are present.

To speed things up, specify **IDL_KW_FAST_SCAN** as the first keyword array element. If **IDL_KW_FAST_SCAN** is the first keyword array element, the keyword array is compiled by **IDL_KWGetParams()** into a more efficient form the first time it is used. Subsequent calls use this efficient version, greatly speeding keyword processing. Usage of **IDL_KW_FAST_SCAN** is optional, and is not worthwhile for small lists. For longer lists, however, the improvement in speed is noticeable. For example, the following list does not use fast scanning:

```
static IDL_KW_PAR kw_pars[] = {
    { "DOUBLE", IDL_TYP_DOUBLE, 1, 0, &d_there, CHARA(d) },
    { "FLOAT", IDL_TYP_FLOAT, 1, IDL_KW_ZERO, 0, CHARA(f) },
    { NULL }
};
```

To use fast scanning, it would be written as:

```
static IDL_KW_PAR kw_pars[] = {
    IDL_KW_FAST_SCAN,
    { "DOUBLE", IDL_TYP_DOUBLE, 1, 0, &d_there, CHARA(d) },
    { "FLOAT", IDL_TYP_FLOAT, 1, IDL_KW_ZERO, 0, CHARA(f) },
    { NULL }
};
```

Cleaning Up

The `IDL_KWCleanup()` function is necessary if the keywords allowed by a system routine include any input-only keywords of type `IDL_TYP_STRING`, or if the `IDL_KW_VIN` flag is used by any of the keyword `IDL_KW_PAR` structures. Such keywords can cause keyword processing to allocate temporary variables that must be cleaned up after they've outlived their usefulness. Call `IDL_KWCleanup()` as follows:

```
void IDL_KWCleanup(int fcn)
```

where `fcn` specifies the operation to be performed, and must be one of the following values:

IDL_KW_MARK

Mark the stack by placing the statement:

```
IDL_KWCleanup(IDL_KW_MARK);
```

above the call to `IDL_KWGetParams()`. In addition, you will need to make a call with `IDL_KW_CLEAN` at the end.

IDL_KW_CLEAN

Clean up from the last call to `IDL_KWGetParams()` by placing the line:

```
IDL_KWCleanup(IDL_KW_CLEAN);
```

just above the `return` statement.

Keyword Examples

The following C function implements `KEYWORD_DEMO`, a system procedure intended to demonstrate how to write the keyword processing code for a routine. It prints the values of its keywords, changes the value of `READWRITE` to 42 if it is present, and returns. Each line is numbered to make discussion easier. These numbers are not part of the actual program.

Note

The following code is designed to demonstrate keyword processing in a simple, uncluttered example. In actual code, you would not use the `printf` mechanism used on lines 35-39.

```

1  #include <stdio.h>
2  #include <idl_export.h>
3
4  void keyword_demo(int argc, IDL_VPTR *argv, char *argk)
5  {
6      int i;
7      IDL_ALLTYPES newval;
8
9      static int d_there, s_there, arr_there;
10     static IDL_LONG l;
11     static float f;
12     static double d;
13     static IDL_STRING s;
14     static IDL_LONG arr_data[10];
15     static IDL_KW_ARRAY_DESC arr_d = {(char *) arr_data,3,10,0};
16     static IDL_VPTR var;
17
18     static IDL_KW_PAR kw_pars[] = { IDL_KW_FAST_SCAN,
19         { "ARRAY", IDL_TYP_LONG, 1, IDL_KW_ARRAY, &arr_there,
20             IDL_CHARA(arr_d) },
21         { "DOUBLE", IDL_TYP_DOUBLE, 1, 0, &d_there, IDL_CHARA(d) },
22         { "FLOAT", IDL_TYP_FLOAT, 1, IDL_KW_ZERO, 0, IDL_CHARA(f) },
23         { "LONG", IDL_TYP_LONG, 1, IDL_KW_ZERO|IDL_KW_VALUE|15, 0,
24             IDL_CHARA(l) },
25         { "READWRITE", IDL_TYP_UNDEF, 1, IDL_KW_OUT|IDL_KW_ZERO,
26             0, IDL_CHARA(var) },
27         { "STRING", TYP_STRING, 1, 0, &s_there, IDL_CHARA(s) },
28         { NULL }
29     };
30
31     IDL_KWcleanup(IDL_KW_MARK);
32
33     (void) IDL_KWGetParams(argc, argv, argk, kw_pars, NULL, 1);
34
35     printf("LONG: <%spresent>\n", l ? "" : "not ");
36     printf("FLOAT: %f\n", f);
37     printf("DOUBLE: <%spresent>\n", d_there ? "" : "not ");
38     printf("STRING: %s\n", s_there ? IDL_STRING_STR(&s) : "<not present>");
39     printf("ARRAY: ");
40

```

Figure 0-1: Obsolete Example

```

41  if (arr_there)
42      for (i = 0; i < arr_d.n; i++)
43          printf(" %d", arr_data[i]);
44  else
45      printf("<not present>");
46  printf("\n");
47
48  printf("READWRITE: ");
49  if (var) {
C  50      IDL_Print(1, &var, (char *) 0);
51      newval.l = 42;
52      IDL_StoreScalar(var, TYP_LONG, &newval);
53  } else {
54      printf("<not present>");
55  }
56  printf("\n");
57
58  IDL_KWcleanup (IDL_KW_CLEAN);
59  }

```

Figure 0-1: Obsolete Example (Continued)

Executing this routine from the IDL command line, by entering:

```
KEYWORD_DEMO
```

gives the output:

```

LONG: <not present>
FLOAT: 0.000000
DOUBLE: <not present>
STRING: <not present>
ARRAY: <not present>
READWRITE: <not present>

```

Executing it again with keywords specified:

```

A = 56
KEYWORD_DEMO, /LONG, FLOAT=2, DOUBLE=34,$
  STRING="hello", ARRAY=FINDGEN(10), READWRITE=A
PRINT, 'Final Value of A: ', A

```

gives the output:

```

LONG: <present>
FLOAT: 2.000000
DOUBLE: <present>
STRING: hello
ARRAY: 0 1 2 3 4 5 6 7 8 9
READWRITE:      56
Final Value of A:      42

```

Those features of this procedure that are interesting in terms of keyword processing are, by line number:

7

The **IDL_StoreScalar()** function used on line 51 requires the scalar to be provided in an **IDL_ALLTYPES** struct.

9

These variables are used to determine if a given keyword is present. Note that all the keyword-related variables are declared static. This is necessary so that the C compiler can build the **IDL_KW_PAR** structure at compile time.

10 – 13

C variables to receive the scalar read-only keyword values.

14

C array to be used for the ARRAY read-only array keyword.

15

The array descriptor used for ARRAY. **arr_data** is the address where the array contents should be copied. The minimum number of elements allowed is 3, the maximum is 10. The value set in the last field (0) is not important, because the keyword processing routine never reads its value. Instead, it puts the number of elements actually seen there.

16

The READWRITE keyword uses the **IDL_KW_OUT** flag, so the routine receives an **IDL_VPTR** instead of a processed value.

18

The keyword definition array. Notice that all of the keywords are ordered lexically (ASCII) and that there is a NULL entry at the end (line 28). Also, all of the mask fields are set to 1, as is the mask argument to **IDL_KWGetParams()** on line 33. This means that all of the keywords in the list are to be considered valid in this routine.

The **IDL_KW_FAST_SCAN** macro is used to define the first keyword array element, speeding the processing of a long **IDL_KW_PAR** list.

19 – 20

ARRAY is defined to be a read-only array keyword of `IDL_TYP_LONG`. The `arr_there` variable will be set to non-zero if the keyword is present. In that case, the array contents will be placed in the variable `arr_data` and the number of elements will be placed into `arr_d.n`.

21

DOUBLE is a scalar keyword of `IDL_TYP_DOUBLE`. It uses the variable `d_there` to know if the keyword is present.

22

FLOAT is an `IDL_TYP_FLOAT` scalar keyword. It does not use the `specified` field of the `IDL_KW_PAR` struct to get notification of whether the keyword is present. Instead, it uses the `IDL_KW_ZERO` flag to make sure that the variable `f` is always zeroed. If the keyword is present, the value will be written into `f`, otherwise it will remain 0. The important point is that the routine can't tell the difference between the keyword being absent, or being present with a user-supplied value of zero. If this distinction doesn't matter, such as when the keyword is to serve as an on/off toggle, use this method. If it does matter, use the `specified` field as demonstrated with the DOUBLE keyword, above.

23 – 24

LONG is a scalar keyword of `IDL_TYP_LONG`. It sets the `IDL_KW_ZERO` flag to get the variable `l` zeroed prior to keyword parsing. The use of the `IDL_KW_VALUE` flag indicates that if the keyword is present, the value 15 (the lower 12 bits of the flags field) will be ORed into the variable `l`.

25 – 26

The `IDL_KW_OUT` flag indicates that the routine wants gets the `IDL_VPTR` for READWRITE if it is present. Since `IDL_KW_ZERO` is also set, the variable `var` will be zero unless the keyword is present. The specification of `IDL_TYP_UNDEF` here indicates that there is no type conversion or processing applied to `IDL_KW_OUT` keywords.

27

This keyword is included here to force the need for `IDL_KWcleanup()` on line 58.

28

Every array of **IDL_KW_PAR** structs must end with a NULL entry.

31

Mark the stack in preparation for the **IDL_KWCleanup()** call on line 58.

33

Do the keyword processing. The first three arguments are simply the arguments the interpreter passed to the routine. The **plain_args** argument is set to NULL because this routine is registered as not accepting any plain arguments. Since no plain arguments will be present, the return value from **IDL_KWGetParams()** is discarded.

35

The **l** variable will be 0 if **LONG** is not present, and 1 if it is.

36

The **f** variable will always have some usable value, but if it is zero there is no way to know if the keyword was actually specified or not.

37 – 38

These keywords use the variables from the specified field of their **IDL_KW_PAR** struct to determine if they were specified or not. Use of the **IDL_STRING_STR** macro is described in “[Accessing IDL_STRING Values](#)” on page 185.

39– 45

Accessing the **ARRAY** keyword is simple. The **arr_there** variable indicates if the keyword is present, and **arr_d.n** gives the number of elements.

47 – 55

Since the **READWRITE** keyword is accessed via the argument’s **IDL_VPTR**, we use the **IDL_Print()** function to print its value. This has the same effect as using the user-level **PRINT** procedure when running IDL. See “[Output of IDL Variables](#)” on page 248. Then, we change its value to 42 using **IDL_StoreScalar()**.

Again, please note that we use this mechanism in order to create a simple example. You will probably want to avoid the use of this type of output (**printf** and **IDL_PRINT()**) in your own code.

57

The use of **IDL_KWCleanup()** is necessitated by the existence of the **STRING** keyword, which is of **IDL_TYP_STRING**.

Interfaces Obsoleted in IDL 5.2.1

Changes were required to implement the ability to enable and disable IDL system routines from runtime and callable IDL. Rather than alter the `IDL_SYSFUN_DEF` structure, and the `IDL_AddSystemRoutine()` function in an incompatible way, a new structure (`IDL_SYSFUN_DEF2`) and new function (`IDL_SysRtnAdd()`) have been created to accomplish the new tasks, and the old structure and function have been obsoleted.

Note

The interfaces described in this section are considered functionally obsolete although they continue to be supported by ITT Visual Information Solutions. This section is supplied to help those maintaining older code. New code should be written using the information found in “[Registering Routines](#)” on page 296.

Registering Routines

The `IDL_AddSystemRoutine()` function adds system routines to IDL’s internal tables of system functions and procedures. As a programmer, you will need to call this function directly if you are linking a version of IDL to which you are adding routines, although this is very rare and not considered to be a good practice for maintainability reasons. More commonly, you use `IDL_AddSystemRoutine()` in the `IDL_Load()` function of a Dynamically Loadable Module (DLM).

Note

LINKIMAGE or DLMs are the preferred way to add system routines to IDL because they do not require building a separate IDL program. These mechanisms are discussed in the following sections of this chapter.

```
int IDL_AddSystemRoutine(IDL_SYSFUN_DEF *defs, int is_function,
int cnt);
```

It returns *True* if it succeeds in adding the routine or *False* in the event of an error:

defs

An array of `IDL_SYSFUN_DEF` structures, one per routine to be declared. This array must be defined with the C language static storage class because IDL keeps pointers to it. **defs** must be sorted by routine name in ascending lexical order.

is_function

Set this parameter to `IDL_TRUE` if the routines in **defs** are functions, and `IDL_FALSE` if they are procedures.

cnt

The number of **IDL_SYSFUN_DEF** structures contained in the **defs** array.

The definition of **IDL_SYSFUN_DEF** is:

```
typedef IDL_VARIABLE *(* IDL_FUN_RET)();

typedef struct {
    IDL_FUN_RET funct_addr;
    char *name;
    UCHAR arg_min;
    UCHAR arg_max;
    UCHAR flags
} IDL_SYSFUN_DEF;
```

IDL_VARIABLE structures are described in [“The IDL_VARIABLE Structure”](#) on page 153.

funct_addr

Address of the function implementing the system routine.

name

The name by which the routine is to be invoked from within IDL. This should be a pointer to a null terminated string. The name should be capitalized. If the routine is an object method, the name should be fully qualified, which means that it should include the class name at the beginning followed by two consecutive colons, followed by the method name (e.g. `CLASS::METHOD`).

arg_min

The minimum number of arguments allowed for the routine.

arg_max

The maximum number of arguments allowed for the routine. If the routine does not place an upper value on the number of arguments, use the value **IDL_MAXPARAMS**.

flags

A bitmask that provides additional information about the routine. Its value can be any combination of the following values (bitwise OR'd together to specify more than one at a time) or zero if no options are necessary:

IDL_SYSFUN_DEF_F_OBSOLETE

IDL should issue a warning message if this routine is called and !WARN.OBS_ROUTINE is set.

IDL_SYSFUN_DEF_F_KEYWORDS

This routine accepts keywords as well as plain arguments.

Simplified Routine Invocation

Note

The functions and techniques described in this section are no longer widely used, and are considered functionally obsolete although they continue to be supported by ITT Visual Information Solutions. This section is supplied to help those maintaining older code. New code should be written using the information found in [Chapter 15, “Adding System Routines”](#).

A great deal of the work involved in writing IDL system routines involves checking positional arguments, screening out illegal combinations of type and structure, and converting them to desired type. The function **IDL_EzCall()** provides a simplified way to handle this task. It processes an array of **IDL_EZ_ARG** structs which describe the processing to be applied to each positional argument.

The **IDL_EzCall()** function is similar to the facility provided for keyword arguments by **IDL_KWGetParams()**:

```
void IDL_EzCall(int argc, IDL_VPTR argv[],
               IDL_EZ_ARG arg_struct[]);
```

where:

argc

The number of positional arguments present.

argv

An array of pointers to the positional arguments.

arg_struct

An array of **IDL_EZ_ARG** structures defining the desired characteristics for each possible argument. Note that this array must have a definition for every possible parameter whether that argument is present in the current call or not. The order of the **IDL_EZ_ARG** structures is the same as the order in which the arguments are specified in the call. (See [“The IDL_EZ_ARG struct”](#) on page 405.)

There are some things you need to be aware of when using **IDL_EzCall()**:

- **IDL_EzCall()** automatically excludes file variables (such as those created by the ASSOC function) so you don't have to take any special action to screen such variables out.

- Every call to **IDL_EzCall()** must have a matching call to **IDL_EzCallCleanup()** before execution returns to the interpreter.
- **IDL_EzCall()** does not handle keyword arguments. If the calling routine allows keyword arguments, it must do its own keyword processing using **IDL_KWGetParams()** (see “[IDL Internals: Keyword Processing](#)” on page 121) and pass an **argv** containing only positional arguments to **IDL_EzCall()**.
- If you mark a variable as being write-only, you shouldn’t count on anything useful being in the **uargv** or **value** fields. This implies that it is not a good idea to set the **IDL_EZ_POST_WRITEBACK** field in the post field. Instead, you will have to allocate a new temporary variable, place the desired value into it, and use the **IDL_VarCopy()** function to write its value back into the original **argv** entry yourself.

Note

IDL_EZ_POST_WRITEBACK is only useful when the access field is set to **IDL_EZ_ACCESS_RW**.

The IDL_EZ_ARG struct

The **IDL_EZ_ARG** struct has the following definition:

```
typedef struct {
    short allowed_dims;
    short allowed_types;
    short access;
    short convert;
    short pre;
    short post;
    IDL_VPTR to_delete;
    IDL_VPTR uargv;
    IDL_ALLTYPES value;
} IDL_EZ_ARG;
```

where:

allowed_dims

A bit mask that specifies the allowed dimensions. Bit 0 means scalar, bit 1 means one-dimensional, etc. The **IDL_EZ_DIM_MASK** macro can be used to specify certain bits. It accepts a single argument that specifies the number of dimensions that are accepted, and returns the bit value that represents that number. For example, to specify that the argument can be scalar or have 2 dimensions:

`IDL_EZ_DIM_MASK(0) | IDL_EZ_DIM_MASK(2)`

In addition, the following constants are defined to simplify the writing of common cases:

IDL_EZ_DIM_ARRAY

Allow all but scalar.

IDL_EZ_DIM_ANY

Allow anything.

allowed_types

This is a bit mask defining the allowed data types for the argument. To convert type codes to the appropriate bits, use the formula $\text{BitMask} = 2^{\text{TypeCode}}$ or use the **IDL_TYP_MASK** macro (see “[Type Masks](#)” on page 115).

Note

If you specify a value for the convert field, it's a good idea to specify **IDL_TYP_B_ALL** or **IDL_TYP_B_SIMPLE** here. The type conversion will catch any problems and your routine will be more flexible.

access

A bitmask that describes the type of access to be allowed to the argument. The following constants should be OR'd together to set the proper value:

IDL_EZ_ACCESS_R

The value of the argument is used by the system routine.

IDL_EZ_ACCESS_W

The value of the argument is changed by the system routine. This means that it must be a named variable (as opposed to a constant or expression).

IDL_EZ_ACCESS_RW

This is equivalent to **IDL_EZ_ACCESS_R | IDL_EZ_ACCESS_W**.

convert

The type code for the type to which the argument will be converted. A value of **IDL_TYP_UNDEF** means that no conversion will be applied.

pre

A bitmask that specifies special purpose processing that should be performed on the variable by **IDL_EzCall()**. These bits are specified with the following constants:

IDL_EZ_PRE_SQMATRIX

The argument must be a square matrix.

IDL_EZ_PRE_TRANSPOSE

Transpose the argument.

Note

This processing occurs after any type conversions specified by **convert**, and is only done if the access field has the **IDL_EZ_ACCESS_R** bit set.

post

A bit mask that specifies special purpose processing that should be performed on the variable by **IDL_EzCallCleanup()**. These bits are specified with the following constants:

IDL_EZ_POST_WRITEBACK

Transfer the contents of the **uargv** field back to the actual argument.

IDL_EZ_POST_TRANSPOSE

Transpose **uargv** prior to transferring its contents back to the actual argument.

Note

This processing occurs only when the **access** field has the **IDL_EZ_ACCESS_W** bit set. If **IDL_EZ_POST_WRITEBACK** is not present, none of the other actions are considered, since that would imply wasted effort.

to_delete

Do not make use of this field. This field is reserved for use by the EZ module. If **IDL_EzCall()** allocated a temporary variable to satisfy the conversion requirements given by the **convert** field, the **IDL_VPTR** to that temporary is saved here for use by **IDL_EzCallCleanup()**.

uargv

After calling `IDL_EzCall()`, `uargv` contains a pointer to the `IDL_VARIABLE` which is the argument. This is the `IDL_VPTR` that your routine should use. Depending on the required type conversions, it might be the actual argument, or a temporary variable containing a converted version of the original. This field won't contain anything useful if the `IDL_EZ_ACCESS_R` bit is not set in the `access` field.

value

This is a copy of the `value` field of the `IDL_VARIABLE` pointed at by `uargv`. For scalar variables, it contains the value, for arrays it points at the array block. This field is here to make reading read-only variables faster. Note that this is only a copy from `uargv`, and changing it will not cause the actual `value` field in `uargv` to be updated.

Cleaning Up

Every call to `IDL_EzCall()` must be bracketed by a call to `IDL_EzCallCleanup()`:

```
void IDL_EzCallCleanup(int argc, IDL_VPTR argv[],
                      IDL_EZ_ARG arg_struct[]);
```

The arguments are exactly the same as those passed to `IDL_EzCall()`.

Example— using `IDL_EzCall()`

The following function skeleton shows how to use the simplified interface to handle argument processing for an older version of the built-in SVD (Singular Value Decomposition) function. SVD accepts the following positional arguments (in order):

A

An m by n matrix (input, required).

w

An n -element vector (output, required).

U

An n by m matrix (output, optional)

V

An n by n matrix (output, optional)

Each line is numbered to make discussion easier. These numbers are not part of the actual program.

```

1 void nr_svdcmp(int argc, IDL_VPTR argv[])
2 {
3     .
4     .
5     .
6     static IDL_EZ_ARG arg_struct[] = {
7         { IDL_EZ_DIM_MASK(2), IDL_TYP_B_SIMPLE, IDL_EZ_ACCESS_R,
8           IDL_TYP_FLOAT, 0, 0 }, /* A */
9         { IDL_EZ_DIM_ANY, IDL_TYP_B_ALL,
10          IDL_EZ_ACCESS_W, 0, 0, 0 }, /* w */
11         { IDL_EZ_DIM_ANY, IDL_TYP_B_ALL,
12          IDL_EZ_ACCESS_W, 0, 0, 0 }, /* U */
13         { IDL_EZ_DIM_ANY, IDL_TYP_B_ALL,
14          IDL_EZ_ACCESS_W, 0, 0, 0 } /* V */
15     };
16
17     IDL_EzCall(argc, argv, arg_struct);
18     .
19     .
20     .
21     /* Do the SVD calculation and prepare temporary
22        variables to be returned as w, U, and V */
23     .
24     .
25     .
26     IDL_EzCallCleanup(argc, argv, arg_struct);
27 }

```

Table A-1: IDL_EzCall() Argument Processing Example

Those features of this procedure that are interesting in terms of plain argument processing are, by line number:

7-8

The settings of the various fields of the **IDL_EZ_ARG** struct for the first positional argument (A) specifies:

allowed_dims

The argument must be 2-dimensional.

allowed_types

It can have any simple type. Types and type codes are discussed in “[IDL Internals: Types](#)” on page 113.

access

The routine will examine the argument’s value, but will not attempt to change it.

convert

The argument should be converted to **IDL_TYP_FLOAT** if necessary.

pre

No pre-processing is required.

post

No post-processing is required.

...

The remaining fields are all set by **IDL_EzCall()** in response to the above.

9-14

Arguments two through four are allowed to have any number of dimensions and are allowed any type. This is because the routine does not intend to examine them, only to change them. For the same reason, a zero (**IDL_TYP_UNDEF**) is specified for the convert field indicating that no type conversion is desired. No pre or post-processing is specified.

17

Process the positional arguments.

26

Clean up.

Obsolete Error Handling API

The following variables can be accessed only on UNIX. These variables have been superseded by the functions listed in “[Functions for Returning System Variables](#)” on page 257, which are available on all platforms. In all cases, these variables should be considered READ-ONLY:.

IDL System Variable	Internal Variable	Type
!DIR	IDL_SysvDir	IDL_STRING
!VERSION.ARCH	IDL_SysvVersion.arch	IDL_STRING
!VERSION.OS	IDL_SysvVersion.os	IDL_STRING
!VERSION.OS_FAMILY	IDL_SysvVersion.os_family	IDL_STRING
!VERSION.RELEASE	IDL_SysvVersion.release	IDL_STRING
!ERR	IDL_SysvErrCode	IDL_LONG
!ERROR	IDL_SysvErrorCode	IDL_LONG
!ORDER	IDL_SysvOrder	IDL_LONG

Table A-2: IDL System Variables Available to User Programs

In addition, the following function has been superseded by the `IDL_SysvErrorCodeValue()` function:

IDL_LONG IDL_SysvErrCodeValue(void)

This function returns the value of !ERR.



Index

Symbols

!DIR system variable, [257](#)

!DLM_PATH system variable

 DLM management, [311](#)

!ERROR_STATE system variable, [257](#), [257](#)
 setting, [257](#)

!ERROR_STATE.CODE system variable, [339](#)

!ORDER system variable, [257](#)

!VERSION.ARCH system variable, [257](#)

!VERSION.OS system variable, [257](#)

!VERSION.OS_FAMILY system variable,
 [257](#)

!VERSION.RELEASE system variable, [257](#)

A

absolute value, [265](#)

accessing structure tags, [161](#)

accessing variable data, [176](#)

action argument, [197](#)

adding

 journal file output, [249](#)

 system routines, [296](#)

adding code to IDL

 overview, [22](#)

 skills required, [23](#)

 system routines, [270](#)

allocating and freeing file units, [243](#)

anonymous structures, [160](#), [160](#)

arguments

 checking, [202](#)

 keywords. *See* keywords

argv argument, [202](#)

array variables, [157](#)

arrays

- creating
 - from existing data, [172](#)
 - temporary, [167](#)
- passing with `CALL_EXTERNAL`, [68](#)
- ASSOC function, [154](#), [158](#)
- associated I/O, [154](#), [158](#)
- `AUTO_GLUE`, [56](#)

B

- bell, ringing with error messages, [199](#)
- blocking timers, [222](#)
- blocking UNIX timers, [226](#)
- buffered data, flushing, [246](#)

C

- `CALL_EXTERNAL` function

- `AUTO_GLUE`, [46](#), [56](#)
- C examples, [58](#)
- calling a C routine, [60](#)
- calling convention, [54](#)
- common errors, [51](#)
- compared to UNIX child process, [45](#)
- compilation and linking, [45](#)
- data types, [47](#)
- Fortran examples, [72](#)
- glue functions, [46](#), [56](#)
- input/output, [47](#)
- memory cleanup, [47](#)
- Microsoft calling conventions, [49](#)
- overview, [16](#), [44](#)
- passing array data, [68](#)
- passing structures, [70](#)
- portable calling convention, [54](#)
- string data, [64](#)
- wrapper routines, [62](#)
- callable IDL
 - about, [18](#)
 - appropriate applications, [327](#)

- appropriate uses, [327](#)
- background mode, [333](#)
- cleanup, [330](#), [341](#), [341](#)
- compiling and linking C programs, [342](#)
- diverting IDL output, [337](#)
- example programs, [343](#), [347](#), [351](#)
- executing IDL statements, [339](#)
- implementation, [324](#)
- initializing IDL, [329](#)
- interactive IDL sessions, [342](#)
- inter-language calling conventions, [327](#)
- licensing issues, [328](#), [333](#)
- no command line, [333](#)
- platform-specific implementation, [324](#)
- program size considerations, [326](#)
- threading, [327](#)
- troubleshooting, [326](#)
- using
 - from C, [343](#)
 - from Fortran, [351](#)
 - overview, [329](#)
 - using the Windows graphics driver, [326](#)
 - when to use, [325](#)
- callable IDL applications
 - simple math function example, [347](#)
- callbacks
 - timer, [223](#)
- calltest program listing
 - C, [343](#)
 - Fortran, [351](#)
- characters
 - reading from the keyboard, [247](#)
- checking arguments, [202](#)
- checking file status, [241](#)
- child processes
 - communicating with, [38](#)
 - spawning, [37](#)
- cleanup
 - callable IDL, [341](#)
- client process, [78](#)
- client variables, [80](#)

- closing
 - files, 239
 - files, preventing, 240
- code argument, 195
- command line
 - initializing IDL without, 333
- communicating with a child process, 37
- compilation and link statements, 357
- complex
 - data type
 - internal, 117
- constants
 - preprocessor, 264
- copying
 - strings, 186
 - variables, 177
- copyrights, 2
- creating
 - structures, 159

D

- data types
 - default output formats, 261
 - internal, 114
 - See also* types, internal.
- debugging
 - DLMs, 270
 - shared library code, 270
 - system routines, 270
- default output formats for data types, 261
- definitions
 - external, 29
- deleting
 - strings, 187
- device
 - special files, 234
- diverting IDL output, 337
- dynamic memory
 - allocating, 174
 - allocation routines, 252

- freed when deleting strings, 187
- freeing, 179
- IDL_MemAlloc(), 253
- IDL_MemAllocPerm(), 254
- IDL_MemFree(), 253

E

- end of file
 - detecting, 245
- errno global variable
 - system level errors, 196
- error messages
 - ringing bell, 199
- errors
 - checking arguments, 202
 - issuing, 195
 - messages
 - format string, 199
 - ringing bell with error message, 199
 - suppressing
 - error messages, 198
 - message prefixes, 198
 - traceback portion of messages, 198
- examples
 - C examples for CALL_EXTERNAL, 58
 - callable IDL
 - from C, 343
 - from Fortran, 351
 - calling a simple math function, 347
 - Fortran CALL_EXTERNAL, 72
 - hello world, 273
 - simple system routine, 274
 - simple_vars.pro, 62
 - using WIDGET_STUB, 377, 379
- exit handlers
 - IDL_ExitRegister(), 255
- export restrictions, 2
- export.h *see* idl_export.h
- external
 - definitions, 29

programs, accessing (SPAWN), 14

F

file

- attributes, verifying, 241
- descriptor, 232
- end of file detection, 245
- IDL_FileOpen(), 236
- prevent closure, 240

file access

- IDL_FILE_STAT struct, 233
- mode, 236

file information

- IDL_FILE_STAT struct, 232

file units

- always open, 238
- LUN table, 232
- special, 238

files

checking

- attributes, 241
- status, 241

closing

- IDL_FileClose, 239
- closing, preventing, 240

FLEXIm floating licence policy, 333

flushing buffered data, 246

Fortran

- binary data, unformatted, 234
- calling, 74
- child processes, 40
- compiler, 342
- complex data types, 117
- external functions, calling, 44
- passing parameters, 24

free() function, 174

FZ_ROOTS function

- example, 277

G

getting dynamic memory, 174

getting file information, 232

H

heap variables, 164

Hello World example, 273

HELP/DLM, 312, 317

I

IDL

about language, 27

combining external code, 22

internal initialization, 329

no command line, 333

spawning child process, 37

IDL output, diverting, 337

IDL portable calling convention, 54

IDL RPC

Client API Example, 81

variable accessor macros, 108

IDL signal API, 211

IDL statements, executing, 339

IDL_ABS() macro, 265

IDL_ALLTYPES union, 153, 156

IDL_ARRAY structure, 153

IDL_BailOut() function, 256

IDL_BasicTypeConversion() function, 207

IDL_CHAR() macro, 265

IDL_CHARA() macro, 265

IDL_Cleanup() function, 330, 341

IDL_CvtByte function, 208

IDL_CvtBytscl function, 208

IDL_CvtComplex function, 208

IDL_CvtDbl function, 208

IDL_CvtDComplex function, 208

IDL_CvtFix function, 208

IDL_CvtFlt function, 208

- IDL_CvtLng function, 208
- IDL_CvtString function, 208
- IDL_Deltmp() function, 171, 175
- IDL_DLM_PATH, 311, 317
- IDL_ENSURE_ARRAY macro, 203
- IDL_ENSURE_OBJREF macro, 203
- IDL_ENSURE_PTR macro, 203
- IDL_ENSURE_SCALAR macro, 203
- IDL_ENSURE_SIMPLE macro, 203
- IDL_ENSURE_STRING macro, 203
- IDL_ENSURE_STRUCTURE macro, 204
- IDL_EXCLUDE_COMPLEX macro, 203
- IDL_EXCLUDE_CONST macro, 202
- IDL_EXCLUDE_EXPR macro, 202
- IDL_EXCLUDE_FILE macro, 203
- IDL_EXCLUDE_SCALAR macro, 203
- IDL_EXCLUDE_STRING macro, 203
- IDL_EXCLUDE_STRUCT macro, 203
- IDL_EXCLUDE_UNDEF macro, 202
- IDL_Execute() function, 339
- IDL_ExecuteStr() function, 339
- IDL_ExitRegister() function, 255
- idl_export.h file, 29
- IDL_FALSE preprocessor constant, 264
- IDL_FILE_STAT struct, 232
- IDL_FileClose() function, 239
- IDL_FileEnsureStatus() function, 241
- IDL_FileEOF() function, 245
- IDL_FileFlushUnit() function, 246
- IDL_FileFreeUnit() function, 243
- IDL_FileGetUnit() function, 243
- IDL_FileOpen() function, 236
- IDL_FileSetClose() function, 240
- IDL_FileStat() function, 232
- IDL_FileTerm global variable, 258
- IDL_FileTermColumns function, 259
- IDL_FileTermIsTty function, 258
- IDL_FileTermLines function, 259
- IDL_FileTermName function, 258
- IDL_FindNamedVariable() function, 182
- IDL_GetKbrd() function, 247
- IDL_GetScratch function, 174
- IDL_Gettmp() function, 166
- IDL_GetUserInfo() function, 263
- IDL_GetVarAddr() function, 181
- IDL_GetVarAddr1() function, 181
- IDL_ImportArray() function, 160, 172
- IDL_ImportNamedArray() function, 160, 172
- IDL_Initialize() function, 329, 331
- IDL_KW_ARR_DESC structure, 129
- IDL_KW_FAST_SCAN macro, 134
- IDL_KW_PAR structure, 123, 126
- IDL_KWCleanup() function, 123
- IDL_KWGetParams() function, 123, 133
- IDL_Load(), 296, 314
- IDL_Logit() function, 249
- IDL_LONG type definition, 116
- IDL_LONG64, 116
- IDL_M_GENERIC message string, 199
- IDL_M_NAMED_GENERIC message code, 199
- IDL_Main() function, 342
- IDL_MakeStruct() function, 159
- IDL_MakeTempArray function, 167
- IDL_MakeTempStruct() function, 168
- IDL_MAX() macro, 265
- IDL_MAX_ARRAY_DIM preprocessor constant, 264
- IDL_MAX_TYPE constant, 114
- IDL_MAXIDLEN preprocessor constant, 264
- IDL_MAXPATH preprocessor constant, 264
- IDL_MBLK_CORE, 192
- IDL_MemAlloc() function, 253
- IDL_MemAllocPerm() function, 254
- IDL_MemFree() function, 253
- IDL_Message() function, 195, 214
- IDL_MessageDefineBlock(), 192, 314
- IDL_MessageNameToCode(), 201
- IDL_MIN() macro, 265
- IDL_MSG_DEF, 192
- IDL_NUM_TYPES constant, 114
- IDL_OutputFormat global variable, 261

- IDL_OutputFormatFunc function, 261
- IDL_OutputFormatLen global variable, 261
- IDL_OutputFormatLenFunc function, 262
- IDL_Print() function, 248
- IDL_Printf() function, 248
- IDL_REGISTER preprocessor constant, 264
- IDL_ROUND_UP() macro, 266
- IDL_RPCCleanup, 86
- IDL_RPCDeltmp, 87
- IDL_RPCExecuteStr, 88
- IDL_RPCGetArrayData, 108
- IDL_RPCGetArrayNumDims, 108
- IDL_RPCGetArrayDimensions, 108
- IDL_RPCGetMainVariable, 89
- IDL_RPCGettmp, 90
- IDL_RPCGetVarByte, 108
- IDL_RPCGetVarComplex, 108
- IDL_RPCGetVarComplexI, 108
- IDL_RPCGetVarComplexR, 108
- IDL_RPCGetVarDComplex, 108
- IDL_RPCGetVarDComplexI, 108
- IDL_RPCGetVarDComplexR, 108
- IDL_RPCGetVarDouble, 109
- IDL_RPCGetVarFloat, 109
- IDL_RPCGetVariable, 91
- IDL_RPCGetVarInt, 109
- IDL_RPCGetVarLong, 109
- IDL_RPCGetVarLong64, 109
- IDL_RPCGetVarString, 109
- IDL_RPCGetVarType, 109
- IDL_RPCGetVarUInt, 109
- IDL_RPCGetVarULong64, 109
- IDL_RPCImportArray, 92
- IDL_RPCInit, 93
- IDL_RPCMakeArray, 94
- IDL_RPCOutputCapture, 96
- IDL_RPCOutputGetStr, 97
- IDL_RPCSetMainVariable, 98
- IDL_RPCSetVariable, 99
- IDL_RPCStoreScalar, 100
- IDL_RPCStrDelete, 101
- IDL_RPCStrDup, 102
- IDL_RPCStrEnsureLength, 103
- IDL_RPCStrStore, 104
- IDL_RPCTimeout, 105
- IDL_RPCVarCopy, 106
- IDL_RPCVarGetData, 107
- IDL_RPCVarIsArray, 109
- IDL_RuntimeExec() function, 340
- IDL_SignalBlock() function, 219
- IDL_SignalMaskBlock() function, 218
- IDL_SignalMaskGet() function, 217
- IDL_SignalMaskSet() function, 218
- IDL_SignalRegister() function, 214
- IDL_SignalSetAdd() function, 216
- IDL_SignalSetDel() function, 217
- IDL_SignalSetInit() function, 216
- IDL_SignalSetIsMember() function, 217
- IDL_SignalSuspend() function, 219
- IDL_SignalUnregister() function, 215
- IDL_SREF structure, 153, 159
- IDL_STDERR_UNIT file unit, 238
- IDL_STDIN_UNIT file unit, 238
- IDL_STDOUT_UNIT file unit, 238
- IDL_StoreScalar() function, 178, 202
- IDL_StoreScalarZero(), 178
- IDL_StrDelete() function, 187
- IDL_StrDup() function, 186
- IDL_StrEnsureLength() function, 189
- IDL_STRING struct, 117
- IDL_STRING structure, 184
- IDL_STRING_STR macro, 185
- IDL_StrStore() function, 188
- IDL_StrToSTRING() function, 188
- IDL_STRUCT_TAG_DEF type definition, 160
- IDL_StructNumTags(), 162
- IDL_StructTagInfoByIndex() function, 161
- IDL_StructTagInfoByName() function, 161
- IDL_StructTagNameByIndex function, 163
- IDL_SYSFUN_DEF, 296
- IDL_SYSFUN_DEF_F_KEYWORDS, 123

- IDL_SYSFUN_DEF2 struct, 123, 296
 - IDL_SysRtnAdd function, 123, 296
 - IDL_SysvDirFunc function, 257
 - IDL_SysvErrorCodeValue function, 257
 - IDL_SysvErrStringFunc function, 257
 - IDL_SysVersionArch function, 257
 - IDL_SysVersionOS function, 257
 - IDL_SysVersionOSFamily function, 257
 - IDL_SysVersionRelease function, 257
 - IDL_SysvOrderValue function, 257
 - IDL_SysvSyserrStringFunc function, 257
 - IDL_TERMINFO struct, 258
 - IDL_TIMER_CONTEXT variable, 224
 - IDL_TimerBlock() function, 226
 - IDL_TimerCancel() function, 225
 - IDL_TimerSet() function, 223
 - IDL_ToutPop() function, 338
 - IDL_ToutPush() function, 338
 - IDL_TRUE preprocessor constant, 264
 - IDL_TTYReset function, 260
 - IDL_TYP_B_ALL constant, 115
 - IDL_TYP_BYTE type code, 114
 - IDL_TYP_COMPLEX type code, 114, 117
 - IDL_TYP_DCOMPLEX type code, 114, 117
 - IDL_TYP_DOUBLE type code, 114
 - IDL_TYP_FLOAT type code, 114
 - IDL_TYP_INT type code, 114
 - IDL_TYP_LONG type code, 114
 - IDL_TYP_LONG64 type code, 115
 - IDL_TYP_MASK preprocessor macro, 115
 - IDL_TYP_OBREF type code, 115
 - IDL_TYP_PTR type code, 115
 - IDL_TYP_STRING type code, 114, 117
 - IDL_TYP_STRUCT type code, 114, 159
 - IDL_TYP_UINT type code, 115
 - IDL_TYP_ULONG type code, 115
 - IDL_TYP_ULONG64 type code, 115
 - IDL_TYP_UNDEF, 114
 - IDL_TYP_UNDEF type code, 114
 - IDL_TypeName global variable, 261
 - IDL_TypeNameFunc function, 262
 - IDL_TypeSize global variable, 261
 - IDL_TypeSizeFunc function, 262
 - IDL_ULONG, 116
 - IDL_ULONG64, 117
 - IDL_USER_INFO struct, 263
 - IDL_VarCopy() function, 177
 - IDL_VarGetData() function, 176
 - IDL_VARIABLE structure, 153
 - IDL_VarName() function, 180
 - IDL_VPTR, 28, 153
 - IDL_WidgetGetStubIds() function, 375, 375
 - IDL_WidgetIssueStubEvent() function, 374
 - IDL_WidgetSetStubIds() function, 375, 375
 - IDL_WidgetStubLock() function, 374
 - IDL_WidgetStubLookup() function, 374
 - IDL_WidgetStubSetSizeFunc() function, 376, 376
 - IDLRPCGetVarULong, 109
 - information on open files
 - IDL_FILE_STAT struct, 232
 - initializing
 - IDL
 - callable IDL, 329
 - IDL_Init() function, 334
 - IDL_Initialize() function, 331
 - no command line, 333
 - input/output
 - internal, 230
 - inter-language
 - calling conventions, 24
 - supported communication techniques, 13
 - internal callback functions (widget stub), 377
 - internal functions for stub widgets, 374
 - interpreted languages, 27
 - interpreter stack, 28
 - interrupt flag, internal, 256
- J**
- journal file, adding to, 249

K

KEYWORD_DEMO procedure, [137](#)

keywords

array, [127](#), [130](#)

Boolean, [127](#)

creating, [123](#)

examples, [137](#)

in external development, [122](#)

input/output, [130](#)

internal input/output, [127](#), [127](#)

processing, [133](#)

processing options, [130](#)

read-only, [129](#)

scalar, [130](#)

speeding processing of, [134](#)

L

language

about IDL, [27](#)

legalities, [2](#)

libraries

IDL portable calling convention, [54](#)

linking to, [81](#)

licensing

callable IDL, [328](#), [333](#)

linking

C programs with callable IDL, [342](#)

client library, [81](#)

external code into IDL, [31](#)

to IDL, [31](#)

logical unit numbers, [158](#)

long integer data type, [116](#)

longjmp() function, [198](#)

LUNs *see* logical unit numbers

M

macros

defined in `idl_export.h`, [265](#)

make file for IDL sharable libraries, [31](#)

malloc() function, [174](#)

mapping

IDL data types to C data types, [116](#)

memory

allocating, [253](#)

allocating permanent, [254](#)

freeing, [253](#)

messages

format string, [199](#)

message blocks, [192](#)

N

names

of variables (external code), [180](#)

O

obtaining names of variables, [180](#)

opening files

IDL_FileOpen(), [236](#)

P

parameters

passing mechanism, [54](#)

preprocessor constants, [264](#)

printf() function, [195](#)

printing

IDL variables, [248](#)

procedure calls, remote, [78](#)

program size considerations

callable IDL, [326](#)

R

registering

exit handlers, [255](#)

routines using `IDL_SysRtnAdd()`, 287

Remote Procedure Calls

about, 15, 78

backward compatibility, 83

example code, 110

IDL as server, 79

library, 85

ringing bell with error messages, 199

rounding

values, 266

RPC server, using IDL as, 79

RPCs *see* Remote Procedure Calls

running

IDL as RPC server, 79

runtime

embedded licensing, 340

S

scalars

values

storing, 178

variables, 156

server ID number, 79

server process, 78

shutting down

IDL, 255

`SIG_DFL`, 210, 212

`SIG_IGN`, 212

`SIGALRM`, 211, 226

`SIGFPE`, 211

`SIGINT`, 256

signal handlers

establishing, 214

removing, 215

signal masks

`IDL_SignalBlock()`, 219

`IDL_SignalMaskBlock()`, 218

`IDL_SignalMaskGet()`, 217

`IDL_SignalMaskSet()`, 218

`IDL_SignalSetAdd()`, 216

`IDL_SignalSetDel()`, 217

`IDL_SignalSetInit()`, 216

`IDL_SignalSetIsMember()`, 217

`IDL_SignalSuspend()`, 219

overview, 216

signals, 210

IDL API, 211

IDL limitations, 211

problems, 210

`SIGTRAP`, 211

`simple_vars.pro`, 62

Skills Required to Add Code to IDL, 23

SPAWN procedure

using, 37

stack, interpreter, 28

standard error, 238

standard input, 238

standard output, 238

stdio buffering, 234

storing

scalar values, 178

string data type, 117

strings

accessing, 185

copying, 186

deleting, 187

ensuring length of, 189

passing with `CALL_EXTERNAL`, 64

processing, 184

setting value of, 188

structure variables, 159

structures, 159

anonymous, 160, 160

creating internal, 159

creating temporary, 168

passing with `CALL_EXTERNAL`, 70

stub widgets

internal functions, 374

overview, 370

`WIDGET_STUB` function, 371

symbol table, 181

system routines
 adding, 296
 examples, 273, 274
 interface, 272
 overview, 270

system variables
 functions for returning, 257

T

temporary array, creating, 167

temporary variables
 about, 165
 internal
 freeing, 171
 getting, 166

Terminal Information, 258

The IDL RPC directory, 79

timer modules in IDL, 222

timers
 blocking, 222, 226
 callbacks, 223
 cancelling requests, 225
 IDL_TimerBlock(), 226
 IDL_TimerCancel(), 225
 IDL_TimerSet(), 223

trademarks, 2

troubleshooting
 callable IDL, 326

type codes
 internal, 114

type information, internal, 261

types, internal
 complex, 117
 global variables, 261
 long integer, 116
 mapping of, 116
 string, 117
 type codes, 114
 type masks, 115
 unsigned byte, 116

U

UCHAR type definition, 116

UNIX
 signal masks, 216

unsigned byte data type, 116

user information (IDL), 263

user interrupts, 256

V

variables
 array, 157
 copying, 177
 in current scope, looking up, 182
 names, 180
 obtaining names of, 180
 returning
 address in main-level program, 181
 current execution scope, 182
 scalar, 156
 setting to scalar values, 178
 structure, 159
 system, 257
 temporary, 165

W

WIDGET_STUB
 examples, 377, 379
 interface, 326, 370
 WIDGET_CONTROL keywords, 372

WIDGET_STUB function
 reference, 371

widgets
 adding custom to IDL, 370
 internal functions, 374
 WIDGET_CONTROL, 372
 WIDGET_STUB, 371

wrapper routines
 CALL_EXTERNAL, 62

X

blocking in background mode, [333](#)

XMANAGER procedure

