

EPICS Application Developer's Guide

EPICS Base Release 3.14.12.8

22 October 2018

Martin R. Kraimer, Janet B. Anderson, Andrew N. Johnson (Argonne National Laboratory)

W. Eric Norum (Lawrence Berkeley Laboratory)

Jeffrey O. Hill (Los Alamos National Laboratory)

Ralph Lange, Benjamin Franksen (Helmholtz-Zentrum Berlin)

Peter Denison (Diamond Light Source)

Contents

EPICS Applications Developer's Guide	1
Table of Contents	7
1 Introduction	9
1.1 Overview	9
1.2 Acknowledgments	11
2 Getting Started	13
2.1 Introduction	13
2.2 Example IOC Application	13
2.3 Channel Access Host Example	15
2.4 iocsh	16
2.5 Building IOC components	16
2.6 makeBaseApp.pl	19
2.7 vxWorks boot parameters	22
2.8 RTEMS boot procedure	23
3 EPICS Overview	25
3.1 What is EPICS?	25
3.2 Basic Attributes	26
3.3 IOC Software Components	26
3.4 Channel Access	28
3.5 OPI Tools	30
3.6 EPICS Core Software	30
4 Build Facility	33
4.1 Overview	33
4.2 Build Requirements	35
4.3 Configuration Definitions	36
4.4 Makefiles	41
4.5 Make	42
4.6 Makefile definitions	43
4.7 Table of Makefile definitions	71
4.8 Configuration Files	78
4.9 Build Documentation Files	81
4.10 Startup Files	82
5 Database Locking, Scanning, And Processing	83
5.1 Overview	83
5.2 Record Links	83
5.3 Database Links	84
5.4 Database Locking	85

5.5	Database Scanning	85
5.6	Record Processing	86
5.7	Guidelines for Creating Database Links	86
5.8	Guidelines for Synchronous Records	89
5.9	Guidelines for Asynchronous Records	89
5.10	Cached Puts	91
5.11	putNotify	91
5.12	Channel Access Links	91
6	Database Definition	95
6.1	Overview	95
6.2	Summary of Database Syntax	95
6.3	General Rules for Database Definition	96
6.4	path addpath – Path Definition	98
6.5	include – Include File	99
6.6	menu – Menu Declaration	99
6.7	recordtype – Record Type Declaration	100
6.8	device – Device Support Declaration	104
6.9	driver – Driver Declaration	104
6.10	registrar – Registrar Declaration	105
6.11	variable – Variable Declaration	105
6.12	function – Function Declaration	106
6.13	breaktable – Breakpoint Table	106
6.14	record – Record Instance	107
6.15	Record Information Item	110
6.16	Record Attributes	111
6.17	Breakpoint Tables – Discussion	111
6.18	Menu and Record Type Include File Generation.	112
6.19	dbExpand	115
6.20	dbLoadDatabase	116
6.21	dbLoadRecords	116
6.22	dbLoadTemplate	117
6.23	dbReadTest	118
7	IOC Initialization	119
7.1	Overview - Environments requiring a main program	119
7.2	Overview - vxWorks	120
7.3	Overview - RTEMS	120
7.4	IOC Initialization	121
7.5	Pausing an IOC	123
7.6	Changing iocCore fixed limits	124
7.7	initHooks	124
7.8	Environment Variables	126
7.9	Initialize Logging	126
8	Access Security	127
8.1	Overview	127
8.2	Quick Start	127
8.3	User's Guide	128
8.4	Design Summary	133
8.5	Access Security Application Programmer's Interface	136
8.6	Database Access Security	141
8.7	Channel Access Security	143
8.8	Trapping Channel Access Writes	145

8.9	Access Control: Implementation Overview	145
8.10	Structures	147
9	IOC Test Facilities	149
9.1	Overview	149
9.2	Database List, Get, Put	149
9.3	Breakpoints	152
9.4	Trace Processing	153
9.5	Error Logging	153
9.6	Hardware Reports	154
9.7	Scan Reports	155
9.8	General Time	155
9.9	Access Security Commands	156
9.10	Channel Access Reports	158
9.11	Interrupt Vectors	159
9.12	Miscellaneous	159
9.13	Database System Test Routines	160
9.14	Record Link Reports	161
9.15	Old Database Access Testing	161
9.16	Routines to dump database information	162
10	IOC Error Logging	165
10.1	Overview	165
10.2	Error Message Routines	166
10.3	errlog Listeners	167
10.4	errlogThread	168
10.5	console output and message queue size	168
10.6	Status Codes	168
10.7	iocLog	169
11	Record Support	171
11.1	Overview	171
11.2	Overview of Record Processing	171
11.3	Record Support and Device Support Entry Tables	172
11.4	Example Record Support Module	173
11.5	Record Support Routines	179
11.6	Global Record Support Routines	183
12	Device Support	187
12.1	Overview	187
12.2	Example Synchronous Device Support Module	188
12.3	Example Asynchronous Device Support Module	189
12.4	Device Support Routines	191
12.5	Extended Device Support	192
13	Driver Support	195
13.1	Overview	195
13.2	Device Drivers	195
14	Static Database Access	199
14.1	Overview	199
14.2	Definitions	199
14.3	Allocating and Freeing DBBASE	201
14.4	DBENTRY Routines	201
14.5	Read and Write Database	202

14.6	Manipulating Record Types	203
14.7	Manipulating Field Descriptions	204
14.8	Manipulating Record Attributes	205
14.9	Manipulating Record Instances	206
14.10	Manipulating Menu Fields	208
14.11	Manipulating Link Fields	209
14.12	Manipulating MenuForm Fields	210
14.13	Manipulating Information Items	211
14.14	Find Breakpoint Table	212
14.15	Dump Routines	212
14.16	Examples	213
15	Runtime Database Access	215
15.1	Overview	215
15.2	Database Include Files	215
15.3	Runtime Database Access Overview	218
15.4	Database Access Routines	221
15.5	Runtime Link Modification	230
15.6	Channel Access Monitors	230
15.7	Lock Set Routines	231
15.8	Channel Access Database Links	232
16	EPICS General Purpose Tasks	237
16.1	Overview	237
16.2	General Purpose Callback Tasks	237
16.3	Task Watchdog	240
17	Database Scanning	243
17.1	Overview	243
17.2	Scan Related Database Fields	243
17.3	Scan Related Software Components	244
17.4	Implementation Overview	247
18	IOC Shell	253
18.1	Introduction	253
18.2	IOC Shell Operation	253
18.3	IOC Shell Programming	256
19	libCom	261
19.1	bucketLib	261
19.2	calc	261
19.3	cppStd	265
19.4	epicsExit	266
19.5	cvtFast	266
19.6	cxxTemplates	267
19.7	dbmf	267
19.8	ellLib	268
19.9	epicsRingBytes	269
19.10	epicsRingPointer	269
19.11	epicsTimer	271
19.12	fdmgr	277
19.13	freeList	277
19.14	gpHash	277
19.15	logClient	278
19.16	macLib	278

19.17misc	280
20 libCom OSI libraries	287
20.1 Overview	287
20.2 epicsAssert	288
20.3 epicsEndian	288
20.4 epicsEvent	289
20.5 epicsFindSymbol	290
20.6 epicsGeneralTime	291
20.7 epicsInterrupt	293
20.8 epicsMath	294
20.9 epicsMessageQueue	294
20.10epicsMutex	296
20.11epicsStdlib	298
20.12epicsStdio	298
20.13epicsStdioRedirect	299
20.14epicsThread	299
20.15epicsTime	304
20.16osiPoolStatus	311
20.17osiProcess	312
20.18Ignoring Posix Signals	312
20.19OS-Independent Socket API	313
20.20 Device Support Library	313
20.21vxWorks Specific routines and Headers	317
21 Registry	319
21.1 Registry.h	319
21.2 registryRecordType.h	319
21.3 registryDeviceSupport.h	320
21.4 registryDriverSupport.h	320
21.5 registryFunction.h	320
21.6 registerRecordDeviceDriver.c	320
21.7 registerRecordDeviceDriver.pl	320
22 Database Structures	321
22.1 Overview	321
22.2 Include Files	321
22.3 Structures	323
Index	325

Chapter 1

Introduction

1.1 Overview

This document describes the core software that resides in an Input/Output Controller (IOC), one of the major components of EPICS. It is intended for anyone developing EPICS IOC databases and/or new record/device/driver support.

The plan of the book is:

Getting Started

A brief description of how to create EPICS support and ioc applications.

EPICS Overview

An overview of EPICS is presented, showing how the IOC software fits into EPICS.

EPICS Build Facility

This chapter describes the EPICS build facility including directory structure, environment and system requirements, configuration files, Makefiles, and related build tools.

Database Locking, Scanning, and Processing

Overview of three closely related IOC concepts. These concepts are at the heart of what constitutes an EPICS IOC.

Database Definition

This chapter gives a complete description of the format of the files that describe IOC databases. This is the format used by Database Configuration Tools and is also the format used to load databases into an IOC.

IOC Initialization

A great deal happens at IOC initialization. This chapter removes some of the mystery about initialization.

Access Security

Channel Access Security is implemented in IOCs. This chapter explains how it is configured and also how it is implemented.

IOC Test Facilities

Epics supplied test routines that can be executed via the epics or vxWorks shell.

IOC Error Logging

IOC code can call routines that send messages to a system wide error logger.

Record Support

The concept of record support is discussed. This information is necessary for anyone who wishes to provide customized record and device support.

Device Support

The concept of device support is discussed. Device support takes care of the hardware specific details of record support, i.e. it is the interface between hardware and a record support module. Device support can directly access hardware or may interface to driver support.

Driver Support

The concepts of driver support is discussed. Drivers, which are not always needed, have no knowledge of records but just take care of interacting with hardware. Guidelines are given about when driver support, instead of just device support, should be provided.

Static Database Access

This is a library that works on both Host and IOC. For IOCs it works on initialized or uninitialized EPICS databases.

Runtime Database Access

The heart of the IOC software is the memory resident database. This chapter describes the interface to this database.

Device Support Library

A set of routines are provided for device support modules that use shared resources such as VME address space.

EPICS General Purpose Tasks

General purpose callback tasks and task watchdog.

Database Scanning

Database scan tasks, i.e. the tasks that request records to process.

IOC Shell

The EPICS IOC shell is a simple command interpreter which provides a subset of the capabilities of the vxWorks shell.

libCom

EPICS base includes a subdirectory src/libCom, which contains a number of c and c++ libraries that are used by the other components of base. This chapter describes most of these libraries.

libCom OSI

This chapter describes the libraries in libCom that provide Operating System Independent (OSI) interfaces used by the rest of EPICS base. LibCom also contains operating system dependent code that implements the OSI interfaces.

Registry

Under vxWorks osiFindGlobalSymbol can be used to dynamically bind to record, device, and driver support. Since on some systems this always returns failure, a registry facility is provided to implement the binding. The basic idea is that any storage meant to be “globally” accessible must be registered before it can be accessed

Database Structures

A description of the internal database structures.

Other than the overview chapter this document describes only core IOC software. Thus it does not describe other EPICS tools which run in an IOC such as the sequencer. It also does not describe Channel Access.

The reader of this manual should also be aware the following additional documentation:

- *EPICS Record Reference Manual*, Philip Stanley, Janet Anderson and Marty Kraimer
- *EPICS R3.14 Channel Access Reference Manual*, Jeffrey O. Hill
- *vxWorks Programmer's Guide*, Wind River Systems
- *vxWorks Reference Manual*, Wind River Systems
- *RTEMS C User's Guide*, Online Applications Research

1.2 Acknowledgments

The basic model of what an IOC should do and how to do it was developed by Bob Dalesio at LANL/GTA. The principle ideas for Channel Access were developed by Jeff Hill at LANL/GTA. Bob and Jeff also were the principle implementers of the original IOC software. This software (called GTACS) was developed over a period of several years with feedback from LANL/GTA users. Without their ideas EPICS would not exist.

During 1990 and 1991, ANL/APS undertook a major revision of the IOC software with the major goal being to provide easily extendible record and device support. Marty Kraimer (ANL/APS) was primarily responsible for designing the data structures needed to support extendible record and device support and for making the changes needed to the IOC resident software. Bob Ziemann (ANL/APS) designed and implemented the UNIX build tools and IOC modules necessary to support the new facilities. Frank Lenkszus (ANL/APS) made extensive changes to the Database Configuration Tool (DCT) necessary to support the new facilities. Janet Anderson developed methods to systematically test various features of the IOC software and is the principal implementer of changes to record support.

During 1993 and 1994, Matt Needes at LANL implemented and supplied the description of fast database links and the database debugging tools.

During 1993 and 1994 Jim Kowalkowski at ANL/APS developed GDCT and also developed the ASCII database instance format now used as the standard format. At that time he also created `dbLoadRecords` and `dbLoadTemplate`.

The `build` utility method resulted in the generation of binary files of UNIX that were loaded into IOCs. As new IOC architectures started being supported this caused problems. During 1995, after learning from an abandoned effort now referred to as `EpicsRX`, the build utilities and binary file (called `default.dct.sdr`) were replaced by all ASCII files. The new method provides architecture independence and a more flexible environment for configuring the record/device/ driver support. This principle implementer was Marty Kraimer with many ideas contributed by John Winans and Jeff Hill. Bob Dalesio made sure that we did not go too far, i.e. 1) make it difficult to upgrade existing applications and 2) lose performance.

In early 1996 Bob Dalesio tackled the problem of allowing runtime link modification. This turned into a cooperative development effort between Bob and Marty Kraimer. The effort included new code for database to Channel Access links, a new library for lock sets, and a cleaner interface for accessing database links.

In early 1999 the port of `iocCore` to non `vxWorks` operating systems was started. The principle developers were Marty Kraimer, Jeff Hill, and Janet Anderson. William Lupton converted the sequencer as well as helping with the `posix` threads implementation of `osiSem` and `osiThread`. Eric Norum provided the port to `RTEMS` and also contributed the shell that is used on non `vxWorks` environments. Ralph Lange provided the port to `HPUX`.

Many other people have been involved with EPICS development, including new record, device, and driver support modules.

Chapter 2

Getting Started

2.1 Introduction

This chapter provides a brief introduction to creating EPICS IOC applications. It contains:

- Instructions for creating, building, and running an example IOC application.
- Instructions for creating, building, and executing example Channel Access clients.
- Briefly describes iocsh, which is a base supplied command shell.
- Describes rules for building IOC components.
- Describes makeBaseApp.pl, which is a perl script that generates files for building applications.
- Briefly discusses vxWorks boot parameters

This chapter will be hard to understand unless you have some familiarity with IOC concepts such as record/device/driver support and have had some experience with creating ioc databases. Once you have this experience, this chapter provides most of the information needed to build applications. The example that follows assumes that EPICS base has already been built.

2.2 Example IOC Application

This section explains how to create an example IOC application in a directory `<top>`, naming the application `myexampleApp` and the ioc directory `iocmyexample`.

2.2.1 Check that EPICS_HOST_ARCH is defined

Execute the command:

```
echo $EPICS_HOST_ARCH           (Unix/Linux)
```

or

```
set EPICS_HOST_ARCH             (Windows)
```

This should display your workstation architecture, for example `linux-x86` or `win32-x86`. If you get an “Undefined variable” error, you should set `EPICS_HOST_ARCH` to your host operating system followed by a dash and then your host architecture, e.g. `solaris-sparc`. The perl script `EpicsHostArch.pl` in the `base/startup` directory has been provided to help set `EPICS_HOST_ARCH`.

2.2.2 Create the example application

The following commands create an example application.

```
mkdir <top>
cd <top>
<base>/bin/<arch>/makeBaseApp.pl -t example myexample
<base>/bin/<arch>/makeBaseApp.pl -i -t example myexample
```

Here, <arch> indicates the operating system architecture of your computer. For example, `solaris-sparc`. The last command will ask you to enter an architecture for the IOC. It provides a list of architectures for which base has been built.

The full path name to <base> (an already built copy of EPICS base) must be given. Check with your EPICS system administrator to see what the path to your <base> is. For example:

```
/home/phoebus/MRK/epics/base/bin/linux-x86/makeBaseApp.pl ...
```

Windows Users Note: Perl scripts must be invoked with the command `perl <scriptname>` on Windows. Perl script names are case sensitive. For example to create an application on Windows:

```
perl C:\epics\base\bin\win32-x86\makeBaseApp.pl -t example myexample
```

2.2.3 Inspect files

Spend some time looking at the files that appear under <top>. Do this *before* building. This allows you to see typical files which are needed to build an application without seeing the files generated by make.

2.2.4 Sequencer Example

The sequencer is now supported as an unbundled product. The example includes an example state notation program, `sncExample.stt`. As created by `makeBaseApp` the example is not built or executed.

Before `sncExample.stt` can be compiled, the sequencer module must have been built using the same version of base that the example uses.

To build `sncExample` edit the following files:

- `configure/RELEASE` – Set `SNCSEQ` to the location of the sequencer.
- `iocBoot/iocmyexample/st.cmd` – Remove the comment character `#` from this line:

```
#seq sncExample, "user=<user>"
```

The Makefile contains commands for building the `sncExample` code both as a component of the example IOC application and as a standalone program called `sncProgram`, an executable that connects through Channel Access to a separate IOC database.

2.2.5 Build

In directory <top> execute the command

```
make
```

NOTE: On systems where GNU make is not the default another command is required, e.g. `gnumake`, `gmake`, etc. See you EPICS system administrator.

2.2.6 Inspect files

This time you will see the files generated by make as well as the original files.

2.2.7 Run the ioc example

The example can be run on vxWorks, RTEMS, or on a supported host.

- On a host, e.g. Linux or Solaris

```
cd <top>/iocBoot/iocmyexample
../bin/linux-x86/myexample st.cmd
```

- vxWorks/RTERMS – Set your boot parameters as described at the end of this chapter and then boot the ioc.

After the ioc is started try some of the shell commands (e.g. `dbl` or `dbpr <recordname>`) described in the chapter “IOC Test Facilities”. In particular run `dbl` to get a list of the records.

The `iocsh` command interpreter used on non-vxWorks IOCs provides a help facility. Just type:

```
help
```

or

```
help <cmd>
```

where `<cmd>` is one of the commands displayed by `help`. The `help` command accepts wildcards, so

```
help db*
```

will provide information on all commands beginning with the characters `db`. On vxWorks the help facility is available by first typing:

```
iocsh
```

2.3 Channel Access Host Example

An example host example can be generated by:

```
cd <mytop>
<base>/bin/<arch>/makeBaseApp.pl -t caClient caClient
make
```

(or `gnumake`, as required by your operating system)

Two channel access examples are provided:

```
caExample
```

This example program expects a `pvname` argument, connects and reads the current value for the `pv`, displays the result and terminates. To run this example just type.

```
<mytop>/bin/<hostarch>/caExample <pvname> where
```

- `<mytop>` is the full path name to your application top directory.
- `<hostarch>` is your host architecture.
- `<pvname>` is one of the record names displayed by the `dbl` ioc shell command.

caMonitor

This example program expects a filename argument which contains a list of pvnames, each appearing on a separate line. It connects to each pv and issues monitor requests. It displays messages for all channel access events, connection events, etc.

2.4 iocsh

Because the vxWorks shell is only available on vxWorks, EPICS base provides iocsh. In the main program it can be invoked as follows:

```
iocsh("filename")
```

or

```
iocsh(0)
```

If the argument is a filename, the commands in the file are executed and iocsh returns. If the argument is 0 then iocsh goes into interactive mode, i.e. it prompts for and executes commands until an exit command is issued.

This shell is described in more detail in Chapter 18, “IOC Shell” on page 249 [FIXPAGEEREF](#).

On vxWorks iocsh is not automatically started. It can be started by just giving the following command to the vxWorks shell.

```
iocsh
```

To get back to the vxWorks shell just say

```
exit
```

2.5 Building IOC components

Detailed build rules are given in chapter “Epics Build Facility”. This section describes methods for building most components needed for IOC applications. It uses excerpts from the `myexampleApp/src/Makefile` that is generated by `makeBaseApp`.

The following two types of applications can be built:

- Support applications

These are applications meant for use by ioc applications. The rules described here install things into one of the following directories that are created just below `<top>`:

`include`

C include files are installed here. Either header files supplied by the application or header files generated from `xxxRecord.dbd` or `xxxMenu.dbd` files.

`dbd`

Each file contains some combination of `include`, `recordtype`, `device`, `driver`, and `registrar` database definition commands. The following are installed:

- `xxxRecord.dbd` and `xxxMenu.dbd` files
- An arbitrary `xxx.dbd` file
- ioc applications install a file `yyy.dbd` generated from file `yyyInclude.dbd`.

`db`

Files containing record instance definitions.


```
lib/<arch>
```

All source modules are compiled and placed in shared or static library (win32 dll)

- IOC applications

These are applications loaded into actual IOCs.

2.5.1 Binding to IOC components

Because many IOC components are bound only during ioc initialization, some method of linking to the appropriate shared and/or static libraries must be provided. The method used for IOCs is to generate, from an `xxxInclude.dbd` file, a C++ program that contains references to the appropriate library modules. The following database definitions keywords are used for this purpose:

```
recordtype
device
driver
function
variable
registrar
```

The method also requires that IOC components contain an appropriate `epicsExport` statement. All components must contain the statement:

```
#include <epicsExport.h>
```

Any component that defines any exported functions must also contain:

```
#include <registryFunction.h>
```

Each record support module must contain a statement like:

```
epicsExportAddress (rset, xxxRSET);
```

Each device support module must contain a statement like:

```
epicsExportAddress (dset, devXxxSoft);
```

Each driver support module must contain a statement like:

```
epicsExportAddress (drvet, drvXxx);
```

Functions are registered using an `epicsRegisterFunction` macro in the C source file containing the function, along with a `function` statement in the application database description file. The `makeBaseApp` example thus contains the following statements to register a pair of functions for use with a subroutine record:

```
epicsRegisterFunction (mySubInit);
epicsRegisterFunction (mySubProcess);
```

The database definition keyword `variable` forces a reference to an integer or double variable, e.g. debugging variables. The `xxxInclude.dbd` file can contain definitions like:

```
variable (asCaDebug, int)
variable (myDefaultTimeout, double)
```

The code that defines the variables must include code like:

```
int asCaDebug = 0;
epicsExportAddress (int, asCaDebug);
```

The keyword `registrar` signifies that the epics component supplies a named registrar function that has the prototype:

```
typedef void (*REGISTRAR)(void);
```

This function normally registers things, as described in Chapter 21, “Registry” on page 319. The `makeBaseApp` example provides a sample `iocsh` command which is registered with the following registrar function:

```
static void helloRegister(void) {
    iocshRegister(&helloFuncDef, helloCallFunc);
}
epicsExportRegistrar(helloRegister);
```

2.5.2 Makefile rules

2.5.2.1 Building a support application.

```
# xxxRecord.h will be created from xxxRecord.dbd
DBDINC += xxxRecord
DBD += myexampleSupport.dbd

LIBRARY_IOC += myexampleSupport

myexampleSupport_SRCS += xxxRecord.c
myexampleSupport_SRCS += devXxxSoft.c
myexampleSupport_SRCS += dbSubExample.c

myexampleSupport_LIBS += $(EPICS_BASE_IOC_LIBS)
```

The `DBDINC` rule looks for a file `xxxRecord.dbd`. From this file a file `xxxRecord.h` is created and installed into `<top>/include`

The `DBD` rule finds `myexampleSupport.dbd` in the source directory and installs it into `<top>/dbd`

The `LIBRARY_IOC` statement states that a shared/static library should be created and installed into `<top>/lib/<arch>`.

The `myexampleSupport_SRCS` statements name all the source files that are compiled and put into the library.

The above statements are all that is needed for building many support applications.

2.5.2.2 Building the IOC application

The following statements build the IOC application:

```
PROD_IOC = myexample

DBD += myexample.dbd

# myexample.dbd will be made up from these files:
myexample_DBD += base.dbd
myexample_DBD += xxxSupport.dbd
myexample_DBD += dbSubExample.dbd

# <name>_registerRecordDeviceDriver.cpp will be created from <name>.dbd
myexample_SRCS += myexample_registerRecordDeviceDriver.cpp
myexample_SRCS_DEFAULT += myexampleMain.cpp
myexample_SRCS_vxWorks += -nil-

# Add locally compiled object code
```

```

myexample_SRCS += dbSubExample.c

#The following adds support from base/src/vxWorks
myexample_OBJS_vxWorks += $(EPICS_BASE_BIN)/vxComLibrary

myexample_LIBS += myexampleSupport
myexample_LIBS += $(EPICS_BASE_IOC_LIBS)

```

PROD_IOC sets the name of the ioc application, here called myexample.

The DBD definition myexample.dbd will cause build rules to create the database definition include file myexampleInclude.dbd from files in the myexample_DBD definition. For each filename in that definition, the created myexampleInclude.dbd will contain an include statement for that filename. In this case the created myexampleInclude.dbd file will contain the following lines.

```

include "base.dbd"
include "xxxSupport.dbd"
include "dbSubExample.dbd"

```

When the DBD build rules find the created file myexampleInclude.dbd, the rules then call dbExpand which reads myexampleInclude.dbd to generate file myexample.dbd, and install it into <top>/dbd.

An arbitrary number of myexample_SRCS statements can be given. Names of the form <name>_registerRecordDeviceDriver.cpp, are special; when they are seen the perl script registerRecordDeviceDriver.pl is executed and given <name>.dbd as input. This script generates the <name>_registerRecordDeviceDriver.cpp file automatically.

2.6 makeBaseApp.pl

makeBaseApp.pl is a perl script that creates application areas. It can create the following:

- <top>/Makefile
- <top>/configure – This directory contains the files needed by the EPICS build system.
- <top>/xxxApp – A set of directories and associated files for a major sub-module.
- <top>/iocBoot – A subdirectory and associated files.
- <top>/iocBoot/iocxxx – A subdirectory and files for a single ioc.

makeBaseApp.pl creates directories and then copies template files into the newly created directories while expanding macros in the template files. EPICS base provides two sets of template files: simple and example. These are meant for simple applications. Each site, however, can create its own set of template files which may provide additional functionality. This section describes the functionality of makeBaseApp itself, the next section provides details about the simple and example templates.

2.6.1 Usage

makeBaseApp has four possible forms of command line:

```
<base>/bin/<arch>/makeBaseApp.pl -h
```

Provides help.

```
<base>/bin/<arch>/makeBaseApp.pl -l [options]
```

List the application templates available. This invocation does not alter the current directory.

```
<base>/bin/<arch>/makeBaseApp.pl [-t type] [options] app ...
```

Create application directories.

```
<base>/bin/<arch>/makeBaseApp.pl -i -t type [options] ioc ...
```

Create ioc boot directories.

Options for all command forms:

`-b base`

Provides the full path to EPICS base. If not specified, the value is taken from the EPICS_BASE entry in config/RELEASE. If the config directory does not exist, the path is taken from the command-line that was used to invoke makeBaseApp

`-T template`

Set the template top directory (where the application templates are). If not specified, the template path is taken from the TEMPLATE_TOP entry in config/RELEASE. If the config directory does not exist the path is taken from the environment variable EPICS_MBA_TEMPLATE_TOP, or if this is not set the templates from EPICS base are used.

`-d`

Verbose output (useful for debugging)

Arguments unique to makeBaseApp.pl [-t type] [options] app ...:

`app`

One or more application names (the created directories will have “App” appended to this name)

`-t type`

Set the template type (use the `-l` invocation to get a list of valid types). If this option is not used, type is taken from the environment variable EPICS_MBA_DEF_APP_TYPE, or if that is not set the values “default” and then “example” are tried.

Arguments unique to makeBaseApp.pl -i [options] ioc ...:

`ioc`

One or more IOC names (the created directories will have “ioc” prepended to this name).

`-a arch`

Set the IOC architecture (e.g. vxWorks-68040). If `-a arch` is not specified, you will be prompted.

2.6.2 Environment Variables:

EPICS_MBA_DEF_APP_TYPE

Application type you want to use as default

EPICS_MBA_TEMPLATE_TOP

Template top directory

2.6.3 Description

To create a new <top> issue the commands:

```
mkdir <top>
cd <top>
<base>/bin/<arch>/makeBaseApp.pl -t <type> <app> ...
<base>/bin/<arch>/makeBaseApp.pl -i -t <type> <ioc> ...
```

makeBaseApp does the following:

- EPICS_BASE is located by checking the following in order:
 - If the `-b` option is specified its value is used.
 - If a `<top>/configure/RELEASE` file exists and defines a value for EPICS_BASE it is used.
 - It is obtained from the invocation of `makeBaseApp`. For this to work, the full path name to the `makeBaseApp.pl` script in the EPICS base release you are using must be given.
- TEMPLATE_TOP is located in a similar fashion:
 - If the `-T` option is specified its value is used.
 - If a `<top>/configure/RELEASE` file exists and defines a value for TEMPLATE_TOP it is used.
 - If EPICS_MBA_TEMPLATE_TOP is defined its value is used.
 - It is set equal to `<epics_base>/templates/makeBaseApp/top`
- If `-l` is specified the list of application types is listed and `makeBaseApp` terminates.
- If `-i` is specified and `-a` is not then the user is prompted for the IOC architecture.
- The application type is determined by checking the following in order:
 - If `-t` is specified it is used.
 - If EPICS_MBA_DEF_APP_TYPE is defined its value is used.
 - If a template `defaultApp` exists, the application type is set equal to default.
 - If a template `exampleApp` exists, the application type is set equal to example.
- If the application type is not found in TEMPLATE_TOP, `makeBaseApp` issues an error and terminates.
- If `Makefile` does not exist, it is created.
- If directory `configure` does not exist, it is created and populated with all the `configure` files.
- If `-i` is specified:
 - If directory `iocBoot` does not exist, it is created and the files from the template boot directory are copied into it.
 - For each `<ioc>` specified on the command line a directory `iocBoot/ioc<ioc>` is created and populated with the files from the template (with `ReplaceLine()` tag replacement, see below).
- If `-i` is NOT specified:
 - For each `<app>` specified on the command line a directory `<app>App` is created and populated with the directory tree from the template (with `ReplaceLine()` tag replacement, see below).

2.6.4 Tag Replacement within a Template

When copying certain files from the template to the new application structure, `makeBaseApp` replaces some predefined tags in the name or text of the files concerned with values that are known at the time. An application template can extend this functionality as follows:

- Two perl subroutines are defined within `makeBaseApp`:

```
ReplaceFilename
```

This substitutes for the following in names of any file taken from the templates.

```
__APPNAME__  
__APPTYPE__
```

ReplaceLine

This substitutes for the following in each line of each file taken from the templates:

```

    _USER_
    _EPICS_BASE_
    _ARCH_
    _APPNAME_
    _APPTYPE_
    _TEMPLATE_TOP_
    _IOC_

```

- If the application type directory has a file named `Replace.pl`, it can:
 - Replace one or both of the above subroutines with its own versions.
 - Add a subroutine `ReplaceFilenameHook($file)` which is called at the end of `ReplaceFilename`.
 - Add a subroutine `ReplaceLineHook($line)` which is called at the end of `ReplaceLine`.
 - Include other code which is run after the command line options are interpreted.

2.6.5 makeBaseApp templates provided with base

2.6.5.1 support

This creates files appropriate for building a support application.

2.6.5.2 ioc

Without the `-i` option, this creates files appropriate for building an ioc application. With the `-i` option it creates an ioc boot directory.

2.6.5.3 example

Without the `-i` option it creates files for running an example. Both a support and an ioc application are built. With the `-i` option it creates an ioc boot directory that can be used to run the example.

2.6.5.4 caClient

This builds two Channel Access clients.

2.6.5.5 caServer

This builds an example Portable Access Server.

2.7 vxWorks boot parameters

The vxWorks boot parameters are set via the console serial port on your IOC. Life is much easier if you can connect the console to a terminal window on your workstation. On Linux the ‘screen’ program lets you communicate through a local serial port; run `screen /dev/ttyS0` if the IOC is connected to `ttyS0`.

The vxWorks boot parameters look something like the following:

```

boot device           : xxx
processor number     : 0
host name           : xxx
file name           : <full path to board support>/vxWorks
inet on ethernet (e) : xxx.xxx.xxx.xxx:<netmask>
host inet (h)       : xxx.xxx.xxx.xxx
user (u)            : xxx
ftp password (pw)   : xxx
flags (f)           : 0x0
target name (tn)    : <hostname for this inet address>
startup script (s)  : <top>/iocBoot/iocmyexample/st.cmd

```

The actual values for each field are site and IOC dependent. Two fields that you can change at will are the vxWorks boot image and the location of the startup script.

Note that the full path name for the correct board support boot image must be specified. If bootp is used the same information will need to be placed in the bootp host's configuration database instead.

When your boot parameters are set properly, just press the reset button on your IOC, or use the @ command to commence booting. You will find it VERY convenient to have the console port of the IOC attached to a scrolling window on your workstation.

2.8 RTEMS boot procedure

RTEMS uses the vendor-supplied bootstrap mechanism so the method for booting an IOC depends upon the hardware in use.

2.8.1 Booting from a BOOTP/DHCP/TFTP server

Many boards can use BOOTP/DHCP to read their network configuration and then use TFTP to read the application program. RTEMS can then use TFTP or NFS to read startup scripts and configuration files. If you are using TFTP to read the startup scripts and configuration files you must install the EPICS application files on your TFTP server as follows:

- Copy all db/xxx files to <tftpbasedir>/epics/<target_hostname\>/db/xxx.
- Copy all dbd/xxx files to <tftpbasedir>/epics/<target_hostname>/dbd/xxx.
- Copy the st.cmd script to <tftpbasedir>/epics/<target_hostname>/st.cmd.

Use DHCP site-specific option 129 to specify the path to the IOC startup script.

2.8.2 Motorola PPCBUG boot parameters

Motorola single-board computers which employ PPCBUG should have their 'NIOT' parameters set up like:

```

Controller LUN =00
Device LUN     =00
Node Control Memory Address =FFE10000
Client IP Address      ='Dotted-decimal' IP address of IOC
Server IP Address      ='Dotted-decimal' IP address of TFTP/NFS server
Subnet IP Address Mask ='Dotted-decimal' IP address of subnet mask (255.255.255.0 for class C subnet)
Broadcast IP Address   ='Dotted-decimal' IP address of subnet broadcast address
Gateway IP Address     ='Dotted-decimal' IP address of network gateway (0.0.0.0 if none)

```

```

Boot File Name           =Path to application bootable image (.../bin/RTEMS-mvme2100/test.boot)
Argument File Name      =Path to application startup script (.../iocBoot/iocTest/st.cmd)
Boot File Load Address  =001F0000 (actual value depends on BSP)
Boot File Execution Address =001F0000 (actual value depends on BSP)
Boot File Execution Delay =00000000
Boot File Length        =00000000
Boot File Byte Offset   =00000000
BOOTP/RARP Request Retry =00
TFTP/ARP Request Retry  =00
Trace Character Buffer Address =00000000

```

2.8.3 Motorola MOTLOAD boot parameters

Motrola single-board computers which employ MOTLOAD should have their network ‘Global Environment Variable’ parameters set up like:

```

mot-/dev/enet0-cipa='Dotted-decimal' IP address of IOC
mot-/dev/enet0-sipa='Dotted-decimal' IP address of TFTP/NFS server
mot-/dev/enet0-snma='Dotted-decimal' IP address of subnet mask (255.255.255.0 for class C subnet)
mot-/dev/enet0-gipa='Dotted-decimal' IP address of network gateway (omit if none)
mot-/dev/enet0-file=Path to application bootable image (.../bin/RTEMS-mvme5500/test.boot)
rtems-client-name=IOC name (mot-/dev/enet0-cipa will be used if this parameter is missing)
rtems-dns-server='Dotted-decimal' IP address of domain name server (omit if none)
rtems-dns-domainname=Domain name (if this parameter is omitted the compiled-in value will be used)
epics-script=Path to application startup script (.../iocBoot/iocTest/st.cmd)

```

The `mot-script-boot` parameter should be set up like:

```

tftpGet -a4000000 -cxxx -sxxx -mxxx -gxxx -d/dev/enet0
        -f....bin/RTEMS-mvme5500/test.boot
netShut
go -a4000000

```

where the `-c`, `-s`, `-m` and `-g` values should match the `cipa`, `sipa`, `snma` and `gipa` values, respectively and the `-f` value should match the file value.

2.8.4 RTEMS NFS access

For IOCs which use NFS for remote file access the EPICS initialization code uses the startup script pathname to determine the parameters for the initial NFS mount. If the startup script pathname begins with a `/` the first component of the pathname is used as both the server path and the local mount point. If the startup script pathname does not begin with a `/` the first component of the pathname is used as the local mount point and the server path is `"/tftpboot/"` followed by the first component of the pathname. This allows the NFS client used for EPICS file access and the TFTP client used for bootstrapping the application to have a similar view of the remote filesystem.

2.8.5 RTEMS ‘Cexp’

The RTEMS ‘Cexp’ add-on package provides the ability to load object modules at application run-time. If your RTEMS build includes this package you can load RTEMS IOC applications in the same fashion as vxWorks IOC applications.

Chapter 3


EPICS Overview

3.1 What is EPICS?

The Experimental Physics and Industrial Control System (EPICS) consists of a set of software components and tools that Application Developers can use to create control systems. The basic components are:

- **OPI:** Operator Interface. This is a workstation which can run various EPICS tools.
- **IOC:** Input/Output Controller. Any platform that can support EPICS run time databases together with the other software components described in the manual. One example is a workstation. Another example is a VME/VXI based system using vxWorks or RTEMS as the realtime operating system.
- **LAN:** Local Area Network. This is the communication network which allows the IOCs and OPIs to communicate. EPICS provides a software component, Channel Access, which provides network transparent communication between a Channel Access client and an arbitrary number of Channel Access servers.

A control system implemented via EPICS has the following physical structure.



The rest of this chapter gives a brief description of EPICS:

- **Basic Attributes:** A few basic attributes of EPICS.
- **Platforms:** The vendor supplied Hardware and Software platforms EPICS supports.
- **IOC Software:** EPICS supplied IOC software components.
- **Channel Access:** EPICS software that supports network independent access to IOC databases.
- **OPI Tools:** EPICS supplied OPI based tools.
- **EPICS Core:** A list of the EPICS core software, i.e. the software components without which EPICS will not work.


3.2 Basic Attributes

The basic attributes of EPICS are:

- **Tool Based:** EPICS provides a number of tools for creating a control system. This minimizes the need for custom coding and helps ensure uniform operator interfaces.
- **Distributed:** An arbitrary number of IOCs and OPIs can be supported. As long as the network is not saturated, no single bottle neck is present. A distributed system scales nicely. If a single IOC becomes saturated, its functions can be spread over several IOCs. Rather than running all applications on a single host, the applications can be spread over many OPIs.
- **Event Driven:** The EPICS software components are all designed to be event driven to the maximum extent possible. For example, rather than having to poll IOCs for changes, a Channel Access client can request that it be notified when a change occurs. This design leads to efficient use of resources, as well as, quick response times.
- **High Performance:** A SPARC based workstation can handle several thousand screen updates a second with each update resulting from a Channel Access event. A 68040 IOC can process more than 6,000 records per second, including generation of Channel Access events.

3.3 IOC Software Components

An IOC contains the following EPICS supplied software components.



- **IOC Database:** The memory resident database plus associated data structures.
- **Database Access:** Database access routines. With the exception of record and device support, all access to the database is via the database access routines.
- **Scanners:** The mechanism for deciding when records should be processed.
- **Record Support:** Each record type has an associated set of record support routines.
- **Device Support:** Each record type can have one or more sets of device support routines.

- **Device Drivers:** Device drivers access external devices. A driver may have an associated driver interrupt routine.
- **Channel Access:** The interface between the external world and the IOC. It provides a network independent interface to database access.
- **Monitors:** Database monitors are invoked when database field values change.
- **Sequencer:** A finite state machine.

Let's briefly describe the major components of the IOC and how they interact.

3.3.1 IOC Database

The heart of each IOC is a memory resident database together with various memory resident structures describing the contents of the database. EPICS supports a large and extensible set of record types, e.g. `ai` (Analog Input), `ao` (Analog Output), etc.

Each record type has a fixed set of fields. Some fields are common to all record types and others are specific to particular record types. Every record has a record name and every field has a field name. The first field of every database record holds the record name, which must be unique across all IOCs that are attached to the same TCP/IP subnet.

Data structures are provided so that the database can be accessed efficiently. Most software components, because they access the database via database access routines, do not need to be aware of these structures.

3.3.2 Database Access

With the exception of record and device support, all access to the database is via the channel or database access routines. See Chapter ?? for details.

3.3.3 Database Scanning

Database scanning is the mechanism for deciding when to process a record. Five types of scanning are possible: Periodic, Event, I/O Event, Passive and Scan Once.

- **Periodic:** A request can be made to process a record periodically. A number of time intervals are supported.
- **Event:** Event scanning is based on the posting of an event by any IOC software component. The actual subroutine call is:
`post_event (event_num)`
- **I/O Event:** The I/O event scanning system processes records based on external interrupts. An IOC device driver interrupt routine must be available to accept the external interrupts.
- **Passive:** Passive records are processed as a result of linked records being processed or as a result of external changes such as Channel Access puts.
- **Scan Once:** In order to provide for caching puts, The scanning system provides a routine `scanOnce` which arranges for a record to be processed one time.

3.3.4 Record Support, Device Support and Device Drivers

Database access needs no record-type specific knowledge, because each record-type has its associated record support module. Therefore, database access can support any number and type of records. Similarly, record support contains no device specific knowledge, giving each record type the ability to have any number of independent device support

modules. If the method of accessing the piece of hardware is more complicated than what can be handled by device support, then a device driver can be developed.

Record types *not* associated with hardware do not have device support or device drivers.

The IOC software is designed so that the database access layer knows nothing about the record support layer other than how to call it. The record support layer in turn knows nothing about its device support layer other than how to call it. Similarly the only thing a device support layer knows about its associated driver is how to call it. This design allows a particular installation and even a particular IOC within an installation to choose a unique set of record types, device types, and drivers. The remainder of the IOC system software is unaffected.

Because an Application Developer can develop record support, device support, and device drivers, these topics are discussed in greater detail in later chapters.

Every record support module must provide a record processing routine to be called by the database scanners. Record processing consists of some combination of the following functions (particular records types may not need all functions):

- **Input:** Read inputs. Inputs can be obtained, via device support routines, from hardware, from other database records via database links, or from other IOCs via Channel Access links.
- **Conversion:** Conversion of raw input to engineering units or engineering units to raw output values.
- **Output:** Write outputs. Output can be directed, via device support routines, to hardware, to other database records via database links, or to other IOCs via Channel Access links.
- **Raise Alarms:** Check for and raise alarms.
- **Monitor:** Trigger monitors related to Channel Access callbacks.
- **Link:** Trigger processing of linked records.

3.3.5 Channel Access

Channel Access is discussed in the next section.

3.3.6 Database Monitors

Database monitors provide a callback mechanism for database value changes. This allows the caller to be notified when database values change without constantly polling the database. A mask can be set to specify value changes, alarm changes, and/or archival changes.

At the present time only Channel Access uses database monitors. No other software should use the database monitors. The monitor routines will not be described because they are of interest only to Channel Access.

3.4 Channel Access

Channel Access provides network transparent access to IOC databases. It is based on a client/ server model. Each IOC provides a Channel Access server which is willing to establish communication with an arbitrary number of clients. Channel Access client services are available on both OPIs and IOCs. A client can communicate with an arbitrary number of servers.

3.4.1 Client Services

The basic Channel Access client services are:

- **Search:** Locate the IOCs containing selected process variables and establish communication with each one.
- **Get:** Get value plus additional optional information for a selected set of process variables.
- **Put:** Change the values of selected process variables.
- **Add Event:** Add a change of state callback. This is a request to have the server send information only when the associated process variable changes state. Any combination of the following state changes can be requested: change of value, change of alarm status and/or severity, and change of archival value. Many record types provide hysteresis factors for value changes.

In addition to requesting process variable values, any combination of the following additional information may be requested:

- **Status:** Alarm status and severity.
- **Units:** Engineering units for this process variable.
- **Precision:** Precision with which to display floating point numbers.
- **Time:** Time when the record was last processed.
- **Enumerated:** A set of ASCII strings defining the meaning of enumerated values.
- **Graphics:** High and low limits for producing graphs.
- **Control:** High and low control limits.
- **Alarm:** The alarm HIHI, HIGH, LOW, and LOLO values for the process variable.

It should be noted that Channel Access does not provide access to database records as records. This is a deliberate design decision. This allows new record types to be added without impacting any software that accesses the database via Channel Access, and it allows a Channel Access client to communicate with multiple IOCs having differing sets of record types.

3.4.2 Search Server

Channel Access provides an IOC resident server which waits for Channel Access search messages. These are generated when a Channel Access client (for example when an Operator Interface task starts) searches for the IOCs containing process variables the client uses. This server accepts all search messages, checks to see if any of the process variables are located in this IOC, and, if any are found, replies to the sender with an “I have it” message.

3.4.3 Connection Request Server

Once the process variables have been located, the Channel Access client issues connection requests for each IOC containing process variables the client uses. The connection request server, in the IOC, accepts the request and establishes a connection to the client. Each connection is managed by two separate tasks: `ca_get` and `ca_put`. The `ca_get` and `ca_put` requests map to `dbGetField` and `dbPutField` database access requests. `ca_add_event` requests result in database monitors being established. Database access and/or record support routines trigger the monitors via a call to `db_post_event`.

3.4.4 Connection Management

Each IOC provides a connection management service. When a Channel Access server fails (e.g. its IOC crashes) the client is notified and when a client fails (e.g. its task crashes) the server is notified. When a client fails, the server breaks the connection. When a server crashes, the client automatically re-establishes communication when the server restarts.

3.5 OPI Tools

EPICS provides a number of OPI based tools. These can be divided into two groups based on whether or not they use Channel Access. Channel Access tools are real time tools, i.e. they are used to monitor and control IOCs.

3.5.1 Examples of Channel Access Tools

A large number of Channel Access tools have been developed. The following are some representative examples.

- **CSS**: Control System Studio, an Eclipse RCP application with many available plug-ins.
- **EDM**: Extensible Display Manager.
- **MEDM**: Motif Editor and Display Manager.
- **StripTool**: A general-purpose stripchart program.
- **ALH**: Alarm Handler. General purpose alarm handler driven by an alarm configuration file.
- **Sequencer**: Runs in an IOC and emulates a finite state machine.
- **Probe**: Allows the user to monitor and/or change a single process variable specified at run time.

3.5.2 Examples of other Tools

- **VDCT**: A Java based database configuration tool which is quickly becoming the recommended database configuration tool.
- **SNC**: State Notation Compiler. It generates a C program that represents the states for the IOC Sequencer tool.

3.6 EPICS Core Software

EPICS consists of a set of core software and a set of optional components. The core software, i.e. the components of EPICS without which EPICS would not function, are:

- Channel Access - Client and Server software
- IOC Database
- Scanners
- Monitors
- Database Definition Tools
- Source/Release

All other software components are optional. Of course, most applications will need equivalent functionality to MEDM (or EDD/DM). Likewise an application developer would not start from scratch developing record and device support. Most OPI tools do not, however, have to be used. Likewise any given record support module, device support module, or driver could be deleted from a particular IOC and EPICS will still function.

Chapter 4

Build Facility

Janet Anderson is the author of this chapter.

4.1 Overview

This chapter describes the EPICS build facility including directory structure, environment and system requirements, configuration files, Makefiles, and related build tools.

4.1.1 <top> Directory structure

EPICS software can be divided into multiple <top> areas. Examples of <top> areas are EPICS base itself, EPICS extensions, and simple or complicated IOC applications. Each <top> may be maintained separately. Different <top> areas can be on different releases of external software such as EPICS base releases.

A <top> directory has the following directory structure:

```
<top>/
  Makefile
  configure/
  dir1/
  dir2/
  ...
```

where `configure` is a directory containing build configuration files and a `Makefile`, where `dir1`, `dir2`, ... are user created subdirectory trees with `Makefiles` and source files to be built. Because the build rules allow make commands like “`make install.vxWorks-68040`”, subdirectory names within a <top> directory structure may not contain a period “.” character.

4.1.2 Install Directories

Files generated during the build are installed into subdirectories of an installation directory which defaults to `$(TOP)`, the <top> directory. For base, extensions, and IOC applications, the default value can be changed in the `configure/CONFIG_SITE` file. The installation directory for the EPICS components is controlled by the definition of the `INSTALL_LOCATION` variable.

Due to a side-effect of the build rules, the *parent* of the installation directory (`$(INSTALL_LOCATION) / . .`) should not contain directories with the same names as the subdirectories listed below.

The following subdirectories may exist in the installation directory. They are created by the build and contain the installed build components.

- `dbd` – A directory into which Database Definition files are installed.
- `include` – A directory into which C and C++ header files are installed. These header files may be generated from menu and record type definitions.
- `bin` – This directory contains a subdirectory for each host and target architecture. These are the directories into which executables, binaries, etc. are installed.
- `lib` – This directory contains a subdirectory for each host and target architecture. These are the directories into which libraries are installed.
- `db` – A directory into which database record instance, template, and substitution files are installed.
- `html` – A directory sub-tree into which html documentation is installed.
- `templates` – A directory sub-tree into which template files are installed.
- `configure` – If the `INSTALL_LOCATION` variable has been explicitly set so it does not equal `TOP`, the configure files are copied from `$(TOP)/configure`.
- `cfg` – A directory into which user created configure files are installed.

4.1.3 Elements of build system

The main ingredients of the build system are:

- A set of configuration files and tools provided in the EPICS `base/configure` directory
- A corresponding set of configuration files in the `<top>/configure` directory of a non-base `<top>` directory structure to be built. The `makeBaseApp.pl` and `makeBaseExt.pl` scripts create these configuration files. Many of these files just include a file of the same name from the `base/configure` directory.
- Makefiles in each directory of the `<top>` directory structure to be built
- User created configuration files in build created `$(INSTALL_LOCATION)/cfg` directories.

4.1.4 Features

The principal features of the build system are:

- Requires a single `Makefile` in each directory of a `<top>` directory structure
- Supports both host or vendor's native compiler and GNU compiler
- Supports building multiple types of software (libraries, executables, databases, java class files, etc.) stored in a single directory tree.
- Supports building EPICS base, extensions, and IOC applications.
- Supports multiple host and target operating system + architecture combinations.
- Allows builds for all hosts and targets within a single `<top>` source directory tree.
- Allows sharing of components such as special record/device/drivers across `<top>` areas.
- `gnumake` is the only command used to build a `<top>` area.

4.1.5 Multiple host and target systems

You can build on multiple host systems and for multiple cross target systems using a single EPICS directory structure. The intermediate and binary files generated by the build will be created in separate `O.*` subdirectories and installed into the appropriate separate host or target install directories. EPICS executables and scripts are installed into the `$(INSTALL_LOCATION)/bin/<arch>` directories. Libraries are installed into `$(INSTALL_LOCATION)/lib/<arch>`. The default definition for `$(INSTALL_LOCATION)` is `$(TOP)` which is the root directory in the directory structure. Architecture dependant created files (e.g. object files) are stored in `O.<arch>` source subdirectories, and architecture independent created files are stored in `O.Common` source subdirectories. This allows objects for multiple cross target architectures to be maintained at the same time.

To build EPICS base for a specific host/target combination you must have the proper host/target `c/c++` cross compiler and target header files, `CROSS_COMPILER_HOST_ARCHS` must empty or include the host architecture in its list value, the `CROSS_COMPILER_TARGET_ARCHS` variable must include the target to be cross-compiled, and the `base/configure/os` directory must have the appropriate configure files.

4.2 Build Requirements

4.2.1 Host Environment Variable

Only one environment variable, `EPICS_HOST_ARCH`, is required to build EPICS `<top>` areas. This variable should be set to be your workstation's operating system - architecture combination to use the os vendor's `c/c++` compiler for native builds or set to the operating system - architecture - alternate compiler combination to use an alternate compiler for native builds if an alternate compiler is supported on your system. The filenames of the `CONFIG.*.Common` files in `base/configure/os` show the currently supported `EPICS_HOST_ARCH` values. Examples are `solaris-sparc`, `solaris-sparc-gnu`, `linux-x86`, `win32-x86`, and `cygwin-x86`.

4.2.2 Software Prerequisites

Before you can build EPICS components your host system must have the following software installed:

- Perl version 5.8 or greater
- GNU make, version 3.81 or greater
- C++ compiler (host operating system vendor's compiler or GNU compiler)

If you will be building EPICS components for vxWorks targets you will also need:

- Tornado II or vxWorks 6.x and one or more board support packages. Consult the vxWorks documentation for details.

If you will be building EPICS components for RTEMS targets you will also need:

- RTEMS development tools and libraries required to run EPICS IOC applications.

4.2.3 Path requirements

You must have the perl executable in your path and you may need C and C++ compilers in your search path. Check definitions of `CC` and `CCC` in `base/configure/os/CONFIG.<host>.<host>` or the definitions for `GCC` and `G++` if `ANSI=GCC` and `CPLUSPLUS=GCC` are specified in `CONFIG_SITE`. For building base you also must have `echo` in your search path. You can override the default settings by defining `PERL`, `CC` and `CCC`, `GCC` and `G++`, `GNU_DIR` ... in the appropriate file (usually `configure/os/CONFIG_SITE.$EPICS_HOST_ARCH.Common`)

4.2.3.1 Unix path

For Unix host builds you also need touch, cpp, cp, rm, mv, and mkdir in your search path and /bin/chmod must exist. On some Unix systems you may also need ar and ranlib in your path, and the c compiler may require ld in your path.

4.2.3.2 Win32 PATH

On WIN32 systems, building shared libraries is the default setting and you will need to add fullpathname to \$ (INSTALL_LOCATION) / to your path so the shared libraries, dlls, can be found during the build.. Building shared libraries is determined by the value of the macro SHARED_LIBRARIES in CONFIG_SITE or os/CONFIG.Common.<host> (either YES or NO).

4.2.4 Directory names

Because the build rules allow make commands like “make <dir>.<action>, <arch>”, subdirectory names within a <top> directory structure may not contain a period.” character.

4.2.5 EPICS_HOST_ARCH environment variable

The startup directory in EPICS base contains a perl script, EpicsHostArch.pl, which can be used to define EPICS_HOST_ARCH. This script can be invoked with a command line parameter defining the alternate compiler (e.g. if invoking EpicsHostArch.pl yields solaris-sparc, then invoking EpicsHostArch.pl gnu will yield solaris-sparc-gnu).

The startup directory also contains scripts to help users set the path and other environment variables.

4.3 Configuration Definitions

4.3.1 Site-specific EPICS Base Configuration

4.3.1.1 Site configuration

To configure EPICS base for your site, you may want to modify the default definitions in the following files:

```
configure/CONFIG_SITE Build choices. Specify target archs.
configure/CONFIG_SITE_ENV Environment variable defaults
```

4.3.1.2 Host configuration

To configure each host system for your site, you may override the default definitions in the configure/os directory by adding a new file with override definitions. The new file should have the same name as the distribution file to be overridden except CONFIG in the name is changed to CONFIG_SITE.

```
configure/os/CONFIG_SITE.<host>.<host> - Host build settings
configure/os/CONFIG_SITE.<host>.Common - Host build settings for all target systems
```

4.3.1.3 Target configuration

To configure each target system, you may override the default definitions in the `configure/os` directory by adding a new file with override definitions. The new file should have the same name as the distribution file to be overridden except `CONFIG` in the name is replaced by `CONFIG_SITE`.

`configure/os/CONFIG_SITE.Common.<target>` - Target cross settings

`configure/os/CONFIG_SITE.<host>.<target>` - Host-target settings

`configure/os/CONFIG_SITE.Common.vxWorksCommon` - vxWorks full paths

4.3.1.4 R3.13 compatibility configuration

To configure EPICS base for building with R3.13 extensions and ioc applications, you must modify the default definitions in the `base/config/CONFIG_SITE*` files to agree with site definitions you made in `base/configure` and `base/configure/os` files. You must also modify the following tow macros in the `base/configure/CONFIG_SITE` file:

`COMPAT_TOOLS_313` - Set to YES to build R3.13 extensions with this base.

`COMPAT_313` - Set to YES to build R3.13 ioc applications and extensions with this base.

4.3.2 Directory definitions

The `configure` files contain definitions for locations in which to install various components. These are all relative to `INSTALL_LOCATION`. The default value for `INSTALL_LOCATION` is `$(TOP)`, and `$(T_A)` is the current build's target architecture. The default value for `INSTALL_LOCATION` can be overridden in the `configure/CONFIG_SITE` file.

```

INSTALL_LOCATION_LIB      = $(INSTALL_LOCATION)/lib
INSTALL_LOCATION_BIN     = $(INSTALL_LOCATION)/bin

INSTALL_HOST_BIN         = $(INSTALL_LOCATION_BIN)/$(EPICS_HOST_ARCH)
INSTALL_HOST_LIB        = $(INSTALL_LOCATION_LIB)/$(EPICS_HOST_ARCH)

INSTALL_INCLUDE          = $(INSTALL_LOCATION)/include
INSTALL_DOC              = $(INSTALL_LOCATION)/doc
INSTALL_HTML             = $(INSTALL_LOCATION)/html
INSTALL_TEMPLATES       = $(INSTALL_LOCATION)/templates
INSTALL_DBD              = $(INSTALL_LOCATION)/dbd
INSTALL_DB                = $(INSTALL_LOCATION)/db
INSTALL_CONFIG           = $(INSTALL_LOCATION)/configure
INSTALL_JAVA             = $(INSTALL_LOCATION)/javalib

INSTALL_LIB              = $(INSTALL_LOCATION_LIB)/$(T_A)
INSTALL_SHRLIB           = $(INSTALL_LOCATION_LIB)/$(T_A)
INSTALL_TCLLIB           = $(INSTALL_LOCATION_LIB)/$(T_A)
INSTALL_BIN              = $(INSTALL_LOCATION_BIN)/$(T_A)

```

4.3.3 Extension and Application Specific Configuration

The `base/configure` directory contains files with the default build definitions and site specific build definitions. The `extensions/configure` directory contains extension specific build definitions (e.g. location of X11 and Motif libraries) and “`include <filename>`” lines for the `base/configure` files. Likewise, the

<application>/configure directory contains application specific build definitions and includes for the application source files. Build definitions such as

CROSS_COMPILER_TARGET_ARCHS can be overridden in an extension or application by placing an override definition in the <top>/configure/CONFIG_SITE file.

4.3.4 RELEASE file

Every <top>/configure directory contains a RELEASE file. RELEASE contains a user specified list of other <top> directory structures containing files needed by the current <top>, and may also include other files to take those definitions from elsewhere. The macros defined in the RELEASE file (or its includes) may reference other defined macros, but cannot rely on environment variables to provide definitions.

When make is executed, macro definitions for include, bin, and library directories are automatically generated for each external <top> definition given in the RELEASE file. Also generated are include statements for any existing RULES_BUILD files, cfg/RULES* files, and cfg/CONFIG* files from each external <top> listed in the RELEASE file.

For example, if configure/RELEASE contains the definition

```
CAMAC = /home/epics/modules/bus/camac
```

then the generated macros will be:

```
CAMAC_HOST_BIN = /home/epics/modules/bus/camac/bin/$(EPICS_HOST_ARCH)
CAMAC_HOST_LIB = /home/epics/modules/bus/camac/lib/$(EPICS_HOST_ARCH)
CAMAC_BIN = /home/epics/modules/bus/camac/bin/$(T_A)
CAMAC_LIB = /home/epics/modules/bus/camac/lib/$(T_A)
RELEASE_INCLUDES += -I/home/epics/modules/bus/camac/include/os
RELEASE_INCLUDES += -I/home/epics/modules/bus/camac/include
RELEASE_DBDFLAGS += -I /home/epics/modules/bus/camac/dbd
RELEASE_DBDFLAGS += -I/home/epics/modules/bus/camac/db
RELEASE_PERL_MODULE_DIRS += /home/epics/modules/bus/camac/lib/perl
```

RELEASE_DBDFLAGS will appear on the command lines for the dbToRecordTypeH, mkmf.pl, and dbExpand tools, and RELEASE_INCLUDES will appear on compiler command lines. CAMAC_LIB and CAMAC_BIN can be used in a Makefile to define the location of needed scripts, executables, object files, libraries or other files.

Definitions in configure/RELEASE can be overridden for a specific host and target architectures by providing the appropriate file or files containing overriding definitions.

```
configure/RELEASE.<epics_host_arch>.Common
configure/RELEASE.Common.<targetarch>
configure/RELEASE.<epics_host_arch>.<targetarch>
```

For <top> directory structures created by makeBaseApp.pl, an EPICS base perl script, convertRelease.pl can perform consistency checks for the external <top> definitions in the RELEASE file and its includes as part of the <top> level build. Consistency checks are controlled by value of CHECK_RELEASE which is defined in <top>/configure/CONFIG_SITE. CHECK_RELEASE can be set to YES, NO or WARN, and if YES (the default value), consistency checks will be performed. If CHECK_RELEASE is set to WARN the build will continue even if conflicts are found.

4.3.5 Modifying configure/RELEASE* files

You should always do a gnumake clean uninstall in the <top> directory BEFORE adding, changing, or removing any definitions in the configure/RELEASE* files and then a gnumake at the top level AFTER making the changes.

The file `<top>/configure/RELEASE` contains definitions for components obtained from outside `<top>`. If you want to link to a new release of anything defined in the file do the following:

```
cd <top>
gnumake clean uninstall
edit configure/RELEASE
```

change the relevant line(s) to point to the new release

```
gnumake
```

All definitions in `<top>/configure/RELEASE` must result in complete path definitions, i.e. relative path names are not permitted. If your site could have multiple releases of base and other support `<top>` components installed at once, these path definitions should contain a release number as one of the components. However as the `RELEASE` file is read by `gnumake`, it is permissible to use macro substitutions to define these pathnames, for example:

```
SUPPORT = /usr/local/iocapps/R3.14.9
EPICS_BASE = $(SUPPORT)/base/3-14-9-asd1
```

4.3.6 OS Class specific definitions

Definitions in a Makefile will apply to the host system (the platform on which `make` is executed) and each system defined by `CROSS_COMPILER_TARGET_ARCHS`.

It is possible to limit the architectures for which a particular definition is used. Most Makefile definition names can be specified with an appended underscore “_” followed by an `osclass` name. If an `_osclass` is not specified, then the definition applies to the host and all `CROSS_COMPILER_TARGET_ARCHS` systems. If an `_osclass` is specified, then the definition applies only to systems with the specified `os class`. A Makefile definition can also have an appended `_DEFAULT` specification. If `_DEFAULT` is appended, then the Makefile definition will apply to all systems that do not have an `_osclass` specification for that definition. If a `_DEFAULT` definition exists but should not apply to a particular system OS Class, the value “-nil-” should be specified in the relevant Makefile definition.

Each system has an `OS_CLASS` definition in its `configure/os/CONFIG.Common.<arch>` file. A few examples are:

For `vxWorks-*` targets `<osclass>` is `vxWorks`.

For `RTEMS-*` targets `<osclass>` is `RTEMS`.

For `solaris-*` targets `<osclass>` is `solaris`.

For `win32-*` targets `<osclass>` is `WIN32`.

For `linux-*` targets `<osclass>` is `Linux`.

For `darwin-*` targets `<osclass>` is `Darwin`.

For `aix-*` targets `<osclass>` is `AIX`.

For example the following Makefile lines specify that product `aaa` should be created for all systems. Product `bbb` should be created for systems that do not have `OS_CLASS` defined as `solaris`.

```
PROD = aaa
PROD_solaris = -nil-
PROD_DEFAULT = bbb
```

4.3.7 Specifying T_A specific definitions

It is possible for the user to limit the systems for which a particular definition applies to specific target systems.

For example the following Makefile lines specify that product aaa should be created for all target architecture which allow IOC type products and product bbb should be created only for the vxWorks-68040 and vxWorks-ppc603 targets. Remember T_A is the build's current target architecture. so PROD_IOC has the bbb value only when the current built target architecture is vxWorks-68040 or vxWorks-ppc603

```
PROD_IOC = aaa
VX_PROD_vxWorks-68040 = bbb
VX_PROD_vxWorks-ppc603 = bbb
PROD_IOC += VX_PROD_$(T_A)
```

4.3.8 Host and Ioc targets

Build creates two type of makefile targets: Host and Ioc. Host targets are executables, object files, libraries, and scripts which are not part of iocCore. Ioc targets are components of ioc libraries, executables, object files, or iocsh scripts which will be run on an ioc.

Each supported target system has a VALID_BUILDS definition which specifies the type of makefile targets it can support. This definition appears in `configure/os/CONFIG.Common.<arch>` or `configure/os/CONFIG.<arch>.<arch>` files.

For vxWorks systems VALID_BUILDS is set to "Ioc".

For Unix type systems, VALID_BUILDS is set to "Host Ioc".

For RTEMS systems, VALID_BUILDS is set to "Ioc".

For WIN32 systems, VALID_BUILDS is set to "Host Ioc".

In a Makefile it is possible to limit the systems for which a particular PROD, TESTPROD, LIBRARY, SCRIPTS, and OBJS is built. For example the following Makefile lines specify that product aaa should be created for systems that support Host type builds. Product bbb should be created for systems that support Ioc type builds. Product ccc should be created for all target systems.

```
PROD_HOST = aaa
PROD_IOC = bbb
PROD = ccc
```

These definitions can be further limited by specifying an appended underscore "_" followed by an oclass or DEFAULT specification.

4.3.9 User specific override definitions

User specific override definitions are allowed in user created files in the user's `<home>/configure` subdirectory. These override definitions will be used for builds in all `<top>` directory structures. The files must have the following names.

```
<home>/configure/CONFIG_USER
<home>/configure/CONFIG_USER.<epics_host_arch>
<home>/configure/CONFIG_USER.Common.<targetarch>
<home>/configure/CONFIG_USER.<epics_host_arch>.<targetarch>
```


4.4 Makefiles

4.4.1 Name

The name of the makefile in each directory must be Makefile.

4.4.2 Included Files

Makefiles normally include files from `<top>/configure`. Thus the makefile “inherits” rules and definitions from `configure`. The files in `<top>/configure` may in turn include files from another `<top>/configure`. This technique makes it possible to share make variables and even rules across `<top>` directories.

4.4.3 Contents of Makefiles

4.4.3.1 Makefiles in directories containing subdirectories

A Makefile in this type of directory must define where `<top>` is relative to this directory, include `<top>/configure` files, and specify the subdirectories in the desired order of make execution. Running `gnumake` in a directory with the following Makefile lines will cause `gnumake` to be executed in `<dir1>` first and then `<dir2>`. The build rules do not allow a Makefile to specify both subdirectories and components to be built.

```
TOP=../..
include $(TOP)/configure/CONFIG
DIRS += <dir1> <dir2>
include $(TOP)/configure/RULES_DIRS
```

4.4.3.2 Makefiles in directories where components are to be built

A Makefile in this type of directory must define where `<top>` is relative to this directory, include `<top>` `configure` files, and specify the target component definitions. Optionally it may contain user defined rules. Running `gnumake` in a directory with this type of Makefile will cause `gnumake` to create an `O.<arch>` subdirectory and then execute `gnumake` to build the defined components in this subdirectory. It contains the following lines:

```
TOP=../..
include $(TOP)/configure/CONFIG
<component definition lines>
include $(TOP)/configure/RULES
<optional rules definitions>
```

4.4.4 Simple Makefile examples

Create an IOC type library named `asIoc` from the source file `asDbLib.c` and install it into the `$(INSTALL_LOCATION)/lib/<arch>` directory.

```
TOP=../..
include $(TOP)/configure/CONFIG
LIBRARY_IOC += asIoc
asIoc_SRCS += asDbLib.c
include $(TOP)/configure/RULES
```

For each Host type target architecture, create an executable named `catest` from the `catest1.c` and `catest2.c` source files linking with the existing EPICS base `ca` and `Com` libraries, and then install the `catest` executable into the `$(INSTALL_LOCATION)/bin/<arch>` directory.

```
TOP=../../..
include $(TOP)/configure/CONFIG
PROD_HOST = catest
catest_SRCS += catest1.c catest2.c
catest_LIBS = ca Com
include $(TOP)/configure/RULES
```

4.5 Make

4.5.1 Make vs. gnumake

EPICS provides an extensive set of make rules. These rules only work with the GNU version of make, `gnumake`, which is supplied by the Free Software Foundation. Thus, on most Unix systems, the native make will not work. On some systems, e.g. Linux, GNU make may be the default. This manual always uses `gnumake` in the examples.

4.5.2 Frequently used Make commands

NOTE: It is possible to invoke the following commands for a single target architecture by appending `<arch>` to the target in the command.

The most frequently used make commands are:

gnumake This rebuilds and installs everything that is not up to date. NOTE: Executing `gnumake` without arguments is the same as “`gnumake install`”

gnumake help This command can be executed from the `<top>` directory only. This command prints a page describing the most frequently used make commands.

gnumake install This rebuilds and installs everything that is not up to date.

gnumake all This is the same as “`gnumake install`”.

gnumake buildInstall This is the same as “`gnumake install`”.

gnumake <arch> This rebuilds and installs everything that is not up to date first for the host `arch` and then (if different) for the specified target `arch`.

NOTE: This is the same as “`gnumake install.<arch>`”

gnumake clean This can be used to save disk space by deleting the `O.<arch>` directories that `gnumake` will create, but does not remove any installed files from the `bin`, `db`, `dbd` etc. directories. “`gnumake clean.<arch>`” can be invoked to clean a single architecture.

gnumake archclean This command will remove the current build’s `O.<arch>` directories but not `O.Common` directory.

gnumake realeclean This command will remove ALL the `O.<arch>` subdirectories (even those created by a `gnumake` from another `EPICS_HOST_ARCH`).

gnumake rebuild This is the same as “`gnumake clean install`”. If you are unsure about the state of the generated files in an application, just execute “`gnumake rebuild`”.

gnumake uninstall This command can be executed from the `<top>` directory only. It will remove everything installed by `gnumake` in the `include`, `lib`, `bin`, `db`, `dbd`, etc. directories.

gnumake realuninstall This command can be executed from the <top> directory only. It will remove all the install directories, include, lib, bin, db, dbd, etc.

gnumake distclean This command can be executed from the <top> directory only. It is the same as issuing both the realclean and realuninstall commands.

gnumake cvsclean This command can be executed from the <top> directory only. It removes cvs .#* files in the make directory tree.

4.5.3 Make targets

The following is a summary of targets that can be specified for gnumake:

- <action>
- <arch>
- <action>.<arch>
- <dir>
- <dir>.<action>
- <dir>.<arch>
- <dir>.<action>.<arch>

where:

<arch> is an architecture such as solaris-sparc, vxWorks-68040, win32-x86, etc.

<action> is help, clean, realclean, distclean, inc, install, build, rebuild, buildInstall, realuninstall, or uninstall

NOTE: help, uninstall, distclean, cvsclean, and realuninstall can only be specified at <top>.

NOTE: realclean cannot be specified inside an O.<arch> subdirectory.

<dir> is subdirectory name

Note: You can build using your os vendors' native compiler and also build using a supported alternate compiler in the same directory structure because the executables and libraries will be created and installed into separate directories (e.g bin/solaris-sparc and bin/solaris-sparc-gnu). You can do this by changing your EPICS_HOST_ARCH, environment variable between builds or by setting EPICS_HOST_ARCH on the gnumake command line.

From Base R3.14.11 the build system will always ensure the host architecture is up to date before building a cross-compiled target, thus Makefiles must be explicit in defining which architectures a componen should be built for.

4.5.4 Header file dependencies

All product, test product, and library source files which appear in one of the source file definitions (e.g. SRCS, PROD_SRCS, LIB_SRCS, <prodname>_SRCS) will have their header file dependencies automatically generated and included as part of the Makefile if HDEPENDS is set to YES in the Makefile and/or in base/configure/CONFIG_SITE.

4.6 Makefile definitions

The following components can be defined in a Makefile:

4.6.1 Source file directories

Normally all product, test product, and library source files reside in the same directory as the Makefile. OS specific source files are allowed and should reside in subdirectories `os/<os_class>` or `os/posix` or `os/default`.

The build rules also allow source files to reside in subdirectories of the current Makefile directory (src directory). For each subdirectory `<dir>` containing source files add the `SRC_DIRS` definition.

```
SRC_DIRS += <dir>
```

where `<dir>` is a relative path definition. An example of `SRC_DIRS` is

```
SRC_DIRS += ../dir1 ../dir2
```

The directory search order for the above definition is

```
.
./os/$(OS_CLASS) ./os/posix ./os/default
./dir1/os/$(OS_CLASS) ./dir1/os/posix ./dir1/os/default
./dir2/os/$(OS_CLASS) ./dir2/os/posix ./dir2/os/default
..
./dir1 ../dir2
```

where the build directory `O.<arch>` is `.` and the `src` directory is `...`

4.6.2 Posix C source code

The epics base config files assume posix source code and define `POSIX` to be `YES` as the default. Individual Makefiles can override this by setting `POSIX` to `NO`. Source code files may have the suffix `.c`, `.cc`, `.cpp`, or `.C`.

4.6.3 Breakpoint Tables

For each breakpoint table `dbd` file, `bpt<table name>.dbd`, to be created from an existing `bpt<table name>.data` file, add the definition

```
DBD += bpt<table name>.dbd
```

to the Makefile. The following Makefile will create a `bptTypeJdegC.dbd` file from an existing `bptTypeJdegC.data` file using the EPICS base utility program `makeBpt` and install the new `dbd` file into the `$(INSTALL_LOCATION)/dbd` directory.

```
TOP=../../..
include $(TOP)/configure/CONFIG
DBD += bptTypeJdegC.dbd
include $(TOP)/configure/RULES
```

4.6.4 Record Type Definitions

For each new record type, the following definition should be added to the makefile:

```
DBDINC += <rectype>Record
```

A `<rectype>Record.h` header file will be created from an existing `<rectype>Record.dbd` file using the EPICS base utility program `dbToRecordTypeH`. This header will be installed into the `$(INSTALL_LOCATION)/include` directory and the `dbd` file will be installed into the `$(INSTALL_LOCATION)/dbd` directory.

The following Makefile will create `xxxRecord.h` from an existing `xxxRecord.dbd` file, install `xxxRecord.h` into `$(INSTALL_LOCATION)/include`, and install `xxxRecord.dbd` into `$(INSTALL_LOCATION)/dbd`.

```
TOP=../../..
include $(TOP)/configure/CONFIG
DBDINC += xxxRecord
include $(TOP)/configure/RULES
```

4.6.5 Menus

If a menu `menu<name>.dbd` file is present, then add the following definition:

```
DBDINC += menu<name>.h
```

The header file, `menu<name>.h` will be created from the existing `menu<name>.dbd` file using the EPICS base utility program `dbToMenuH` and installed into the `$(INSTALL_LOCATION)/include` directory and the menu `dbd` file will be installed into `$(INSTALL_LOCATION)/dbd`.

The following Makefile will create a `menuConvert.h` file from an existing `menuConvert.dbd` file and install `menuConvert.h` into `$(INSTALL_LOCATION)/include` and `menuConvert.dbd` into `$(INSTALL_LOCATION)/dbd`.

```
TOP=../../..
include $(TOP)/configure/CONFIG
DBDINC = menuConvert.h
include $(TOP)/configure/RULES
```

4.6.6 Expanded Database Definition Files

Database definition include files named `<name>Include.dbd` containing includes for other database definition files can be expanded by the EPICS base utility program `dbExpand` into a created `<name>.dbd` file and the `<name>.dbd` file installed into `$(INSTALL_LOCATION)/dbd`. The following variables control the process:

```
DBD += <name>.dbd
USR_DBDFLAGS += -I <include path>
USR_DBDFLAGS += -S <macro substitutions>
<name>_DBD += <file1>.dbd <file2>.dbd ...
```

where

```
DBD += <name>.dbd
```

is the name of the output `dbd` file to contain the expanded definitions. It is created by expanding an existing or build created `<name>Include.dbd` file and then copied into `$(INSTALL_LOCATION)/dbd`.

An example of a file to be expanded is `exampleInclude.dbd` containing the following lines

```
include "base.dbd"
include "xxxRecord.dbd"
device(xxx, CONSTANT, devXxxSoft, "SoftChannel")
```

`USR_DBDFLAGS` defines optional flags for `dbExpand`. Currently only an include path (`-I <path>`) and macro substitution (`-S <substitution>`) are supported. The include paths for EPICS base/`dbd`, and other `<top>/dbd` directories will automatically be added during the build if the `<top>` names are specified in the `configure/RELEASE` file.

A database definition include file named `<name>Include.dbd` containing includes for other database definition files can be created from a `<name>_DBD` definition. The lines

```
DBD += <name>.dbd
<name>_DBD += <file1>.dbd <file2>.dbd ...
```

will create an expanded dbd file `<name>.dbd` by first creating a `<name>Include.dbd`. For each filename in the `<name>_DBD` definition, the created `<name>Include.dbd` will contain an include statement for that filename. Then the expanded DBD file is generated from the created `<name>Include.dbd` file and installed into `$(INSTALL_LOCATION)/dbd`.

The following Makefile will create an expanded dbd file named `example.dbd` from an existing `exampleInclude.dbd` file and then install `example.dbd` into the `$(INSTALL_LOCATION)/dbd` directory.

```
TOP=../../..
include $(TOP)/configure/CONFIG
DBD += exampleApp.dbd
include $(TOP)/configure/RULES
```

The following Makefile will create an `exampleInclude.dbd` file from the `example_DBD` definition then expand it to create an expanded dbd file, `example.dbd`, and install `example.dbd` into the `$(INSTALL_LOCATION)/dbd` directory.

```
TOP=../../..
include $(TOP)/configure/CONFIG
DBD += example.dbd
example_DBD += base.dbd xxxRecord.dbd xxxSupport.dbd
include $(TOP)/configure/RULES
```

The created `exampleInclude.dbd` file will contain the following lines

```
include "base.dbd"
include "xxxRecord.dbd"
include "xxxSupport.dbd"
```

4.6.7 Registering Support Routines for Expanded Database Definition Files

A source file which registers simple static variables and record/device/driver support routines with `iocsh` can be created. The list of variables and routines to register is obtained from lines in an existing dbd file.

The following line in a Makefile will result in `<name>_registerRecordDeviceDriver.cpp` being created, compiled, and linked into `<prodname>`. It requires that the file `<name>.dbd` exist or can be created using other make rules.

```
<prodname>_SRCS += <name>_registerRecordDeviceDriver.cpp
```

An example of registering the variable `mySubDebug` and the routines `mySubInit` and `mySubProcess` is `<name>.dbd` containing the following lines

```
variable (mySubDebug)
function (mySubInit)
function (mySubProcess)
```

4.6.8 Database Definition Files

The following line installs the existing named dbd files into `$(INSTALL_LOCATION)/dbd` without expansion.

```
DBD += <name>.dbd
```

4.6.9 DBD install files

Definitions of the form:

```
DBD_INSTALLS += <name>
```

result in files being installed to the \$(INSTALL_LOCATION)/dbd directory. The file <name> can appear with or without a directory prefix. If the file has a directory prefix e.g. \$(APPNAME)/dbd/, it is copied from the specified location. If a directory prefix is not present, make will look in the current source directory for the file.

4.6.10 Database Files

For most databases just the name of the database has to be specified. Make will figure out how to generate the file:

```
DB += xxx.db
```

generates xxx.db depending on which source files exist and installs it into \$(INSTALL_LOCATION)/db.

A <name>.db database file will be created from an optional <name>.template file and/or an optional <name>.substitutions file. If the substitution file exists but the template file is not named <name>.template, the template file name can be specified as

```
<name>_TEMPLATE = <template file name>
```

A *<nn>.db database file will be created from a *.template and a *<nn>.substitutions file, (where nn is an optional index number).

If a <name> substitutions file contains “file” references to other input files, these referenced files are made dependencies of the created <name>.db by the makeDbDepends.pl perl tool.

The Macro Substitutions and Include tool, msi, will be used to generate the database, and msi must either be in your path or you must redefine MSI as the full path name to the msi binary in a RELEASE file or Makefile. An example MSI definition is

```
MSI = /usr/local/epics/extensions/bin/${EPICS_HOST_ARCH}/msi
```

Template files <name>.template, and db files, <name>.db, will be created from an edf file <name>.edf and an <name>.edf file will be created from a <name>.sch file.

Template and substitution files can be installed.

```
DB += xxx.template xxx.substitutions
```

generates and installs these files. If one or more xxx.substitutions files are to be created by script, the script name must be placed in the CREATSUBSTITUTIONS variable (e.g. CREATSUBSTITUTIONS=mySubst.pl). This script will be executed by gnumake with the prefix of the substitution file name to be generated as its argument. If (and only if) there are script generated substitutions files, the prefix of any inflated database’s name may not equal the prefix of the name of any template used within the directory.

4.6.11 DB install files

Definitions of the form:

```
DB_INSTALLS += <name>
```

result in files being installed to the \$(INSTALL_LOCATION)/db directory. The file <name> can appear with or without a directory prefix. If the file has a directory prefix e.g. \$(APPNAME)/db/, it is copied from the specified location. If a directory prefix is not present, make will look in the current source directory for the file.

4.6.12 Compile and link command options

Any of the following can be specified:

4.6.12.1 Options for all compile/link commands.

These definitions will apply to all compiler and linker targets.

```
USR_INCLUDES += -I<name>
```

header file directories each prefixed by a “-I”.

```
USR_INCLUDES_<osclass> += -I<name>
```

os specific header file directories each prefixed by a “-I”.

```
USR_INCLUDES_DEFAULT += -I<name>
```

header file directories each prefixed by “-I” for any arch that does not have a `USR_INCLUDE_<osclass>` definition

```
USR_CFLAGS += <c flags>
```

C compiler options.

```
USR_CFLAGS_<osclass> += <c flags>
```

os specific C compiler options.

```
USR_CFLAGS_<arch> += <c flags>
```

target architecture specific C compiler options.

```
USR_CFLAGS_DEFAULT += <c flags>
```

C compiler options for any arch that does not have a `USR_CFLAGS_<osclass>` definition

```
USR_CXXFLAGS += <c++ flags>
```

C++ compiler options.

```
USR_CXXFLAGS_<osclass> += <c++ flags>
```

C++ compiler options for the specified osclass.

```
USR_CXXFLAGS_<arch> += <c++ flags>
```

C++ compiler options for the specified target architecture.

```
USR_CXXFLAGS_DEFAULT += <c++ flags>
```

C++ compiler options for any arch that does not have a `USR_CXXFLAGS_<osclass>` definition

```
USR_CPPFLAGS += <preprocessor flags>
```

C preprocessor options.

```
USR_CPPFLAGS_<osclass> += <preprocessor flags>
```

os specific C preprocessor options.

```
USR_CPPFLAGS_<arch> += <preprocessor flags>
```

target architecture specific C preprocessor options.

```
USR_CPPFLAGS_DEFAULT += <preprocessor flags>
```

C preprocessor options for any arch that does not have a `USR_CPPFLAGS_<osclass>` definition

```
USR_LDFLAGS += <linker flags>
```

linker options.

USR_LDFLAGS_<osclass> += <linker flags>

os specific linker options.

USR_LDFLAGS_DEFAULT += <linker flags>

linker options for any arch that does not have a USR_LDFLAGS_<osclass> definition

4.6.12.2 Options for a target specific compile/link command.

<name>_INCLUDES += -I<name>

header file directories each prefixed by a “-I”.

<name>_INCLUDES_<osclass> += -I<name>

os specific header file directories each prefixed by a “-I”.

<name>_INCLUDES_<T_A> += -I<name>

target architecture specific header file directories each prefixed by a “-I”.

<name>_CFLAGS += <c flags>

c compiler options.

<name>_CFLAGS_<osclass> += <c flags>

os specific c compiler options.

<name>_CFLAGS_<T_A> += <c flags>

target architecture specific c compiler options.

<name>_CXXFLAGS += <c++ flags>

c++ compiler options.

<name>_CXXFLAGS_<osclass> += <c++ flags>

c++ compiler options for the specified osclass.

<name>_CXXFLAGS_<T_A> += <c++ flags>

c++ compiler options for the specified target architecture.

<name>_CPPFLAGS += <preprocessor flags>

c preprocessor options.

<name>_CPPFLAGS_<osclass> += <preprocessor flags>

os specific c preprocessor options.

<name>_CPPFLAGS_<T_A> += <preprocessor flags>

target architecture specific c preprocessor options.

<name>_LDFLAGS += <linker flags>

linker options.

<name>_LDFLAGS_<osclass> += <linker flags>

os specific linker options.

4.6.13 Libraries

A library is created and installed into `$(INSTALL_LOCATION)/lib/<arch>` by specifying its name and the name of the object and/or source files containing code for the library. An object or source file name can appear with or without a directory prefix. If the file name has a directory prefix e.g. `$(EPICS_BASE_BIN)`, it is taken from the specified location. If a directory prefix is not present, make will first look in the source directories for a file with the specified name and next try to create the file using existing configure rules. A library filename prefix may be prepended to the library name when the file is created. For Unix type systems and vxWorks the library prefix is `lib` and there is no prefix for WIN32. Also a library suffix appropriate for the library type and target arch (e.g. `.a`, `.so`, `.lib`, `.dll`) will be appended to the filename when the file is created.

vxWorks and RTEMS Note: Only archive libraries are created.

Shared libraries Note: Shared libraries can be built for any or all HOST type architectures. The definition of `SHARED_LIBRARIES` (YES/NO) in `base/configure/CONFIG_SITE` determines whether shared or archive libraries will be built. When `SHARED_LIBRARIES` is YES, both archive and shared libraries are built. This definition can be overridden for a specific arch in an `configure/os/CONFIG_SITE.<arch>.Common file.` The default definition for `SHARED_LIBRARIES` in the EPICS base distribution file is YES for all host systems.

win32 Note: An object library file is created when `SHARED_LIBRARIES=NO`, `<name>.lib` which is installed into `$(INSTALL_LOCATION)/lib/<arch>`. Two library files are created when `SHARED_LIBRARIES=YES`, `<name>.lib`, an import library for DLLs, which is installed into `$(INSTALL_LOCATION)/lib/<arch>`, and `<name>.dll` which is installed into `$(INSTALL_LOCATION)/bin/<arch>`. (Warning: The file `<name>.lib` will only be created by the build if there are exported symbols from the library.) If `SHARED_LIBRARIES=YES`, the directory

`$(INSTALL_LOCATION)/bin/<arch>` must be in the user's path during builds to allow invoking executables which were linked with shared libraries. NOTE: the `<name>.lib` files are different for shared and nonshared builds.

4.6.13.1 Specifying the library name.

Any of the following can be specified:

```
LIBRARY += <name>
```

A library will be created for every target arch.

```
LIBRARY_<osclass> += <name>
```

Library `<name>` will be created for all archs of the specified `osclass`.

```
LIBRARY_DEFAULT += <name>
```

Library `<name>` will be created for any arch that does not have a `LIBRARY_<osclass>` definition

```
LIBRARY_IOC += <name>
```

Library `<name>` will be created for IOC type archs.

```
LIBRARY_IOC_<osclass> += <name>
```

Library `<name>` will be created for all IOC type archs of the specified `osclass`.

```
LIBRARY_IOC_DEFAULT += <name>
```

Library `<name>` will be created for any IOC type arch that does not have a `LIBRARY_IOC_<osclass>` definition

```
LIBRARY_HOST += <name>
```

Library `<name>` will be created for HOST type archs.

`LIBRARY_HOST_<osclass> += <name>`

Library <name> will be created for all HOST type archs of the specified osclass.

`LIBRARY_HOST_DEFAULT += <name>`

Library <name> will be created for any HOST type arch that does not have a `LIBRARY_HOST_<osclass>` definition

4.6.13.2 Specifying library source file names

Source file names, which must have a suffix, are defined as follows:

`SRCS += <name>`

Source files will be used for all defined libraries and products.

`SRCS_<osclass> += <name>`

Source files will be used for all defined libraries and products for all archs of the specified osclass.

`SRCS_DEFAULT += <name>`

Source files will be used for all defined libraries and products for any arch that does not have a `SRCS_<osclass>` definition

`LIBSRCS` and `LIB_SRCS` have the same meaning. `LIBSRCS` is deprecated, but retained for R3.13 compatibility.

`LIBSRCS += <name>`

Source files will be used for all defined libraries.

`LIBSRCS_<osclass> += <name>`

Source files will be used for all defined libraries for all archs of the specified osclass.

`LIBSRCS_DEFAULT += <name>`

Source files will be used for all defined libraries for any arch that does not have a `LIBSRCS_<osclass>` definition

`USR_SRCS += <name>`

Source files will be used for all defined products and libraries.

`USR_SRCS_<osclass> += <name>`

Source files will be used for all defined products and libraries for all archs of the specified osclass.

`USR_SRCS_DEFAULT += <name>`

Source files will be used for all defined products and libraries for any arch that does not have a `USR_SRCS_<osclass>` definition

`LIB_SRCS += <name>`

Source files will be used for all libraries.

`LIB_SRCS_<osclass> += <name>`

Source files will be used for all defined libraries for all archs of the specified osclass.

`LIB_SRCS_DEFAULT += <name>`

Source files will be used for all defined libraries for any arch that does not have a `LIB_SRCS_<osclass>` definition

```
<libname>_SRCS += <name>
```

Source files will be used for the named library.

```
<libname>_SRCS_<osclass> += <name>
```

Source files will be used for named library for all archs of the specified osclass.

```
<libname>_SRCS_DEFAULT += <name>
```

Source files will be used for named library for any arch that does not have a `<libname>_SRCS_<osclass>` definition

4.6.13.3 Specifying library object file names

Library object file names should only be specified for object files which will not be built in the current directory. For object files built in the current directory, library source file names should be specified. See Specifying Library Source File Names above.

Object files which have filename with a “.o” or “.obj” suffix are defined as follows and can be specified without the suffix but should have the directory prefix

```
USR_OBJS += <name>
```

Object files will be used in builds of all products and libraries

```
USR_OBJS_<osclass> += <name>
```

Object files will be used in builds of all products and libraries for archs with the specified osclass.

```
USR_OBJS_DEFAULT += <name>
```

Object files will be used in builds of all products and libraries for archs without a `USR_OBJS_<osclass>` definition specified.

```
LIB_OBJS += <name>
```

Object files will be used in builds of all libraries.

```
LIB_OBJS_<osclass> += <name>
```

Object files will be used in builds of all libraries for archs of the specified osclass.

```
LIB_OBJS_DEFAULT += <name>
```

Object files will be used in builds of all libraries for archs without a `LIB_OBJS_<osclass>` definition specified.

```
<libname>_OBJS += <name>
```

Object files will be used for all builds of the named library)

```
<libname>_OBJS_<osclass> += <name>
```

Object files will be used in builds of the library for archs with the specified osclass.

```
<libname>_OBJS_DEFAULT += <name>
```

Object files will be used in builds of the library for archs without a `<libname>_OBJS_<osclass>` definition specified.

Combined object files, from R3.13 built modules and applications which have file names that do not include a “.o” or “.obj” suffix (e.g. xyzLib) are defined as follows:

```
USR_OBJLIBS += <name>
```

Combined object files will be used in builds of all libraries and products.

```
USR_OBJLIBS_<osclass> += <name>
```

Combined object files will be used in builds of all libraries and products for archs of the specified osclass.

```
USR_OBJLIBS_DEFAULT += <name>
```

Combined object files will be used in builds of all libraries and products for archs without a USR_OBJLIBS_<osclass> definition specified.

```
LIB_OBJLIBS += <name>
```

Combined object files will be used in builds of all libraries.

```
LIB_OBJLIBS_<osclass> += <name>
```

Combined object files will be used in builds of all libraries for archs of the specified osclass.

```
LIB_OBJLIBS_DEFAULT += <name>
```

Combined object files will be used in builds of all libraries for archs without a LIB_OBJLIBS_<osclass> definition specified.

```
<libname>_OBJLIBS += <name>
```

Combined object files will be used for all builds of the named library.

```
<libname>_OBJLIBS_<osclass> += <name>
```

Combined object files will be used in builds of the library for archs with the specified osclass.

```
<libname>_OBJLIBS_DEFAULT += <name>
```

Combined object files will be used in builds of the library for archs without a <libname>_OBJLIBS_<osclass> definition specified.

```
<libname>_LDOBJJS += <name>
```

Combined object files will be used for all builds of the named library. (deprecated)

```
<libname>_LDOBJJS_<osclass> += <name>
```

Combined object files will be used in builds of the library for archs with the specified osclass. (deprecated)

```
<libname>_LDOBJJS_DEFAULT += <name>
```

Combined object files will be used in builds of the library for archs without a <libname>_LDOBJJS_<osclass> definition specified. (deprecated)

4.6.13.4 LIBOBJJS definitions

Previous versions of epics (3.13 and before) accepted definitions like:

```
LIBOBJJS += $(<support>_BIN)/xxx.o
```

These are gathered together in files such as baseLIBOBJJS. To use such definitions include the lines:

```
-include ../baseLIBOBJJS
```

```
<libname>_OBS += $(LIBOBS)
```

Note: vxWorks applications created by makeBaseApp.pl from 3.14 Base releases no longer have a file named baseLIBOBS. Base record and device support now exists in archive libraries.

4.6.13.5 Specifying dependant libraries to be linked when creating a library

For each library name specified which is not a system library nor a library from an EPICS top defined in the configure/RELEASE file, a <name>_DIR definition must be present in the Makefile to specify the location of the library.

Library names, which must not have a directory and “lib” prefix nor a suffix, are defined as follows:

```
LIB_LIBS += <name>
```

Libraries to be used when linking all defined libraries.

```
LIB_LIBS_<osclass> += <name>
```

Libraries to be used or all archs of the specified osclass when linking all defined libraries.

```
LIB_LIBS_DEFAULT += <name>
```

Libraries to be used for any arch that does not have a LIB_LIBS_<osclass> definition when linking all defined libraries.

```
USR_LIBS += <name>
```

Libraries to be used when linking all defined products and libraries.

```
USR_LIBS_<osclass> += <name>
```

Libraries to be used or all archs of the specified osclass when linking all defined products and libraries.

```
USR_LIBS_DEFAULT += <name>
```

Libraries to be used for any arch that does not have a USR_LIBS_<osclass> definition when linking all defined products and libraries.

```
<libname>_LIBS += <name>
```

Libraries to be used for linking the named library.

```
<libname>_LIBS_<osclass> += <name>
```

Libraries will be used for all archs of the specified osclass for linking named library.

```
<libname>_LIBS_DEFAULT += <name>
```

Libraries to be used for any arch that does not have a <libname>_LIBS_<osclass> definition when linking named library.

```
<libname>_SYS_LIBS += <name>
```

System libraries to be used for linking the named library.

```
<libname>_SYS_LIBS_<osclass> += <name>
```

System libraries will be used for all archs of the specified osclass for linking named library.

```
<libname>_SYS_LIBS_DEFAULT += <name>
```

System libraries to be used for any arch that does not have a `<libname>_LIBS_<osclass>` definition when linking named library.

4.6.13.6 The order of dependant libraries

Dependant library names appear in the following order on a library link line:

1. `<libname>_LIBS`
2. `<libname>_LIBS_<osclass>` or `<libname>_LIBS_DEFAULT`
3. `LIB_LIBS`
4. `LIB_LIBS_<osclass>` or `LIB_LIBS_DEFAULT`
5. `USR_LIBS`
6. `USR_LIBS_<osclass>` or `USR_LIBS_DEFAULT`
7. `<libname>_SYS_LIBS`
8. `<libname>_SYS_LIBS_<osclass>` or `<libname>_SYS_LIBS_DEFAULT`
9. `LIB_SYS_LIBS`
10. `LIB_SYS_LIBS_<osclass>` or `LIB_SYS_LIBS_DEFAULT`
11. `USR_SYS_LIBS`
12. `USR_SYS_LIBS_<osclass>` or `USR_SYS_LIBS_DEFAULT`

4.6.13.7 Specifying library DLL file names (deprecated)

WIN32 libraries require all external references to be resolved, so if a library contains references to items in other DLL libraries, these DLL library names must be specified (without directory prefix and without “.dll” suffix) as follows:

```
DLL_LIBS += <name>
```

These DLLs will be used for all libraries.

```
<libname>_DLL_LIBS += <name>
```

These DLLs will be used for the named library.

Each `<name>` must have a corresponding `<name>_DIR` definition specifying its directory location.

4.6.13.8 Specifying shared library version number

A library version number can be specified when creating a shared library as follows:

```
SHRLIB_VERSION = <version>
```

On WIN32 this results in `/version:$(SHRLIB_VERSION)` link option. On Unix type hosts `.$(SHRLIB_VERSION)` is appended to the shared library name and a symbolic link is created for the unversioned library name. `$(EPICS_VERSION) .$(EPICS_REVISION)` is the default value for `SHRLIB_VERSION`.

4.6.13.9 Library example:

```

LIBRARY_vxWorks += vxWorksOnly
LIBRARY_IOC += iocOnly
LIBRARY_HOST += hostOnly
LIBRARY += all
vxWorksOnly_OBJS += $(LINAC_BIN)/vxOnly1
vxWorksOnly_SRCS += vxOnly2.c
iocOnly_OBJS += $(LINAC_BIN)/iocOnly1
iocOnly_SRCS += iocOnly2.cpp
hostOnly_OBJS += $(LINAC_BIN)/host1
all_OBJS += $(LINAC_BIN)/all1
all_SRCS += all2.cpp

```

If the architectures defined in <top>/configure are solaris-sparc and vxWorks-68040 and LINAC is defined in the <top>/configure/RELEASE file, then the following libraries will be created:

- \$(INSTALL_LOCATION)/bin/vxWork-68040/libvxWorksOnly.a : \$(LINAC_BIN)/vxOnly1.o vxOnly2.o
- \$(INSTALL_LOCATION)/bin/vxWork-68040/libiocOnly.a : \$(LINAC_BIN)/iocOnly1.o iocOnly2.o
- \$(INSTALL_LOCATION)/lib/solaris-sparc/libiocOnly.a : \$(LINAC_BIN)/iocOnly1.o iocOnly2.o
- \$(INSTALL_LOCATION)/lib/solaris-sparc/libhostOnly.a : \$(LINAC_BIN)/host1.o
- \$(INSTALL_LOCATION)/bin/vxWork-68040/liball.a : \$(LINAC_BIN)/all1.o all2.o
- \$(INSTALL_LOCATION)/lib/solaris-sparc/liball.a : \$(LINAC_BIN)/all1.o all2.o

4.6.14 Loadable libraries

Loadable libraries are regular libraries which are not required to have all symbols resolved during the build. The intent is to create dynamic plugins so no archive library is created. Source file, object files, and dependant libraries are specified in exactly the same way as for regular libraries.

Any of the following can be specified:

```
LOADABLE_LIBRARY += <name>
```

The <name> loadable library will be created for every target arch.

```
LOADABLE_LIBRARY_<osclass> += <name>
```

Loadable library <name> will be created for all archs of the specified osclass.

```
item LOADABLE_LIBRARY_DEFAULT += <name>
```

Loadable library <name> will be created for any arch that does not have a LOADABLE_LIBRARY_<osclass> definition

```
LOADABLE_LIBRARY_HOST += <name>
```

Loadable library <name> will be created for HOST type archs.

```
LOADABLE_LIBRARY_HOST_<osclass> += <name>
```

Loadable library <name> will be created for all HOST type archs of the specified osclass.

```
LOADABLE_LIBRARY_HOST_DEFAULT += <name>
```


Loadable library <name> will be created for any HOST type arch that does not have a LOADABLE_LIBRARY_HOST_<osclass> definition

4.6.15 Combined object libraries (VxWorks only)

Combined object libraries are regular combined object files which have been created by linking together multiple object files. OBJLIB specifications in the Makefile create a combined object file and a corresponding munch file for vxWorks target architectures only. Combined object libraries have a Library.o suffix. It is possible to generate and install combined object libraries by using definitions:

```
OBJLIB += <name>
OBJLIB_vxWorks += <name>
OBJLIB_SRCS += <srcname1> <srcname2> ...
OBJLIB_OBJJS += <objname1> <objname2> ...
```

These definitions result in the combined object file <name>Library.o and its corresponding <name>Library.munch munch file being built for each vxWorks architecture from source/object files in the OBJLIB_SRCS/OBJLIB_OBJJS definitions. The combined object file and the munch file are installed into the \$(INSTALL_LOCATION)/bin/<arch> directory.

4.6.16 Object Files

It is possible to generate and install object files by using definitions:

```
OBJJS += <name>
OBJJS_<osclass> += <name>
OBJJS_DEFAULT += <name>
OBJJS_IOC += <name>
OBJJS_IOC_<osclass> += <name>
OBJJS_IOC_DEFAULT += <name>
OBJJS_HOST += <name>
OBJJS_HOST_<osclass> += <name>
OBJJS_HOST_DEFAULT += <name>
```

These will cause the specified file to be generated from an existing source file for the appropriate target arch and installed into \$(INSTALL_LOCATION)/bin/<arch>.

The following Makefile will create the abc object file for all target architectures, the def object file for all target archs except vxWorks, and the xyz object file only for the vxWorks target architecture and install them into the appropriate \$(INSTALL_LOCATION)/bin/<arch> directory.

```
TOP=../../../../
include $(TOP)/configure/CONFIG
OBJJS += abc
OBJJS_vxWorks += xyz
OBJJS_DEFAULT += def
include $(TOP)/configure/RULES
```

4.6.17 State Notation Programs

A state notation program file can be specified as a source file in any SRC definition. For example:

```
<prodname>_SRCS += <name>.stt
```

The state notation compiler `snc` will generate the file `<name>.c` from the state notation program file `<name>.stt`. This C file is compiled and the resulting object file is linked into the `<prodname>` product.

A state notation source file must have the extension `.st` or `.stt`. The `.st` file is passed through the C preprocessor before it is processed by `snc`.

If you have state notation language source files (`.stt` and `.st` files), the module `seq` must be built and `SNCSEQ` defined in the `RELEASE` file. If the state notation language source files require C preprocessing before conversion to C source (`.st` files), `gcc` must be in your path.

4.6.18 Scripts, etc.

Any of the following can be specified:

```
SCRIPTS += <name>
```

A script will be installed from the `src` directory to the `$(INSTALL_LOCATION)/bin/<arch>` directories.

```
SCRIPTS_<osclass> += <name>
```

Script `<name>` will be installed for all archs of the specified `osclass`.

```
SCRIPTS_DEFAULT += <name>
```

Script `<name>` will be installed for any arch that does not have a `SCRIPTS_<osclass>` definition

```
SCRIPTS_IOC += <name>
```

Script `<name>` will be installed for IOC type archs.

```
SCRIPTS_IOC_<osclass> += <name>
```

Script `<name>` will be installed for all IOC type archs of the specified `osclass`.

```
SCRIPTS_IOC_DEFAULT += <name>
```

Script `<name>` will be installed for any IOC type arch that does not have a `SCRIPTS_IOC_<osclass>` definition

```
SCRIPTS_HOST += <name>
```

Script `<name>` will be installed for HOST type archs.

```
SCRIPTS_HOST_<osclass> += <name>
```

Script `<name>` will be installed for all HOST type archs of the specified `osclass`.

```
SCRIPTS_HOST_DEFAULT += <name>
```

Script `<name>` will be installed for any HOST type arch that does not have a `SCRIPTS_HOST_<osclass>` definition

Definitions of the form:

```
SCRIPTS_<osclass> += <name1>
SCRIPTS_DEFAULT += <name2>
```

results in the `<name1>` script being installed from the `src` directory to the `$(INSTALL_LOCATION)/bin/<arch>` directories for all target archs of the specified os class `<osclass>` and the `<name2>` script installed into the `$(INSTALL_LOCATION)/bin/<arch>` directories of all other target archs.

4.6.19 Include files

A definition of the form:

```
INC += <name>.h
```

results in file `<name>.h` being installed or created and installed to the `$(INSTALL_LOCATION)/include` directory.

Definitions of the form:

```
INC_DEFAULT += <name>.h
INC_<osclass> += <name>.h
```

results in file `<name>.h` being installed or created and installed into the appropriate `$(INSTALL_LOCATION)/include/os/<osclass>` directory.

4.6.20 Html and Doc files

A definition of the form:

```
HTMLS_DIR = <dirname>
HTMLS += <name>
```

results in file `<name>` being installed from the `src` directory to the `$(INSTALL_LOCATION)/html/<dirname>` directory.

A definition of the form:

```
DOCS += <name>
```

results in file `<name>` being installed from the `src` directory to the `$(INSTALL_LOCATION)/doc` directory.

4.6.21 Templates

Adding definitions of the form

```
TEMPLATES_DIR = <dirname>
TEMPLATES += <name>
```

results in the file `<name>` being installed from the `src` directory to the `$(INSTALL_LOCATION)/templates/<dirname>` directory. If a directory structure of template files is to be installed, the template file names may include a directory prefix.

4.6.22 Lex and yac

If a `<name>.c` source file specified in a Makefile definition is not found in the source directory, gnumake will try to build it from `<name>.y` and `<name>_lex.l` files in the source directory. Lex converts a `<name>.l` Lex code file to a `lex.yy.c` file which the build rules renames to `<name>.c`. Yacc converts a `<name>.y` yacc code file to a `y.tab.c` file, which the build rules renames to `<name>.c`. Optionally yacc can create a `y.tab.h` file which the build rules renames to `<name>.h`.

4.6.23 Products

A product executable is created for each `<arch>` and installed into `$(INSTALL_LOCATION)/bin/<arch>` by specifying its name and the name of either the object or source files containing code for the product. An object or source file name can appear with or without a directory prefix. Object files should contain a directory prefix. If the file has a directory prefix e.g. `$(EPICS_BASE_BIN)`, the file is taken from the specified location. If a directory prefix is not present, make will look in the source directories for a file with the specified name or try build it using existing rules. An executable filename suffix appropriate for the target arch (e.g. `.exe`) may be appended to the filename when the file is created.

PROD specifications in the Makefile for vxWorks target architectures create a combined object file with library references resolved and a corresponding `.munch` file.

```
PROD_HOST += <name>
<name>_SRC += <srcname>.c
```

results in the executable `<name>` being built for each HOST architecture, `<arch>`, from a `<srcname>.c` file. Then `<name>` is installed into the `$(INSTALL_LOCATION)/bin/<arch>` directory.

4.6.23.1 Specifying the product name.

Any of the following can be specified:

```
PROD += <name>
```

Product `<name>` will be created for every target arch.

```
PROD_<osclass> += <name>
```

Product `<name>` will be created for all archs of the specified `osclass`.

```
PROD_DEFAULT += <name>
```

Product `<name>` will be created for any arch that does not have a `PROD_<osclass>` definition

```
PROD_IOC += <name>
```

Product `<name>` will be created for IOC type archs.

```
PROD_IOC_<osclass> += <name>
```

Product `<name>` will be created for all IOC type archs of the specified `osclass`.

```
PROD_IOC_DEFAULT += <name>
```

Product `<name>` will be created for any IOC type arch that does not have a `PROD_IOC_<osclass>` definition

```
PROD_HOST += <name>
```

Product `<name>` will be created for HOST type archs.

```
PROD_HOST_<osclass> += <name>
```

Product `<name>` will be created for all HOST type archs of the specified `osclass`.

```
PROD_HOST_DEFAULT += <name>
```

Product `<name>` will be created for any HOST type arch that does not have a `PROD_HOST_<osclass>` definition

4.6.23.2 Specifying product object file names

Object files which have filenames with a “.o” or “.obj” suffix are defined as follows and can be specified without the suffix but should have the directory prefix

```
USR_OBJS += <name>
```

Object files will be used in builds of all products and libraries

```
USR_OBJS_<osclass> += <name>
```

Object files will be used in builds of all products and libraries for archs with the specified osclass.

```
USR_OBJS_DEFAULT += <name>
```

Object files will be used in builds of all products and libraries for archs without a USR_OBJS_<osclass> definition specified.

```
PROD_OBJS += <name>
```

Object files will be used in builds of all products

```
PROD_OBJS_<osclass> += <name>
```

Object files will be used in builds of all products for archs with the specified osclass.

```
PROD_OBJS_DEFAULT += <name>
```

Object files will be used in builds of all products for archs without a PROD_OBJS_<osclass> definition specified.

```
<prodname>_OBJS += <name>
```

Object files will be used for all builds of the named product

```
<prodname>_OBJS_<osclass> += <name>
```

Object files will be used in builds of the named product for archs with the specified osclass.

```
<prodname>_OBJS_DEFAULT += <name>
```

Object files will be used in builds of the named product for archs without a <prodname>_OBJS_<osclass> definition specified.

Combined object files, from R3.13 built modules and applications which have file names that do not include a “.o” or “.obj” suffix (e.g. xyzLib) are defined as follows:

```
USR_OBGLIBS += <name>
```

Combined object files will be used in builds of all libraries and products.

```
USR_OBGLIBS_<osclass> += <name>
```

Combined object files will be used in builds of all libraries and products for archs of the specified osclass.

```
USR_OBGLIBS_DEFAULT += <name>
```

Combined object files will be used in builds of all libraries and products for archs without a USR_OBGLIBS_<osclass> definition specified.

```
PROD_OBGLIBS += <name>
```

Combined object files will be used in builds of all products.

```
PROD_OBGLIBS_<osclass> += <name>
```

Combined object files will be used in builds of all products for archs of the specified osclass.

PROD_OBJLIBS_DEFAULT += <name>

Combined object files will be used in builds of all products for archs without a PROD_OBJLIBS_<osclass> definition specified.

<prodname>_OBJLIBS += <name>

Combined object files will be used for all builds of the named product.

<prodname>_OBJLIBS_<osclass> += <name>

Combined object files will be used in builds of the named product for archs with the specified osclass.

<prodname>_OBJLIBS_DEFAULT += <name>

Combined object files will be used in builds of the named product for archs without a <prodname>_OBJLIBS_<osclass> definition specified.

<prodname>_LDOBJJS += <name>

Object files will be used for all builds of the named product. (deprecated)

<prodname>_LDOBJJS_<osclass> += <name>

Object files will be used in builds of the name product for archs with the specified osclass. (deprecated)

<prodname>_LDOBJJS_DEFAULT += <name>

Object files will be used in builds of the product for archs without a <prodname>_LDOBJJS_<osclass> definition specified. (deprecated)

4.6.23.3 Specifying product source file names

Source file names, which must have a suffix, are defined as follows:

SRCS += <name>

Source files will be used for all defined libraries and products.

SRCS_<osclass> += <name>

Source files will be used for all defined libraries and products for all archs of the specified osclass.

SRCS_DEFAULT += <name>

Source files will be used for all defined libraries and products for any arch that does not have a SRCS_<osclass> definition

USR_SRCS += <name>

Source files will be used for all products and libraries.

USR_SRCS_<osclass> += <name>

Source files will be used for all defined products and libraries for all archs of the specified osclass.

USR_SRCS_DEFAULT += <name>

Source files will be used for all defined products and libraries for any arch that does not have a USR_SRCS_<osclass> definition

```
PROD_SRCS += <name>
```

Source files will be used for all products.

```
PROD_SRCS_<osclass> += <name>
```

Source files will be used for all defined products for all archs of the specified osclass.

```
PROD_SRCS_DEFAULT += <name>
```

Source files will be used for all defined products for any arch that does not have a `PROD_SRCS_<osclass>` definition

```
<prodname>_SRCS += <name>
```

Source file will be used for the named product.

```
<prodname>_SRCS_<osclass> += <name>
```

Source files will be used for named product for all archs of the specified osclass.

```
<prodname>_SRCS_DEFAULT += <name>
```

Source files will be used for named product for any arch that does not have a `<prodname>_SRCS_<osclass>` definition

4.6.23.4 Specifying libraries to be linked when creating the product

For each library name specified which is not a system library nor a library from `EPICS_BASE`, a `<name>_DIR` definition must be present in the Makefile to specify the location of the library.

Library names, which must not have a directory and “lib” prefix nor a suffix, are defined as follows:

```
PROD_LIBS += <name>
```

Libraries to be used when linking all defined products.

```
PROD_LIBS_<osclass> += <name>
```

Libraries to be used or all archs of the specified osclass when linking all defined products.

```
PROD_LIBS_DEFAULT += <name>
```

Libraries to be used for any arch that does not have a `PROD_LIBS_<osclass>` definition when linking all defined products.

```
USR_LIBS += <name>
```

Libraries to be used when linking all defined products.

```
USR_LIBS_<osclass> += <name>
```

Libraries to be used or all archs of the specified osclass when linking all defined products.

```
USR_LIBS_DEFAULT += <name>
```

Libraries to be used for any arch that does not have a `USR_LIBS_<osclass>` definition when linking all defined products.

```
<prodname>_LIBS += <name>
```

Libraries to be used for linking the named product.

`<prodname>_LIBS_<osclass> += <name>`

Libraries will be used for all archs of the specified osclass for linking named product.

`<prodname>_LIBS_DEFAULT += <name>`

Libraries to be used for any arch that does not have a `<prodname>_LIBS_<osclass>` definition when linking named product.

`SYS_PROD_LIBS += <name>`

System libraries to be used when linking all defined products.

`SYS_PROD_LIBS_<osclass> += <name>`

System libraries to be used for all archs of the specified osclass when linking all defined products.

`SYS_PROD_LIBS_DEFAULT += <name>`

System libraries to be used for any arch that does not have a `PROD_LIBS_<osclass>` definition when linking all defined products.

`<prodname>_SYS_LIBS += <name>`

System libraries to be used for linking the named product.

`<prodname>_SYS_LIBS_<osclass> += <name>`

System libraries will be used for all archs of the specified osclass for linking named product.

`<prodname>_SYS_LIBS_DEFAULT += <name>`

System libraries to be used for any arch that does not have a `<prodname>_LIBS_<osclass>` definition when linking named product.

4.6.23.5 The order of dependant libraries

Dependant library names appear in the following order on a product link line:

1. `<prodname>_LIBS`
2. `<prodname>_LIBS_<osclass>` or `<prodname>_LIBS_DEFAULT`
3. `PROD_LIBS`
4. `PROD_LIBS_<osclass>` or `PROD_LIBS_DEFAULT`
5. `USR_LIBS`
6. `USR_LIBS_<osclass>` or `USR_LIBS_DEFAULT`
7. `<prodname>_SYS_LIBS`
8. `<prodname>_SYS_LIBS_<osclass>` or `<prodname>_SYS_LIBS_DEFAULT`
9. `PROD_SYS_LIBS`
10. `PROD_SYS_LIBS_<osclass>` or `PROD_SYS_LIBS_DEFAULT`
11. `USR_SYS_LIBS`
12. `USR_SYS_LIBS_<osclass>` or `USR_SYS_LIBS_DEFAULT`

4.6.23.6 Specifying product version number

On WIN32 only a product version number can be specified as follows:

```
PROD_VERSION += <version>
```

This results in “/version:\$(PROD_VERSION)” link option.

4.6.23.7 Product static builds

Product executables can be linked with either archive versions or shared versions of EPICS libraries. Shared versions of system libraries will always be used in product linking. The definition of `STATIC_BUILD` (YES/NO) in `base/configure/CONFIG_SITE` determines which EPICS libraries to use. When `STATIC_BUILD` is NO, shared libraries will be used. (`SHARED_LIBRARIES` must be set to YES.) The default definition for `STATIC_BUILD` in the EPICS base `CONFIG_SITE` distribution file is NO. A `STATIC_BUILD` definition in a Makefile will override the definition in `CONFIG_SITE`. Static builds may not be possible on all systems. For static builds, all nonsystem libraries must have an archive version, and this may not be true for all libraries.

4.6.24 Test Products

Test products are product executables that are created but not installed into `$(INSTALL_LOCATION)/bin/<arch>` directories. Test product libraries, source, and object files are specified in exactly the same way as regular products.

Any of the following can be specified:

```
TESTPROD += <name>
```

Test product `<name>` will be created for every target arch.

```
TESTPROD_<osclass> += <name>
```

Test product `<name>` will be created for all archs of the specified `osclass`.

```
TESTPROD_DEFAULT += <name>
```

Test product `<name>` will be created for any arch that does not have a `TESTPROD_<osclass>` definition

```
TESTPROD_IOC += <name>
```

Test product `<name>` will be created for IOC type archs.

```
TESTPROD_IOC_<osclass> += <name>
```

Test product `<name>` will be created for all IOC type archs of the specified `osclass`.

```
TESTPROD_IOC_DEFAULT += <name>
```

Test product `<name>` will be created for any IOC type arch that does not have a `TESTPROD_IOC_<osclass>` definition

```
TESTPROD_HOST += <name>
```

Test product `<name>` will be created for HOST type archs.

```
TESTPROD_HOST_<osclass> += <name>
```

Test product `<name>` will be created for all HOST type archs of the specified `osclass`.

```
TESTPROD_HOST_DEFAULT += <name>
```

Test product <name> will be created for any HOST type arch that does not have a TESTPROD_HOST_<osclass> definition

4.6.25 Test Scripts

Test scripts are perl scripts whose names end in `.t` that get executed to satisfy the `runtests` make target. They are run by the perl `Test::Harness` library, and should send output to `stdout` following the Test Anything Protocol. Any of the following can be specified, although only `TESTSCRIPTS_HOST` is currently useful:

```
TESTSCRIPTS += <name>
```

Test script <name> will be created for every target arch.

```
TESTSCRIPTS_<osclass> += <name>
```

Test script <name> will be created for all archs of the specified osclass.

```
TESTSCRIPTS_DEFAULT += <name>
```

Test script <name> will be created for any arch that does not have a TESTSCRIPTS_<osclass> definition

```
TESTSCRIPTS_IOC += <name>
```

Test script <name> will be created for IOC type archs.

```
TESTSCRIPTS_IOC_<osclass> += <name>
```

Test script <name> will be created for all IOC type archs of the specified osclass.

```
TESTSCRIPTS_IOC_DEFAULT += <name>
```

Test script <name> will be created for any IOC type arch that does not have a TESTSCRIPTS_IOC_<osclass> definition

```
TESTSCRIPTS_HOST += <name>
```

Test script <name> will be created for HOST type archs.

```
TESTSCRIPTS_HOST_<osclass> += <name>
```

Test script <name> will be created for all HOST type archs of the specified osclass.

```
TESTSCRIPTS_HOST_DEFAULT += <name>
```

Test script <name> will be created for any HOST type arch that does not have a TESTSCRIPTS_HOST_<osclass> definition.

If a name in one of the above variables matches a regular executable program name (normally generated as a test product) with `.t` appended, a suitable perl script will be generated that will execute that program directly; this makes it simple to run programs that use the `epicsUnitTest` routines in `libCom`. A test script written in Perl with a name ending `.plt` will be copied into the `O.<arch>` directory with the ending changed to `.t`; such scripts will usually use the perl `Test::Simple` or `Test::More` libraries.

4.6.26 Miscellaneous Targets

A definition of the form:

```
TARGETS += <name>
```

results in the file <name> being built in the `O.<arch>` directory from existing rules and files in the source directory. These target files are not installed.

4.6.27 Installing Other Binaries

Definitions of the form:

```
BIN_INSTALLS += <name>
BIN_INSTALLS += <dir>/<name>
BIN_INSTALLS_DEFAULT += <name>
BIN_INSTALLS_<osclass> += <name>
```

will result in the named files being installed to the appropriate `$(INSTALL_LOCATION)/bin/<arch>` directory. The file <name> can appear with or without a directory prefix. If the file has a directory prefix e.g. `$(EPICS_BASE_BIN)`, it is copied from the specified location. If a directory prefix is not present, make will look in the source directory for the file.

4.6.28 Installing Other Libraries

Definitions of the form:

```
LIB_INSTALLS += <name>
LIB_INSTALLS += <dir>/<name>
LIB_INSTALLS_DEFAULT += <name>
LIB_INSTALLS_<osclass> += <name>
```

result in files being installed to the appropriate `$(INSTALL_LOCATION)/lib/<arch>` directory. The file <name> can appear with or without a directory prefix. If the file has a directory prefix e.g. `$(EPICS_BASE_LIB)`, it is copied from the specified location. If a directory prefix is not present, make will look in the source directory for the file.

4.6.29 Win32 resource files

Definitions of the form:

```
RCS += <name>    Resource definition script files for all products and libraries.
RCS_<osclass> += <name>
```

```
PROD_RCS += <name> Resource definition script files for all products.
PROD_RCS_<osclass> += <name>
PROD_RCS_DEFAULT += <name>
```

```
LIB_RCS += <name> Resource definition script files for all libraries.
LIB_RCS_<osclass> += <name>
LIB_RCS_DEFAULT += <name>
```

```
<name>_RCS += <name> Resource definition script files for specified product or library.
<name>_RCS_<osclass> += <name>
<name>_RCS_DEFAULT += <name>
```

result in resource files (*.res files) being created from the specified *.rc resource definition script files and linked into the prods and/or libraries.

4.6.30 TCL libraries

Definitions of the form:

```
TCLLIBNAME += <name>
TCLINDEX += <name>
```

result in the specified tcl files being installed to the $\$(INSTALL_LOCATION)/lib/<arch>$ directory.

4.6.31 Java class files

Java class files can be created by the javac tool into $\$(INSTALL_JAVA)$ or into the `O.Common` subdirectory, by specifying the name of the java class file in the Makefile. Command line options for the javac tool can be specified. The configuration files set the java c option “-sourcepath/..”.

Any of the following can be specified:

```
JAVA += <name>.java
```

The `<name>.java` file will be used to create the `<name>.class` file in the $\$(INSTALL_JAVA)$ directory.

```
TESTJAVA += <name>.java
```

The `<name>.java` files will be used to create the `<name>.class` file in the `O.Common` subdirectory.

```
USR_JAVACFLAGS += <name>
```

The javac option `<name>` will be used on the javac command lines.

4.6.31.1 Example 1

In this example, three class files are created in $\$(INSTALL_LOCATION)/javalib/mytest$. The javac deprecation flag is used to list the description of each use or override of a deprecated member or class.

```
JAVA = mytest/one.java
```

```
JAVA = mytest/two.java
```

```
JAVA = mytest/three.java
```

```
USR_JAVACFLAGS = -deprecation
```

4.6.31.2 Example 2

In this example, the test.class file is created in the `O.Common` subdirectory.

```
TESTJAVA = test.java
```

4.6.32 Java jar file

A single java jar file can be created using the java jar tool and installed into $\$(INSTALL_JAVA)$ (i.e. $\$(INSTALL_LOCATION)/javalib$) by specifying its name, and the names of its input files to be included in the created jar file. The jar input file names must appear with a directory prefix.

Any of the following can be specified:

```
JAR += <name>
```

The $\langle name \rangle$ jar file will be created and installed into the $\$(INSTALL_JAVA)$ directory.

```
JAR_INPUT += <name>
```

Names of images, audio files and classes files to be included in the jar file.

```
JAR_MANIFEST += <name>
```

The preexisting manifest file will be used for the created jar file.

```
JAR_PACKAGES += <name>
```

Names of java packages to be installed and added to the created jar file.

4.6.32.1 Example 1

In this example, all the class files created by the current Makefile's "JAVA+=" definitions, are placed into a file named mytest1.jar. A manifest file will be automatically generated for the jar.

Note: $\$(INSTALL_CLASSES)$ is set to $\$(addprefix \$(INSTALL_JAVA)/,\$(CLASSES))$ in the EPICS base configure files.

```
JAR = mytest1.jar
```

```
JAR_INPUT = \$(INSTALL\_CLASSES)
```

4.6.32.2 Example 2

In this example, three class files are created and placed into a new jar archive file named mytest2.jar. An existing manifest file, mytest2.mf is put into the new jar file.

```
JAR = mytest2.jar
```

```
JAR_INPUT = \$(INSTALL\_JAVA)/mytest/one.class
```

```
JAR_INPUT = \$(INSTALL\_JAVA)/mytest/two.class
```

```
JAR_INPUT = \$(INSTALL\_JAVA)/mytest/three.class
```

```
JAR_MANIFEST = mytest2.mf
```

4.6.33 Java native method C header files

A C header files for use with java native methods will be created by the javah tool in the `O.Common` subdirectory by specifying the name of the header file to be created. The name of the java class file used to generate the header is derived from the name of the header file. Underscores (`_`) are used as a header file name delimiter. Command line options for the javah tool can be specified.

Any of the following can be specified:

```
JAVAINC += <name>.h
```

The `<name>.h` header file will be created in the `O.Common` subdirectory.

```
USR_JAVAHFLAGS += <name>
```

The `javah` option `<name>` will be used on the `javah` tool command line.

4.6.33.1 Example

In this example, the C header `xx_yy_zz.h` will be created in the `$(COMMON_DIR)` subdirectory from the class `xx.yy.zz` (i.e. the java class file `$(INSTALL_JAVA)/xx/yy/zz.class`). The option “-old” will tell `javah` to create old JDK1.0 style header files.

```
JAVAINC = xx_yy_zz.h
```

```
USR_JAVAHFLAGS = -old
```

4.6.34 User Created CONFIG* and RULES* files

Module developers can now create new `CONFIG` and `RULES*` files in a `<top>` application source directory. These new `CONFIG*` or `RULES*` files will be installed into the directory `$(INSTALL_LOCATION)/cfg` by including lines like the following Makefile line:

```
CFG += CONFIG_MY1 RULES_MY1
```

The build will install the new files `CONFIG_MY1` and `RULES_MY1` into the `$(INSTALL_LOCATION)/cfg` directory.

Files in a `$(INSTALL_LOCATION)/cfg` directory are now included during a build by so that the definitions and rules in them are available for use by later `src` directory Makefiles in the same module or by other modules with a `RELEASE` line pointing to the `TOP` of this module.

4.6.35 User Created File Types

Module developers can now define a new type of file, e.g. `ABC`, so that files of type `ABC` will be installed into a directory defined by `INSTALL_ABC`. This is done by creating a new `CONFIG_<name>` file, e.g. `CONFIG_ABC`, with the following lines:

```
FILE_TYPE += ABC
```

```
INSTALL_ABC = $(INSTALL_LOCATION)/abc
```

The `INSTALL_ABC` directory should be a subdirectory of `$(INSTALL_LOCATION)`. The file type `ABC` should be target architecture independent (all files, medm files, edm files).

Optional rules necessary for files of type `ABC` should be put in a `RULES_ABC` file.

The module developer installs new `CONFIG_ABC` and `RULES_ABC` files for the new file type into the directory `$(INSTALL_LOCATION)/cfg` by including the following Makefile line:

```
CFG += CONFIG_ABC RULES_ABC
```

Files of type `ABC` are installed into `INSTALL_ABC` directory by adding a line like the following to a Makefile.

```
ABC += <filename1> <filename2> <filename3>
```

Since the files in `$(INSTALL_LOCATION)/cfg` directory are now included by the base config files, the `ABC +=` definition lines are available for use by later `src` directory Makefiles in the same module or by other modules with a `RELEASE` line pointing to the `TOP` of this module.

4.7 Table of Makefile definitions

Definitions given below containing `<osclass>` are used when building for target archs of a specific osclass, and the `<osclass>` part of the name should be replaced by the desired osclass, e.g. solaris, vxWorks, etc. If a `_DEFAULT` setting is given but a particular `<osclass>` requires that the default not apply and there are no items in the definition that apply for that `<osclass>`, the value “-nil-” should be specified in the relevant Makefile definition.

Build Option	Description
Products to be built (host type archs only)	
PROD	products (names without execution suffix) to build and install. Specify xyz to build executable xyz on Unix and xyz.exe on WIN32
PROD_<osclass>	os class specific products to build and install for <osclass> archs only
PROD_DEFAULT	products to build and install for archs with no PROD_<osclass> specified
PROD_IOC	products to build and install for ioc type archs
PROD_IOC_<osclass>	os specific products to build and install for ioc type archs
PROD_IOC_DEFAULT	products to build and install for ioc type arch systems with no PROD_IOC_<osclass> specified
PROD_HOST	products to build and install for host type archs.
PROD_HOST_<osclass>	os class specific products to build and install for <osclass> type archs
PROD_HOST_DEFAULT	products to build and install for arch with no PROD_HOST_<osclass> specified
Test products to be built	
TESTPROD	test products (names without execution suffix) to build but not install
TESTPROD_<osclass>	os class specific test products to build but not install
TESTPROD_DEFAULT	test products to build but not install for archs with no TESTPROD_<osclass> specified
TESTPROD_IOC	test products to build and install for ioc type archs
TESTPROD_IOC_<osclass>	os specific test products to build and install for ioc type archs
TESTPROD_IOC_DEFAULT	test products to build and install for ioc type arch systems with no TESTPROD_IOC_<osclass> specified
TESTPROD_HOST	testproducts to build and install for host type archs.
TESTPROD_HOST_<osclass>	os class specific testproducts to build and install for <osclass> type archs
TESTPROD_HOST_DEFAULT	test products to build and install for arch with no TESTPROD_HOST_<osclass> specified
Test scripts to be built	
TESTSCRIPTS	test scripts (names with .t suffix) to build but not install
TESTSCRIPTS_<osclass>	os class specific test scripts to build but not install
TESTSCRIPTS_DEFAULT	test scripts to build but not install for archs with no TESTSCRIPTS_<osclass> specified
TESTSCRIPTS_IOC	test scripts to build and install for ioc type archs
TESTSCRIPTS_IOC_<osclass>	os specific test scripts to build and install for ioc type archs
TESTSCRIPTS_IOC_DEFAULT	test scripts to build and install for ioc type arch systems with no TESTSCRIPTS_IOC_<osclass> specified
TESTSCRIPTS_HOST	test scripts to build and install for host type archs.
TESTSCRIPTS_HOST_<osclass>	os class specific testscripts to build and install for <osclass> type archs
TESTSCRIPTS_HOST_DEFAULT	test scripts to build and install for arch with no TESTSCRIPTS_HOST_<osclass> specified
Libraries to be built	

LIBRARY	name of library to build and install. The name should NOT include a prefix or extension e.g. specify Ca to build libCa.a on Unix, Ca.lib or Ca.dll on WIN32
LIBRARY_<osclass> LIBRARY_DEFAULT	os specific libraries to build and install libraries to build and install for archs with no LIBRARY_<osclass> specified
LIBRARY_IOC	name of library to build and install for ioc type archs. The name should NOT include a prefix or extension e.g. specify Ca to build libCa.a on Unix, Ca.lib or Ca.dll on WIN32
LIBRARY_IOC_<osclass> LIBRARY_IOC_DEFAULT	os specific libraries to build and install for ioc type archs libraries to build and install for ioc type arch systems with no LIBRARY_IOC_<osclass> specified
LIBRARY_HOST	name of library to build and install for host type archs. The name should NOT include a prefix or extension, e.g. specify Ca to build libCa.a on Unix, Ca.lib or Ca.dll on WIN32
LIBRARY_HOST_<osclass> LIBRARY_HOST_DEFAULT	os class specific libraries to build and install for host type archs libraries to build and install for host type arch systems with no LIBRARY_HOST_<osclass> specified
SHARED_LIBRARIES	build shared libraries? Must be YES or NO
SHRLIB_VERSION	shared library version number
Loadable libraries to be built	
LOADABLE_LIBRARY	name of loadable library to build and install. The name should NOT include a prefix or extension e.g. specify Ca to build libCa.so on Unix and Ca.dll on WIN32
LOADABLE_LIBRARY_<osclass> LOADABLE_LIBRARY_DEFAULT	os specific loadable libraries to build and install loadable libraries to build and install for archs with no LOADABLE_LIBRARY_<osclass> specified
LOADABLE_LIBRARY_HOST	name of loadable library to build and install for host type archs. The name should NOT include a prefix or extension, e.g. specify test to build libtest.so on Unix and test.dll on WIN32
LOADABLE_LIBRARY_HOST_<osclass>	os class specific loadable libraries to build and install for host type archs
LOADABLE_LIBRARY_HOST_DEFAULT	loadable libraries to build and install for host type arch systems with no LOADABLE_LIBRARY_HOST_<osclass> specified
Combined object files (vxWorks only)	
OBJLIB	name of a combined object file library and corresponding munch file to build and install. The name will have a Library suffix appended same as OBJLIB
OBJLIB_vxWorks	source files to build the OBJLIB
OBJLIB_SRCS	object files to include in OBJLIB
OBJLIB_OBJS	
Product and library source files	
SRCS	source files to build all PRODs and LIBRARYs
SRCS_<osclass> SRCS_DEFAULT	osclass specific source files to build all PRODs and LIBRARYs source file to build all PRODs and LIBRARYs for archs with no SRCS_<osclass> specified
USR_SRCS	source files to build all PRODs and LIBRARYs
USR_SRCS_<osclass> USR_SRCS_DEFAULT	osclass specific source files to build all PRODs and LIBRARYs source file to build all PRODs and LIBRARYs for archs with no SRCS_<osclass> specified
PROD_SRCS	source files to build all PRODs
PROD_SRCS_<osclass> PROD_SRCS_DEFAULT	osclass specific source files to build all PRODs source files needed to build PRODs for archs with no SRCS_<osclass> specified

LIB_SRCS	source files for building LIBRARY (e.g. LIB_SRCS=la.c lb.c lc.c)
LIB_SRCS_<osclass>	os-specific library source files
LIB_SRCS_DEFAULT	library source files for archs with no LIB_SRCS_<osclass> specified
LIBSRCS	source files for building LIBRARY (deprecated)
LIBSRCS_<osclass>	os-specific library source files (deprecated)
LIBSRCS_DEFAULT	library source files for archs with no LIBSRCS_<osclass> specified (deprecated)
<name>_SRCS	source files to build a specific PROD or LIBRARY
<name>_SRCS_<osclass>	os specific source files to build a specific PROD or LIBRARY
<name>_SRCS_DEFAULT	source files needed to build a specific PROD or LIBRARY for archs with no <prod>_SRCS_<osclass> specified

Product and library object files

USR_OBJS	object files, specified without suffix, to build all PRODs and LIBRARYs
USR_OBJS_<osclass>	osclass specific object files, specified without suffix, to build all PRODs and LIBRARYs
USR_OBJS_DEFAULT	object files, specified without suffix, needed to build PRODs and LIBRARYs for archs with no OBJs_<osclass> specified
PROD_OBJS	object files, specified without suffix, to build all PRODs
PROD_OBJS_<osclass>	osclass specific object files, specified without suffix, to build all PRODs
PROD_OBJS_DEFAULT	object files, specified without suffix, needed to build PRODs for archs with no OBJs_<osclass> specified
LIB_OBJS	object files, specified without suffix, for building all LIBRARYs (e.g. LIB_OBJS+=\$(AB_BIN)/la \$(AB_BIN)/lb)
LIB_OBJS_<osclass>	os-specific library object files, specify without suffix,
LIB_OBJS_DEFAULT	library object files, specified without suffix, for archs with no LIB_OBJS_<osclass> specified
<name>_OBJS	object files, specified without suffix, to build a specific PROD or LIBRARY
<name>_OBJS_<osclass>	os specific object files, specified without suffix, to build a specific PROD or LIBRARY
<name>_OBJS_DEFAULT	object files, without suffix, needed to build a specific PROD or LIBRARY for archs with no <prod>_OBJS_<osclass> specified

Product and library R3.13 combined object files

USR_OBJLIBS	combined object files with filenames that do not have a suffix, needed for building all PRODs and LIBRARYs (e.g. USR_OBJLIBS+=\$(XYZ_BIN)/xyzLib)
USR_OBJLIBS_<osclass>	os-specific combined object files with filenames that do not have a suffix for building all PRODs and LIBRARYs
USR_OBJLIBS_DEFAULT	combined object files with filenames that do not have a suffix, for archs with no USR_OBJLIBS_<osclass> specified for building all PRODs and LIBRARYs
PROD_OBJLIBS	combined object files with filenames that do not have a suffix, needed for building all PRODs (e.g. PROD_OBJLIBS+=\$(XYZ_BIN)/xyzLib)
PROD_OBJLIBS_<osclass>	os-specific combined object files with filenames that do not have a suffix for building all PRODs
PROD_OBJLIBS_DEFAULT	combined object files with filenames that do not have a suffix, for archs with no PROD_OBJLIBS_<osclass> specified for building all PRODs

LIB_OBJLIBS	combined object files with filenames that do not have a suffix, needed for building all LIBRARYs (e.g. LIB_OBJLIBS+=\$(XYZ_BIN)/xyzLib)
LIB_OBJLIBS_<osclass>	os-specific combined object files with filenames that do not have a suffix for building all LIBRARYs
LIB_OBJLIBS_DEFAULT	combined object files with filenames that do not have a suffix, for archs with no LIB_OBJLIBS_<osclass> specified for building all LIBRARYs
<name>_OBJLIBS	combined object files with filenames that do not have a suffix, needed to build a specific PROD or LIBRARY
<name>_OBJLIBS_<osclass>	os specific combined object files with filenames that do not have a suffix, to build a specific PROD or LIBRARY
<name>_OBJLIBS_DEFAULT	combined object files with filenames that do not have a suffix, needed to build a specific PROD or LIBRARY for archs with no <name>_OBJLIBS_<osclass> specified
<name>_LDOBJS	combined object files with filenames that do not have a suffix, needed to build a specific PROD or LIBRARY (deprecated)
<name>_LDOBJS_<osclass>	os specific combined object files with filenames that do not have a suffix, to build a specific PROD or LIBRARY (deprecated)
<name>_LDOBJS_DEFAULT	combined object files with filenames that do not have a suffix, needed to build a specific PROD or LIBRARY for archs with no <name>_LDOBJS_<osclass> specified (deprecated)

Product and library dependant libraries

<name>_DIR	directory to search for the specified lib. (For libs listed in all PROD_LIBS, LIB_LIBS, <name>_LIBS and USR_LIBS listed below) System libraries do not need a <name>_dir definition.
USR_LIBS	load libraries (e.g. Xt X11) for all products and libraries
USR_LIBS_<osclass>	os specific load libraries for all makefile links
USR_LIBS_DEFAULT	load libraries for systems with no USR_LIBS_<osclass> specified libs
<name>_LIBS	named prod or library specific ld libraries (e.g. probe_LIBS=X11 Xt)
<name>_LIBS_<osclass>	os-specific libs needed to link named prod or library
<name>_LIBS_DEFAULT	libs needed to link named prod or library for systems with no <name>_LIBS_<osclass> specified
PROD_LIBS	libs needed to link every PROD
PROD_LIBS_<osclass>	os-specific libs needed to link every PROD
PROD_LIBS_DEFAULT	libs needed to link every PROD for archs with no PROD_LIBS_<osclass> specified
LIB_LIBS	libraries to be linked with every library being created
LIB_LIBS_<osclass>	os class specific libraries to be linked with every library being created
LIB_LIBS_DEFAULT	libraries to be linked with every library being created for archs with no LIB_LIBS_<osclass> specified
USR_SYS_LIBS	system libraries (e.g. Xt X11) for all products and libraries
USR_SYS_LIBS_<osclass>	os class specific system libraries for all makefile links
USR_SYS_LIBS_DEFAULT	system libraries for archs with no USR_SYS_LIBS_<osclass> specified
<name>_SYS_LIBS	named prod or library specific system ld libraries
<name>_SYS_LIBS_<osclass>	os class specific system libs needed to link named prod or library
<name>_SYS_LIBS_DEFAULT	system libs needed to link named prod or library for systems with no <name>_SYS_LIBS_<osclass> specified
PROD_SYS_LIBS	system libs needed to link every PROD
PROD_SYS_LIBS_<osclass>	os class specific system libs needed to link every PROD

PROD_SYS_LIBS_DEFAULT	system libs needed to link every PROD for archs with no PROD_SYS_LIBS_<osclass> specified
LIB_SYS_LIBS	system libraries to be linked with every library being created
LIB_SYS_LIBS_<osclass>	os class specific system libraries to be linked with every library being created
LIB_SYS_LIBS_DEFAULT	system libraries to be linked with every library being created for archs with no LIB_SYS_LIBS_<osclass> specified
SYS_PROD_LIBS	system libs needed to link every PROD for all systems (deprecated)
SYS_PROD_LIBS_<osclass>	os class specific system libs needed to link every PROD (deprecated)
SYS_PROD_LIBS_DEFAULT	system libs needed to link every PROD for systems with no SYS_PROD_LIBS_<osclass> specified (deprecated)

Compiler flags

USR_CFLAGS	C compiler flags for all systems
USR_CFLAGS_<T_A>	target architecture specific C compiler flags
USR_CFLAGS_<osclass>	os class specific C compiler flags
USR_CFLAGS_DEFAULT	C compiler flags for archs with no USR_CFLAGS_<osclass> specified
<name>_CFLAGS	file specific C compiler flags (e.g. xxxRecord.CFLAGS=-g)
<name>_CFLAGS_<T_A>	file specific C compiler flags for a specific target architecture
<name>_CFLAGS_<osclass>	file specific C compiler flags for a specific os class
USR_CXXFLAGS	C++ compiler flags for all systems (e.g. xyxMain.CFLAGS=-DSDDS)
USR_CXXFLAGS_<T_A>	target architecture specific C++ compiler flags
USR_CXXFLAGS_<osclass>	os-specific C++ compiler flags
USR_CXXFLAGS_DEFAULT	C++ compiler flags for systems with no USR_CXXFLAGS_<osclass> specified
<name>_CXXFLAGS	file specific C++ compiler flags
<name>_CXXFLAGS_<T_A>	file specific C++ compiler flags for a specific target architecture
<name>_CXXFLAGS_<osclass>	file specific C++ compiler flags for a specific osclass
USR_CPPFLAGS	C pre-processor flags (for all makefile compiles)
USR_CPPFLAGS_<T_A>	target architecture specific cpp flags
USR_CPPFLAGS_<osclass>	os specific cpp flags
USR_CPPFLAGS_DEFAULT	cpp flags for systems with no USR_CPPFLAGS_<osclass> specified
<name>_CPPFLAGS	file specific C pre-processor flags(e.g. xxxRecord.CPPFLAGS=-DDEBUG)
<name>_CPPFLAGS_<T_A>	file specific cpp flags for a specific target architecture
<name>_CPPFLAGS_<osclass>	file specific cpp flags for a specific os class
USR_INCLUDES	directories, with -I prefix, to search for include files(e.g. -I\$(EPICS_EXTENSIONS_INCLUDE))
USR_INCLUDES_<osclass>	directories, with -I prefix, to search for include files for a specific os class
USR_INCLUDES_DEFAULT	directories, with -I prefix, to search for include files for systems with no <name>_INCLUDES_<osclass> specified
<name>_INCLUDES	directories, with -I prefix, to search for include files when building a specific object file (e.g. -I\$(MOTIF_INC))
<name>_INCLUDES_<T_A>	file specific directories, with -I prefix, to search for include files for a specific target architecture
<name>_INCLUDES_<osclass>	file specific directories, with -I prefix, to search for include files for a specific os class
HOST_WARN	Are compiler warning messages desired for host type builds? (YES or NO) (default is YES)
CROSS_WARN	C cross-compiler warning messages desired (YES or NO) (default YES)

HOST_OPT	Is host build compiler optimization desired (default is NO optimization)
CROSS_OPT	Is cross-compiler optimization desired (YES or NO) (default is NO optimization)
CMPLR	C compiler selection, TRAD, ANSI or STRICT (default is STRICT)
CXXCMPLR	C++ compiler selection, NORMAL or STRICT (default is STRICT)
Linker options	
USR_LDFLAGS	linker options (for all makefile links)
USR_LDFLAGS_<osclass>	os specific linker options (for all makefile links)
USR_LDFLAGS_DEFAULT	linker options for systems with no USR_LDFLAGS_<osclass> specified
PROD_LDFLAGS	prod linker options
PROD_LDFLAGS_<osclass>	os specific prod linker options
PROD_LDFLAGS_DEFAULT	prod linker options for systems with no PROD_LDFLAGS_<osclass> specified
LIB_LDFLAGS	library linker options
LIB_LDFLAGS_<osclass>	os specific library linker options
LIB_LDFLAGS_DEFAULT	library linker options for systems with no LIB_LDFLAGS_<osclass> specified
<name>_LDFLAGS	prod or library specific linker options
<name>_LDFLAGS_<osclass>	prod or library specific linker flags for a specific os class
<name>_LDFLAGS_DEFAULT	linker options for systems with no <name>_LDFLAGS_<osclass> specified
STATIC_BUILD	Is static build desired (YES or NO) (default is NO). On win32 if STATIC_BUILD=YES then set SHARED_LIBRARIES=NO
Header files to be installed	
INC	list of include files to install into \$(INSTALL_DIR)/include
INC_<osclass>	os specific includes to installed under \$(INSTALL_DIR)/include/os/<osclass>
INC_DEFAULT	include files to install where no INC_<osclass> is specified
Perl, csh, tcl etc. script installation	
SCRIPTS	scripts to install for all systems
SCRIPTS_<osclass>	os-specific scripts to install
SCRIPTS_DEFAULT	scripts to install for systems with no SCRIPTS_<osclass> specified
SCRIPTS_IOC	scripts to install for ioc type archs.
SCRIPTS_IOC_<osclass>	os specific scripts to install for ioc type archs
SCRIPTS_IOC_DEFAULT	scripts to install for ioc type arch systems with no SCRIPTS_IOC_<osclass> specified
SCRIPTS_HOST	scripts to install for host type archs. T
SCRIPTS_HOST_<osclass>	os class specific scripts to install for host type archs
SCRIPTS_HOST_DEFAULT	scripts to install for host type arch systems with no OBJS_HOST_<osclass> specified
TCLLIBNAME	list of tcl scripts to install into \$(INSTALL_DIR)/lib/<osclass> (Unix hosts only)
TCLINDEX	name of tcl index file to create from TCLLIBNAME scripts
Object files	
OBJS	object files to build and install for all system.
OBJS_<osclass>	os-specific object files to build and install.
OBJS_DEFAULT	object files to build and install for systems with no OBJS_<osclass> specified.
OBJS_IOC	object files to build and install for ioc type archs.
OBJS_IOC_<osclass>	os specific object files to build and install for ioc type archs

OBJS_IOC_DEFAULT	object files to build and install for ioc type arch systems with no OBJS_IOC_<osclass> specified
OBJS_HOST	object files to build and install for host type archs. T
OBJS_HOST_<osclass>	os class specific object files to build and install for host type archs
OBJS_HOST_DEFAULT	object files to build and install for host type arch systems with no OBJS_HOST_<osclass> specified
Documentation	
DOCS	text files to be installed into the \$(INSTALL_DIR)/doc directory
HTMLS_DIR	name install Hypertext directory name i.e. \$(INSTALL_DIR)/html/\$(HTMLS_DIR)
HTMLS	hypertext files to be installed into the \$(INSTALL_DIR)/html/\$(HTMLS_DIR) directory
TEMPLATES_DIR	template directory to be created as \$(INSTALL_DIR)/templates/\$(TEMPLATE_DIR)
TEMPLATES	template files to be installed into \$(TEMPLATE_DIR)
Database Definition files	
DBD	database definition files to be installed or created and installed into \$(INSTALL_DBD).
DBDINC	names, without suffix, of menus or record database definitions and headers to be installed or created and installed.
USR_DBDFLAGS	optional flags for dbExpand. Currently only include path (-I <path>) and macro substitution (-S <substitution>) are supported.
DBD_INSTALLS	files from specified directory to install into \$(INSTALL_DBD) (e.g. DBD_INSTALLS = \$(APPNAME)/dbd/test.dbd)
Database Files	
DB	database files to be installed or created and installed into \$(INSTALL_DB).
DB_INSTALLS	files from specified directory to install into \$(INSTALL_DB) (e.g. DB_INSTALLS = \$(APPNAME)/db/test.db)
USR_DBFLAGS	optional flags for msi (EPICS Macro Substitution Tool)
Options for other programs	
YACCOPT	yacc options
LEXOPT	lex options
SNCFLAGS	state notation language, snc, options
<name>_SNCFLAGS	product specific state notation language options
E2DB_FLAGS	e2db options
SCH2EDIF_FLAGS	sch2edif options
RANLIBFLAGS	ranlib options
USR_ARFLAGS	ar options
Facilities for building Java programs	
JAVA	names of Java source files to be built and installed
TESTJAVA	names of Java source files to be built
JAVAINC	names of C header file to be created in O.Common subdirectory
JAR	name of Jar file to be built
JAR_INPUT	names of files to be included in JAR
JAR_MANIFEST	name of manifest file for JAR
USR_JAVACFLAGS	javac tool options
USR_JAVAHFLAGS	javah tool options
Facilities for Windows 95/NT resource (.rc) files	
RCS	resource files (<name>.rc) needed to build every PROD and LIBRARY
RCS_<osclass>	resource files (<name>.rc) needed to build every PROD and LIBRARY for ioc type archs

RCS_DEFAULT	resource files needed to build every PROD and LIBRARY for ioc type arch systems with no RCS_<osclass> specified
<name>_RCS	resource files needed to build a specific PROD or LIBRARY
<name>_RCS_<osclass>	os specific resource files to build a specific PROD or LIBRARY
<name>_RCS_DEFAULT	resource files needed to build a specific PROD or LIBRARY for ioc type arch systems with no RCS_<osclass> specified

Other definitions:

USR_VPATH	list of directories
BIN_INSTALLS	files from specified directories to be installed into \$(INSTALL_BIN) (e.g. BIN_INSTALLS = \$(EPICS.BASE_BIN)/aiRecord\$(OBJ))
BIN_INSTALLS_<osclass>	os class specific files from specified directories to be installed into \$(INSTALL_BIN)
BIN_INSTALLS_DEFAULT	files from specified directories to be installed into \$(INSTALL_BIN) for target archs with no BIN_INSTALLS_<osclass> specified
LIB_INSTALLS	files from specified directories to be installed into \$(INSTALL_LIB)
LIB_INSTALLS_<osclass>	os class specific files from specified directories to be installed into \$(INSTALL_LIB)
LIB_INSTALLS_DEFAULT	files from specified directories to be installed into \$(INSTALL_LIB) for target archs with no LIB_INSTALLS_<osclass> specified
TARGETS	files to create but not install
INSTALL_LOCATION	installation directory (defaults to \$(TOP))

4.8 Configuration Files

4.8.1 Base Configure Directory

The base/configure directory has the following directory structure:

```
base/
  configure/
    os/
    tools/
```

4.8.2 Base Configure File Descriptions

The configure files contain definitions and make rules to be included in the various makefiles.

CONFIG.CrossCommon

Definitions for all hosts and all targets for a cross build (host different than target).

CONFIG.gnuCommon

Definitions for all hosts and all targets for builds using the gnu compiler.

CONFIG_ADDONS

Definitions which setup the variables that have <osclass> and DEFAULT options.

CONFIG_APP_INCLUDE

Definitions to generate include, bin, lib, perl module, db, and dbd directory definitions for RELEASE <top>s.

CONFIG_BASE

EPICS base specific definitions.

CONFIG_BASE_VERSION

Definitions for the version number of EPICS base. This file is used for creating epicsVersion.h which is installed into base/include.

CONFIG_COMMON

Definitions common to all builds.

CONFIG_ENV

Default definitions of the EPICS environment variables. This file is used for creating envData.c which is included in the Com library.

CONFIG_FILE_TYPE

Definitions to allow user created file types.

CONFIG_SITE

File in which you add to or modify make variables in EPICS base. A definition commonly overridden is CROSS_COMPILER_TARGET_ARCHS

CONFIG_SITE_ENV

Defaults for site specific definitions of EPICS environment variables. This file is used for creating envData.c which is included in the Com library.

CONFIG

Include statements for all the other configure files. You can override any definitions in other CONFIG* files by placing override definitions at the end of this file.

RELEASE

Specifies the location of external products such as Tornado II and external <tops> such as EPICS base.

RULES

This file just includes the appropriate rules configuration file.

RULES.Db

Rules for building and installing database and database definition files. Databases generated from templates and/or CapFast schematics are supported.

RULES.ioc

Rules which allow building in the iocBoot/<iocname> directory of a makeBaseApp created ioc application.

RULES_ARCHS

Definitions and rules which allow building the make target for each target architecture.

RULES_BUILD

Build rules for the Makefiles

RULES_DIRS

Definitions and rules which allow building the make targets in each subdirectory. This file is included by Makefiles in directories with subdirectories to be built.

RULES_EXPAND

Definitions and rules to use `expandVars.pl` to expand `@VAR@` variables in a file.

RULES_FILE_TYPE

Definitions and rules to allow user created `CONFIG*` and `RULES*` files and rules to allow user created file types.

`RULES_JAVA` Definitions and rules which allow building java class files and java jar files.

RULES_TARGET

Makefile code to create target specific dependency lines for libraries and product targets.

RULES_TOP

Rules specific to a `<top>` level directory e.g. `uninstall` and `tar`. It also includes the `RULES_DIRS` file.

Makefile Definitions to allow creation of `CONFIG_APP_INCLUDE` and installation of the `CONFIG*` files into the `$(INSTALL_LOCATION)` directory.

4.8.3 Base configure/os File Descriptions

The `configure/os` directory contains os specific make definitions. The naming convention for the files in this directory is `CONFIG.<host>.<target>` where `<host>` is either the arch for a specific host system or `Common` for all supported host systems and `<target>` is either the arch for a specific target system or `Common` for all supported target systems.

For example, the file `CONFIG.Common.vxWorks-pentium` will contain make definitions to be used for builds on all host systems when building for a `vxWorks-pentium` target system.

Also, if a group of host or target files have the same make definitions these common definitions can be moved to a new file which is then included in each host or target file. An example of this is all Unix hosts which have common definitions in a `CONFIG.UnixCommon.Common` file and all `vxWorks` targets with definitions in `CONFIG.Common.vxWorksCommon`.

The `base/configure/os` directory contains the following os-arch specific definitions

`CONFIG.<host>.<target>`

Specific host-target build definitions

`CONFIG.Common.<target>`

Specific target definitions for all hosts

`CONFIG.<host>.Common`

Specific host definitions for all targets

`CONFIG.UnixCommon.Common`

Definitions for Unix hosts and all targets

`CONFIG.<host>.vxWorksCommon`

Specific host definitions for all vx targets

`CONFIG_COMPAT`

R3.13 arch compatibility definitions

`CONFIG_SITE.<host>.<target>`

Site specific host-target definitions

CONFIG_SITE.Common.<target>

Site specific target definitions for all hosts

CONFIG_SITE.<host>.Common

Site specific host definitions for all targets

4.8.4 Base src/tools File Descriptions

The src/tools directory contains Perl script tools used for the build. The are installed by the build into \$(INSTALL_LOCATION)/bin/\$(T_A) for Host type target archs. The tools currently in this directory are:

convertRelease.pl This Perl script does consistency checks for the external <top> definitions in the RELEASE file.

This script also creates envPaths, cdCommands, and dllPath.bat files for vxWorks and other IOCs.

cvsclean.pl This perl script finds and deletes cvs .#* files in all directories of the directory tree.

dos2unix.pl This perl script converts text file in DOS CR/LF format to unix ISO format.

expandVars.pl This perl tool expands @VAR@ variables while copying a file.

filterWarnings.pl This is a perl script that filters compiler warning output (for HP-UX).

fullpathname.pl This perl script returns the fullpathname of a file.

installEpics.pl This is a Perl script that installs build created files into the install directories.

makeDbDepends.pl This perl script searches .substitutions and .template files for entries to create a DEPENDS file.

makeIncludeDbd.pl This perl script creates an include dbd file from file names

makeMakefile.pl This is a perl script that creates a Makefile in the created O.<arch> directories.

makeTestfile.pl This perl script generates a file \$target.t which executes a real test program in the same directory.

mkmf.pl This perl script generates include file dependencies for targets from source file include statements.

munch.pl This is a perl script that creates a ctdt.c file for vxWorks target arch builds which lists the c++ static constructors and destructors. See munching in the vxWorks documentation for more information.

replaceVAR.pl This is a perl script that changes VAR(XXX) style macros in CapFast generated databases into the \$(XXX) notation used in EPICS databases.

useManifestTool.pl This tools uses MS Visual C++ compiler version number to determine if we want to use the Manifest Tool (status=1) or not (status=0).

4.9 Build Documentation Files

4.9.1 Base Documentation Directory

The base/documentation directory contains README files to help users setup and build epics/base.

4.9.2 Base Documentation File Descriptions

The files currently in the base/documentation directory are:

README.1st Instructions for setup and building epics base

README.html html version of README.1st

README.MS_WINDOWS Microsoft WIN32 specific instructions

README.niCpu030 NI cpu030 specific instructions

README.hpux HPUX 11 (hpux-parisc) specific instructions

README.cris Cris architecture specific instructions

README.tru64unix Tru64Unix/Alpha specific instructions

README.darwin.html Installation notes for Mac OS X (Darwin)

BuildingR3.13AppsWithR3.14.html Describes how to modify a R3.13 vxWorks application so that it builds with release R3.14.1.

ConvertingR3.13AppstoR3.14.html Describes how to convert a R3.13 vxWorks application so that it contains a R3.14 configure directory and R3.14 Makefiles and builds with R3.14.1.

ConvertingR3.14.0alpha2AppstoBeta1.html Describes how to modify a R3.14.0alpha1 application so that it builds with release R3.14.0beta1.

ConvertingR3.14.0beta1AppstoBeta2.html Describes how to modify a R3.14.0beta1 application so that it builds with release R3.14.0beta2.

ConvertingR3.14.0beta2AppstoR3.14.1.html Describes how to modify a R3.14.0beta2 application so that it builds with release R3.14.1.

ConvertingR3.14.*AppstoR3.14.*.html Describes how to modify a R3.14.* application so that it builds with next release after R3.14.*.

BuildingR3.13ExtensionsWithR3.14.html Describes how to modify a R3.13 extension so that it builds with release R3.14.1.

RELEASE_NOTES.html Describes changes in the R3.14.1 release

KnownProblems.html List of known problems in EPICS base R3.14.1.

ReleaseChecklist.html Checklist of things that must be done when creating a new release of EPICS Base.

4.10 Startup Files

4.10.1 Base Startup Directory

The base/startup directory contains scripts to help users set the required environment variables and path. The appropriate startup files should be executed before any EPICS builds.

4.10.2 Base Startup File Descriptions

The scripts currently in the base/startup directory are:

EpicsHostArch c shell script to set EPICS_HOST_ARCH environment variable

EpicsHostArch.pl perl script to set EPICS_HOST_ARCH environment variable

Site.profile Unix bourne shell script to set path and environment variables

Site.cshrc Unix c shell script to set path and environment variables

cygwin.bat WIN32 bat file to set path and environment variables for building with cygwin gcc/g++ compilers

win32.bat WIN32 bat file to set path and environment variables for building with MS Visual C++ compilers

Chapter 5

Database Locking, Scanning, And Processing

5.1 Overview

Before describing particular components of the IOC software, it is helpful to give an overview of three closely related topics: Database locking, scanning, and processing. Locking is done to prevent two different tasks from simultaneously modifying related database records. Database scanning is the mechanism for deciding when records should be processed. The basics of record processing involves obtaining the current value of input fields and outputting the current value of output fields. As records become more complex so does the record processing.

One powerful feature of the DATABASE is that records can contain links to other records. This feature also causes considerable complication. Thus, before discussing locking, scanning, and processing, record links are described.

5.2 Record Links

A database record may contain links to other records. Each link is one of the following types:

- INLINK
- OUTLINK

INLINKs and OUTLINKs can be one of the following:

- constant link

Not discussed in this chapter

- database link

A link to another record in the same IOC.

- channel access link

A link to a record in another IOC. It is accessed via a special IOC client task. It is also possible to force a link to be a channel access link even it references a record in the same IOC.

- hardware link

Not discussed in this chapter

- **FWDLINK**

A forward link refers to a record that should be processed whenever the record containing the forward link is processed. The following types are supported:

- constant link

Ignored.

- database link

A link to another record in the same IOC.

- channel access link

A link to a record in another IOC or a link forced to be a channel access link. Unless the link references the PROC field it is ignored. If it does reference the PROC field a channel access put with a value of 1 is issued.

Links are defined in file `link.h`.

NOTE: This chapter discusses mainly database links.

5.3 Database Links

Database links are referenced by calling one of the following routines:

- `dbGetLink`: The value of the field referenced by the input link retrieved.
- `dbPutLink`: The value of the field referenced by the output link is changed.
- `dbScanPassive`: The record referred to by the forward link is processed if it is passive.

A forward link only makes sense when it refers to a passive record that should be processed after the record that contains the link.

For input and output links, two additional attributes can be specified by the application developer: `process passive`, and `maximize severity`.

5.3.1 Process Passive

The Process Passive attribute takes the value `NPP` (Non-Process Passive) or `PP` (Process Passive). It determines if the linked record should be processed before getting a value from an input link or after writing a value to an output link. The linked record will be processed only if link's Process Passive attribute is `PP` and the target record's `SCAN` field is `Passive`.

NOTE: Three other options may also be specified: `CA`, `CP`, and `CPP`. These options force the link to be handled like a Channel Access Link. See last section of this chapter for details.

5.3.2 Maximize Severity

The Maximize Severity attribute is one of `NMS` (Non-Maximize Severity), `MS` (Maximize Severity), `MSS` (Maximize Status and Severity) or `MSI` (Maximize Severity if Invalid). It determines whether alarm severity is propagated across links. If the attribute is `MSI` only a severity of `INVALID_ALARM` is propagated; settings of `MS` or `MSS` propagate all alarms that are more severe than the record's current severity. For input links the alarm severity of the record referred to by the link is propagated to the record containing the link. For output links the alarm severity of the record containing the link is propagated to the record referred to by the link. If the severity is changed the associated alarm status is set to `LINK_ALARM`, except if the attribute is `MSS` when the alarm status will be copied along with the severity.

The method of determining if the alarm status and severity should be changed is called “maximize severity”. In addition to its actual status and severity, each record also has a new status and severity. The new status and severity are initially 0, which means `NO_ALARM`. Every time a software component wants to modify the status and severity, it first checks the new severity and only makes a change if the severity it wants to set is greater than the current new severity. If it does make a change, it changes the new status and new severity, not the current status and severity. When database monitors are checked, which is normally done by a record processing routine, the current status and severity are set equal to the new values and the new values reset to zero. The end result is that the current alarm status and severity reflect the highest severity outstanding alarm. If multiple alarms of the same severity are present the alarm status reflects the first one detected.

5.4 Database Locking

The purpose of database locking is to prevent a record from being processed simultaneously by two different tasks. In addition, it prevents “outside” tasks from changing any field while the record is being processed.

The following routines are provided for database locking.

```
dbScanLock (precord) ;
dbScanUnlock (precord) ;
```

The basic idea is to call `dbScanLock` before accessing database records and calling `dbScanUnlock` afterwards. Because of database links (Input, Output, and Forward) a modification to one record can cause modification to other records. Records linked together with database links are placed in the same lock set. `dbScanLock` locks the entire lock set, not just the record requested. `dbScanUnlock` unlocks the entire set.

The following rules determine when the lock routines must be called:

1. The periodic, I/O event, and event tasks lock before and unlock after processing:
2. `dbPutField` locks before modifying a record and unlocks afterwards.
3. `dbGetField` locks before reading and unlocks afterwards.
4. Any asynchronous record support completion routine must lock before modifying a record and unlock afterwards.

All records connected by any kind of database link are placed in the same lock set. Versions of EPICS Base prior to R3.14 allowed an NPP NMS input link to span two different lock sets, but this was not safe where the read and write operations on the field value were not atomic in nature and is no longer available to break a lockset.

5.5 Database Scanning

Database scanning refers to requests that database records be processed. Four types of scanning are possible:

1. Periodic - Records are scanned at regular intervals.
2. I/O event - A record is scanned as the result of an I/O interrupt.
3. Event - A record is scanned as the result of any task issuing a `post_event` request.
4. Passive - A record is scanned as a result of a call to `dbScanPassive`. `dbScanPassive` will issue a record processing request if and only if the record is passive and is not already being processed.

A `dbScanPassive` request results from a task calling one of the following routines:

- `dbScanPassive`: Only record processing routines, `dbGetLink`, `dbPutLink`, and `dbPutField` call `dbScanPassive`. Record processing routines call it for each forward link in the record.

- **dbPutField:** This routine changes the specified field and then, if the field has been declared `process_passive`, calls `dbScanPassive`. Each field of each record type has the attribute `process_passive` declared `TRUE` or `FALSE` in the definition file. The attribute is a global property which is set by the record type. This use of `process_passive` only affects calls to the `dbPutField` routine. If `dbPutField` finds the record already active (this can happen to asynchronous records) and it is supposed to cause it to process, it arranges for it to be processed again once the current processing completes.
- **dbGetLink:** If the link specifies `process passive`, this routine calls `dbScanPassive`. Whether or not `dbScanPassive` is called, it then obtains the specified value.
- **dbPutLink:** This routine changes the specified field. Then, if the link specifies `process passive`, it calls `dbScanPassive`. `dbPutLink` is only called from record processing routines. Note that this usage of `process_passive` is under the control of the application developer. If `dbPutLink` finds the record already active because of a `dbPutField` directed to this record then it arranges for the record to be processed again, once the current processing completes.


All non-record processing tasks (Channel Access, Sequence Programs, etc.) call `dbGetField` to obtain database values. `dbGetField` just reads values without asking that a record be processed.

5.6 Record Processing

A record is processed as a result of a call to `dbProcess`. Each record support module must supply a routine `process`. This routine does most of the work related to record processing. Since the details of record processing are record type specific this topic is discussed in greater detail in the Chapter “Record Support”.

5.7 Guidelines for Creating Database Links

The ability to link records together is an extremely powerful feature of the IOC software. In order to use links properly it is important that the Application Developer understand how they are processed. As an introduction consider the following example:



Assume that A, B, and C are all passive records. The notation states that A has a forward link to B and B to C. C has an input link obtaining a value from A. Assume, for some reason, A gets processed. The following sequence of events occurs:

1. A begins processing. While processing a request is made to process B.
2. B starts processing. While processing a request is made to process C.
3. C starts processing. One of the first steps is to get a value from A via the input link.
4. At this point a question occurs. Note that the input link specifies `process passive` (signified by the `PP` after `InLink`). But `process passive` states that A should be processed before the value is retrieved. Are we in an infinite loop? The answer is no. Every record contains a field `PACT` (`processing active`), which is set `TRUE` when record processing begins and is not set `FALSE` until all processing completes. When C is processed A still has `PACT TRUE` and will not be processed again.
5. C obtains the value from A and completes its processing. Control returns to B.
6. B completes returning control to A

7. A completes processing.


This brief example demonstrates that database links need more discussion.

5.7.1 Rules Relating to Database Links

5.7.1.1 Processing Order

The processing order follows the following rules:

1. Forward links are processed in order from left to right and top to bottom. For example the following records are processed in the order FLNK1, FLNK2, FLNK3, FLNK4 .



2. If a record has multiple input links (such as the calculation or select records) the input values are normally fetched in the natural order. For example for link fields named INPA, INPB, ..., INPL, the links would be read in the order A, B, C etc. Thus if obtaining an input results in a record being processed, the processing order is guaranteed. Some record types may not follow this rule however.
3. All input and output links are processed before the forward link.

5.7.1.2 Lock Sets


All records, except for the conditions listed in the next paragraph, linked together directly or indirectly are placed in the same lock set. When `dbScanLock` is called the entire set, not just the specified record, is locked. This prevents two different tasks from simultaneously modifying records in the same lock set.

5.7.1.3 PACT - Process Active

Every record contains a field `PACT`. This field is set `TRUE` at the beginning of record processing and is not set `FALSE` until the record is completely processed. To prevent infinite processing loops, whenever a record gets processed through a forward link, or a database link with the `PP` link option, the linking record's `PACT` field is saved and set to `TRUE`, then restored again afterwards. The example given at the beginning of this section gives an example. It will be seen in the next two sections that `PACT` has other uses.

5.7.1.4 Process Passive: Link option


Input and output links have an option called process passive. For each such link the application developer can specify process passive `TRUE` (`PP`) or process passive `FALSE` (`NPP`). Consider the following example:



Assume that all records except fanout are passive. When the fanout record is processed the following sequence of events occur:

1. Fanout starts processing and asks that B be processed.
2. B begins processing. It calls `dbGetLink` to obtain data from A.
3. Because the input link has `process passive true`, a request is made to process A.
4. A is processed, the data value fetched, and control is returned to B
5. B completes processing and control is returned to fanout. Fanout asks that C be processed.
6. C begins processing. It calls `dbGetLink` to obtain data from A.
7. Because the input link has `process passive TRUE`, a request is made to process A.
8. A is processed, the data value fetched, and control is returned to C.
9. C completes processing and returns to fanout
10. The fanout completes

Note that A was processed twice. This is unnecessary. If the input link to C were declared `No Process Passive` then A would only be processed once. Thus a better solution would be:



5.7.1.5 Process Passive: Field attribute

All record type field definitions have an attribute called `process_passive` which is specified in the record definition file. It cannot be changed by an IOC application developer. This attribute is used only by `dbPutField`. It determines if a passive record will be processed after `dbPutField` sets a field in the record. Consult the record specific information in the record reference manual for the setting of individual fields.

5.7.1.6 Maximize Severity: Link option

Input and output links have an option called `maximize severity`. For each such link the application developer can specify the option as `MS` (Maximize Severity), `NMS` (Non-Maximize Severity), `MSS` (Maximize Status and Severity) or `MSI` (Maximize Severity if Invalid).

When database input or output links are defined, the application developer can use this option to specify whether and how alarm severities should be propagated across links with the data. The alarm severity is transferred only if the new severity will be greater than the current severity of the destination record. If the severity is propagated the alarm status is set equal to `LINK_ALARM` (unless the link option is `MSS` when the alarm status will also be copied from the source record).

5.8 Guidelines for Synchronous Records

A synchronous record is a record that can be completely processed without waiting. Thus the application developer never needs to consider the possibility of delays when he defines a set of related records. The only consideration is deciding when records should be processed and in what order a set of records should be processed.

The following reviews the methods available to the application programmer for deciding when to process a record and for enforcing the order of record processing.

1. A record can be scanned periodically (at one of several rates), via I/O event, or via Event.
2. For each periodic group and for each Event group the PHAS field can be used to specify processing order.
3. The application programmer has no control over the record processing order of records in different groups.
4. The disable fields (SDIS, DISA, and DISV) can be used to disable records from being processed. By letting the SDIS field of an entire set of records refer to the same input record, the entire set can be enabled or disabled simultaneously. See the Record Reference Manual for details.
5. A record (periodic or other) can be the root of a set of passive records that will all be processed whenever the root record is processed. The set is formed by input, output, and forward links.
6. The `process_passive` attribute of each record field determines if a passive record will be processed when a `dbPutField` is directed to the field. The application developer must be aware of the possibility of record processing being triggered by external sources using this mechanism.
7. The `process_passive` option for input and output links provides the application developer control over how a set of records are scanned.
8. General link structures can be defined. The application programmer should be wary, however, of defining arbitrary structures without carefully analyzing the processing order.

5.9 Guidelines for Asynchronous Records

The previous discussion does not cover asynchronous device support. An example might be a GPIB input record. When the record is processed the GPIB request is started and the processing routine returns. Processing, however, is not really complete until the GPIB request completes. This is handled via an asynchronous completion routine. Let's state a few attributes of asynchronous record processing.

During the initial processing for all asynchronous records the following is done:

1. PACT is set TRUE
2. Data is obtained for all input links
3. Record processing is started
4. The record processing routine returns


The asynchronous completion routine performs the following algorithm:

1. Record processing continues
2. Record specific alarm conditions are checked
3. Monitors are raised
4. Forward links are processed
5. PACT is set FALSE.

A few attributes of the above rules are:

1. Asynchronous record processing does not delay the scanners.
2. Between the time that record processing begins and the asynchronous completion routine completes, no attempt will be made to again process the record. This is because `PACT` is `TRUE`. The routine `dbProcess` checks `PACT` and does not call the record processing routine if it is `TRUE`. Note, however, that if `dbProcess` finds the record active 10 times in succession, it raises a `SCAN_ALARM`.
3. Forward and output links are triggered only when the asynchronous completion routine completes record processing.

With these rules the following works just fine:




When `dbProcess` is called for record `ASYN`, processing will be started but `dbScanPassive` will not be called. Until the asynchronous completion routine executes any additional attempts to process `ASYN` are ignored. When the asynchronous callback is invoked the `dbScanPassive` is performed.

Problems still remain. A few examples are:

5.9.1 Infinite Loop

Infinite processing loops are possible.




Assume both `A` and `B` are asynchronous passive records and a request is made to process `A`. The following sequence of events occur.

1. `A` starts record processing and returns leaving `PACT TRUE`.
2. Sometime later the record completion for `A` occurs. During record completion a request is made to process `B`. `B` starts processing and control returns to `A` which completes leaving its `PACT` field `FALSE`.
3. Sometime later the record completion for `B` occurs. During record completion a request is made to process `A`. `A` starts processing and control returns to `B` which completes leaving its `PACT` field `FALSE`.

Thus an infinite loop of record processing has been set up. It is up to the application developer to prevent such loops.

5.9.2 Obtain Old Data


A `dbGetLink` to a passive asynchronous record can get old data.



If `A` is a passive asynchronous record then the `dbGetLink` request forces `dbProcess` to be called for `A`. `dbProcess` starts the processing and returns. `dbGetLink` then reads the field value which is still old because processing will only be completed at a later time.

5.9.3 Delays

Consider the following:



The second ASYN record will not begin processing until the first completes, etc. This is not really a problem except that the application developer must be aware of delays caused by asynchronous records. Again, note that scanners are not delayed, only records downstream of asynchronous records.

5.10 Cached Puts

The rules followed by `dbPutLink` and `dbPutField` provide for “cached” puts. This is necessary because of asynchronous records. Two cases arise.

The first results from a `dbPutField`, which is a put coming from outside the database, i.e. Channel Access puts. If this is directed to a record that already has `PACT TRUE` because the record started processing but asynchronous completion has not yet occurred, then a value is written to the record but nothing will be done with the value until the record is again processed. In order to make this happen `dbPutField` arranges to have the record reprocessed when the record finally completes processing.

The second case results from `dbPutLink` finding a record already active because of a `dbPutField` directed to the record. In this case `dbPutLink` arranges to have the record reprocessed when the record finally completes processing. If the record is already active because it appears twice in a chain of record processing, it is not reprocessed because the chain of record processing would constitute an infinite loop.

Note that the term caching not queuing is used. If multiple requests are directed to a record while it is active, each new value is placed in the record but it will still only be processed once, i.e. last value wins.

5.11 putNotify

`dbPutNotify` (called when a Channel Access client calls `ca_put_callback`) is a request to notify the caller when all records processed as a result of the put are complete. Because of asynchronous records this can be complicated and the set of records that are processed because of a put may not be deterministic. The result of a `dbPutNotify` is the same as a `dbPutField` except for the following:

- `dbPutNotify` requests are queued rather than cached. Thus when additional requests are directed to a record that already has an active `dbPutNotify`, they are queued. As each one finishes it releases the next one in the queue.
- If a `dbPutNotify` links to a record that is not active but has a `dbPutNotify` attached to it, no attempt is made to process the record.

5.12 Channel Access Links

A channel access link is:

1. A record link that references a record in a different IOC.
2. A link that the application developer forces to be a channel access link.

A channel access client task (`dbCa`) handles all I/O for channel access links. It does the following:

- At IOC initialization, `dbCa` issues channel access search requests for each channel access link.
- For each input link it establishes a channel access monitor. It uses `ca_field_type` and `ca_element_count` when it establishes the monitor. It also monitors the alarm status. Whenever the monitor is invoked the new data

is stored in a buffer belonging to dbCa. When iocCore or the record support module asks for data the data is taken from the buffer and converted to the requested type.

- For each output link, a buffer is allocated the first time iocCore/record support issues a put and a channel access connection has been made. This buffer is allocated according to `ca_field_type` and `ca_element_count`. Each time iocCore/record support issues a put, the data is converted and placed in the buffer and a request is made to dbCa to issue a new `ca_put`.

Even if a link references a record in the same IOC it can be useful to force it to act like a channel access link. In particular the records will not be forced to be in the same lock set. As an example consider a scan record that links to a set of unrelated records, each of which can cause a lot of records to be processed. It is often NOT desirable to force all these records into the same lock set. Forcing the links to be handled as channel access links solves the problem.

CA links which connect between IOCs incur the extra overhead associated with message passing protocols, operating system calls, and network activity. In contrast, CA links which connect records in the same IOC are executed more efficiently by directly calling database access functions such as `dbPutField` and `dbGetField`, or by receiving callbacks directly from a database monitor subscription event queue.

Because channel access links interact with the database only via `dbPutField`, `dbGetField`, and a database monitor subscription event queue, their interaction with the database is fundamentally different from database links which are tightly integrated within the code that executes database records. For this reason and because channel access does not support the passing of a process passive flag, the semantics of channel access links are not the same as database links. Let's discuss the channel access semantics of INLINK, OUTLINK, and FWDLINK separately.

5.12.1 INLINK

The options for process passive are:

- Input links always act like NPP.
- CA - Forces the link to be a channel access link.
- CP - Forces the link to be a channel access link and also requests that the record containing the link be processed whenever a monitor occurs.
- CPP - Force the link to be a channel access link and also requests that the record containing the link, if it is passive, be processed whenever a monitor occurs.

Maximize Severity is honored.

5.12.2 OUTLINK

The options for process passive are:

- It is not possible to honor PP or NPP options; the put operation completes immediately but whether the destination record will process depends on the process passive attribute of the destination field.
- CA - Force the link to be a channel access link.

Maximize Severity is not honored.

5.12.3 FWDLINK

A channel access forward link is honored only if it references the `PROC` field of a record. In that case a `ca_put` with a value of 1 is written each time a forward link request is issued.

The available options are:

- CA - Force the link to be a channel access link.

Maximize Severity is not honored.

Chapter 6

Database Definition

6.1 Overview

This chapter describes database definitions. The following definitions are described:

- Menu
- Record Type
- Device
- Driver
- Registrar
- Variable
- Function
- Breakpoint Table
- Record Instance

Record Instances are fundamentally different from the other definitions. A file containing record instances should never contain any of the other definitions and vice-versa. Thus the following convention is followed:

Database Definition File A file that contains any type of definition except record instances.

Record Instance File A file that contains only record instance definitions.

This chapter also describes utility programs which operate on these definitions

Any combination of definitions can appear in a single file or in a set of files related to each other via include files.

6.2 Summary of Database Syntax

The following summarizes the Database Definition syntax:

```
path "path"
addpath "path"
include "filename"
#comment
menu (name) {
    include "filename"
```

```

        choice(choice_name, "choice_value")
        ...
    }

recordtype(record_type) {
    include "filename"
    field(field_name, field_type) {
        asl(asl_level)
        initial("init_value")
        promptgroup(gui_group)
        prompt("prompt_value")
        special(special_value)
        pp(pp_value)
        interest(interest_level)
        base(base_type)
        size(size_value)
        extra("extra_info")
        menu(name)
    }
    %C_declaration
    ...
}

device(record_type, link_type, dset_name, "choice_string")

driver(drvet_name)

registrar(function_name)

variable(variable_name)

breaktable(name) {
    raw_value eng_value
    ...
}

```

The Following defines a Record Instance

```

record(record_type, record_name) {
    include "filename"
    field(field_name, "value")
        alias(alias_name)
        info(info_name, "value")
        ...
}
alias(record_name, alias_name)

```

6.3 General Rules for Database Definition

6.3.1 Keywords

The following are keywords, i.e. they may not be used as values unless they are enclosed in quotes:


```

path
addpath
include
menu
choice
recordtype
field
device
driver
registrar
function
variable
breaktable
record
grecord
info
alias

```

6.3.2 Unquoted Strings

In the summary section, some values are shown as quoted strings and some unquoted. The actual rule is that any string consisting of only the following characters does not have to be quoted unless it contains one of the above keywords:

```
a-z A-Z 0-9 _ + - : . [ ] < > ;
```

These are all legal characters for process variable names, although `.` is not allowed in a record name since it separates the record from the field name in a PV name. Thus in many cases quotes are not needed around record or field names in database files. Any string containing a macro does need to be quoted though.

6.3.3 Quoted Strings

A quoted string can contain any ascii character except the quote character `"`. The quote character itself can given by using `\` as an escape. For example `"\"` is a quoted string containing the single character `"`.

6.3.4 Macro Substitution

Macro substitutions are permitted inside quoted strings. Macro instances take the form:

```
$(name)
```

or

```
${name}
```

There is no distinction between the use of parentheses or braces for delimiters, although the two must match for a given macro instance. The macro name can be made up from other macros, for example:

```
$(name_$(sel))
```

A macro instance can also provide a default value that is used when no macro with the given name is defined. The default value can be defined in terms of other macros if desired, but cannot contain any unescaped comma characters. The syntax for specifying a default value is as follows:

```
$(name=default)
```

Finally macro instances can also contain definitions of other macros, which can (temporarily) override any existing values for those macros but are in scope only for the duration of the expansion of this macro instance. These definitions consist of `name=value` sequences separated by commas, for example:

```
$ (abcd=$ (a) $ (b) $ (c) $ (d) , a=A, b=B, c=C, d=D)
```

6.3.5 Escape Sequences

The database routines translate standard C escape sequences inside database field value strings only. The standard C escape sequences supported are:

```
\a \b \f \n \r \t \v \\ \? \' \" \ooo \xhh
```

\ooo represents an octal number with 1, 2, or 3 digits. \xhh represents a hexadecimal number with 1 or 2 digits.

6.3.6 Comments

The comment symbol is “#”. Whenever the comment symbol appears, it and all characters through the end of the line are ignored.

6.3.7 Define before referencing

No item can be referenced until it is defined. For example a `recordtype` menu field can not reference a menu unless that menu definition has already been defined. Another example is that a record instance can not appear until the associated record type has been defined.

6.3.8 Multiple Definitions

If a menu, `recordtype`, device, driver, or breakpoint table is defined more than once, then only the first instance is used. Record instance definitions however are (normally) cumulative, so multiple instances of the same record may be loaded and each time a field value is encountered it replaces the previous value.

6.3.9 Filename Extensions

By convention:

- Record instances files have the extension “.db” or “.vdb” if the file also contains visual layout information
- Database definition files have the extension “.dbd”

6.4 path addpath – Path Definition

6.4.1 Format

```
path "dir:dir...:dir"
addpath "dir:dir...:dir"
```

The path string follows the standard convention for the operating system, i.e. directory names are separated by a colon “:” on Unix and a semicolon “;” on Windows.

The `path` command specifies the current search path for use when loading database and database definition files. The `addpath` appends directory names to the current path. The path is used to locate the initial database file and included files. An empty `dir` at the beginning, middle, or end of a non-empty path string means the current directory. For example:

```
nnn::mmm      # Current directory is between nnn and mmm
:nnn          # Current directory is first
nnn:         # Current directory is last
```

Utilities which load database files (`dbExpand`, `dbLoadDatabase`, etc.) allow the user to specify an initial path. The `path` and `addpath` commands can be used to change or extend the initial path.

The initial path is determined as follows:

If an initial path is specified, it is used. Else:

If the environment variable `EPICS_DB_INCLUDE_PATH` is defined, it is used. Else:

the default path is “.”, i.e. the current directory.

The path is used unless the filename contains a / or \. The first directory containing the specified filename is used.

6.5 include – Include File

6.5.1 Format

```
include "filename"
```

An include statement can appear at any place shown in the summary. It uses the path as specified above.

6.6 menu – Menu Declaration

6.6.1 Format

```
menu (name) {
    choice(choice_name, "choice_string")
    ...
}
```

6.6.2 Definitions

name Name for menu. This is the unique name identifying the menu. If duplicate definitions are specified, only the first is used.

choice_name The name used in the `enum` generated by `dbToMenuH` or `dbToRecordtypeH`. This must be a legal C/C++ identifier.

choice_string The text string associated with this particular choice.

6.6.3 Example

```
menu(menuYesNo) {
    choice(menuYesNoNO, "NO")
    choice(menuYesNoYES, "YES")
}
```

6.7 recordtype – Record Type Declaration

6.7.1 Format

```
recordtype(record_type) {
    field(field_name, field_type) {
        asl(as_level)
        initial("init_value")
        promptgroup(gui_group)
        prompt("prompt_value")
        special(special_value)
        pp(pp_value)
        interest(interest_level)
        base(base_type)
        size(size_value)
        extra("extra_info")
        menu("name")
    }
    %C_declaration
    ...
}
```

6.7.2 Field Definition Rules

asl Sets the Access Security Level for the field. Access Security is discussed in chapter 8.

initial Provides an initial (default) value for the field.

promptgroup The group to which the field belongs, for database configuration tools.

prompt A prompt string for database configuration tools. Optional if `promptgroup` is not defined.

special If specified, special processing is required for this field at run time.

pp Whether a passive record should be processed when Channel Access writes to this field.

interest Interest level for the field.

base For integer fields, the number base to use when converting the field value to a string.

size Must be specified for `DBF_STRING` fields.

extra Must be specified for `DBF_NOACCESS` fields.

menu Must be specified for `DBF_MENU` fields. It is the name of the associated menu.

6.7.3 Definitions

record_type The unique name of the record type. If duplicates are specified, only the first definition is used.

field_name The field name, which must be a valid C identifier. When include files are generated, the field name is converted to lower case. Previous versions of EPICS required the field name be a maximum of four characters, but this restriction no longer applies.

field_type This must be one of the following values:

- DBF_STRING
- DBF_CHAR, DBF_UCHAR
- DBF_SHORT, DBF_USHORT
- DBF_LONG, DBF_ULONG
- DBF_FLOAT, DBF_DOUBLE
- DBF_ENUM, DBF_MENU, DBF_DEVICE
- DBF_INLINK, DBF_OUTLINK, DBF_FWDLINK
- DBF_NOACCESS

as_level This must be one of the following values:

- ASL0
- ASL1 (default value)

Fields which operators normally change are assigned ASL0. Other fields are assigned ASL1. For example, the VAL field of an analog output record is assigned ASL0 and all other fields ASL1. This is because only the VAL field should be modified during normal operations.

init_value A legal value for data type.

prompt_value A prompt value for database configuration tools.

gui_group This must be one of the following:

- GUI_COMMON
- GUI_ALARMS
- GUI_BITS1
- GUI_BITS2
- GUI_CALC
- GUI_CLOCK
- GUI_COMPRESS
- GUI_CONVERT
- GUI_DISPLAY
- GUI_HIST
- GUI_INPUTS
- GUI_LINKS
- GUI_MBB
- GUI_MOTOR

- GUI_OUTPUT
- GUI_PID
- GUI_PULSE
- GUI_SELECT
- GUI_SEQ1
- GUI_SEQ2
- GUI_SEQ3
- GUI_SUB
- GUI_TIMER
- GUI_WAVE
- GUI_SCAN

This information is for use by Database Configuration Tools. This is defined only for fields that can be given values by database configuration tools. File `guigroup.h` contains all possible definitions. This allows database configuration tools to group fields together by functionality, not just order them by name. This feature has seldom been used, so many record types do not have appropriate values assigned to some fields.

special_value Must be one of the following:

- SPC_MOD – Notify record support when modified. The record support `special` routine will be called whenever the field is modified by the database access routines.
- SPC_NOMOD – No external modifications allowed. This value disables external writes to the field, so it can only be set by the record or device support module.
- SPC_DBADDR – Use this if the record support `cvt_dbaddr` routine should be called by `dbNameToAddr`, i.e. when code outside record/device support is connecting to the field.

The following values are for database common fields. They must *not* be used for record specific fields:

- SPC_SCAN – Scan related field.
- SPC_ALARMACK – Alarm acknowledgment field.
- SPC_AS – Access security field.

The following values are deprecated, use `SPC_MOD` instead:

- An integer value greater than 103.
- SPC_RESET – a reset field is being modified.
- SPC_LINCONV – A linear conversion field is being modified.
- SPC_CALC – A calc field is being modified.

pp_value Should a passive record be processed when Channel Access writes to this field? The allowed values are:

- NO (default)
- YES

interest_level An interest level for the `dbpr` command.

base For integer type fields, the default base. The legal values are:

- DECIMAL (Default)

- HEX

size_value The number of characters for a DBF_STRING field.

extra_info For DBF_NOACCESS fields, this is the C language definition for the field. The definition must end with the fieldname in lower case.

%C.declaration A percent sign % inside the record body introduces a line of code that is to be included in the generated C header file.

6.7.4 Example

The following is the definition of the event record type:

```
recordtype(event) {
    include "dbCommon.dbd"
    field(VAL,DBF_USHORT) {
        prompt("Event Number To Post")
        promptgroup(GUI_INPUTS)
        asl(ASL0)
    }
    field(INP,DBF_INLINK) {
        prompt("Input Specification")
        promptgroup(GUI_INPUTS)
        interest(1)
    }
    field(SIOL,DBF_INLINK) {
        prompt("Sim Input Specifctn")
        promptgroup(GUI_INPUTS)
        interest(1)
    }
    field(SVAL,DBF_USHORT) {
        prompt("Simulation Value")
    }
    field(SIML,DBF_INLINK) {
        prompt("Sim Mode Location")
        promptgroup(GUI_INPUTS)
        interest(1)
    }
    field(SIMM,DBF_MENU) {
        prompt("Simulation Mode")
        interest(1)
        menu(menuYesNo)
    }
    field(SIMS,DBF_MENU) {
        prompt("Sim mode Alarm Svrty")
        promptgroup(GUI_INPUTS)
        interest(2)
        menu(menuAlarmSevr)
    }
}
```

6.8 device – Device Support Declaration

6.8.1 Format

```
device(record_type, link_type, dset_name, "choice_string")
```

6.8.2 Definitions

record_type Record type. The combination of `record_type` and `choice_string` must be unique. If the same combination appears more than once, only the first definition is used.

link_type Link type. This must be one of the following:

- CONSTANT
- PV_LINK
- VME_IO
- CAMAC_IO
- AB_IO
- GPIB_IO
- BITBUS_IO
- INST_IO
- BBGPIB_IO
- RF_IO
- VXI_IO

dset_name The name of the device support entry table for this device support.

choice_string The DTYP choice string for this device support. A `choice_string` value may be reused for different record types, but must be unique for each specific record type.

6.8.3 Examples

```
device(ai, CONSTANT, devAiSoft, "Soft Channel")
device(ai, VME_IO, devAiXy566Se, "XYCOM-566 SE Scanned")
```

6.9 driver – Driver Declaration

6.9.1 Format

```
driver(drvet_name)
```

6.9.2 Definitions

drvet_name If duplicates are defined, only the first is used.

6.9.3 Examples

```
driver(drvVxi)
driver(drvXy210)
```

6.10 registrar – Registrar Declaration

6.10.1 Format

```
registrar(function_name)
```

6.10.2 Definitions

function_name The name of an C function that accepts no arguments, returns `void` and has been marked in its source file with an `epicsExportRegistrar` declaration, e.g.

```
static void myRegistrar(void);
epicsExportRegistrar(myRegistrar);
```

This can be used to register functions for use by subroutine records or that can be invoked from `iocsh`. The example application described in Section 2.2, “Example IOC Application” gives an example of how to register functions for subroutine records.

6.10.3 Example

```
registrar(myRegistrar)
```

6.11 variable – Variable Declaration

6.11.1 Format

```
variable(variable_name[, type])
```

6.11.2 Definitions

variable_name The name of a C variable which has been marked in its source file with an `epicsExportAddress` declaration.

type The C variable’s type. If not present, `int` is assumed. Currently only `int` and `double` variables are supported.

This registers a diagnostic/configuration variable for device or driver support or a subroutine record subroutine. This variable can be read and set with the `iocsh var` command (see Section 18.2.5). The example application described in Section 2.2 shows how to register a debug variable for use in a subroutine record.

6.11.3 Example

In an application C source file:

```
#include <epicsExport.h>

static double myParameter;
epicsExportAddress(double, myParameter);
```

In an application database definition file:

```
variable(myParameter, double)
```

6.12 function – Function Declaration

6.12.1 Format

```
function(function_name)
```

6.12.2 Definitions

function_name The name of a C function which has been exported from its source file with an `epicsRegisterFunction` declaration.

This registers a function so that it can be found in the function registry for use by record types such as `sub` or `aSub` which refer to the function by name. The example application described in Section 2.2 shows how to register functions for a subroutine record.

6.12.3 Example

In an application C source file:

```
#include <epicsExport.h>
#include <registryFunction.h>

static long myFunction(void *argp) {
    /* my code ... */
}
epicsRegisterFunction(myFunction);
```

In an application database definition file:

```
function(myFunction)
```

6.13 breaktable – Breakpoint Table

6.13.1 Format

```
breaktable(name) {
    raw_value eng_value
    ...
}
```

6.13.2 Definitions

name Name, which must be alpha-numeric, of the breakpoint table. If duplicates are specified the first is used.

raw_value The raw value, i.e. the actual ADC value associated with the beginning of the interval.

eng_value The engineering value associated with the beginning of the interval.

6.13.3 Example

```
breaktable(typeJdegC) {
    0.000000 0.000000
    365.023224 67.000000
    1000.046448 178.000000
    3007.255859 524.000000
    3543.383789 613.000000
    4042.988281 692.000000
    4101.488281 701.000000
}
```

6.14 record – Record Instance

6.14.1 Format

```
record(record_type, record_name) {
    alias(alias_name)
    field(field_name, "field_value")
    info(info_name, "info_value")
    ...
}
alias(record_name, alias_name)
```

6.14.2 Definitions

record_type The record type.

record_name The record name. This must be composed of the following characters:

```
a-z A-Z 0-9 _ - + : [ ] < > ;
```

NOTE: If macro substitutions are used the name must be quoted.

Duplicate definitions are normally allowed for a record as long as the record type is the same. The last value given for each field is the value used.

The variable `dbRecordsOnceOnly` can be set to any non-zero value using the `iocsh var` command to make loading duplicate record definitions into the IOC illegal.

alias_name An alternate name for the record, following the same rules as the record name.

field_name A field name.

field_value A value for the named field, depending on the particular field type. Inside double quotes the field value string may contain escaped C89 characters such as `\`, `\t`, `\n`, `\064` and `\x7e`, and these will be translated appropriately when loading the database. Permitted values are as follows:

- `DBF_STRING`
Any ASCII string. If it exceeds the field length, it will be truncated.
- `DBF_CHAR`, `DBF_UCHAR`, `DBF_SHORT`, `DBF_USHORT`, `DBF_LONG`, `DBF_ULONG`
A string that represents a valid integer. The standard C conventions are applied, i.e. a leading 0 means the value is given in octal and a leading 0x means that value is given in hex.
- `DBF_FLOAT`, `DBF_DOUBLE`
The string must represent a valid floating point number.
- `DBF_MENU`
The string must be one of the valid choices for the associated menu.
- `DBF_DEVICE`
The string must be one of the valid device choice strings.
- `DBF_INLINK`, `DBF_OUTLINK`, `DBF_FWDLINK`

NOTES:

- If the field name is `INP` or `OUT` then this field is associated with `DTYP`, and the permitted values are determined by the link type of the device support selected by the current `DTYP` choice string. Other `DBF_INLINK` and `DBF_OUTLINK` fields must be either `CONSTANT` or `PV_LINKS`.
- A device support that specifies a link type of `CONSTANT` can be given either a constant or a `PV_LINK`.

The allowed values for the field depend on the device support's link type as follows:

- `CONSTANT`
A numeric literal, valid for the field type it is to be read into.
- `PV_LINK`
A value of the form:

```
record.field process maximize
```

`record` is the name of a record that exists in this or another IOC.

The `.field`, `process`, and `maximize` parts are all optional.

The default value for `.field` is `.VAL`.

`process` can have one of the following values:

- `NPP` – No Process Passive (Default)
- `PP` – Process Passive
- `CA` – Force link to be a channel access link
- `CP` – CA and process on monitor
- `CPP` – CA and process on monitor if record is passive

NOTES:

`CP` and `CPP` are valid only for `DBF_INLINK` fields.

`DBF_FWDLINK` fields can use `PP` or `CA`. If a `DBF_FWDLINK` is a channel access link it must reference the target record's `PROC` field.

`maximize` can have one of the following values:

- `NMS` – No Maximize Severity (Default)
- `MS` – Maximize Severity
- `MSS` – Maximize Severity and Status

- MSI – Maximize Severity if Invalid
- VME_IO
 - #Ccard Ssignal @parm
 - card – the card number of associated hardware module
 - signal – signal on card
 - parm – An arbitrary character string of up to 31 characters. This field is optional and is device specific.
- CAMAC_IO
 - #Bbranch Ccrate Nstation Asubaddress Ffunction @parm
 - branch, crate, station, subaddress, and function should be obvious to camac users. subaddress and function are optional (0 if not given). parm is also optional and is device specific (25 characters max).
- AB_IO
 - #Llink Aadapter Ccard Ssignal @parm
 - link – Scanner, i.e. vme scanner number
 - adapter – Adapter. Allen Bradley also calls this rack
 - card – Card within Allen Bradley Chassis
 - signal – signal on card
 - parm – optional device-specific character string (27 char max)
- GPIB_IO
 - #Llink Aaddr @parm
 - link – gpib link, i.e. interface
 - addr – GPIB address
 - parm – device-specific character string (31 char max)
- BITBUS_IO
 - #Llink Nnode Pport Ssignal @parm
 - link – link, i.e. vme bitbus interface
 - node – bitbus node
 - port – port on the node
 - signal – signal on port
 - parm – device specific-character string (31 char max)
- INST_IO @parm
 - parm – Device dependent character string
- BBGPIB_IO
 - #Llink Bbbaddr Ggpibaddr @parm
 - link – link, i.e. vme bitbus interface
 - bbaddr – bitbus address
 - gpibaddr – gpib address
 - parm – optional device-specific character string (31 char max)
- RF_IO
 - #Rcryo Mmicro Ddataset Eelement
- VXI_IO
 - #Vframe Cslot Ssignal @parm (Dynamic addressing)
 - or
 - #Vla Signal @parm (Static Addressing)

frame – VXI frame number
 slot – Slot within VXI frame
 la – Logical Address
 signal – Signal Number
 parm – device specific character string(25 char max)

info_name The name of an Information Item related to this record. See section 6.15 below for more on Information Items.

info_value Any ASCII string. IOC applications using this information item may place additional restrictions on the contents of the string.

6.14.3 Examples

```

record(ai, STS_AbAiMaS0) {
  field(SCAN, ".1 second")
  field(DTYP, "AB-1771IFE-4to20MA")
  field(INP, "#L0 A2 C0 S0 F0 @")
  field(PREC, "4")
  field(LINR, "LINEAR")
  field(EGUF, "20")
  field(EGUL, "4")
  field(EGU, "MilliAmps")
  field(HOPR, "20")
  field(LOPR, "4")
}
record(ao, STS_AbAoMaC1S0) {
  field(DTYP, "AB-1771OFE")
  field(OUT, "#L0 A2 C1 S0 F0 @")
  field(LINR, "LINEAR")
  field(EGUF, "20")
  field(EGUL, "4")
  field(EGU, "MilliAmp")
  field(DRVH, "20")
  field(DRVL, "4")
  field(HOPR, "20")
  field(LOPR, "4")
  info(autosaveFields, "VAL")
}
record(bi, STS_AbDiA0C0S0) {
  field(SCAN, "I/O Intr")
  field(DTYP, "AB-Binary Input")
  field(INP, "#L0 A0 C0 S0 F0 @")
  field(ZNAM, "Off")
  field(ONAM, "On")
}

```

6.15 Record Information Item

Information items provide a way to attach named string values to individual record instances that are loaded at the same time as the record definition. They can be attached to any record without having to modify the record type, and can be retrieved by programs running on the IOC (they are not visible via Channel Access at all). Each item attached

to a single record must have a unique name by which it is addressed, and database access provides routines to allow a record's info items to be scanned, searched for, retrieved and set. At runtime a `void*` pointer can also be associated with each item, although only the string value can be initialized from the record definition when the database is loaded.

6.16 Record Attributes

Each record type can have any number of record attributes. Each attribute is a psuedo field that can be accessed via database and channel access. Each attribute has a name that acts like a field name but returns the same value for all instances of the record type. Two attributes are generated automatically for each record type: `RTYP` and `VERS`. The value for `RTYP` is the record type name. The default value for `VERS` is "none specified", which can be changed by record support. Record support can call the following routine to create new attributes or change existing attributes:

```
long dbPutAttribute(char *recordTypename,
                  char *name, char*value)
```

The arguments are:

- `recordTypename` – The name of recordtype.
- `name` – The attribute name, i.e. the psuedo field name.
- `value` – The value assigned to the attribute.

6.17 Breakpoint Tables – Discussion

The menu `menuConvert` is used for field `LINR` of the `ai` and `ao` records. These records allow raw data to be converted to/from engineering units via one of the following:

1. No Conversion.
2. Slope Conversion.
3. Linear Conversion.
4. Breakpoint table.

Other record types can also use this feature. The first choice specifies no conversion; the second and third are both linear conversions, the difference being that for Slope conversion the user specifies the conversion slope and offset values directly, whereas for Linear conversions these are calculated by the device support from the requested Engineering Units range and the device support's knowledge of the hardware conversion range. The remaining choices are assumed to be the names of breakpoint tables. If a breakpoint table is chosen, the record support modules calls `cvtRawToEngBpt` or `cvtEngToRawBpt`. You can look at the `ai` and `ao` record support modules for details.

If a user wants to add additional breakpoint tables, then the following should be done:

- Copy the `menuConvert.dbd` file from EPICS `base/src/bpt`
- Add definitions for new breakpoint tables to the end
- Make sure modified `menuConvert.dbd` is loaded into the IOC instead of EPICS version.

It is only necessary to load a breakpoint file if a record instance actually chooses it. It should also be mentioned that the Allen Bradley IXE device support misuses the `LINR` field. If you use this module, it is very important that you do not change any of the EPICS supplied definitions in `menuConvert.dbd`. Just add your definitions at the end.

If a breakpoint table is chosen, then the corresponding breakpoint file must be loaded into the IOC before `iocInit` is called.

Normally, it is desirable to directly create the breakpoint tables. However, sometimes it is desirable to create a breakpoint table from a table of raw values representing equally spaced engineering units. A good example is the Thermo-couple tables in the OMEGA Engineering, INC Temperature Measurement Handbook. A tool `makeBpt` is provided to convert such data to a breakpoint table.

The format for generating a breakpoint table from a data table of raw values corresponding to equally spaced engineering values is:

```
!comment line
<header line>
<data table>
```

The header line contains the following information:

Name An alphanumeric ascii string specifying the breakpoint table name

Low Value Eng Engineering Units Value for first breakpoint table entry

Low Value Raw Raw value for first breakpoint table entry

High Value Eng Engineering Units: Highest Value desired

High Value Raw Raw Value for High Value Eng

Error Allowed error (Engineering Units)

First Table Engineering units corresponding to first data table entry

Last Table Engineering units corresponding to last data table entry

Delta Table Change in engineering units per data table entry

An example definition is:

```
"TypeKdegF" 32 0 1832 4095 1.0 -454 2500 1
<data table>
```

The breakpoint table can be generated by executing

```
makeBpt bptXXX.data
```

The input file must have the extension of data. The output filename is the same as the input filename with the extension of `.dbd`.

Another way to create the breakpoint table is to include the following definition in a `Makefile`:

```
BPTS += bptXXX.dbd
```

NOTE: This requires the naming convention that all data tables are of the form `bpt<name>.data` and a breakpoint table `bpt<name>.dbd`.

6.18 Menu and Record Type Include File Generation.

6.18.1 Introduction

Given a file containing menus, `dbToMenuH` generates an include file that can be used by any code which uses the associated menus. Given a file containing any combination of menu definitions and record type definitions, `dbToRecordtypeH` generates an include file that can be used by any code which uses the menus and record type.

EPICS base uses the following conventions for managing menu and recordtype definitions. Users generating local record types are encouraged to do likewise.

- Each menu that is either for fields in database common (for example `menuScan`) or is of global use (for example `menuYesNo`) is defined in a separate file. The name of the file is the same as the menu name with an extension of `.dbd`. The name of the generated include file is the menu name with an extension of `.h`. Thus `menuScan` is defined in a file `menuScan.dbd` and the generated include file is named `menuScan.h`
- Each record type definition is defined in a separate file. In addition, this file contains any menu definitions that are used only by that record type. The name of the file is the same as the recordtype name followed by `Record.dbd`. The name of the generated include file is the same name with an extension of `.h`. Thus `aoRecord` is defined in a file `aoRecord.dbd` and the generated include file is named `aoRecord.h`. Since `aoRecord` has a private menu called `aoOIF`, the `dbd` file and the generated include file have definitions for this menu. Thus for each record type, there are two source files (`xxxRecord.dbd` and `xxxRecord.c`) and one generated file (`xxxRecord.h`).

Before continuing, it should be mentioned that developers don't normally execute `dbToMenuH` or `dbToRecordtypeH` themselves. If the proper naming conventions are used, it is only necessary to add definitions to the appropriate `Makefile`. Consult the chapter on the EPICS Build Facility for details.

6.18.2 dbToMenuH

This tool is executed as follows:

```
dbToMenuH -I dir -Smacsub menuXXX.dbd
```

It generates a file which has the same name as the input file but with an extension of `.h`. Multiple `-I` options can be specified for an include path and multiple `-S` options for macro substitution.

For example `menuPriority.dbd`, which contains the definitions for processing priority contains:

```
menu(menuPriority) {
    choice(menuPriorityLOW, "LOW")
    choice(menuPriorityMEDIUM, "MEDIUM")
    choice(menuPriorityHIGH, "HIGH")
}
```

The include file, `menuPriority.h`, generated by `dbToMenuH` contains:

```
#ifndef INCMenuPriorityH
#define INCMenuPriorityH
typedef enum {
    menuPriorityLOW,
    menuPriorityMEDIUM,
    menuPriorityHIGH,
}menuPriority;
#endif /*INCMenuPriorityH*/
```

Any code that needs to use the priority menu values should use these definitions.

6.18.3 dbToRecordtypeH

This tool is executed as follows:

```
dbToRecordtypeH -I dir -Smacsub xxxRecord.dbd
```

It generates a file which has the same name as the input file but with an extension of `.h`. Multiple `-I` options can be specified for an include path and multiple `-S` options for macro substitution.

For example `aoRecord.dbd`, which contains the definitions for the analog output record contains:

```

menu(aoOIF) {
    choice(aoOIF_Full, "Full")
    choice(aoOIF_Incremental, "Incremental")
}
recordtype(ao) {
    include "dbCommon.dbd"
    field(VAL, DBF_DOUBLE) {
        prompt("Desired Output")
        asl(ASL0)
        pp(TRUE)
    }
    field(OVAL, DBF_DOUBLE) {
        prompt("Output Value")
    }
    ... (Many more field definitions)
}
}

```

The include file, `aoRecord.h`, generated by `dbToRecordtypeH` contains:

```

#include "ellLib.h"
#include "epicsMutex.h"
#include "link.h"
#include "epicsTime.h"
#include "epicsTypes.h"

#ifndef INCaoOIFH
#define INCaoOIFH
typedef enum {
    aoOIF_Full,
    aoOIF_Incremental,
}aoOIF;
#endif /*INCaoOIFH*/
#ifndef INCaoH
#define INCaoH
typedef struct aoRecord {
    char          name[29]; /*Record Name*/
    ... Remaining fields in database common
    double        val;     /*Desired Output*/
    double        oval;    /*Output Value*/
    ... remaining record specific fields
} aoRecord;
#define aoRecordNAME    0
... defines for remaining fields in database common
#define aoRecordVAL     42
#define aoRecordOVAL    43
... defines for remaining record specific fields
#ifdef GEN_SIZE_OFFSET
int aoRecordSizeOffset(dbRecordType *pdbRecordType)
{
    aoRecord *prec = 0;
    pdbRecordType->papFldDes[0]->size=sizeof(prec->name);
    pdbRecordType->papFldDes[0]->offset=
(short)((char *)&prec->name -- (char *)prec);
    ... code to compute size&offset for other fields in dbCommon
}

```

```

    pdbRecordType->papFldDes[42]->size=sizeof(prec->val);
    pdbRecordType->papFldDes[42]->offset=
(short)((char *)&prec->val -- (char *)prec);
    pdbRecordType->papFldDes[43]->size=sizeof(prec->oval);
    pdbRecordType->papFldDes[43]->offset=
(short)((char *)&prec->oval -- (char *)prec);
    ... code to compute size&offset for remaining fields
    pdbRecordType->rec_size = sizeof(*prec);
    return(0);
}
#endif /*GEN_SIZE_OFFSET*/

```

The analog output record support module and all associated device support modules should use this include file. No other code should use it. Let's discuss the various parts of the file.:

- The `enum` generated from the menu definition should be used to reference the value of the field associated with the menu.
- The `typedef` and `structure` defining the record are used by record support and device support to access fields in an analog output record.
- A `#define` is present for each field within the record. This is useful for the record support routines that are passed a pointer to a `DBADDR` structure. They can have code like the following:

```

switch (dbGetFieldIndex(pdbAddr)) {
    case aoRecordVAL :
        ...
        break;
    case aoRecordXXX:
        ...
        break;
    default:
        ...
}

```

The C source routine `aoRecordSizeOffset` is automatically called when a record type file is loaded into an IOC. Thus user code does not have to be aware of this routine except for the following convention: The associate record support module **MUST** include the statements:

```

#define GEN_SIZE_OFFSET
#include "xxxRecord.h"
#undef GEN_SIZE_OFFSET

```

This convention ensures that the routine is defined exactly once.

6.19 dbExpand

```
dbExpand -Idir -Smacsub -outfile file1 file2 ...
```

Multiple `-I` options can be specified for an include path, and multiple `-S` options for macro substitution. If no output filename is specified with `-outfile` then the output will go to `stdout`. Note that the environment variable `EPICS_DB_INCLUDE_PATH` can also be used in place of the `-I` options.

NOTE: This is supported only on the host.

This command reads all the input files and then writes a file containing the definitions for all information described by the input files. The output content differs from the input in that comment lines do not appear and all include files are expanded.

This routine is extremely useful if an IOC is not using NFS for the `dbLoadDatabase` commands. It takes more than 2 minutes to load the `base/rec/base.dbd` file into an IOC if NFS is not used. If `dbExpand` creates a local `base.dbd` file, it takes about 7 seconds to load (25 MHZ 68040 IOC).

6.20 dbLoadDatabase

```
dbLoadDatabase(char *dbdfile, char *path, char *substitutions)
```

NOTES:

- IOC Only
- Using a path on a vxWorks ioc does not work very well.
- Both path and substitutions can be null, and are usually omitted.
- `dbdfile` may contain environment variable macros of the form `${MOTOR}` which will be expanded before the file is opened.

This command loads a database file containing any of the definitions given in the summary at the beginning of this chapter. Note that `dbLoadDatabase` should only be used to load a Database Definition (`.dbd`) file, although it is currently possible to use it for loading Record Instance (`.db`) files as well.

As each line of `dbdfile` is read, the substitutions specified in `substitutions` are performed. Substitutions are specified as follows:

```
"var1=sub1,var2=sub3,..."
```

Variables are specified in the `dbfile` as `$(var)`. If the substitution string

```
"a=1,b=2,c=\"this is a test\""
```

were used, any variables `$(a)`, `$(b)`, `$(c)` in the database file would have the appropriate values substituted during parsing.

6.21 dbLoadRecords

```
dbLoadRecords(char* dbfile, char* substitutions)
```

NOTES:

- IOC Only.
- The file named in `dbfile` should contain only record instances, record aliases and/or breakpoint tables.
- The `dbfile` string may itself contain environment variable macros of the form `${MOTOR}` which will be expanded before the file is opened.

6.21.1 Example

For example, let the file `test.db` contain:

```

record(ai, "$(pre)testrec1")
record(ai, "$(pre)testrec2")
record(stringout, "$(pre)testrec3") {
    field(VAL, "$(STR)")
    field(SCAN, "$(SCAN)")
}

```

Then issuing the command:

```
dbLoadRecords("test.db", "pre=TEST,STR=test,SCAN=Passive")
```

gives the same results as loading:

```

record(ai, "TESTtestrec1")
record(ai, "TESTtestrec2")
record(stringout, "TESTtestrec3") {
    field(VAL, "test")
    field(SCAN, "Passive")
}

```

6.22 dbLoadTemplate

```
dbLoadTemplate(char* template_def)
```

NOTES:

- IOC Only.
- MSI can be used to expand templates on the host instead of using this.

`dbLoadTemplate` reads a template definition file. This file contains rules about loading database instance files, which contain `$(xxx)` macros, and performing substitutions.

`template_def` contains the rules for performing substitutions on the instance files. For convenience two formats are provided. The format is either:

```

file name.template {
    { var1=sub1_for_set1, var2=sub2_for_set1, var3=sub3_for_set1, ... }
    { var1=sub1_for_set2, var2=sub2_for_set2, var3=sub3_for_set2, ... }
    { var1=sub1_for_set3, var2=sub2_for_set3, var3=sub3_for_set3, ... }
}

```

or:

```

file name.template {
pattern { var1, var2, var3, ... }
    { sub1_for_set1, sub2_for_set1, sub3_for_set1, ... }
    { sub1_for_set2, sub2_for_set2, sub3_for_set2, ... }
    { sub1_for_set3, sub2_for_set3, sub3_for_set3, ... }
}

```

The first line (`file name.template`) specifies the record instance input file. The file name may appear inside double quotation marks; these are required if the name contains any characters that are not in the following set, or if it contains environment variable macros of the form `${ENV_VAR_NAME}` which are to be expanded before the file is opened:

```
a-z A-Z 0-9 _ + - . / \ : ; [ ] < >
```

Each set of definitions enclosed in `{ }` is variable substitution for the input file. The input file has each set applied to it to produce one composite file with all the completed substitutions in it. Version 1 should be obvious. In version 2, the

variables are listed in the `pattern{ }` line, which must precede the braced substitution lines. The braced substitution lines contains sets which match up with the `pattern{ }` line.

6.22.1 Example

Two simple template file examples are shown below. The examples specify the same substitutions to perform: `this=sub1` and `that=sub2` for a first set, and `this=sub3` and `that=sub4` for a second set.

```
file test.template {
  { this=sub1,that=sub2 }
  { this=sub3,that=sub4 }
}

file test.template {
  pattern{this,that}
  {sub1,sub2}
  {sub3,sub4 }
}
```

Assume that the file `test.template` contains:

```
record(ai, "$ (this) record") {
  field(DESC, "this = $ (this) ")
}
record(ai, "$ (that) record") {
  field(DESC, "this = $ (that) ")
}
```

Using `dbLoadTemplate` with either input is the same as defining the records:

```
record(ai, "sub1record") {
  field(DESC, "this = sub1")
}
record(ai, "sub2record") {
  field(DESC, "this = sub2")
}

record(ai, "sub3record") {
  field(DESC, "this = sub3")
}
record(ai, "sub4record") {
  field(DESC, "this = sub4")
}
```

6.23 dbReadTest

```
dbReadTest -I-dir -S-macsub file.dbd ... file.db ...
```

This utility can be used to check for correct syntax in database definition and database instance files. It just reads all the specified files

Multiple `-I`, and `-S` options can be specified. An arbitrary number of database definition and database instance files can be specified.

Chapter 7

IOC Initialization

7.1 Overview - Environments requiring a main program

If a main program is required (most likely on all environments except vxWorks and RTEMS), then initialization is performed by statements residing in startup scripts which are executed by `iocsh`. An example main program is:

```
int main(int argc, char *argv[])
{
    if (argc >= 2) {
        iocsh(argv[1]);
        epicsThreadSleep(.2);
    }
    iocsh(NULL);
    epicsExit(0)
    return 0;
}
```

The first call to `iocsh` executes commands from the startup script filename which must be passed as an argument to the program. The second call to `iocsh` with a `NULL` argument puts `iocsh` into interactive mode. This allows the user to issue the commands described in the chapter on “IOC Test Facilities” as well as some additional commands like `help`.

The command file passed is usually called the startup script, and contains statements like these:

```
< envPaths
cd ${TOP}
dbLoadDatabase "dbd/appname.dbd"
appname_registerRecordDeviceDriver pdbname
dbLoadRecords "db/file.db", "macro=value"
cd ${TOP}/iocBoot/${IOC}
iocInit
```

The `envPaths` file is automatically generated in the IOC’s boot directory and defines several environment variables that are useful later in the startup script. The definitions shown below are always provided; additional entries will be created for each support module referenced in the application’s `configure/RELEASE` file:

```
epicsEnvSet("ARCH", "linux-x86")
epicsEnvSet("IOC", "iocname")
epicsEnvSet("TOP", "/path/to/application")
epicsEnvSet("EPICS_BASE", "/path/to/base")
```

7.2 Overview - vxWorks

After vxWorks is loaded at IOC boot time, commands like the following, normally placed in the vxWorks startup script, are issued to load and initialize the application code:

```
# Many vxWorks board support packages need the following:
#cd <full path to IOC boot directory>
< cdCommands
cd topbin
ld 0,0, "appname.munch"

cd top
dbLoadDatabase "dbd/appname.dbd"
appname_registerRecordDeviceDriver pdbname
dbLoadRecords "db/file.db", "macro=value"

cd startup
iocInit
```

The `cdCommands` script is automatically generated in the IOC boot directory and defines several vxWorks global variables that allow `cd` commands to various locations, and also sets several environment variables. The definitions shown below are always provided; additional entries will be created for each support module referenced in the application's `configure/RELEASE` file:

```
startup = "/path/to/application/iocBoot/iocname"
putenv "ARCH=vxWorks-68040"
putenv "IOC=iocname"
top = "/path/to/application"
putenv "TOP=/path/to/application"
topbin = "/path/to/application/bin/vxWorks-68040"
epics_base = "/path/to/base"
putenv "EPICS_BASE=/path/to/base"
epics_basebin = "/path/to/base/bin/vxWorks-68040"
```

The `ld` command in the startup script loads EPICS core, the record, device and driver support the IOC needs, and any application specific modules that have been linked into it.

`dbLoadDatabase` loads database definition files describing the record/device/driver support used by the application..

`dbLoadRecords` loads record instance definitions.

`iocInit` initializes the various epics components and starts the IOC running.

7.3 Overview - RTEMS

RTEMS applications can start up in many different ways depending on the board-support package for a particular piece of hardware. Systems which use the Cexp package can be treated much like vxWorks. Other systems first read initialization parameters from non-volatile memory or from a BOOTP/DHCP server. The exact mechanism depends upon the BSP. TFTP or NFS filesystems are then mounted and the IOC shell is used to read commands from a startup script. The location of this startup script is specified by a initialization parameter. This script is often similar or identical to the one used with vxWorks. The RTEMS startup code calls

```
epicsRtemsInitPreSetBootConfigFromNVRAM(struct rtems_bsdnet_config *);
```

just before setting the initialization parameters from non-volatile memory, and


```
epicsRtemsInitPostSetBootConfigFromNVRAM(struct rtems_bsdnet_config *);
```

just after setting the initialization parameters. An application may provide either or both of these routines to perform any custom initialization required. These function prototypes and some useful external variable declarations can be found in the header file `epicsRtemsInitHooks.h`

7.4 IOC Initialization

An IOC is normally started with the `iocInit` command as shown in the startup scripts above, which is actually implemented in two distinct parts. The first part can be run separately as the `iocBuild` command, which puts the IOC into a quiescent state without allowing the various internal threads it starts to actually run. From this state the second command `iocRun` can be used to bring it online very quickly. A running IOC can be quiesced using the `iocPause` command, which freezes all internal operations; at this point the `iocRun` command can restart it from where it left off, or the IOC can be shut down (exit the program, or reboot on vxWorks/RTEMS). Most device support and drivers have not yet been written with the possibility of pausing an IOC in mind though, so this feature may not be safe to use on an IOC which talks to external devices or software.

IOC initialization using the `iocBuild` and `iocRun` commands then consists of the following steps:

7.4.1 Configure Main Thread

Providing the IOC has not already been initialized, `initHookAtIocBuild` is announced first.

The main thread's `epicsThreadIsOkToBlock` flag is set, the message "Starting `iocInit`" is logged and `epicsSignalInstallSigHupIgnore` called, which on Unix architectures prevents the process from shutting down if it later receives a HUP signal.

At this point, `initHookAtBeginning` is announced.

7.4.2 General Purpose Modules

Calls `coreRelease` which prints a message showing which version of `iocCore` is being run.

Calls `taskwdInit` to start the task watchdog. This accepts requests to watch other tasks. It runs periodically and checks to see if any of the tasks is suspended. If so it issues an error message, and can also invoke callback routines registered by the task itself or by other software that is interested in the state of the IOC. See "Task Watchdog" on page 240 for details.

Starts the general purpose callback tasks by calling `callbackInit`. Three tasks are started at different scheduling priorities.

`initHookAfterCallbackInit` is announced.

7.4.3 Channel Access Links

Calls `dbCaLinkInit`. This initializes the module that handles database channel access links, but does not allow its task to run yet.

`initHookAfterCaLinkInit` is announced.

7.4.4 Driver Support

`initDrvSup` locates each device driver entry table and calls the `init` routine of each driver.

`initHookAfterInitDrvSup` is announced.

7.4.5 Record Support

`initRecSup` locates each record support entry table and calls the `init` routine for each record type.

`initHookAfterInitRecSup` is announced.

7.4.6 Device Support

`initDevSup` locates each device support entry table and calls its `init` routine specifying that this is the initial call.

`initHookAfterInitDevSup` is announced.

7.4.7 Database Records

`initDatabase` is called which makes three passes over the database performing the following functions:

1. Initializes the fields `RSET`, `RDES`, `MLOK`, `MLIS`, `PACT` and `DSET` for each record.
Calls record support's `init_record` (first pass).
2. Convert each `PV_LINK` into a `DB_LINK` or `CA_LINK`
Calls any extended device support's `add_record` routine.
3. Calls record support's `init_record` (second pass).

Finally it registers an `epicsAtExit` routine to shut down the database when the IOC application exits.

Next `dbLockInitRecords` is called to create the lock sets.

Then `dbBkptInit` is run to initialize the database debugging module.

`initHookAfterInitDatabase` is announced.

7.4.8 Device Support again

`initDevSup` locates each device support entry table and calls its `init` routine specifying that this is the final call.

`initHookAfterFinishDevSup` is announced.

7.4.9 Scanning and Access Security

The periodic, event, and I/O event scanners are initialized by calling `scanInit`, but the scan threads created are not allowed to process any records yet.

A call to `asInit` initializes access security. If this reports failure, the IOC initialization is aborted.

`dbPutNotifyInit` initializes support for put notification.

After a short delay to allow settling, `initHookAfterScanInit` is announced.

7.4.10 Initial Processing

`initialProcess` processes all records that have PINI set to YES.

`initHookAfterInitialProcess` is announced.

7.4.11 Channel Access Server

The Channel Access server is started by calling `rsrv_init`, but its tasks are not allowed to run so it does not announce its presence to the network yet.

`initHookAfterCaServerInit` is announced.

At this point, the IOC has been fully initialized but is still quiescent. `initHookAfterIocBuilt` is announced. If started using `iocBuild` this command completes here.

7.4.12 Enable Record Processing

If the `iocRun` command is used to bring the IOC out of its initial quiescent state, it starts here.

`initHookAtIocRun` is announced.

The routines `scanRun` and `dbCaRun` are called in turn to enable their associated tasks and set the global variable `interruptAccept` to TRUE (this now happens inside `scanRun`). Until this is set all I/O interrupts should have been ignored.

`initHookAfterDatabaseRunning` is announced. If the `iocRun` command (or `iocInit`) is being executed for the first time, `initHookAfterInterruptAccept` is announced.

7.4.13 Enable CA Server

The Channel Access server tasks are allowed to run by calling `rsrv_run`.

`initHookAfterCaServerRunning` is announced. If the IOC is starting for the first time, `initHookAtEnd` is announced.

A command completion message is logged, and `initHookAfterIocRunning` is announced.

7.5 Pausing an IOC

The command `iocPause` brings a running IOC to a quiescent state with all record processing frozen (other than possibly the completion of asynchronous I/O operations). A paused IOC may be able to be restarted using the `iocRun` command, but whether it will fully recover or not can depend on how long it has been quiescent and the status of any device drivers which have been running. The operations which make up the pause operation are as follows:

1. `initHookAtIocPause` is announced.
2. The Channel Access Server tasks are paused by calling `rsrv_pause`
3. `initHookAfterCaServerPaused` is announced.
4. The routines `dbCaPause` and `scanPause` are called to pause their associated tasks and set the global variable `interruptAccept` to FALSE.
5. `initHookAfterDatabasePaused` is announced.

6. After logging a pause message, `initHookAfterIocPaused` is announced.

7.6 Changing iocCore fixed limits

The following commands can be issued after `iocCore` is loaded to change `iocCore` fixed limits. The commands should be given before any `dbLoadDatabase` commands.

```
callbackSetQueueSize(size)
dbPvdTableSize(size)
scanOnceSetQueueSize(size)
errlogInit(bufferSize)
errlogInit2(bufferSize, maxMessageSize)
```

7.6.1 callbackSetQueueSize

Requests for the general purpose callback tasks are placed in a ring buffer. This command can be used to set the size for the ring buffers. The default is 2000. A message is issued when a ring buffer overflows. It should rarely be necessary to override this default. Normally the ring buffer overflow messages appear when a callback task fails.

7.6.2 dbPvdTableSize

Record instance names are stored in a process variable directory, which is a hash table. The default number of hash entries is 512. `dbPvdTableSize` can be called to change the size. It must be called before any `dbLoad` commands and must be a power of 2 between 256 and 65536. If an IOC contains very large databases (several thousand records) then a larger hash table size speeds up searches for records.

7.6.3 scanOnceSetQueueSize

`scanOnce` requests are placed in a ring buffer. This command can be used to set the size for the ring buffer. The default is 1000. It should rarely be necessary to override this default. Normally the ring buffer overflow messages appear when the `scanOnce` task fails.

7.6.4 errlogInit or errlogInit2

These commands can increase (but not decrease) the default buffer and maximum message sizes for the `errlog` message queue. The default buffer size is 1280 bytes, the maximum message size defaults to 256 bytes.

7.7 initHooks

The `inithooks` facility allows application functions to be called at various states during `ioc` initialization. The states are defined in `inithooks.h`, which contains the following definitions:

```
typedef enum {
    initHookAtIocBuild = 0,          /* Start of iocBuild/iocInit commands */
    initHookAtBeginning,
    initHookAfterCallbackInit,
    initHookAfterCaLinkInit,
    initHookAfterInitDrvSup,
```

```

    initHookAfterInitRecSup,
    initHookAfterInitDevSup,
    initHookAfterInitDatabase,
    initHookAfterFinishDevSup,
    initHookAfterScanInit,
    initHookAfterInitialProcess,
    initHookAfterCaServerInit,
    initHookAfterIocBuilt,           /* End of iocBuild command */

    initHookAtIocRun,               /* Start of iocRun command */
    initHookAfterDatabaseRunning,
    initHookAfterCaServerRunning,
    initHookAfterIocRunning,       /* End of iocRun/iocInit commands */

    initHookAtIocPause,           /* Start of iocPause command */
    initHookAfterCaServerPaused,
    initHookAfterDatabasePaused,
    initHookAfterIocPaused,       /* End of iocPause command */

/* Deprecated states, provided for backwards compatibility.
 * These states are announced at the same point they were before,
 * but will not be repeated if the IOC gets paused and restarted.
 */
    initHookAfterInterruptAccept, /* After initHookAfterDatabaseRunning */
    initHookAtEnd,               /* Before initHookAfterIocRunning */
}initHookState;

typedef void (*initHookFunction)(initHookState state);
int initHookRegister(initHookFunction func);
const char *initHookName(int state);

```

Any functions that are registered before `iocInit` reaches the desired state will be called when it reaches that state. The `initHookName` function returns a static string representation of the state passed into it which is intended for printing. The following skeleton code shows how to use this facility:

```

static initHookFunction myHookFunction;

int myHookInit(void)
{
    return(initHookRegister(myHookFunction));
}

static void myHookFunction(initHookState state)
{
    switch(state) {
        case initHookAfterInitRecSup:
            ...
            break;
        case initHookAfterInterruptAccept:
            ...
            break;
        default:
            break;
    }
}

```

An arbitrary number of functions can be registered.

7.8 Environment Variables

Various environment variables are used by iocCore:

```
EPICS_CA_ADDR_LIST
EPICS_CA_AUTO_ADDR_LIST
EPICS_CA_CONN_TMO
EPICS_CAS_BEACON_PERIOD
EPICS_CA_REPEATER_PORT
EPICS_CA_SERVER_PORT
EPICS_CA_MAX_ARRAY_BYTES
EPICS_TS_NTP_INET
EPICS_IOC_LOG_PORT
EPICS_IOC_LOG_INET
```

For an explanation of the EPICS_CA_... and EPICS_CAS_... variables see the EPICS Channel Access Reference Manual. For an explanation of the EPICS_IOC_LOG_... variables see "iocLogClient" on page 170 of this manual. EPICS_TS_NTP_INET is used only on vxWorks and RTEMS, where it sets the address of the Network Time Protocol server. If it is not defined the IOC uses the boot server as its NTP server.

These variables can be set through iocsh via the epicsEnvSet command, or on vxWorks using putenv. For example:

```
epicsEnvSet("EPICS_CA_CONN_TMO", "10")
```

All epicsEnvSet commands should be issued after iocCore is loaded and before any dbLoad commands.

The following commands can be issued to iocsh:

epicsPrtEnvParams – This shows just the environment variables used by iocCore.

epicsEnvShow – This shows all environment variables on your system.

7.9 Initialize Logging

Initialize the logging system. See the chapter on "IOC Error Logging" for details. The following can be used to direct the log client to use a specific host log server.

```
epicsEnvSet("EPICS_IOC_LOG_PORT", "<port>")
epicsEnvSet("EPICS_IOC_LOG_INET", "<inet addr>")
```

These command must be given immediately after iocCore is loaded.

To start logging you must issue the command:

```
iocLogInit
```

Chapter 8

Access Security

8.1 Overview

This chapter describes access security, i.e. the system that limits access to IOC databases. It consists of the following sections:

Overview This section

Quick start A summary of the steps necessary to start access security.

User's Guide This explains what access security is and how to use it.

Design Summary Functional Requirements and Design Overview.

Application Programmer's Interface

Database Access Security Access Security features for EPICS IOC databases.

Channel Access Security Access Security features in Channel Access

Trapping Channel Access Writes This allows trapping of all writes from external channel access clients.

Implementation Overview

The requirements for access security were generated at ANL/APS in 1992. The requirements document is:

EPICS: Channel Access Security - Functional Requirements, Ned D. Arnold, 03/-9/92.

This document is available through the EPICS website.

8.2 Quick Start

In order to “turn on” access security for a particular IOC the following must be done:

- Create the access security file.
- IOC databases may have to be modified
 - Record instances may have to have values assigned to field ASG. If ASG is null the record is in group DEFAULT.
 - Access security files can be reloaded after `iocInit` via a subroutine record with `asSubInit` and `asSubProcess` as the associated subroutines. Writing the value 1 to this record will cause a reload.
- The startup script must contain the following command before `iocInit`.

```
asSetFilename("/full/path/to/accessSecurityFile")
```

- The following is an optional command.

```
asSetSubstitutions("var1=sub1,var2=sub2,...")
```

The following rules decide if access security is turned on for an IOC:

- If `asSetFilename` is not executed before `iocInit`, access security will *never* be started.
- If `asSetFile` is given and any error occurs while first initializing access security, then *all* access to that ioc is denied.
- If after successfully starting access security, an attempt is made to restart and an error occurs then the previous access security configuration is maintained.

After an IOC has been booted with access security enabled, the access security rules can be changed by issuing the `asSetFilename`, `asSetSubstitutions`, and `asInit`. The functions `asInitialize`, `asInitFile`, and `asInitFP`, which are described below, can also be used.

8.3 User's Guide

8.3.1 Features

Access security protects IOC databases from unauthorized Channel Access Clients. Access security is based on the following:

Who Userid of the channel access client.

Where Hostid where the user is logged on. This is the host on which the channel access client exists. Thus no attempt is made to see if a user is local or is remotely logged on to the host.

What Individual fields of records are protected. Each record has a field containing the Access Security Group (ASG) to which the record belongs. Each field has an access security level, either `ASL0` or `ASL1`. The security level is defined in the record definition file. Thus the access security level for a field is the same for all record instances of a record type.

When Access rules can contain input links and calculations similar to the calculation record.

8.3.2 Limitations

An IOC database can be accessed only via Channel Access or via the `vxWorks` or `ioc` shell. It is assumed that access to the local IOC console is protected via physical security, and that network access is protected via normal networking and physical security methods.

No attempt has been made to protect against the sophisticated saboteur. Network and physical security methods must be used to limit access to the subnet on which the iocs reside.

8.3.3 Definitions

This document uses the following terms:

ASL Access Security Level (Called access level in Req Doc)

ASG Access Security Group (Called PV Group in Req Doc)

UAG User Access Group

HAG Host Access Group

8.3.4 Access Security Configuration File

This section describes the format of a file containing definitions of the user access groups, host access groups, and access security groups. An IOC creates an access configuration database by reading an access configuration file (the extension `.acf` is recommended). Lets first give a simple example and then a complete description of the syntax.

8.3.4.1 Simple Example

```
UAG(uag) {user1,user2}
HAG(hag) {host1,host2}
ASG(DEFAULT) {
  RULE(1,READ)
  RULE(1,WRITE) {
    UAG(uag)
    HAG(hag)
  }
}
```

These rules provide read access to anyone located anywhere and write access to `user1` and `user2` if they are located at `host1` or `host2`.

8.3.4.2 Syntax Definition

In the following description:

[] surrounds optional elements

| separates alternatives

... means that an arbitrary number of definitions may be given.

introduces a comment line

The elements `<name>`, `<user>`, `<host>`, `<pvname>` and `<calculation>` can be given as quoted or unquoted strings. The rules for unquoted strings are the same as for database definitions.

```
UAG(<name>) [{ <user> [, <user> ...] }]
...
HAG(<name>) [{ <host> [, <host> ...] }]
...
ASG(<name>) [{
  [INP<index>(<pvname>)
  ...]
  RULE(<level>,NONE | READ | WRITE [, NOTRAPWRITE | TRAPWRITE]) {
    [UAG(<name> [, <name> ...])]
    [HAG(<name> [, <name> ...])]
    CALC(<calculation>)
  }
  ...
}]
...
```

8.3.4.3 Discussion

- UAG: User Access Group. This is a list of user names. The list may be empty. A user name may appear in more than one UAG. To match, a user name must be identical to the user name read by the CA client library running on the client machine. For vxWorks clients, the user name is usually taken from the user field of the boot parameters.
- HAG: Host Access Group. This is a list of host names. It may be empty. The same host name can appear in multiple HAGs. To match, a host name must match the host name read by the CA client library running on the client machine; both names are converted to lower case before comparison however. For vxWorks clients, the host name is usually taken from the target name of the boot parameters.
- ASG: An access security group. The group `DEFAULT` is a special case. If a member specifies a null group or a group which has no ASG definition then the member is assigned to the group `DEFAULT`.
- INP<index> Index must have one of the values A to L. These are just like the INP fields of a calculation record. It is necessary to define INP fields if a CALC field is defined in any RULE for the ASG.
- RULE This defines access permissions. <level> must be 0 or 1. Permission for a level 1 field implies permission for level 0 fields. The permissions are NONE, READ, and WRITE. WRITE permission implies READ permission. The standard EPICS record types have all fields set to level 1 except for VAL, CMD (command), and RES (reset). An optional argument specifies if writes should be trapped. See the section below on trapping Channel Access writes for how this is used. If not given the default is NOTRAPWRITE.
 - UAG specifies a list of user access groups that can have the access privilege. If UAG is not defined then all users are allowed.
 - HAG specifies a list of host access groups that have the access privilege. If HAG is not defined then all hosts are allowed.
 - CALC is just like the CALC field of a calculation record except that the result must evaluate to TRUE or FALSE. The rule only applies if the calculation result is TRUE, where the actual test for TRUE is $(0.99 < result < 1.01)$. Anything else is regarded as FALSE and will cause the rule to be ignored. Assignment statements are not permitted in CALC expressions here.

Each IOC record contains a field ASG, which specifies the name of the ASG to which the record belongs. If this field is null or specifies a group which is not defined in the access security file then the record is placed in group `DEFAULT`.

The access privilege for a channel access client is determined as follows:

1. The ASG associated with the record is searched.
2. Each RULE is checked for the following:
 - (a) The field's level must be less than or equal to the level for this RULE.
 - (b) If UAG is defined, the user must belong to one of the specified UAGs. If UAG is not defined all users are accepted.
 - (c) If HAG is defined, the user's host must belong to one one of the HAGs. If HAG is not defined all hosts are accepted.
 - (d) If CALC is specified, the calculation must yield the value 1, i.e. TRUE. If any of the INP fields associated with this calculation are in INVALID alarm severity the calculation is considered false. The actual test for TRUE is $.99 < result < 1.01$.
3. The maximum access allowed by step 2 is the access chosen.

Multiple RULEs can be defined for a given ASG, even RULEs with identical levels and access permissions. The TRAPWRITE setting used for a client is determined by the first WRITE rule that passes the rule checks.

8.3.5 ascheck - Check Syntax of Access Configuration File

After creating or modifying an access configuration file it can be checked for syntax errors by issuing the command:

```
ascheck -S "xxx=yyy,..." < "filename"
```

This is a Unix command. It displays errors on `stdout`. If no errors are detected it prints nothing. Only syntax errors not logic errors are detected. Thus it is still possible to get your self in trouble. The flag `-S` means a set of macro substitutions may appear. This is just like the macro substitutions for `dbLoadDatabase`.

8.3.6 IOC Access Security Initialization

In order to have access security turned on during IOC initialization the following command must appear in the startup file before `iocInit` is called:

```
asSetFilename("/full/path/to/access/security/file.acf")
```

If this command is not used then access security will not be started by `iocInit`. If an error occurs when `iocInit` calls `asInit` than all access to the ioc is disabled, i.e. no channel access client will be able to access the ioc. Note that this command does not read the file itself, it just saves the argument string for use later on, nor does it save the current working directory, which is why the use of an absolute path-name for the file is recommended (a path name could be specified relative to the current directory at the time when `iocInit` is run, but this is not recommended if the IOC also loads the subroutine record support as a later reload of the file might happen after the current directory had been changed).

Access security also supports macro substitution just like `dbLoadDatabase`. The following command specifies the desired substitutions:

```
asSetSubstitutions("var1=sub1,var2=sub2,...")
```

This command must be issued before `iocInit`.

After an IOC is initialized the access security database can be changed. The preferred way is via the subroutine record described in the next section. It can also be changed by issuing the following command to the `vxWorks` shell:

```
asInit
```

It is also possible to reissue `asSetFilename` and/or `asSetSubstitutions` before `asInit`. If any error occurs during `asInit` the old access security configuration is maintained. It is NOT permissible to call `asInit` before `iocInit` is called.

Restarting access security after ioc initialization is an expensive operation and should not be used as a regular procedure.

8.3.7 Database Configuration

8.3.7.1 Access Security Group

Each database record has a field `ASG` which holds a character string. Any database configuration tool can be used to give a value to this field. If the `ASG` of a record is not defined or is not equal to a `ASG` in the configuration file then the record is placed in `DEFAULT`.

8.3.7.2 Subroutine Record Support

Two subroutines, which can be attached to a subroutine record, are available (provided with `iocCore`):

```
asSubInit
asSubProcess
```

NOTE: These subroutines are automatically registered thus do NOT put a `registrar` definition in your database definition file.

If a record is created that attaches to these routines, it can be used to force the IOC to load a new access configuration database. To change the access configuration:

1. Modify the file specified by the last call to `asSetFilename` so that it contains the new configuration desired.
2. Write a 1 to the subroutine record `VAL` field. Note that this can be done via channel access.

The following action is taken:

1. When the value is found to be 1, `asInit` is called and the value set back to 0.
2. The record is treated as an asynchronous record. Completion occurs when the new access configuration has been initialized or a time-out occurs. If initialization fails the record is placed into alarm with a severity determined by `BRSV`.

8.3.7.3 Record Type Description

Each field of each record type has an associated access security level of `ASL0` or `ASL1`. See the chapter “Database Definition” for details.

8.3.8 Example:

Lets design a set of rules for a Linac. Assume the following:

1. Anyone can have read access to all fields at anytime.
2. Linac engineers, located in the injection control or control room, can have write access to most level 0 fields only if the Linac is not in operational mode.
3. Operators, located in the injection control or control room, can have write access to most level 0 fields anytime.
4. The operations supervisor, linac supervisor, and the application developers can have write access to all fields but must have some way of not changing something inadvertently.
5. Most records use the above rules but a few (high voltage power supplies, etc.) are placed under tighter control. These will follow rules 1 and 4 but not 2 or 3.
6. IOC channel access clients always have level 1 write privilege.

Most Linac IOC records will not have the `ASG` field defined and will thus be placed in `ASG DEFAULT`. The following records will have an `ASG` defined:

- `LI:OPSTATE` and any other records that need tighter control have `ASG="critical"`. One such record could be a subroutine record used to cause a new access configuration file to be loaded. `LI:OPSTATE` has the value (0,1) if the Linac is (not operational, operational).
- `LI:lev1permit` has `ASG="permit"`. In order for the `opSup`, `linacSup`, or an `appDev` to have write privilege to everything this record must be set to the value 1.

The following access configuration satisfies the above rules.

```
UAG(op) {op1,op2, superguy}
UAG(opSup) {superguy}
UAG(linac) {waw,nassiri,grelick,berg,fuja,gsm}
UAG(linacSup) {gsm}
```

```

UAG(appDev) {nda,kko}
HAG(icr) {silver,phebos,gaea}
HAG(cr) {mars,hera,gold}
HAG(ioc) {ioclic1,ioclic2,ioclid1,ioclid2,ioclid3,ioclid4,ioclid5}
ASG(DEFAULT) {
  INPA(LI:OPSTATE)
  INPB(LI:levlpermit)
  RULE(0,WRITE) {
    UAG(op)
    HAG(icr,cr)
    CALC("A=1")
  }
  RULE(0,WRITE) {
    UAG(op,linac,appdev)
    HAG(icr,cr)
    CALC("A=0")
  }
  RULE(1,WRITE) {
    UAG(opSup,linacSup,appdev)
    CALC("B=1")
  }
  }
  RULE(1,READ)
  RULE(1,WRITE) {
    HAG(ioc)
  }
}
ASG(permit) {
  RULE(0,WRITE) {
    UAG(opSup,linacSup,appDev)
  }
  RULE(1,READ)
  RULE(1,WRITE) {
    HAG(ioc)
  }
}
ASG(critical) {
  INPB(LI:levlpermit)
  RULE(1,WRITE) {
    UAG(opSup,linacSup,appdev)
    CALC("B=1")
  }
  RULE(1,READ)
  RULE(1,WRITE) {
    HAG(ioc)
  }
}
}

```

8.4 Design Summary

8.4.1 Summary of Functional Requirements

A brief summary of the Functional Requirements is:

1. Each field of each record type is assigned an access security level.
2. Each record instance is assigned to a unique access security group.
3. Each user is assigned to one or more user access groups.
4. Each node is assigned to a host access group.
5. For each access security group a set of access rules can be defined. Each rule specifies:
 - (a) Access security level
 - (b) READ or READ/WRITE access.
 - (c) An optional list of User Access Groups or * meaning anyone.
 - (d) An optional list of Host Access Groups or * meaning anywhere.
 - (e) Conditions based on values of process variables

8.4.2 Additional Requirements

8.4.2.1 Performance

Although the functional requirements doesn't mention it, a fundamental goal is performance. The design provides almost no overhead during normal database access and moderate overhead for the following: channel access client/server connection, ioc initialization, a change in value of a process variable referenced by an access calculation, and dynamically changing a records access control group. Dynamically changing the user access groups, host access groups, or the rules, however, can be a time consuming operation. This is done, however, by a low priority IOC task and thus does not impact normal ioc operation.

8.4.2.2 Generic Implementation

Access security should be implemented as a stand alone system, i.e. it should not be imbedded tightly in database or channel access.

8.4.2.3 No Access Security within an IOC

Within an IOC no access security is invoked. This means that database links and local channel access clients calls are not subject to access control. Also test routines such as dbgf should not be subject to access control.

8.4.2.4 Defaults

It must be possible to easily define default access rules.

8.4.2.5 Access Security is Optional

When an IOC is initialized, access security is optional.

8.4.3 Design Overview

The implementation provides a library of routines for accessing the security system. This library has no knowledge of channel access or IOC databases, i.e. it is generic. Database access, which is responsible for protecting an IOC database, calls library routines to add each IOC record to one of the access control groups.

Lets briefly discuss the access security system and how database access and channel access interact with it.

8.4.3.1 Configuration File

User access groups, host access groups, and access security groups are configured via an ASCII file.

8.4.3.2 Access Security Library

The access security library consists of the following groups of routines: initialization, group manipulation, client manipulation, access computation, and diagnostic. The initialization routine reads a configuration file and creates a memory resident access control database. The group manipulation routines allow members to be added and removed from access groups. The client routines provide services for clients attached to members.

8.4.3.3 IOC Database Access Security

The interface between an IOC database and the access security system.

8.4.3.4 Channel Access Security

Whenever the Channel Access broadcast server receives a `ca_search` request and finds the process variable, it calls `asAddClient`. Whenever it disconnects it calls `asRemoveClient`. Whenever it issues a get or put to the database it must call `asCheckGet` or `asCheckPut`.

8.4.4 Comments

It is likely that the access rules will be defined such that many IOCs will attach to a common process variable. As a result the IOC containing the PV will have many CA clients.

What about password protection and encryption? I maintain that this is a problem to be solved in a level above the access security described in this document. This is the issue of protecting against the sophisticated saboteur.

8.4.5 Performance and Memory Requirements

Performance has not yet been measured but during the tests to measure memory usage no noticeable change in performance during ioc initialization or during Channel Access clients connection was noticed. Unless access privilege is violated the overhead during channel access gets and puts is only an extra comparison.

In order to measure memory usage, the following test was performed:

1. A database consisting of 5000 soft analog records was created.
2. A channel access client (`caput`) was created that performs `ca_puts` on each of the 5000 channels. Each time it begins a new set of puts the value increments by 1.
3. A channel access client (`caget`) was created that has monitors on each of the 5000 channels.

The memory consumption was measured before `iocInit`, after `iocInit`, after `caput` connected to all channels, and after `caget` connected to all 5000 channels. This was done for APS release 3.11.5 (before access security) and the first version which included access security. The results were:

	R3.11.5	After
Before <code>iocInit</code>	4,244,520	4,860,840
After <code>iocInit</code>	4,995,416	5,964,904
After <code>caput</code>	5,449,780	6,658,868
After <code>caget</code>	8,372,444	9,751,796

Before the database was loaded the memory used was 1,249,692 bytes. Thus most of the memory usage before `iocInit` resulted from storage for records. The increase since R3.11.5 results from added fields to `dbCommon`. Fields were added for access security, synchronous time support and for the new caching put support. The other increases in memory usage result from the control blocks needed to support access control. The entire design was based on maximum performance. This resulted in increased memory usage.

8.5 Access Security Application Programmer's Interface

8.5.1 Introduction

File `asLib.h` describes the access security data structures and the last section of this chapter has a diagram describing the relationship between the structures. The structures are:

- `ASBASE` - Contains the list head for lists of `UAGs`, `HAGs`, and `ASGs`
- `UAG` - A user access group.
- `HAG` - A host access group
- `ASG` - An access security group. It contains the list head for `ASGINPs`, `ASGRULEs`, and `ASGMEMBERS`
- `ASGINP` - Contains the information for an `INPx`.
- `ASGRULE` - Contains the information for a rule
- `ASGMEMBER` - Contains the information for a member of an access security group. It contains the list head for `ASGCLIENTs`.

All structures except `ASGMEMBER` and `ASGCLIENT` are created by the access security library itself when it reads an access security file. An `ASGMEMBER` is created each time `asAddMember` is called by code that interfaces to the database. An `ASGCLIENT` is created each time `asAddClient` is called by a channel access server.

8.5.2 Definitions

The following are descriptions of arguments of routines described later.

```
typedef struct asgMember *ASMEMBERPVT;
typedef struct asgClient *ASCLIENTPVT;
typedef int (*ASINPUTFUNCPtr)(char *buf, int max_size);
typedef enum{
    asClientCOAR /*Change of access rights*/
    /*For now this is all*/
} asClientStatus;
```



```
typedef void (*ASCLIENTCALLBACK) (ASCLIENTPVT, asClientStatus);
```

8.5.3 Initialization

```
long asInitialize(ASINPUTFUNPTR inputFunction)
long asInitFile(const char *filename, const char *substitutions)
long asInitFP(FILE *fp, const char *substitutions)
```

These routines read an access definition file and perform all initialization necessary. The caller must provide a routine to provide input lines for `asInitialize`. `asInitFile` and `asInitFP` do their own input and also perform macro substitutions.

The initialization routines can be called multiple times. If an access system already exists the old definitions are removed and the new one initialized. Existing members are placed in the new ASGs.

8.5.4 Group manipulation

The routines are called by code that knows how to associate ASG names with the database. In the case of IOC databases, `dbCommon` has a field `ASG`. At IOC initialization a call is made to `asAddMember` for every record instance in the IOC database.

8.5.4.1 add Member

```
long asAddMember(ASMEMBERPVT *ppvt, const char *asgName);
```

This routine adds a new member to ASG `asgName`. The calling routine must provide storage for `ASMEMBERPVT`. Upon successful return `*ppvt` will be equal to the address of storage used by the access control system. The access system keeps an orphan list for all `asgNames` not defined in the access configuration.

The caller must provide permanent storage for `asgName`.

This routine returns `S_asLib_asNotActive` without doing anything if access control is not active.

8.5.4.2 remove Member

```
long asRemoveMember(ASMEMBERPVT *ppvt);
```

This routine removes a member from an access control group. If any clients are still present it returns an error status of `S_asLib_clientExists` without removing the member.

This routine returns `S_asLib_asNotActive` without doing anything if access control is not active.

8.5.4.3 get Member Pvt

```
void *asGetMemberPvt(ASMEMBERPVT pvt);
```

For each member, the access system keeps a pointer that can be used by the caller. This routine returns the value of the pointer.

This routine returns `NULL` if access security is not active

8.5.4.4 put Member Pvt

```
long asPutMemberPvt (ASMEMBERPVT pvt, void *userPvt);
```

This routine is used to set the pointer returned by `asGetMemberPvt`.

This routine returns `S_asLib_asNotActive` without doing anything if access control is not active.

8.5.4.5 change Group

```
long asChangeGroup (ASMEMBERPVT *ppvt, const char *newAsgName);
```

This routine changes the group for an existing member. The access rights of all clients of the member are recomputed.

The caller must provide permanent storage for `newAsgName`.

This routine returns `S_asLib_asNotActive` without doing anything if access control is not active.

8.5.5 Client Manipulation

This code is called by a channel access server.

8.5.5.1 add Client

```
long asAddClient (ASCLIENTPVT *ppvt, ASMEMBERPVT pvt, int asl,
                 const char *user, char*host);
```

This routine adds a client to an ASG member. The calling routine must provide storage for `ASCLIENTPVT`. `ASMEMBERPVT` is the value that was set by calling `asAddMember`. The database code and the server code must develop a convention that allows the server code to locate the `ASMEMBERPVT`. For IOC databases, `ASMEMBERPVT` is kept in `dbCommon`. `asl` is the access security level.

The caller must provide permanent storage for `user` and `host`. Note that `user` is “const char *” but `host` is just “char *”. The reason is the host names are converted to lower case.

This routine returns `S_asLib_asNotActive` without doing anything if access control is not active.

8.5.5.2 change Client

```
long asChangeClient (ASCLIENTPVT ppvt, int asl,
                    const char *user, char*host);
```

This routine changes one or more of the values `asl`, `user`, and `host` for an existing client. Again the caller must provide permanent storage for `user` and `host`. It is permissible to use the same `user` and `host` used in the call to `asAddClient` with different values.

This routine returns `S_asLib_asNotActive` without doing anything if access control is not active.

8.5.5.3 remove Client

```
long asRemoveClient (ASCLIENTPVT *pvt);
```

This call removes a client.

This routine returns `S_asLib_asNotActive` without doing anything if access control is not active.

8.5.5.4 get Client Pvt

```
void *asGetClientPvt (ASCLIENTPVT pvt);
```

For each client, the access system keeps a pointer that can be used by the caller. This routine returns the value of the pointer.

This routine returns NULL if access security is not active.

8.5.5.5 put Client Pvt

```
void asPutClientPvt (ASCLIENTPVT pvt, void *userPvt);
```

This routine is used to set the pointer returned by `asGetClientPvt`.

8.5.5.6 register Callback

```
long asRegisterClientCallback (ASCLIENTPVT pvt,  
                               ASCLIENTCALLBACK pcallback);
```

This routine registers a callback that will be called whenever the access privilege of the client changes.

This routine returns `S_asLib_asNotActive` without doing anything if access control is not active.

8.5.5.7 check Get

```
long asCheckGet (ASCLIENTPVT pvt);
```

This routine, actually a macro, returns (TRUE,FALSE) if the client (has, doesn't have) get access rights.

8.5.5.8 check Put

```
long asCheckPut (ASCLIENTPVT pvt);
```

This routine, actually a macro, returns (TRUE,FALSE) if the client (has, doesn't have) put access rights.

8.5.5.9 asTrapWriteBefore and asTrapWriteAfter

```
void *asTrapWriteBefore (ASCLIENTPVT clientPvt,  
                        const char *userid, const char *hostid, void *serverSpecific);  
void *asTrapWriteAfter (void *trapPvt);
```

These routines must be called before and after any write performed for a client. The value returned by `asTrapWriteBefore` must be the value passed to `asTrapWriteAfter`. The `serverSpecific` argument is assigned to the `serverSpecific` field of the `asTrapWriteMessage` described below.

8.5.6 Access Computation

8.5.6.1 compute all Asg

```
long asComputeAllAsg (void);
```

This routine calls `asComputeAsg` for each access security group.

This routine returns `S_asLib_asNotActive` without doing anything if access control is not active.

8.5.6.2 compute Asg

```
long asComputeAsg(ASG *pasg);
```

This routine calculates all `CALC` entries for the `ASG` and calls `asCompute` for each client of each member of the specified access security group.

This routine returns `S_asLib_asNotActive` without doing anything if access control is not active.

8.5.6.3 compute access

rights

```
long asCompute(ASCLIENTPVT pvt);
```

This routine computes the access rights of a client. This routine is normally called by the access library itself rather than use code.

This routine returns `S_asLib_asNotActive` without doing anything if access control is not active.

8.5.7 Diagnostics

8.5.7.1 Dump

```
int asDump(void (*member)(ASMEMBERPVT),
void (*client)(ASCLIENTPVT), int verbose);
int asDumpFP(FILE *fp, void (*member)(ASMEMBERPVT),
void (*client)(ASCLIENTPVT), int verbose);
```

These routines print the current access security database. If `verbose` is 0 (`FALSE`), then only the information obtained from the access security file is printed.

If `verbose` is `TRUE` then additional information is printed. The value of each `INP` is displayed. The list of members belonging to each `ASG` and the clients belonging to each member are displayed. If member callback is specified as an argument, then it is called for each member. If client callback is specified, it is called for each access security client.

8.5.7.2 Dump UAG

```
int asDumpUag(char *uagname)
int asDumpUagFP(FILE *fp, char *uagname)
```

These routines display the specified `UAG` or if `uagname` is `NULL` each `UAG` defined in the access security database.

8.5.7.3 Dump HAG

```
int asDumpHag(char *hagname)
int asDumpHagFP(FILE *fp, char *hagname)
```

These routines display the specified `UAG` or if `uagname` is `NULL` each `UAG` defined in the access security database.

8.5.7.4 Dump Rules

```
int asDumpRules(char *asgname)
int asDumpRulesFP(FILE *fp, char *asgname)
```

These routines display the rules for the specified ASG or if `asgname` is NULL the rules for each ASG defined in the access security database.

8.5.7.5 Dump member

```
int asDumpMem(char *asgname,
              void (*memcallback)(ASMEMBERPVT), int clients)
int asDumpMemFP(FILE *fp, char *asgname,
                void (*memcallback)(ASMEMBERPVT), int clients)
```

This routine displays the member and, if `clients` is TRUE, client information for the specified ASG or if `asgname` is NULL the member and client information for each ASG defined in the access security database. It also calls `memcallback` for each member if this argument is not NULL.

8.5.7.6 Dump hash table

```
int asDumpHash(void)
int asDumpHash(FILE *fp, void)
```

These show the contents of the hash table used to locate UAGs and HAGs,

8.6 Database Access Security

8.6.1 Access Level definition

The definition of access level means that a level is defined for each field of each record type.

1. Structure `fldDes` (`dbBase.h`), which describes the attributes of each field, contains a field `access_security_level`. In addition definitions exist for the symbols: `ASL0` and `ASL1`.
2. Each field description in a record description contains a field with the value `ASLx`.

The meanings of the Access Security Level definitions are as follows:

- `ASL0` Assigned to fields used during normal operation
- `ASL1` Assigned to fields that may be sensitive to change. Permission to access this level implies permission for `ASL0`.

Most record types assign ASL as follows: The fields `VAL`, `RES` (Reset), and `CMD` use the value `ASL0`. All other fields use `ASL1`.

8.6.2 Access Security Group definition

`dbCommon` contains the fields `ASG` and `ASP`. `ASG` (Access Security Group) is a character string. The value can be assigned via a database configuration tool or else a utility could be provided to assign values during `ioc` initialization. `ASP` is an access security private field. It contains the address of an `ASGMEMBER`.

8.6.3 Access Client Definition

Struct `dbAddr` contains a field `asPvt`, which contains the address of an `ASGCLIENT`. This definition is also added to struct `db_addr` so that old database access also supports access security.

8.6.4 Database Access Library

Two files `asDbLib.c` and `asCa.c` implement the interface between IOC databases and access control. They contain the following routines:

8.6.4.1 Initialization

```
int asSetFilename(char *acf)
```

Calling this routine sets the filename of an access configuration file, but does not save the current working directory, so the use of an absolute pathname is strongly recommended. The next call to `asInit` uses this filename. `asSetFilename` must be called before `iocInit` otherwise access configuration is disabled. If access security is disabled during `iocInit` it will never be turned on.

```
int asSetSubstitutions(char *substitutions)
```

This routine specifies macro substitutions.

```
int asInit()
int asInitAsyn(ASDBCALLBACK *pcallback)
```

This routines call `asInitialize`. If the current access configuration file, as specified by `asSetFilename`, is `NULL` then the routine just returns, otherwise the configuration file is used to create the access configuration database. After initialization all records in the database are made members of the appropriate access control group.

`asInit` is called by `iocInit`, and can also be called after `iocInit` to change the access configuration information.

`asInitAsyn` spawns a task `asInitTask` to perform the initialization. This allows `asInitAsyn` to be called from a subroutine called by the process entry of a subroutine record. `asInitTask` calls `taskwdInsert` so that if it suspends for some reason `taskwd` can detect the failure.

If the caller provides an `ASDBCALLBACK` then when either initialization completes or `taskwd` detects a failure the users callback routine is called via one of the standard callback tasks.

`asInitAsyn` will return a value of `-1` if access initialization is already active. It returns `0` if `asInitTask` is successfully spawned.

8.6.4.2 Routines used by Channel Access Server

```
int asDbGetAsl(void *paddr)
```

Get Access Security level for the field referenced by a database access structure. The argument is defined as a `void*` so that both old and new database access can be used.

```
void * asDbGetMemberPvt(void *paddr)
```

Get `ASMEMBERPVT` for the field referenced by a database access structure. The argument is defined as a `void*` so that both old and new database access can be used.

8.6.4.3 Routine to test asAddClient

```
int astac(char *pname, char *user, char *host)
```

This is a routine to test `asAddClient`. It simulates the calls that are made by Channel Access.

8.6.4.4 Subroutines attached to a subroutine record

These routines are provided so that a channel access client can force an ioc to load a new access configuration database.

```
long asSubInit(struct subRecord *prec, int pass)
long asSubProcess(struct subRecord *prec)
```

These are routines that can be attached to a subroutine record. Whenever a 1 is written to the record, `asSubProcess` calls `asInit`. If `asInit` returns success, it returns with asynchronously. When `asInitTask` calls the completion routine supplied by `asSubProcess`, the return status is used to place the record in alarm.

8.6.4.5 Diagnostic Routines

These routines provide interfaces to the `asDump` routines described in the previous chapter. They do NOT lock before calling the associated routine. Thus they may fail if the access security configuration is changing while they are running. However the danger of the user accidentally aborting a command and leaving the access security system locked is considered a risk that should be avoided.

```
asdbdump(void)
asdbdumpFP(FILE *fp)
```

These routines call `asDumpFP` with a member callback and with `verbose TRUE`.

```
aspuag(char *uagname)
aspuagFP(FILE *fp, char *uagname)
```

These routines call `asDumpUagFP`.

```
asphag(char *hagname)
asphagFP(FILE *fp, char *hagname)
```

These routines call `asDumpHagFP`.

```
asprules(char *asgname)
asprulesFP(FILE *fp, char *asgname)
```

These routines call `asDumpRulesFP`.

```
aspmem(char *asgname, int clients)
aspmemFP(FILE *fp, char *asgname, int clients)
```

These routines call `asDumpMemFP`.

8.7 Channel Access Security

EPICS Access Security was originally designed to protect Input Output Controllers (IOCs) from unauthorized access via the Channel Access (CA) network protocol. It can also be used by any Channel Access Server (CAS) tool. For example the Channel Access PV Gateway implements its own access security. This section describes the interaction between a CA server and the Access Security system. It also briefly describes how the current access rights state is communicated to clients of the EPICS control system via the CA client interface.

8.7.1 CA Server Interfaces to the Access Security System

The CA server calls `asAddClient()` and `asRegisterClientCallback()` for each of the channels that a client connects to the server. The routine `asRemoveClient()` is called whenever the client clears (removes) a channel or when the client disconnects.

The server maintains storage for the clients host and user names. The initial value of these strings are supplied to the server when the client connects and can be updated at any time by the client. When these strings change then `asChangeClient()` is called for each of the channels maintained by the server for the client.

The server checks for read access when processing gets and for write access when processing puts. If access is denied then an exception message is sent to the client. The macros `asCheckGet()` and `asCheckPut()` perform the checks.

The server checks for read access when processing requests to register an event callback (monitor) for the client. If there is read access the server always sends an initial update indicating the current value. If there isn't read access the server sends one update indicating no read access and disables subsequent updates.

The server receives asynchronous notification of access rights change via the callback registered with `asRegisterClientCallback()`. When a channel's access rights change the server communicates the current state to the client library. If read access to a channel is lost and there are events (monitors) registered on the channel then the server sends an update to the client for each of them indicating no access and disables future updates for each event. If read access is reestablished to a channel and there are events (monitors) registered on the channel then the server re-enables updates and sends an initial update message to the client for each of them.

The server must also call `asTrapWriteBefore()` and `asTrapWriteAfter()` before and after a put request from a client is performed.

8.7.2 Client Interfaces

Additional details on the channel access client side callable interfaces to access security can be obtained from the "Channel Access Reference Manual".

The client library stores and maintains the current state of the access rights for each channel that it has established. The client library receives asynchronous updates of the current access rights state from the server. It uses this state to check for read access when processing gets and for write access when processing puts. If a program issues a channel access request that is inconsistent with the client library's current knowledge of the access rights state then access is denied and an error code is returned to the application. The current access rights state as known by the client library can be tested by an applications program with the C macros `ca_read_access()` and `ca_write_access()`.

An application program can also receive asynchronous notification of changes to the access rights state by registering a function to be called back when the client library updates its storage of the access rights state. The application's call back function is installed for this purpose by calling `ca_replace_access_rights_event()`.

If the access rights state changes in the server after a request is queued in the client library but before the request is processed by the server then it is possible that the request will fail in the server. Under these circumstances then an exception will be raised in the client.

The server always sends one update to the client when the event (monitor) is initially registered. If there isn't read access then the status in the arguments to the application program's event call back function indicates no read access and the value in the arguments to the clients event call back is set to zero. If the read access right changes after the event is initially registered then another update is supplied to the application programs call back function.

8.8 Trapping Channel Access Writes

Access security provides a facility `asTrapWrite` that can trap write requests and pass them to any facility that registers a listener. In order to use this facility three things are necessary:

1. The facility, e.g. the channel access server, using access security must make two calls: `asTrapWriteBefore()` and `asTrapWriteAfter()`. These are defined in `asLib.h`. The channel access server on the ioc makes these calls.
2. `asTrapWrite()` gets called by `asTrapWriteBefore()` and `asTrapWriteAfter()` and uses the `TRAPWRITE` option specified with the `RULEs` given in the access configuration file to decide if listeners should be called. `asTrapWrite` also includes a routine `asTrapWriteRegisterListener()`.
3. Some facility not included with access security must call `asTrapWriteRegisterListener()`. If nothing calls `asTrapWriteRegisterListener`, `asTrapWrite` does nothing.

The remainder of this section describes how a facility can use `asTrapWrite.h`, which is defined as:

```
typedef struct asTrapWriteMessage {
    const char *userid;
    const char *hostid;
    void *serverSpecific;
    void *userPvt;
} asTrapWriteMessage;

typedef void *asTrapWriteId;
typedef void (*asTrapWriteListener)(asTrapWriteMessage *pmessage, int after);

asTrapWriteId asTrapWriteRegisterListener(asTrapWriteListener func);
void asTrapWriteUnregisterListener(asTrapWriteId id);
```

After a facility calls `asTrapWriteRegisterListener()` its `asTrapWriteListener()` will get called before and after each write with an associated `RULE` that has the option `TRAPWRITE` set.

`asTrapWriteRegisterListener()` is passed the address of an `asTrapWriteMessage`. This message contains the following fields:

- `userid` - Userid of whoever originated the request.
- `hostid` - Hostid of whoever originated the request.
- `serverSpecific` - The meaning of this field is server specific. If the listener uses this field it must know what type of server is supplying the messages. It is the value the server provides to `asTrapWriteBefore`.
- `userPvt` - This field is for use by the `asTrapWriteListener`. When the listener is called before the write, `userPvt` has the value 0. The listener can give it any value it desires and `userPvt` will have the same value when the listener gets called after the write.

`asTrapWriteListener` delays the associated server thread so it must not do anything that causes it to block.

The IOC's `RSRV` server calls `asTrapWriteBefore` with `serverSpecific` set to the `dbAddr` describing the database location.

8.9 Access Control: Implementation Overview

This section provides a few aids for reading the access security code. Include file `asLib.h` describes the control blocks used by the access security library.

8.9.1 Implementation Overview

The following files form the access security system:

asLib.h Definitions for the portion of access security that is independent of IOC databases.

asDbLib.h Definitions for access routines that interface to an IOC database.

asLib.lex.l `Lex` and `Yacc` (actually `EPICS flex` and `antelope`) are used to parse the access configuration file. This is the `lex` input file.

asLib.y This is the `yacc` input file. Note that it includes `asLibRoutines.c`, which do most of the work.

asLibRoutines.c These are the routines that implement access security. This code has no knowledge of the database or channel access. It is a general purpose access security implementation.

asDbLib.c This contains the code for interfacing access security to the IOC database.

asCa.c This code contains the channel access client code that implements the `INP` and `CALC` definitions in an access security database.


ascheck.c The Unix program which performs a syntax check on a configuration file.

8.9.2 Locking

Because it is possible for multiple tasks to simultaneously modify the access security database it is necessary to provide locking. Rather than try to provide low level locking, the entire access security database is locked during critical operations. The only things this should hold up are access initialization, CA searches, CA clears, and diagnostic routines. It should NEVER cause record processing to wait. In addition CA gets and puts should never be delayed. One exception exists. If the ASG field of a record is changed then `asChangeGroup` is called which locks.

All operations invoked from outside the access security library that cause changes to the internal structures of the access security database.routines lock.

8.10 Structures



Chapter 9

IOC Test Facilities

9.1 Overview

This chapter describes a number of IOC test routines that are of interest to both application developers and system developers. The routines are available from either `iocsh` or the `vxWorks` shell. In both shells the parentheses around arguments are optional. On `vxWorks` all character string arguments must be enclosed in double quote characters `" "` and all arguments must be separated by commas. For `iocsh` single or double quotes must be used around string arguments that contain spaces or commas but are otherwise optional, and arguments may be separated by either commas or spaces. For example:

```
dbpf("aiTest", "2")
dbpf "aiTest", "2"
```

are both valid with both `iocsh` and with the `vxWorks` shell.

```
dbpf aiTest 2
```

Is valid for `iocsh` but not for the `vxWorks` shell.

Both `iocsh` and `vxWorks` shells allow output redirection, i.e. the standard output of any command can be redirected to a file. For example

```
dbl > dbl.lst
```

will send the output of the `dbl` command to the file `dbl.lst`

If `iocsh` is being used it provides help for all commands that have been registered. Just type

```
help
```

or

```
help pattern*
```

9.2 Database List, Get, Put

9.2.1 dbl

Database List:

```
dbl("<record type>", "<field list>")
```

Examples

```
dbl
dbl ("ai")
dbl ("*")
dbl ("")
```

This command prints the names of records in the run time database. If `<record type>` is empty (`""`), `"*"`, or not specified, all records are listed. If `<record type>` is specified, then only the names of the records of that type are listed.

If `<field list>` is given and not empty then the values of the fields specified are also printed.

9.2.2 dbgrep

List Record Names That Match a Pattern:

```
dbgrep ("<pattern>")
```

Examples

```
dbgrep ("S0*")
dbgrep ("*gpibAi*")
```

Lists all record names that match a pattern. The pattern can contain any characters that are legal in record names as well as `"*"`, which matches 0 or more characters.

9.2.3 dbla

List Record Alias Names with optional pattern:

```
dbla
dbla ("<pattern>")
```

Lists the names of all aliases (which match the pattern if given) and the records they refer to. Examples:

```
dbla
dbla "alia*"
```

9.2.4 dba

Database Address:

```
dba ("<record_name.field_name>")
```

Example

```
dba ("aitest")
dba ("aitest.VAL")
```

This command calls `dbNameToAddr` and then prints the value of each field in the `dbAddr` structure describing the field. If the field name is not specified then `VAL` is assumed (the two examples above are equivalent).

9.2.5 dbgf

Get Field:

```
dbgf("<record_name.field_name>")
```

Example:

```
dbgf("aitest")  
dbgf("aitest.VAL")
```

This performs a `dbNameToAddr` and then a `dbGetField`. It prints the field type and value. If the field name is not specified then `VAL` is assumed (the two examples above are equivalent). Note that `dbGetField` locks the record lockset, so `dbgf` will not work on a record with a stuck lockset; use `dbpr` instead in this case.

9.2.6 dbpf

Put Field:

```
dbpf("<record_name.field_name>", "<value>")
```

Example:

```
dbpf("aitest", "5.0")
```

This command performs a `dbNameToAddr` followed by a `dbPutField` and `dbgf`. If `<field_name>` is not specified `VAL` is assumed.

9.2.7 dbpr

Print Record:

```
dbpr("<record_name>", <interest level>)
```

Example

```
dbpr("aitest", 2)
```

This command prints all fields of the specified record up to and including those with the indicated interest level. Interest level has one of the following values:

- 0: Fields of interest to an Application developer and that can be changed as a result of record processing.
- 1: Fields of interest to an Application developer and that do not change during record processing.
- 2: Fields of major interest to a System developer.
- 3: Fields of minor interest to a System developer.
- 4: Fields of no interest.

9.2.8 dbtr

Test Record:

```
dbtr("<record_name>")
```

This calls `dbNameToAddr`, then `dbProcess` and finally `dbpr` (interest level 3). Its purpose is to test record processing.

9.2.9 dbnr

Print number of records:

```
dbnr (<all_recordtypes>)
```

This command displays the number of records of each type and the total number of records. If `all_record_types` is 0 then only record types with record instances are displayed otherwise all record types are displayed.

9.3 Breakpoints

A breakpoint facility that allows the user to step through database processing on a per lockset basis. This facility has been constructed in such a way that the execution of all locksets other than ones with breakpoints will not be interrupted. This was done by executing the records in the context of a separate task.

The breakpoint facility records all attempts to process records in a lockset containing breakpoints. A record that is processed through external means, e.g.: a scan task, is called an entrypoint into that lockset. The `dbstat` command described below will list all detected entrypoints to a lockset, and at what rate they have been detected.

9.3.1 dbb

Set Breakpoint:

```
dbb ("<record_name>")
```

Sets a breakpoint in a record. Automatically spawns the `bkptCont`, or breakpoint continuation task (one per lockset). Further record execution in this lockset is run within this task's context. This task will automatically quit if two conditions are met, all breakpoints have been removed from records within the lockset, and all breakpoints within the lockset have been continued.

9.3.2 dbd

Remove Breakpoint:

```
dbd ("<record_name>")
```

Removes a breakpoint from a record.

9.3.3 dbs

Single Step:

```
dbs ("<record_name>")
```

Steps through execution of records within a lockset. If this command is called without an argument, it will automatically step starting with the last detected breakpoint.

9.3.4 dbc

Continue:

```
dbc ("<record_name>")
```

Continues execution until another breakpoint is found. This command may also be called without an argument.

9.3.5 dbp

Print Fields Of Suspended Record:

```
dbp("<record_name>,<interest_level>")
```

Prints out the fields of the last record whose execution was suspended.

9.3.6 dbap

Auto Print:

```
dbap("<record_name>")
```

Toggles the automatic record printing feature. If this feature is enabled for a given record, it will automatically be printed after the record is processed.

9.3.7 dbstat

Status:

```
dbstat
```

Prints out the status of all locksets that are suspended or contain breakpoints. This lists all the records with breakpoints set, what records have the autoprint feature set (by `dbap`), and what entrypoints have been detected. It also displays the vxWorks task ID of the breakpoint continuation task for the lockset. Here is an example output from this call:

```
LSet: 00009   Stopped at: so#B: 00001   T: 0x23cafac
              Entrypoint: so#C: 00001   C/S:      0.1
              Breakpoint: so(ap)
LSet: 00008#B: 00001   T: 0x22fee4c
              Breakpoint: output
```

The above indicates that two locksets contain breakpoints. One lockset is stopped at record “so.” The other is not currently stopped, but contains a breakpoint at record “output.” “LSet:” is the lockset number that is being considered. “#B:” is the number of breakpoints set in records within that lockset. “T:” is the vxWorks task ID of the continuation task. “C:” is the total number of calls to the entrypoint that have been detected. “C/S:” is the number of those calls that have been detected per second. (ap) indicates that the autoprint feature has been turned on for record “so.”

9.4 Trace Processing

The user should also be aware of the field `TPRO`, which is present in every database record. If it is set `TRUE` then a message is printed each time its record is processed and a message is printed for each record processed as a result of it being processed.

9.5 Error Logging

9.5.1 eltc

Display error log messages on console:

```
eltc(int noYes)
```

This determines if error messages are displayed on the IOC console. 0 means no and any other value means yes.

9.5.2 `errlogInit`, `errlogInit2`

Initialize error log client buffering

```
errlogInit(int bufSize)
errlogInit2(int bufSize, int maxMsgSize)
```

The error log client maintains a circular buffer of messages that are waiting to be sent to the log server. If not set using one or other of these routines the default value for `bufSize` is 1280 bytes and for `maxMsgSize` is 256 bytes.

9.5.3 `errlog`

Send a message to the log server

```
errlog("<message>")
```

This command is provided for use from the `ioc` shell only. It sends its string argument and a new-line to the log server, without displaying it on the IOC console. Note that the `iocsh` will have expanded any environment variable macros in the string (if it was double-quoted) before passing it to `errlog`.

9.6 Hardware Reports

9.6.1 `dbior`

I/O Report:

```
dbior("<driver_name>", <interest level>)
```

This command calls the report entry of the indicated driver. If `<driver_name>` is "" or "*", then a report for all drivers is generated. The command also calls the report entry of all device support modules. Interest level is one of the following:

- 0: Print a short report for each module.
- 1: Print additional information.
- 2: Print even more info. The user may be prompted for options.

9.6.2 `dbhcr`

Hardware Configuration Report:

```
dbhcr()
```

This command produces a report of all hardware links. To use it on the IOC, issue the command:

```
dbhcr > report
```

The report will probably not be in the sort order desired. The Unix command:

```
sort report > report.sort
```

should produce the sort order you desire.

9.7 Scan Reports

9.7.1 scanppl

Print Periodic Lists:

```
scanppl(double rate)
```

This routine prints a list of all records in the periodic scan list of the specified rate. If rate is 0.0 all period lists are shown.

9.7.2 scanpel

Print Event Lists:

```
scanpel(int event_number)
```

This routine prints a list of all records in the event scan list for the specified event number. If event_number is 0 all event scan lists are shown.

9.7.3 scanpiol

Print I/O Event Lists:

```
scanpiol
```

This routine prints a list of all records in the I/O event scan lists.

9.8 General Time

The built-in time providers depend on the IOC's target architecture, so some of the specific subsystem report commands listed below are only available on the architectures that use that particular provider.

9.8.1 generalTimeReport

Format:

```
generalTimeReport(int level)
```

This routine displays the time providers and their priority levels that have registered with the General Time subsystem for both current and event times. At level 1 it also shows the current time as obtained from each provider.

9.8.2 installLastResortEventProvider

Format:

```
installLastResortEventProvider
```

Installs the optional Last Resort event provider at priority 999, which returns the current time for every event number.

9.8.3 NTPTime_Report

Format:

```
NTPTime_Report (int level)
```

Only vxWorks and RTEMS targets use this time provider. The report displays the provider's synchronization state, and at interest level 1 it also gives the synchronization interval, when it last synchronized, the nominal and measured system tick rates, and on vxWorks the NTP server address.

9.8.4 NTPTime_Shutdown

Format:

```
NTPTime_Shutdown
```

On vxWorks and RTEMS this command shuts down the NTP time synchronization thread. With the thread shut down, the driver will no longer act as a current time provider.

9.8.5 ClockTime_Report

Format:

```
ClockTime_Report (int level)
```

This time provider is used on several target architectures, registered as the time provider of last resort. On vxWorks and RTEMS the report displays the synchronization state, when it last synchronized the system time with a higher priority provider, and the synchronization interval. On workstation operating systems the synchronization task is not started on the assumption that some other process is taking care of synchronizing the OS clock as appropriate, so the report is minimal.

9.8.6 ClockTime_Shutdown

Format:

```
ClockTime_Shutdown
```

Some sites may prefer to provide their own implementation of a system clock time provider to replace the built-in one. On vxWorks and RTEMS this command stops the OS Clock synchronization thread, allowing the OS clock to free-run. The time provider *will* continue to return the current system time after this command is used however.

9.9 Access Security Commands

9.9.1 asSetSubstitutions

Format:

```
asSetSubstitutions ("substitutions")
```

Specifies macro substitutions used when access security is initialized.

9.9.2 asSetFilename

Format:

```
asSetFilename("<filename>")
```

This command defines a new access security file.

9.9.3 asInit

Format:

```
asInit
```

This command reinitializes the access security system. It rereads the access security file in order to create the new access security database. This command is useful either because the `asSetFilename` command was used to change the file or because the file itself was modified. Note that it is also possible to reinitialize the access security via a subroutine record. See the access security document for details.

9.9.4 asdbdump

Format:

```
asdbdump
```

This provides a complete dump of the access security database.

9.9.5 aspuag

Format:

```
aspuag("<user access group>")
```

Print the members of the user access group. If no user access group is specified then the members of all user access groups are displayed.

9.9.6 asphag

Format:

```
asphag("<host access group>")
```

Print the members of the host access group. If no host access group is specified then the members of all host access groups are displayed.

9.9.7 asprules

Format:

```
asprules("<access security group>")
```

Print the rules for the specified access security group or if no group is specified for all groups.

9.9.8 aspmem

Format:

```
aspmem("<access security group>", <print clients>)
```

Print the members (records) that belong to the specified access security group, for all groups if no group is specified. If <print clients> is (0, 1) then Channel Access clients attached to each member (are not, are) shown.

9.10 Channel Access Reports

9.10.1 casr

Channel Access Server Report

```
casr(<level>)
```

Level can have one of the following values:

0

Prints server's protocol version level and a one line summary for each client attached. The summary lines contain the client's login name, client's host name, client's protocol version number, and the number of channel created within the server by the client.

1

Level one provides all information in level 0 and adds the task id used by the server for each client, the client's IP protocol type, the file number used by the server for the client, the number of seconds elapsed since the last request was received from the client, the number of seconds elapsed since the last response was sent to the client, the number of unprocessed request bytes from the client, the number of response bytes which have not been flushed to the client, the client's IP address, the client's port number, and the client's state.

2

Level two provides all information in levels 0 and 1 and adds the number of bytes allocated by each client and a list of channel names used by each client. Level 2 also provides information about the number of bytes in the server's free memory pool, the distribution of entries in the server's resource hash table, and the list of IP addresses to which the server is sending beacons. The channel names are shown in the form:

```
<name>(nrw)
```

where

n is number of ca_add_events the client has on this channel

r is (-,R) if client (does not, does) have read access to the channel.

w is(-, W) if client (does not, does) have write access to the channel.

9.10.2 dbel

Format:

```
dbel ("<record_name>")
```

This routine prints the Channel Access event list for the specified record.

9.10.3 dbcar

Database to Channel Access Report - See "Record Link Reports"

9.10.4 ascar

Format:

```
ascar (level)
```

Prints a report of the channel access links for the INP fields of the access security rules. Level 0 produces a summary report. Level 1 produces a summary report plus details on any unconnect channels. Level 2 produces the summary report plus a detail report on each channel.

9.11 Interrupt Vectors**9.11.1 veclist**

Format:

```
veclist
```

NOTE: This routine is only available on vxWorks. On PowerPC CPUs it requires BSP support to work, and even then it cannot display chained interrupts using the same vector.

Print Interrupt Vector List

9.12 Miscellaneous**9.12.1 epicsParamShow**

Format:

```
epicsParamShow
or
epicsPrtEnvParams
```

Print the environment variables that are created with epicsEnvSet. These are defined in <base>/config/CONFIG_ENV and <base>/config/CONFIG_SITE_ENV or else by user applications calling epicsEnvSet.

9.12.2 epicsEnvShow

Format:

```
epicsEnvShow ("<name>")
```

Show Environment variables. On vxWorks it shows the variables created via calls to putenv.

9.12.3 coreRelease

Format:

```
coreRelease
```

Print release information for iocCore.

9.13 Database System Test Routines

These routines are normally only of interest to EPICS system developers NOT to Application Developers.

9.13.1 dbtgf

Test Get Field:

```
dbtgf("<record_name.field_name>")
```

Example:

```
dbtgf("aitest")
dbtgf("aitest.VAL")
```

This performs a `dbNameToAddr` and then calls `dbGetField` with all possible request types and options. It prints the results of each call. This routine is of most interest to system developers for testing database access.

9.13.2 dbtpf

Test Put Field:

```
dbtpf("<record_name.field_name>", "<value>")
```

Example:

```
dbtpf("aitest", "5.0")
```

This command performs a `dbNameToAddr`, then calls `dbPutField`, followed by `dbgf` for each possible request type. This routine is of interest to system developers for testing database access.

9.13.3 dbtpn

Test Put Notify:

```
dbtpn("<record_name.field_name>", "<value>")
```

Example:

```
dbtpn("aitest", "5.0")
```

This command performs a `dbNameToAddr`, then calls `dbPutNotify` and has a callback routine that prints a message when it is called. This routine is of interest to system developers for testing database access.

9.14 Record Link Reports

9.14.1 dblsr

Lock Set Report:

```
dblsr (<recordname>, <level>)
```

This command generates a report showing the lock set to which each record belongs. If `recordname` is 0, "", or "*" all records are shown, otherwise only records in the same lock set as `recordname` are shown.

`level` can have the following values:

- 0 - Show lock set information only.
- 1 - Show each record in the lock set.
- 2 - Show each record and all database links in the lock set.

9.14.2 dbLockShowLocked

Show locked locksets:

```
dbLockShowLocked (<level>)
```

This command generates a report showing all locked locksets, the records they contain, the lockset state and the thread that currently owns the lockset. The `level` argument is passed to `epicsMutexShow` to adjust the information reported about each locked `epicsMutex`.

9.14.3 dbcar

Database to channel access report

```
dbcar (<recordname>, <level>)
```

This command generates a report showing database channel access links. If `recordname` is "*" then information about all records is shown otherwise only information about the specified record.

`level` can have the following values:

- 0 - Show summary information only.
- 1 - Show summary and each CA link that is not connected.
- 2 - Show summary and status of each CA link.

9.14.4 dbhcr

Report hardware links. See "Hardware Reports".

9.15 Old Database Access Testing

These routines are of interest to EPICS system developers. They are used to test the old database access interface, which is still used by Channel Access.

9.15.1 gft

Get Field Test:

```
gft("<record_name.field_name>")
```

Example:

```
gft("aitest")
gft("aitest.VAL")
```

This performs a `db_name_to_addr` and then calls `db_get_field` with all possible request types. It prints the results of each call. This routine is of interest to system developers for testing database access.

9.15.2 pft

Put Field Test:

```
pft("<record_name.field_name>", "<value>")
```

Example:

```
pft("aitest", "5.0")
```

This command performs a `db_name_to_addr`, `db_put_field`, `db_get_field` and prints the result for each possible request type. This routine is of interest to system developers for testing database access.

9.15.3 tpn

Test Put Notify:

```
tpn("<record_name.field_name>", "<value>")
```

Example:

```
tpn("aitest", "5.0")
```

This routine tests `dbPutNotify` via the old database access interface.

9.16 Routines to dump database information

9.16.1 dbDumpPath

Dump Path:

```
dbDumpPath(pdbname)
dbDumpPath(pdbname)
```

The current path for database includes is displayed.

9.16.2 dbDumpMenu

Dump Menu:

```
dbDumpMenu(pdbname, "<menu>")
dbDumpMenu(pdbname, "menuScan")
```

If the second argument is 0 then all menus are displayed.

9.16.3 dbDumpRecordType

Dump Record Description:

```
dbDumpRecordType (pdbbase, "<record type>")
dbDumpRecordType (pdbbase, "ai")
```

If the second argument is 0 then all descriptions of all records are displayed.

9.16.4 dbDumpField

Dump Field Description:

```
dbDumpField (pdbbase, "<record type>", "<field name>")
dbDumpField (pdbbase, "ai", "VAL")
```

If the second argument is 0 then the field descriptions of all records are displayed. If the third argument is 0 then the description of all fields are displayed.

9.16.5 dbDumpDevice

Dump Device Support:

```
dbDumpDevice (pdbbase, "<record type>")
dbDumpDevice (pdbbase, "ai")
```

If the second argument is 0 then the device support for all record types is displayed.

9.16.6 dbDumpDriver

Dump Driver Support:

```
dbDumpDriver (pdbbase)
dbDumpDriver (pdbbase)
```

9.16.7 dbDumpRecord

Dump Record Instances:

```
dbDumpRecord (pdbbase, "<record type>", level)
dbDumpRecords (pdbbase, "ai")
```

If the second argument is 0 then the record instances for all record types are displayed. The third argument determines which fields are displayed just like for the command `dbpr`.

9.16.8 dbDumpBreaktable

Dump breakpoint table

```
dbDumpBreaktable (pdbname, name)
dbDumpBreaktable (pdbname, "typeKdegF")
```

This command dumps a breakpoint table. If the second argument is 0 all breakpoint tables are dumped.

9.16.9 dbPvdDump

Dump the Process variable Directory:

```
dbPvdDump (pdbname, verbose)
dbPvdDump (pdbname, 0)
```

This command shows how many records are mapped to each hash table entry of the process variable directory. If verbose is not 0 then the command also displays the names which hash to each hash table entry.

Chapter 10

IOC Error Logging

10.1 Overview

Errors detected by an IOC can be divided into classes: Errors related to a particular client and errors not attributable to a particular client. An example of the first type of error is an illegal Channel Access request. For this type of error, a status value should be passed back to the client. An example of the second type of error is a device driver detecting a hardware error. This type of error should be reported to a system wide error handler.

Dividing errors into these two classes is complicated by a number of factors.

- In many cases it is not possible for the routine detecting an error to decide which type of error occurred.
- Normally, only the routine detecting the error knows how to generate a fully descriptive error message. Thus, if a routine decides that the error belongs to a particular client and merely returns an error status value, the ability to generate a fully descriptive error message is lost.
- If a routine always generates fully descriptive error messages then a particular client could cause error message storms.
- While developing a new application the programmer normally prefers fully descriptive error messages. For a production system, however, the system wide error handler should not normally receive error messages cause by a particular client.

If used properly, the error handling facilities described in this chapter can process both types of errors.

This chapter describes the following:

- Error Message Generation Routines - Routines which pass messages to the errlog Task.
- Error Log Listeners - Any code can register to receive errlog messages.
- errlogThread - A thread that passes the messages to all registered listeners.
- console output and message buffer size - Messages can also be written to the console. The storage for the message queue can be specified by the user.
- status codes - EPICS status codes.
- iocLog- A system wide error logger supplied with base. It writes all messages to a system wide file.

NOTE: Many sites use CMLOG instead of iocLog.

NOTE: `recGbl` error routines are also provided. They in turn call one of the error message routines.

10.2 Error Message Routines

10.2.1 Basic Routines

```
int errlogPrintf(const char *pformat, ...);
int errlogVprintf(const char *pformat, va_list pvar);
int errlogMessage(const char *message);
void errlogFlush(void);
```

`errlogPrintf` and `errlogVprintf` are like `printf` and `vprintf` provided by the standard C library, except that their output is sent to the `errlog` task; unless configured not to, the output will appear on the console as well. Consult any book that describes the standard C library such as “The C Programming Language ANSI C Edition” by Kernighan and Ritchie.

`errlogMessage` sends message to the `errlog` task.

`errlogFlush` wakes up the `errlog` task and then waits until all messages are flushed from the queue.

10.2.2 Log with Severity

```
typedef enum {
    errlogInfo, errlogMinor, errlogMajor, errlogFatal
}errlogSevEnum;

int errlogSevPrintf(const errlogSevEnum severity,
    const char *pformat, ...);
int errlogSevVprintf(const errlogSevEnum severity,
    const char *pformat, va_list pvar);

char *errlogGetSevEnumString(const errlogSevEnum severity);

void errlogSetSevToLog(const errlogSevEnum severity );
errlogSevEnum errlogGetSevToLog(void);
```

`errlogSevPrintf` and `errlogSevVprintf` are like `errlogPrintf` and `errlogVprintf` except that they add the severity to the beginning of the message in the form “sevr=<value>” where value is one of “info, minor, major, fatal”. Also the message is suppressed if severity is less than the current severity to suppress. If `epicsThreadIsOkToBlock` is true, which is true during `iocInit`, `errlogSevVprintf` does NOT send output to the `errlog` task.

`errlogGetSevEnumString` gets the string value of severity.

`errlogSetSevToLog` sets the severity to log. `errlogGetSevToLog` gets the current severity to log.

10.2.3 Status Routines

```
void errMessage(long status, char *message);

void errPrintf(long status, const char *pFileName,
    int lineno, const char *pformat, ...);
```

Routine `errMessage` (actually a macro that calls `errPrintf`) has the following format:

```
void errMessage(long status, char *message);
```

Where `status` is defined as:

- 0: Find latest vxWorks or Unix error.
- -1: Don't report status.
- Other: See "Return Status Values" above.

`errMessage`, via a call to `errPrintf`, prints the message, the status symbol and string values, and the name of the task which invoked `errMessage`. It also prints the name of the source file and the line number from which the call was issued.

The calling routine is expected to pass a descriptive message to this routine. Many subsystems provide routines built on top of `errMessage` which generate descriptive messages.

An IOC global variable `errVerbose`, defined as an external in `errMdef.h`, specifies verbose messages. If `errVerbose` is `TRUE` then `errMessage` should be called whenever an error is detected even if it is known that the error belongs to a specific client. If `errVerbose` is `FALSE` then `errMessage` should be called only for errors that are not caused by a specific client.

Routine `errPrintf` is normally called as follows:

```
errPrintf(status, __FILE__, __LINE__, "<fmt>", ...);
```

Where `status` is defined as:

- 0: Find latest vxWorks or Unix error.
- -1: Don't report status.
- Other: See "Return Status Values", above.

`FILE` and `LINE` are defined as:

- `__FILE__` As shown or `NULL` if the file name and line number should not be printed.
- `__LINE__` As shown

The remaining arguments are just like the arguments to the C `printf` routine. `errVerbose` determines if the filename and line number are shown.

An EPICS status code can also be converted to a string. If the supplied status code isn't registered in the status code database then the raw status code number is converted into a string in the destination buffer.

```
#include "errMdef.h"
void errSymLookup(long status, char *pBuf, unsigned bufLength);
```

10.2.4 Obsolete Routines

```
int epicsPrintf(const char *pformat, ...);
int epicsVprintf(const char *pformat, va_list pvar);
```

These are macros that call `erlogPrintf` and `erlogVprintf`. They are provided for compatibility.

10.3 errlog Listeners

Any code can receive `erlog` message. The following are the calls to add and remove a listener.

```
typedef void(*errlogListener) (void *pvt, const char *message);
void errlogAddListener(errlogListener listener, void *pPrivate);
void errlogRemoveListener(errlogListener listener);
```

These routines add/remove a callback that receives each error message. These routines are the interface to the actual system wide error handlers.

10.4 `errlogThread`

The error message routines can be called by any non-interrupt level code. These routines pass the message to the `errlogThread`. If any of the error message routines are called at interrupt level, `epicsInterruptContextMessage` is called with the message “`errlogPrintf called from interrupt level`”.

`errlogThread` manages the messages. Messages are placed in a message queue, which is read by `errlogThread`. The message queue uses a fixed block of memory to hold all messages. When the message queue is full additional messages are rejected but a count of missed messages is kept. The next time the message queue empties an extra message about the missed messages is generated.

The maximum message size is by default 256 characters. If a message is longer, the message is truncated and a message explaining that it was truncated is appended. There is a chance that long messages corrupt memory. This only happens if client code is defective. Long messages most likely result from “`%s`” formats with a bad string argument.

`errlogThread` passes each message to any registered listener.

10.5 console output and message queue size

The `errlog` system can also display messages on the ioc console. It calls `epicsThreadIsOkToBlock` to decide when to display the message. If it is OK to block, the message is displayed by the same thread that calls one of the `errlog` print routines. If it is not OK to block, `errlogThread` displays the messages.

Normally the `errlog` system displays all messages on the console. `eltc` can be used to suppress these messages.

```
int eltc(int yesno); /* error log to console (0 or 1) */
int errlogInit(int bufsize);
int errlogInit2(int bufsize, int maxMsgSize);
```

`eltc` determines if `errlog` task writes message to the console. During error message storms this command can be used to suppress console messages. A argument of 0 suppresses the messages, any other value lets messages go to the console.

`errlogInit` or `errlogInit2` can be used to initialize the error logging system with a larger buffer and maximum message size. The default buffer size is 1280 bytes, and the default maximum message size is 256.

10.6 Status Codes

EPICS defined status values provide the following features:

- Whenever possible, IOC routines return a status value: 0 means OK, non-0 means an error.
- The header files for most IOC subsystems contain macros defining error status symbols and strings.
- Routines are provided for run time access of the error status symbols and strings.
- A global variable `errVerbose` helps code decide if error messages should be generated.

IOC routines often return a long integer status value, encoded similar to the `vxWorks` error status encoding. On some modern architectures a long integer is more than 32 bits wide, but in order to keep the API compatible the status values are still passed as long integers, even though only 32 bits are ever used. The most significant 16 bits indicate the subsystem or module where the error occurred. The least significant 16 bits contain a subsystem-specific status value.

In order that status values do not conflict with the vxWorks error status values, all subsystem numbers are greater than 500.

A header file `errMdef.h` defines macros for all the subsystem numbers. For example the database access routines use this module number:

```
#define M_dbAccess (501 << 16) /*Database Access Routines*/
```

There are header files for every IOC subsystem that returns standard status values. The status values are encoded with lines of the following format:

```
#define S_xxxxxxx value /*string value*/
```

For example:

```
#define S_dbAccessBadDBR (M_dbAccess|3) /*Invalid Database Request*/
```

For example, when `dbGetField` detects a bad database request type, it executes the statement:

```
return (S_dbAccessBadDBR);
```

The calling routine checks the return status as follows:

```
status = dbGetField(...);
if(status) { /* Call was not successful */ }
```

10.7 iocLog

NOTE: Many sites use CMLOG instead of iocLog. See the CMLOG documentation for details.

This consists of two modules: `iocLogServer` and `iocLogClient`. The client code runs on each ioc and listens for the messages generated locally by the `errlog` system. It also reports the messages from the vxWorks `logMsg` facility.

10.7.1 iocLogServer

This runs on a host. It receives messages for all enabled `iocLogClients` in the local area network. The messages are written to a file. Epics base provides a startup file “`base/src/util/rc2.logServer`”, which is a SystemV init script to start the server. Consult this script for details.

To start a log server on a UNIX or PC workstation you must first set the following environment variables and then run the executable “`iocLogServer`” on your PC or UNIX workstation.

`EPICS_IOC_LOG_FILE_NAME`

The name and path to the log file.

`EPICS_IOC_LOG_FILE_LIMIT`

The maximum size in characters for the log file (after which it becomes a circular file and writes new messages over old messages at the beginning of the file). If the value is zero then there is no limit on the size of the log file.

`EPICS_IOC_LOG_FILE_COMMAND`

A shell command string used to obtain the log file path name during initialization and in response to `SIGHUP`. The new path name will replace any path name supplied in `EPICS_IOC_LOG_FILE_NAME`.

Thus, if `EPICS_IOC_LOG_FILE_NAME` is “`a/b/c.log`” and `EPICS_IOC_LOG_FILE_COMMAND` returns “`A/B`” or “`A/B/`” the log server will be stored at “`A/B/c.log`”

If `EPICS_IOC_LOG_FILE_COMMAND` is empty then this behavior is disabled. This feature is used at some sites for switching the server to a new directory at a fixed time each day. This variable is currently used only by the UNIX version of the log server.

`EPICS_IOC_LOG_PORT`

THE TCP/IP port used by the log server.

To configure an IOC so that its messages are placed in the log you must set the environment variable `EPICS_IOC_LOG_INET` to the IP address of the host that is running the log server, and `EPICS_IOC_LOG_PORT` to the TCP/IP port used by the log server.

Defaults for all of the above parameters are specified in the files `$(EPICS_BASE)/config/CONFIG_SITE_ENV` and `$(EPICS_BASE)/config/CONFIG_ENV`.

10.7.2 `iocLogClient`

This runs on each ioc. It is started by calling:

```
iocLogInit();
```

The global variable `iocLogDisable` can be used to enable/disable the messages from being sent to the server. Setting this variable to (0,1) (enables, disables) the messages generation. If `iocLogDisable` is set to 1 before calling `iocLogInit` then `iocLogClient` will not even initialize itself. `iocLogDisable` can also be changed to turn logging on or off.

`iocLogClient` calls `errlogAddListener` and sends each message to the `iocLogServer`.

10.7.3 Configuring a Private Log Server

In a testing environment it is desirable to use a private log server. This can be done as follows:

- Add a `epicsEnvSet` command to your IOC startup file. For example

```
ld < iocCore
epicsEnvSet ("EPICS_IOC_LOG_INET=xxx.xxx.xxx.xxx")
```

- The inet address is that of your host workstation.
- On your host workstation, start the log server.

Chapter 11

Record Support

11.1 Overview

The purpose of this chapter is to describe record support in sufficient detail such that a C programmer can write new record support modules. Before attempting to write new support modules, you should carefully study a few of the existing support modules. If an existing support module is similar to the desired module most of the work will already be done.

From previous chapters, it should be clear that many things happen as a result of record processing. The details of what happens are dependent on the record type. In order to allow new record types and new device types without impacting the core IOC system, the concept of record support and device support is used. For each record type, a record support module exists. It is responsible for all record specific details. In order to allow a record support module to be independent of device specific details, the concept of device support has been created.

A record support module consists of a standard set of routines which are called by database access routines. These routines implement record specific code. Each record type can define a standard set of device support routines specific to that record type.

By far the most important record support routine is `process`, which `dbProcess` calls when it wants to process a record. This routine is responsible for the details of record processing. In many cases it calls a device support I/O routine. The next section gives an overview of what must be done in order to process a record. Next is a description of the entry tables that must be provided by record and device support modules. The remaining sections give example record and device support modules and describe some global routines useful to record support modules.

The record and its device support modules are the only source files that should include the record specific header files. Thus they will be the only routines that access record specific fields without going through database access.

11.2 Overview of Record Processing

The most important record support routine is `process`. This routine determines what record processing means. Before the record specific “`process`” routine is called, the following has already been done:

- Decision to process a record.
- Check that record is not active, i.e. `pact` must be `FALSE`.
- Check that the record is not disabled.

The `process` routine, together with its associated device support, is responsible for the following tasks:

- Set record active while it is being processed

- Perform I/O (with aid of device support)
- Check for record specific alarm conditions
- Raise database monitors
- Request processing of forward links

A complication of record processing is that some devices are intrinsically asynchronous. It is NEVER permissible to wait for a slow device to complete. Asynchronous records perform the following steps:

1. Initiate the I/O operation and set `pact = TRUE`
2. Determine a method for again calling process when the operation completes
3. Return immediately without completing record processing
4. When process is called after the I/O operation complete record processing
5. Set `pact = FALSE` and return

The examples given below show how this can be done.

11.3 Record Support and Device Support Entry Tables

Each record type has an associated set of record support routines. These routines are located via the `struct rset` data structure declared in `recSup.h` and defined by the specific record type. This use of a record support vector table isolates the `iocCore` software from the implementation details of each record type. Thus new record types can be defined without having to modify the IOC core software.

Each record type also has zero or more sets of device support routines. Record types without associated hardware, e.g. calculation records, normally do not have any associated device support. Record types with associated hardware normally have a device support module for each device type. The concept of device support isolates IOC core software and even record support from device specific details.

Corresponding to each record type is a set of record support routines. The set of routines is the same for every record type. These routines are located via a Record Support Entry Table (RSET), which has the following structure:

```

struct rset { /* record support entry table */
    long      number;          /* number of support routine */
    RECSUPFUN report;         /* print report */
    RECSUPFUN init;          /* init support */
    RECSUPFUN init_record;   /* init record */
    RECSUPFUN process;       /* process record */
    RECSUPFUN special;       /* special processing */
    RECSUPFUN get_value;     /* OBSOLETE: Just leave NULL */
    RECSUPFUN cvt_dbaddr;    /* cvt dbAddr */
    RECSUPFUN get_array_info;
    RECSUPFUN put_array_info;
    RECSUPFUN get_units;
    RECSUPFUN get_precision;
    RECSUPFUN get_enum_str;  /* get string from enum */
    RECSUPFUN get_enum_strs; /* get all enum strings */
    RECSUPFUN put_enum_str;  /* put enum from string */
    RECSUPFUN get_graphic_double;
    RECSUPFUN get_control_double;
    RECSUPFUN get_alarm_double;
};

```

Each record support module must define its RSET. The external name must be of the form:

```
<record_type>RSET
```

Any routines not needed for the particular record type should be initialized to the value NULL. Look at the example below for details.

Device support routines are located via a Device Support Entry Table (DSET), which has the following structure:

```
struct dset { /* device support entry table */
    long      number; /* number of support routines */
    DEVSUPFUN report; /* print report */
    DEVSUPFUN init; /* init support */
    DEVSUPFUN init_record; /* init record instance*/
    DEVSUPFUN get_ioint_info; /* get io interrupt info*/
    /* other functions are record dependent*/
};
```

Each device support module must define its associated DSET. The external name must be the same as the name which appears in devSup.ascii.

Any record support module which has associated device support must also include definitions for accessing its associated device support modules. The field dset, which is declared in dbCommon, contains the address of the DSET. It is given a value by iocInit.

11.4 Example Record Support Module

This section contains the skeleton of a record support package. The record type is xxx and the record has the following fields in addition to the dbCommon fields: VAL, PREC, EGU, HOPR, LOPR, HIHI, LOLO, HIGH, LOW, HHSV, LLSV, HSV, LSV, HYST, ADEL, MDEL, LALM, ALST, MLST. These fields will have the same meaning as they have for the ai record. Consult the Record Reference manual for a description.

11.4.1 Declarations

```
/* Create RSET - Record Support Entry Table*/
#define report NULL
#define initialize NULL
static long init_record();
static long process();
#define special NULL
#define get_value NULL
#define cvt_dbaddr NULL
#define get_array_info NULL
#define put_array_info NULL
static long get_units();
static long get_precision();
#define get_enum_str NULL
#define get_enum_strs NULL
#define put_enum_str NULL
static long get_graphic_double();
static long get_control_double();
static long get_alarm_double();

rset xxxRSET={
```

```

    RSETNUMBER,
    report,
    initialize,
    init_record,
    process,
    special,
    get_value,
    cvt_dbaddr,
    get_array_info,
    put_array_info,
    get_units,
    get_precision,
    get_enum_str,
    get_enum_strs,
    put_enum_str,
    get_graphic_double,
    get_control_double,
    get_alarm_double
};
epicsExportAddress (rset, xxxRSET);

/* declarations for associated DSET */
typedef struct xxxdset { /* analog input dset */
    long    number;
    DEVSUPFUN    dev_report;
    DEVSUPFUN    init;
    DEVSUPFUN    init_record; /* returns: (1,0)=> (failure, success)*/
    DEVSUPFUN    get_ioint_info;
    DEVSUPFUN    read_xxx;
}xxxdset;

/* forward declaration for internal routines*/
static void checkAlarams(xxxRecord *pxxx);
static void monitor(xxxRecord *pxxx);

```

The above declarations define the Record Support Entry Table (RSET), a template for the associated Device Support Entry Table (DSET), and forward declarations to private routines.

The RSET must be declared with an external name of xxxRSET. It defines the record support routines supplied for this record type. Note that forward declarations are given for all routines supported and a NULL declaration for any routine not supported.

The template for the DSET is declared for use by this module.

11.4.2 init_record

```

static long init_record(void *precord, int pass)
{
    xxxRecord*pxxx = (xxxRecord *)precord;
    xxxdset*pdset;
    longstatus;

    if(pass==0) return(0);

```

```

    if((pdset = (xxx_dset *) (pxxx->dset)) == NULL) {
        recGblRecordError(S_dev_noDSET, pxxx, "xxx: init_record");
        return(S_dev_noDSET);
    }
    /* must have read_xxx function defined */
    if( (pdset->number < 5) || (pdset->read_xxx == NULL) ) {
        recGblRecordError(S_dev_missingSup, pxxx,
            "xxx: init_record");
        return(S_dev_missingSup);
    }
    if( pdset->init_record ) {
        if((status=(*pdset->init_record)(pxxx))) return(status);
    }
    return(0);
}

```

This routine, which is called by `iocInit` twice for each record of type `xxx`, checks to see if it has a proper set of device support routines and, if present, calls the `init_record` entry of the DSET.

During the first call to `init_record` (`pass=0`) only initializations relating to this record can be performed. During the second call (`pass=1`) initializations that may refer to other records can be performed. Note also that during the second pass, other records may refer to fields within this record. A good example of where these rules are important is a waveform record. The `VAL` field of a waveform record actually refers to an array. The waveform record support module must allocate storage for the array. If another record has a database link referring to the waveform `VAL` field then the storage must be allocated before the link is resolved. This is accomplished by having the waveform record support allocate the array during the first pass (`pass=0`) and having the link reference resolved during the second pass (`pass=1`).

11.4.3 process

```

static long process(void *precord)
{
    xxxRecord* pxxx = (xxxRecord *)precord;
    xxx_dset* pdset = (xxx_dset *)pxxx->dset;
    long status;
    unsigned char pact=pxxx->pact;

    if( (pdset==NULL) || (pdset->read_xxx==NULL) ) {
        /* leave pact true so that dbProcess doesnt call again*/
        pxxx->pact=TRUE;
        recGblRecordError(S_dev_missingSup, pxxx, "read_xxx");
        return(S_dev_missingSup);
    }

    /* pact must not be set true until read_xxx completes*/
    status=(*pdset->read_xxx)(pxxx); /* read the new value */
    /* return if beginning of asynch processing*/
    if(!pact && pxxx->pact) return(0);
    pxxx->pact = TRUE;
    recGblGetTimeStamp(pxxx);

    /* check for alarms */
    alarm(pxxx);
    /* check event list */
}

```

```

    monitor(pxxx);
    /* process the forward scan link record */
    recGblFwdLink(pxxx);

    pxxx->pact=FALSE;
    return(status);
}

```

The record processing routines are the heart of the IOC software. The record specific process routine is called by `dbProcess` whenever it decides that a record should be processed. Process decides what record processing really means. The above is a good example of what should be done. In addition to being called by `dbProcess` the process routine may also be called by asynchronous record completion routines.

The above model supports both synchronous and asynchronous device support routines. For example, if `read_xxx` is an asynchronous routine, the following sequence of events will occur:

- process is called with `pact FALSE`
- `read_xxx` is called. Since `pact` is `FALSE` it starts I/O, arranges callback, and sets `pact TRUE`
- `read_xxx` returns
- because `pact` went from `FALSE` to `TRUE` process just returns
- Any new call to `dbProcess` is ignored because it finds `pact TRUE`
- Sometime later the callback occurs and process is called again.
- `read_xxx` is called. Since `pact` is `TRUE` it knows that it is a completion request.
- `read_xxx` returns
- process completes record processing
- `pact` is set `FALSE`
- process returns

At this point the record has been completely processed. The next time process is called everything starts all over from the beginning.

11.4.4 Miscellaneous Utility Routines

```

static long get_units(DBADDR *paddr, char *units)
{
    xxxRecord *pxxx=(xxxRecord *)paddr->precord;

    strncpy(units,pxxx->egu,sizeof(pxxx->egu));
    return(0);
}

static long get_graphic_double(DBADDR *paddr,
    struct dbr_grDouble *pgd)
{
    xxxRecord *pxxx=(xxxRecord *)paddr->precord;
    int fieldIndex = dbGetFieldIndex(paddr);

    if(fieldIndex == xxxRecordVAL) {
        pgd->upper_disp_limit = pxxx->hopr;
        pgd->lower_disp_limit = pxxx->lopr;
    }
}

```



```

        } else recGblGetGraphicDouble (paddr,pgd);
        return(0);
    }
    /* similar routines would be provided for */
    /* get_control_double and get_alarm_double*/

```

These are a few examples of various routines supplied by a typical record support package. The functions that must be performed by the remaining routines are described in the next section.

11.4.5 Alarm Processing

```

static void checkAlarms (xxxRecord *pxxx)
{
    double          val;
    float           hyst, lalm, hihi, high, low, lolo;
    unsigned short  hhsv, llsv, hsv, lsv;

    if (pxxx->udf == TRUE ) {
        recGblSetSevr (pxxx, UDF_ALARM, VALID_ALARM);
        return;
    }

    hihi=pxxx->hihi; lolo=pxxx->lolo;
    high=pxxx->high; low=pxxx->low;
    hhsv=pxxx->hhsv; llsv=pxxx->llsv;
    hsv=pxxx->hsv; lsv=pxxx->lsv;
    val=pxxx->val; hyst=pxxx->hyst; lalm=pxxx->lalm;

    /* alarm condition hihi */
    if (hhsv && (val >= hihi
    || ((lalm==hihi) && (val >= hihi-hyst)))) {
        if (recGblSetSevr (pxxx, HIHI_ALARM, pxxx->hhsv)
            pxxx->lalm = hihi;
        return;
    }
    /* alarm condition lolo */
    if (llsv && (val <= lolo
    || ((lalm==lolo) && (val <= lolo+hyst)))) {
        if (recGblSetSevr (pxxx, LOLO_ALARM, pxxx->llsv)
            pxxx->lalm = lolo;
        return;
    }
    /* alarm condition high */
    if (hsv && (val >= high
    || ((lalm==high) && (val >= high-hyst)))) {
        if (recGblSetSevr (pxxx, HIGH_ALARM, pxxx->hsv)
            pxxx->lalm = high;
        return;
    }
    /* alarm condition low */
    if (lsv && (val <= low
    || (lalm==low) && (val <= low+hyst)))) {
        if (recGblSetSevr (pxxx, LOW_ALARM, pxxx->lsv)

```

```

        pxxx->lalm = low;
    return;
}
/*we get here only if val is out of alarm by at least hyst*/
pxxx->lalm=val;
return;
}

```

This is a typical set of code for checking alarms conditions for an analog type record. The actual set of code can be very record specific. Note also that other parts of the system can raise alarms. The algorithm is to always maximize alarm severity, i.e. the highest severity outstanding alarm will be reported.

The above algorithm also honors a hysteresis factor for the alarm. This is to prevent alarm storms from occurring in the event that the current value is very near an alarm limit and noise makes it continually cross the limit. It honors the hysteresis only when the value is going to a lower alarm severity.

Note the test:

```

    if (pxxx->udf == TRUE ) {
        recGblSetSevr (pxxx, UDF_ALARM, VALID_ALARM) ;
        return;
    }

```

Database common defines the field UDF, which means that field VAL is undefined. The STAT and SEVR fields are initialized as though `recGblSetSevr (pxxx, UDF_ALARM, VALID_ALARM)` was called. Thus if the record is never processed the record will be in an INVALID UNDEFINED alarm state. Field UDF is initialized to the value 1, i.e. TRUE. Thus the above code will keep the record in the INVALID UNDEFINED alarm state as long as UDF is not given the value 0.

The UDF field means Undefined, i.e. the VAL field has never been given a value. When records are loaded into an ioc this is the initial state of records. Whenever code gives a value to the VAL field it is also supposed to set UDF false. Unless a particular record type has unusual semantics no code should set UDF true. UDF normally means that the field was never given a value.

For input records device support is responsible for obtaining an input value. If no input value can be obtained neither record support nor device support sets UDF false. If device support reads a raw value it returns a value telling record support to perform a conversion. After the record support sets VAL equal to the converted value, it sets UDF false. If device support obtains a converted value that it writes to VAL, it sets UDF false.

For output records either something outside record/device support writes to the VAL field or else VAL is given a value because record support obtains a value via the OMSL field. In either case the code that writes to the VAL field sets UDF false.

Whenever database access writes to the VAL field it sets UDF false.

Routine `recGblSetSevr` is called to raise alarms. It can be called by `iocCore`, record support, or device support. The code that detects an alarm is responsible for raising the alarm.

11.4.6 Raising Monitors

```

static void monitor(xxxRecord *pxxx)
{
    unsigned short    monitor_mask;
    float             delta;

    monitor_mask = recGblResetAlarms (pxxx);
    /* check for value change */
    delta = pxxx->mlst - pxxx->val;

```

```

    if(delta<0.0) delta = -delta;
    if (delta > pxxx->mdel) {
        /* post events for value change */
        monitor_mask |= DBE_VALUE;
        /* update last value monitored */
        pxxx->mlst = pxxx->val;
    }
    /* check for archive change */
    delta = pxxx->alst - pxxx->val;
    if(delta<0.0) delta = 0.0;
    if (delta > pxxx->adel) {
        /* post events on value field for archive change */
        monitor_mask |= DBE_LOG;
        /* update last archive value monitored */
        pxxx->alst = pxxx->val;
    }
    /* send out monitors connected to the value field */
    if (monitor_mask){
        db_post_events (pxxx, &pxxx->val, monitor_mask);
    }
    return;
}

```

All record types should call `recGblResetAlarms` as shown. Note that `nsta` and `nsev` will have the value 0 after this routine completes. This is necessary to ensure that alarm checking starts fresh after processing completes. The code also takes care of raising alarm monitors when a record changes from an alarm state to the no alarm state. It is essential that record support routines follow the above model or else alarm processing will not follow the rules.

Analog type records should also provide monitor and archive hysteresis fields as shown by this example.

`db_post_events` results in channel access issuing monitors for clients attached to the record and field. The call is

```

int db_post_events(void *precord, void *pfield,
    unsigned int monitor_mask)

```

where:

`precord` - The address of the record

`pfield` - The address of the field

`monitor_mask` - A bit mask that can be any combinations of the following:

DBE_ALARM - A change of alarm state has occurred. This is set by `recGblResetAlarms`.

DBE_LOG - Archive change of state.

DBE_VAL - Value change of state

IMPORTANT: The record support module is responsible for calling `db_post_event` for any fields that change as a result of record processing. Also it should NOT call `db_post_event` for fields that do not change.

11.5 Record Support Routines

This section describes the routines defined in the RSET. Any routine that does not apply to a specific record type must be declared `NULL`.

11.5.1 Generate Report of Each Field in Record

```
report(void *precord); /* addr of record*/
```

This routine is not used by most record types. Any action is record type specific.

11.5.2 Initialize Record Processing

```
initialize(void);
```

This routine is called once at IOC initialization time. Any action is record type specific. Most record types do not need this routine.

11.5.3 Initialize Specific Record

```
init_record(
    void *precord, /* addr of record*/
    int    pass);
```

`iocInit` calls this routine twice (`pass=0` and `pass=1`) for each database record of the type handled by this routine. It must perform the following functions:

- Check and/or issue initialization calls for the associated device support routines.
- Perform any record type specific initialization.
- During the first pass it can only perform initializations that affect the record referenced by `precord`.
- During the second pass it can perform initializations that affect other records.

11.5.4 Process Record

```
process(void *precord); /* addr of record*/
```

This routine must follow the guidelines specified previously.

11.5.5 Special Processing

```
special(
    struct dbAddr *paddr,
    int    after); /* (FALSE, TRUE) => (Before, After) Processing*/
```

This routine implements the record type specific special processing for the field referred to by `dbAddr`. Note that it is called twice. Once before any changes are made to the associated field and once after. File `special.h` defines special types. This routine is only called for user special fields, i.e. fields with `SPC_XXX >= 100`. A field is declared special in the ASCII record definition file. New values should not be added to `special.h`, instead use `SPC_MOD`.

The database access routine, `dbGetFieldIndex` can be used to determine which field is being modified.

11.5.6 Get Value

This routine is no longer used. It should be left as a NULL procedure in the record support entry table.

11.5.7 Convert dbAddr Definitions

```
cvt_dbaddr(struct dbAddr *paddr);
```

This routine is called by `dbNameToAddr` if the field has `special` set equal to `SPC_DBADDR`. A typical use is when a field refers to an array. This routine can change any combination of the `dbAddr` fields: `no_elements`, `field_type`, `field_size`, `special`, `pfield`, and `dbr_type`. For example if the `VAL` field of a waveform record is passed to `dbNameToAddr`, `cvt_dbaddr` would change `dbAddr` so that it refers to the actual array rather than `VAL`.

The database access routine, `dbGetFieldIndex` can be used to determine which field is being modified.

NOTES:

- Channel access calls `db_name_to_addr`, which is part of old database access. `db_name_to_addr` calls `dbNameToAddr`. This is done when a client connects to the record.
- `no_elements` must be set to the maximum number of elements that will ever be stored in the array.

11.5.8 Get Array Information

```
get_array_info(
    struct dbAddr *paddr,
    long *no_elements,
    long *offset);
```

This routine returns the current number of elements and the offset of the first value of the specified array. The offset field is meaningful if the array is actually a circular buffer.

The database access routine, `dbGetFieldIndex` can be used to determine which field is being modified. It is permissible for `get_array_info` to change `pfield`. This feature can be used to implement double buffering.

When an array field is being written `get_array_info` is called before the field values are changed.

11.5.9 Put Array Information

```
put_array_info(
    struct dbAddr *paddr,
    long nNew);
```

This routine is called after new values have been placed in the specified array.

The database access routine, `dbGetFieldIndex` can be used to determine which field is being modified.

11.5.10 Get Units

```
get_units(
    struct dbAddr *paddr,
    char *punits);
```

This routine sets units equal to the engineering units for the field.

The database access routine, `dbGetFieldIndex` can be used to determine which field is being modified.

11.5.11 Get Precision

```
get_precision(
    struct dbAddr *paddr,
    long *precision);
```

This routine gets the precision, i.e. number of decimal places, which should be used to convert the field value to an ASCII string. `recGblGetPrec` should be called for fields not directly related to the value field.

The database access routine, `dbGetFieldIndex` can be used to determine which field is being modified.

11.5.12 Get Enumerated String

```
get_enum_str(
    struct dbAddr *paddr,
    char *p);
```

This routine sets `*p` equal to the ASCII string for the field value. The field must have type `DBF_ENUM`.

Look at the code for the `bi` or `mbbi` records for examples.

The database access routine, `dbGetFieldIndex` can be used to determine which field is being modified.

11.5.13 Get Strings for Enumerated Field

```
get_enum_strs(
    struct dbAddr *paddr,
    struct dbr_enumStrs *p);
```

This routine gives values to all fields of structure `dbr_enumStrs`.

Look at the code for the `bi` or `mbbi` records for examples.

The database access routine, `dbGetFieldIndex` can be used to determine which field is being modified.

11.5.14 Put Enumerated String

```
put_enum_str(
    struct dbAddr *paddr,
    char *p);
```

Given an ASCII string, this routine updates the database field. It compares the string with the string values associated with each enumerated value and if it finds a match sets the database field equal to the index of the string which matched.

Look at the code for the `bi` or `mbbi` records for examples.

The database access routine, `dbGetFieldIndex` can be used to determine which field is being modified.

11.5.15 Get Graphic Double Information

```
get_graphic_double(
    struct dbAddr *paddr,
    struct dbr_grDouble *p); /* addr of return info*/
```

This routine fills in the graphics related fields of structure `dbr_grDouble`. `recGblGetGraphicDouble` should be called for fields not directly related to the value field.

The database access routine, `dbGetFieldIndex` can be used to determine which field is being modified.

11.5.16 Get Control Double Information

```
get_control_double(
    struct dbAddr *paddr,
    struct dbr_ctrlDouble *p); /* addr of return info*/
```

This routine gives values to all fields of structure `dbr_ctrlDouble`. `recGblGetControlDouble` should be called for fields not directly related to the value field.

The database access routine, `dbGetFieldIndex` can be used to determine which field is being modified.

11.5.17 Get Alarm Double Information

```
get_alarm_double(
    struct dbAddr *paddr,
    struct dbr_alDouble *p); /* addr of return info*/
```

This routine gives values to all fields of structure `dbr_alDouble`.

The database access routine, `dbGetFieldIndex` can be used to determine which field is being modified.

11.6 Global Record Support Routines

A number of global record support routines are available. These routines are intended for use by the record specific processing routines but can be called by any routine that wishes to use their services.

The name of each of these routines begins with “`recGbl`”.

11.6.1 Alarm Status and Severity

Alarms may be raised in many different places during the course of record processing. The algorithm is to maximize the alarm severity, i.e. the highest severity outstanding alarm is raised. If more than one alarm of the same severity is found then the first one is reported. This means that whenever a code fragment wants to raise an alarm, it does so only if the alarm severity it will declare is greater than that already existing. Four fields (in database common) are used to implement alarms: `sevr`, `stat`, `nsev`, and `nsta`. The first two are the status and severity after the record is completely processed. The last two fields (`nsta` and `nsev`) are the status and severity values to set during record processing. Two routines are used for handling alarms. Whenever a routine wants to raise an alarm it calls `recGblSetSevr`. This routine will only change `nsta` and `nsev` if it will result in the alarm severity being increased. At the end of processing, the record support module must call `recGblResetAlarms`. This routine sets `stat = nsta`, `sevr = nsev`, `nsta = 0`, and `nsev = 0`. If `stat` or `sevr` has changed value since the last call it calls `db_post_event` for `stat` and `sevr` and returns a value of `DBE_ALARM`. If no change occurred it returns 0. Thus after calling `recGblResetAlarms` everything is ready for raising alarms the next time the record is processed. The example record support module presented above shows how these macros are used.

```
recGblSetSevr(
    void *precord,
    short nsta,
    short nsevr);
```

Returns TRUE if it changed `nsta` and/or `nsev`, FALSE if it did not change them.

```
unsigned short recGblResetAlarms(void *precord);
```

Returns: Initial value for `monitor_mask`

11.6.2 Alarm Acknowledgment

Database common contains two additional alarm related fields: `acks` (Highest severity unacknowledged alarm) and `ackt` (does transient alarm need to be acknowledged). These field are handled by `iocCore` and `recGblResetAlarms` and are not the responsibility of record support. These fields are intended for use by the alarm handler.

11.6.3 Generate Error: Process Variable Name, Caller, Message

SUGGESTION: use `errlogPrintf` instead of this for new code.

```
recGblDbaddrError(
    long    status,
    struct dbAddr *paddr,
    char    *pcaller_name); /* calling routine name */
```

This routine interfaces with the system wide error handling system to display the following information: Status information, process variable name, calling routine.

11.6.4 Generate Error: Status String, Record Name, Caller

SUGGESTION: use `errlogPrintf` instead of this for new code.

```
recGblRecordError(
    long    status,
    void    *precord, /* addr of record */
    char    *pcaller_name); /* calling routine name */
```

This routine interfaces with the system wide error handling system to display the following information: Status information, record name, calling routine.

11.6.5 Generate Error: Record Name, Caller, Record Support Message

SUGGESTION: use `errlogPrintf` instead of this for new code.

```
recGblRecsupError(
    long    status,
    struct dbAddr *paddr,
    char    *pcaller_name, /* calling routine name */
    char    *psupport_name); /* support routine name*/
```

This routine interfaces with the system wide error handling system to display the following information: Status information, record name, calling routine, record support entry name.

11.6.6 Get Graphics Double

```
recGblGetGraphicDouble(
    struct dbAddr *paddr,
    struct dbr_grDouble *pgd);
```

This routine can be used by the `get_graphic_double` record support routine to obtain graphics values for fields that it doesn't know how to set.

11.6.7 Get Control Double

```
recGblGetControlDouble(
    struct dbAddr *paddr,
    struct dbr_ctrlDouble *pcd);
```

This routine can be used by the `get_control_double` record support routine to obtain control values for fields that it doesn't know how to set.

11.6.8 Get Alarm Double

```
recGblGetAlarmDouble(
    struct dbAddr *paddr,
    struct dbr_alDouble *pcd);
```

This routine can be used by the `get_alarm_double` record support routine to obtain control values for fields that it doesn't know how to set.

11.6.9 Get Precision

```
recGblGetPrec(
    struct dbAddr *paddr,
    long *pprecision);
```

This routine can be used by the `get_precision` record support routine to obtain the precision for fields that it doesn't know how to set the precision.

11.6.10 Get Time Stamp

```
recGblGetTimeStamp(void *precord)
```

This routine gets the current time stamp and puts it in the record. It does the following:

- If TSEL is not a constant link and TSEL refers to the TIME field of a record, the time is obtained from the record reference by TSEL and this put in field TIME. The routine then returns.
- If TSEL is not a constant link `dbGetLink` is called and the value put in field TSE.
- If TSE is equal to `epicsTimeEventDeviceTime (-2)` then nothing is done, i.e. the routine just returns.
- `epicsTimeGetEvent` is called.

11.6.11 Forward link

```
recGblFwdLink(  
    void *precord);
```

This routine can be used by process to request processing of forward links.

11.6.12 Initialize Constant Link

```
int recGblInitConstantLink(  
    struct link *plink,  
    short dbfType,  
    void *pdest);
```

Initialize a constant link. This routine is usually called by `init_record` (or by associated device support) to initialize the field associated with a constant link. It returns (FALSE, TRUE) if it (did not, did) modify the destination.

Chapter 12

Device Support

12.1 Overview

In addition to a record support module, each record type can have an arbitrary number of device support modules. The purpose of device support is to hide hardware specific details from record processing routines. Thus support can be developed for a new device without changing the record support routines.

A device support routine has knowledge of the record definition. It also knows how to talk to the hardware directly or how to call a device driver which interfaces to the hardware. Thus device support routines are the interface between hardware specific fields in a database record and device drivers or the hardware itself.

Release 3.14.8 introduced the concept of extended device support, which provides an optional interface that a device support can implement to obtain notification when a record's address is changed at runtime. This permits records to be reconnected to a different kind of I/O device, or just to a different signal on the same device. Extended device support is described in more detail in Section ?? below.

Database common contains two device related fields:

- `dtyp`: Device Type.
- `dset`: Address of Device Support Entry Table.

The field `DTYPE` contains the index of the menu choice as defined by the device ASCII definitions. `iocInit` uses this field and the device support structures defined in `devSup.h` to initialize the field `DSET`. Thus record support can locate its associated device support via the `DSET` field.

Device support modules can be divided into two basic classes: synchronous and asynchronous. Synchronous device support is used for hardware that can be accessed without delays for I/O. Many register based devices are synchronous devices. Other devices, for example all GPIB devices, can only be accessed via I/O requests that may take large amounts of time to complete. Such devices must have associated asynchronous device support. Asynchronous device support makes it more difficult to create databases that have linked records.

If a device can be accessed with a delay of less than a few microseconds then synchronous device support is appropriate. If a device causes delays of greater than 100 microseconds then asynchronous device support is appropriate. If the delay is between these values your guess about what to do is as good as mine. Perhaps you should ask the hardware designer why such a device was created.

If a device takes a long time to accept requests there is another option than asynchronous device support. A driver can be created that periodically polls all its attached input devices. The device support just returns the latest polled value. For outputs, device support just notifies the driver that a new value must be written. the driver, during one of its polling phases, writes the new value. The EPICS Allen Bradley device/driver support is a good example.

12.2 Example Synchronous Device Support Module

```

/* Create the dset for devAiSoft */
long init_record();
long read_ai();
struct {
    long    number;
    DEVSUPFUN    report;
    DEVSUPFUN    init;
    DEVSUPFUN    init_record;
    DEVSUPFUN    get_ioint_info;
    DEVSUPFUN    read_ai;
    DEVSUPFUN    special_linconv;
}devAiSoft={
    6,
    NULL,
    NULL,
    init_record,
    NULL,
    read_ai,
    NULL
};
epicsExportAddress(dset,devAiSoft);

static long init_record(void *precord)
{
    aiRecord *pai = (aiRecord *)precord;
    long status;

    /* ai.inp must be a CONSTANT, PV_LINK, DB_LINK or CA_LINK*/
    switch (pai->inp.type) {
        case (CONSTANT) :
            if(recGblInitConstantLink(&pai->inp,DBF_DOUBLE,&pai->val))
                pai->udf = FALSE;
            break;

        case (PV_LINK) :
        case (DB_LINK) :
        case (CA_LINK) :
            break;
        default :
            recGblRecordError(S_db_badField, (void *)pai,
                "devAiSoft (init_record) Illegal INP field");
            return(S_db_badField);
    }
    /* Make sure record processing routine does not perform any conversion*/
    pai->linr=0;
    return(0);
}

static long read_ai(void *precord)
{
    aiRecord*pai =(aiRecord *)precord;
    long status;

```

```

    status=dbGetGetLink (& (pai->inp.value.db_link),
    (void *)pai,DBR_DOUBLE, & (pai->val), 0, 1);
    if (pai->inp.type!=CONSTANT && RTN_SUCCESS(status)) pai->udf = FALSE;
    return(2); /*don't convert*/
}

```

The example is `devAiSoft`, which supports soft analog inputs. The `INP` field can be a constant or a database link or a channel access link. Only two routines are provided (the rest are declared `NULL`). The `init_record` routine first checks that the link type is valid. If the link is a constant it initializes `VAL`. If the link is a Process Variable link it calls `dbCaGetLink` to turn it into a Channel Access link. The `read_ai` routine obtains an input value if the link is a database or Channel Access link, otherwise it doesn't have to do anything.

12.3 Example Asynchronous Device Support Module

This example shows how to write an asynchronous device support routine. It does the following sequence of operations:

1. When first called `PACT` is `FALSE`. It arranges for a callback (`myCallback`) routine to be called after a number of seconds specified by the `DISV` field.
2. It prints a message stating that processing has started, sets `PACT` to `TRUE`, and returns. The record processing routine returns without completing processing.
3. When the specified time elapses `myCallback` is called. It calls `dbScanLock` to lock the record, calls `process`, and calls `dbScanUnlock` to unlock the record. It calls the `process` entry of the record support module, which it locates via the `RSET` field in `dbCommon`, directly rather than `dbProcess`. `dbProcess` would not call `process` because `PACT` is `TRUE`.
4. When `process` executes, it again calls `read_ai`. This time `PACT` is `TRUE`.
5. `read_ai` prints a message stating that record processing is complete and returns a status of 2. Normally a value of 0 would be returned. The value 2 tells the record support routine not to attempt any conversions. This is a convention (a bad convention!) used by the analog input record.
6. When `read_ai` returns the record processing routine completes record processing.

At this point the record has been completely processed. The next time `process` is called everything starts all over.

Note that this is somewhat of an artificial example since real code of this form would more likely use the `callbackcallbackRequestProcessCallbackDelayed` function to perform the required processing.

```

static void myCallback(CALLBACK *pcallback)
{
    struct dbCommon *precord;
    struct rset      *prset;

    callbackGetUser (precord, pcallback);
    prset=(struct rset *) (precord->rset);
    dbScanLock (precord);
    (*prset->process) (precord);
    dbScanUnlock (precord);
}

static long init_record(struct aiRecord *pai)
{
    CALLBACK *pcallback;

```

```

switch (pai->inp.type) {
case (CONSTANT) :
    pcallback = (CALLBACK *) (calloc(1, sizeof(CALLBACK)));
    callbackSetCallback(myCallback, pcallback);
    callbackSetUser(pai, pcallback);
    pai->dpvt = (void *)pcallback;
    break;
default :
    recGblRecordError(S_db_badField, (void *)pai,
        "devAiTestAsyn (init_record) Illegal INP field");
    return(S_db_badField);
}
return(0);
}

static long read_ai(struct aiRecord *pai)
{
    CALLBACK *pcallback = (CALLBACK *)pai->dpvt;
    if(pai->pact) {
        pai->val += 0.1; /* Change VAL just to show we've done something. */
        pai->udf = FALSE; /* We modify VAL so we are responsible for UDF too. */
        printf("Completed asynchronous processing: %s\n", pai->name);
        return(2); /* don't convert*/
    }
    printf("Starting asynchronous processing: %s\n", pai->name);
    pai->pact=TRUE;
    callbackRequestDelayed(pcallback, pai->disv);
    return(0);
}

/* Create the dset for devAiTestAsyn */
struct {
    long          number;
    DEVSUPFUN     report;
    DEVSUPFUN     init;
    DEVSUPFUN     init_record;
    DEVSUPFUN     get_ioint_info;
    DEVSUPFUN     read_ai;
    DEVSUPFUN     special_linconv;
}devAiTestAsyn={
    6,
    NULL,
    NULL,
    init_record,
    NULL,
    read_ai,
    NULL
};
epicsExportAddress(dset, devAiTestAsyn);

```

12.4 Device Support Routines

This section describes the routines defined in the DSET. Any routine that does not apply to a specific record type must be declared `NULL`.

12.4.1 Generate Device Report

```
long report(
    int    interest);
```

This routine is responsible for reporting all I/O cards it has found. If `interest` is (0,1) then generate a (short, long) report. If a device support module is using a driver, it normally does not have to implement this routine because the driver generates the report.

12.4.2 Initialize Device Processing

```
long init(
    int    after);
```

This routine is called twice at IOC initialization time. Any action is device specific. This routine is called twice: once before any database records are initialized and once after all records are initialized but before the scan tasks are started. `after` has the value (0,1) (before, after) record initialization.

12.4.3 Initialize Specific Record

```
long init_record(
    void *precord); /* addr of record*/
```

The record support `init_record` routine calls this routine.

12.4.4 Get I/O Interrupt Information

```
long get_ioint_info(
    int    cmd,
    struct dbCommon *precord,
    IOSCANPVT *ppvt);
```

This is called by the I/O interrupt scan task. If `cmd` is (0,1) then this routine is being called when the associated record is being (placed in, taken out of) an I/O scan list. See the chapter on scanning for details.

It should be noted that a previous type of I/O event scanning is still supported. It is not described in this document because, hopefully, it will go away in the near future. When it calls this routine the arguments have completely different meanings.

12.4.5 Other Device Support Routines

All other device support routines are record type specific.

12.5 Extended Device Support

This section describes the additional behaviour and routines required for a device support layer to support online changes to a record's hardware address.

12.5.1 Rationale

In releases prior to R3.14.8 it is possible to change the value of the INP or OUT field of a record but (unless a soft device support is in use) this generally has no effect on the behaviour of the device support at all. Some device supports have been written that check this hardware address field for changes every time they process, but they are in the minority and in any case they do not provide any means to switch between different device support layers at runtime since no software is present that can lookup a new value for the DSET field after `iocInit`.

The extended device interface has been carefully designed to retain maximal backwards compatibility with existing device and record support layers, and as a result it cannot just introduce new routines into the DSET:

- Different record types have different numbers of DSET routines
- Every device support layer defines its own DSET structure layout
- Some device support layers add their own routines to the DSET (GPIB, BitBus)

Since both basic and extended device support layers have to co-exist within the same IOC, some rules are enforced concerning whether the device address of a particular record is allowed to be changed:

1. Records that were connected at `iocInit` to a device support layer that does not implement the extended interface are never allowed to have address fields changed at runtime.
2. Extended device support layers are not required to implement both the `add_record` and `del_record` routines, thus some devices may only allow one-way changes.
3. The old device support layer is informed and allowed to refuse an address change before the field change is made (it does not get to see the new address).
4. The new device support layer is informed after the field change has been made, and it may refuse to accept the record. In this case the record will be set as permanently busy (`PACT=true`) until an address is accepted.
5. Record support layers can also get notified about this process by making their address field special, in which case the record type's special routine can refuse to accept the new address before it is presented to the device support layer. Special cannot prevent the old device support from being disconnected however.

If an address change is refused, the change to the INP or OUT field will cause an error or exception to be passed to the software performing the change. If this was a Channel Access client the result is to generate an exception callback.

To switch to a different device support layer, it is necessary to change the DTYP field before the INP or OUT field. The change to the DTYP field has no effect until the latter field change takes place.

If a record is set to I/O Interrupt scan but the new layer does not support this, the scan will be changed to Passive.

12.5.2 Initialization/Registration

Device support that implements the extended behaviour must provide an `init` routine in the Device Support Entry Table (see Section ??). In the first call to this routine (pass 0) it registers the address of its Device Support eXtension Table (DSXT) in a call to `devExtend`.

The only exception to this registration requirement is when the device support uses a link type of CONSTANT. In this circumstance the system will automatically register an empty DSXT for that particular support layer (both the `add_record` and `del_record` routines pointed to by this DSXT do nothing and return zero). This exception

allows existing soft channel device support layers to continue to work without requiring any modification, since the iocCore software already takes care of changes to PV_LINK addresses.

The following is an example of a DSXT and the initialization routine that registers it:

```
static struct dsxt myDsxt = {
    add_record, del_record
};

static long init(int pass) {
    if (pass==0) devExtend(&myDsxt);
    return 0;
}
```

A call to devExtend can only be made during the first pass of the device support initialization process, and registers the DSXT for that device support layer; if called at any other time it will log an error message and immediately return.

12.5.3 Device Support eXtension Table

The full definition of struct dsxt is found in devSup.h and currently looks like this:

```
typedef struct dsxt {
    long (*add_record)(struct dbCommon *precord);
    long (*del_record)(struct dbCommon *precord);
} dsxt;
```

There may be future additions to this table to support additional functionality; such extensions may only be made by changing the devSup.h header file and rebuilding EPICS Base and all support modules, thus neither record types nor device support are permitted to make any private use of this table.

The two function pointers are the means by which the extended device support is notified about the record instances it is being given or that are being moved away from its control. In both cases the only parameter is a pointer to the record concerned, which the code will have to cast to the appropriate pointer for the record type. The return value from the routines should be zero for success, or an EPICS error status code.

12.5.4 Add Record Routine

```
long add_record(
    struct dbCommon *precord);
```

This function is called to offer a new record to the device support. It is also called during iocInit, in between the pass 0 and pass 1 calls to the regular device support init_record routine (described in Section ?? above). When converting an existing device support layer, this routine will usually be very similar to the old init_record routine, although in some cases it may be necessary to do a little more work depending on the particular record type involved. The extra code required in these cases can generally be copied straight from the record type implementation itself. This is necessary because the record type has no knowledge of the address change that is taking place, so the device support must perform any bitmask generation and/or readback value conversions itself. This document does not attempt to describe all the necessary processing for the various different standard record types, although the following (incomplete) list is presented as an aid to device support authors:

- mbbi/mbbo record types: Set SHFT, convert NOBT and SHFT into MASK
- bi/bo record types: Set SHFT, convert SHFT to MASK
- analog record types: Calculate ESLO and EOFF

- Output record types: Possibly read the current value from hardware and back-convert to VAL, or send the current record output value to the hardware. *This behaviour is not required or defined, and it's not obvious what should be done. There may be complications here with ao records using OROC and/or OIF=Incremental; solutions to this issue have yet to be considered by the community.*

If the `add_record` routine discovers any errors, say in the link address, it should return a non-zero error status value to reject the record. This will cause the record's PACT field to be set, preventing any further processing of this record until some other address change to it gets accepted.

12.5.5 Delete Record Routine

```
long del_record(
    struct dbCommon *precord);
```

This function is called to notify the device support of a request to change the hardware address of a record, and allow the device support to free up any resources it may have dedicated to this particular record.

Before this routine is called, the record will have had its SCAN field changed to Passive if it had been set to I/O Interrupt. This ensures that the device support's `get_ioint_info` routine is never called after the the call to `del_record` has returned successfully, although it may also lead to the possibility of missed interrupts if the address change is rejected by the `del_record` routine.

If the device support is unable to disconnect from the hardware for some reason, this routine should return a non-zero error status value, which will prevent the hardware address from being changed. In this event the SCAN field will be restored if it was originally set to I/O Interrupt.

After a successful call to `del_record`, the record's DPVT field is set to NULL and PACT is cleared, ready for use by the new device support.

12.5.6 Init Record Routine

The `init_record` routine from the DSET (see Section ??) is called by the record type, and must still be provided since the record type's per-record initialization is run some time *after* the initial call to the DSXT's `add_record` routine. Most record types perform some initialization of record fields at this point, and an extended device support layer may have to fix anything that the record overwrites. The following (incomplete) list is presented as an aid to device support authors:

- mbbi/mbbo record types: Calculate MASK from SHFT
- analog record types: Calculate ESLO and EOFF
- Output record types: Perform readback of the initial raw value from the hardware.

Chapter 13

Driver Support

13.1 Overview

It is not necessary to create a driver support module in order to interface EPICS to hardware. For simple hardware device support is sufficient. At the present time most hardware support has both. The reason for this is historical. Before EPICS there was GTACS. During the change from GTACS to EPICS, record support was changed drastically. In order to preserve all existing hardware support the GTACS drivers were used without change. The device support layer was created just to shield the existing drivers from the record support changes.

Since EPICS now has both device and driver support the question arises: When do I need driver support and when don't I? Lets give a few reasons why drivers should be created.

- The hardware is actually a subnet, e.g. GPIB. In this case a driver should be provided for accessing the subnet. There is no reason to make the driver aware of EPICS except possibly for issuing error messages.
- The hardware is complicated. In this case supplying driver support helps modularized the software. The Allen Bradley driver, which is also an example of supporting a subnet, is a good example.
- An existing driver, maintained by others, is available. I don't know of any examples.
- The driver should be general purpose, i.e. not tied to EPICS. The CAMAC driver is a good example. It is used by other systems, such as CODA. This is perhaps the most important reason for driver support.
- For common devices, e.g. GPIB, CAN, CAMAC, etc. a generic driver layer should be created. This generic layer should be independent of EPICS and independent of low level interfaces. It should also define an interface for low level drivers. This allows low level interfaces to be replaced without impacting IOC records, record support, or device support.

The only thing needed to interface a driver to EPICS is to provide a driver support module, which can be layered on top of an existing driver, and provide a database definition for the driver. The driver support module is described in the next section. The database definition is described in chapter "Database Definition".

13.2 Device Drivers

Device drivers are modules that interface directly with the hardware. They are provided to isolate device support routines from details of how to interface to the hardware. Device drivers have no knowledge of the internals of database records. Thus there is no necessary correspondence between record types and device drivers. For example the Allen Bradley driver provides support for many different types of signals including analog inputs, analog outputs, binary inputs, and binary outputs.

In general only device support routines know how to call device drivers. Since device support varies widely from device to device, the set of routines provided by a device driver is almost completely driver dependent. The only requirement is that routines `report` and `init` must be provided. Device support routines must, of course, know what routines are provided by a driver.

File `drvSup.h` describes the format of a driver support entry table. The driver support module must supply a driver entry table. An example definition is:

```
LOCAL long report();
LOCAL long init();
struct {
    long    number;
    DRVSUPFUN  report;
    DRVSUPFUN  init;
} drvAb={
    2,
    report,
    init
};
epicsExportAddress(drvet, drvGpib);
```

The above example is for the Allen Bradley driver. It has an associated ascii definition of:

```
driver(drvGpib)
```

Thus it is seen that the driver support module should supply two EPICS callable routines: `int` and `report`.

13.2.0.1 init

This routine, which has no arguments, is called by `iocInit`. The driver is expected to look for and initialize the hardware it supports. As an example the `init` routine for Allen Bradley is:

```
LOCAL long init()
{
    return(ab_driver_init());
}
```

13.2.0.2 report

The `report` routine is called by the `dbior` IOC test routine. It is responsible for producing a report describing the hardware it found at `init` time. It is passed one argument, `level`, which is a hint about how much information to display. An example, taken from Allen Bradley, is:

```
LOCAL long report(int level)
{
    return(ab_io_report(level));
}
```

Guidelines for `level` are as follows:

Level=0 Display a one line summary for each device

Level=1 Display more information

Level=2 Display a lot of information. It is even permissible to prompt for what is wanted.

13.2.0.3 Hardware Configuration

Hardware configuration includes the following:

- VME/VXI address space
- VME Interrupt Vectors and levels
- Device Specific Information

The information contained in hardware links supplies some but not all configuration information. In particular it does not define the VME/VXI addresses and interrupt vectors. This additional information is what is meant by hardware configuration in this chapter.

The problem of defining hardware configuration information is an unsolved problem for EPICS. At one time configuration information was defined in `module_types.h`. Many existing device/driver support modules still uses this method. It should NOT be used for any new support for the following reasons:

- There is no way to manage this file for the entire EPICS community.
- It does not allow arbitrary configuration information.
- It is hard for users to determine what the configuration information is.

The fact that it is now easy to include ASCII definitions for only the device/driver support used in each IOC makes the configuration problem much more manageable than previously. Previously if you wanted to support a new VME modules it was necessary to pick addresses that nothing in `module_types.h` was using. Now you only have to check modules you are actually using.

Since there are no EPICS defined rules for hardware configuration, the following minimal guidelines should be used:

- Never use `#define` to specify things like VME addresses. Instead use variables and assign default values. Allow the default values to be changed before `iocInit` is executed. The best way is to supply a global routine that can be invoked from the IOC startup file. Note that all arguments to such routines should be one of the following:

```
int
char *
double
```

- Call the routines described in chapter “Device Support Library” whenever possible.

Chapter 14

Static Database Access

14.1 Overview

An IOC database is created on a Unix system via a Database Configuration Tool and stored in a Unix file. EPICS provides two sets of database access routines: Static Database Access and Runtime Database Access. Static database access can be used on Unix or IOC database files. Runtime database requires an initialized IOC databases. Static database access is described in this chapter and runtime database access in the next chapter.

Static database access provides a simplified interface to a database, i.e. much of the complexity is hidden. `DBF_MENU` and `DBF_DEVICE` fields are accessed via a common type called `DCT_MENU`. A set of routines are provided to simplify access to link fields. All fields can be accessed as character strings. This interface is called static database access because it can be used to access an uninitialized as well as an initialized database.

Before accessing database records, the files describing menus, record types, and devices must be read via `dbReadDatabase` or `dbReadDatabaseFP`. These routines, which are also used to load record instances, can be called multiple times.

Database Configuration Tools (DCTs) should manipulate an EPICS database only via the static database access interface. An IOC database is created on a host system via a database configuration tool and stored in a host file with a file extension of “.db”. Three routines (`dbReadDatabase`, `dbReadDatabaseFP` and `dbWriteRecord`) access the host database file. These routines read/write a database file to/from a memory resident EPICS database. All other access routines manipulate the memory resident database.

An include file `dbStaticLib.h` contains all the definitions needed to use the static database access library. Two structures (`DBBASE` and `DBENTRY`) are used to access a database. The fields in these structures should not be accessed directly. They are used by the static database access library to keep state information for the caller.

14.2 Definitions

14.2.1 DBBASE

Multiple memory resident databases can be accessed simultaneously. The user must provide definitions in the form:

```
DBBASE *pdbname;
```

NOTE: On an IOC `pdbname` is a global variable, which is accessible if you include `dbAccess.h`

14.2.2 DBENTRY

A typical declaration for a database entry structure is:

```
DBENTRY *pdbentry;  
pdbentry=dbAllocEntry(pdbbase);
```

Most static access to a database is via a DBENTRY structure. As many DBENTRYS as desired can be allocated.

The user should NEVER access the fields of DBENTRY directly. They are meant to be used by the static database access library.

Most static access routines accept an argument which contains the address of a DBENTRY. Each routine uses this structure to locate the information it needs and gives values to as many fields in this structure as possible. All other fields are set to NULL.

14.2.3 Field Types

Each database field has a type as defined in the next chapter. For static database access a new and simpler set of field types are defined. In addition, at runtime, a database field can be an array. With static database access, however, all fields are scalars. Static database access field types are called DCT field types.

The DCT field types are:

- DCT_STRING: Character string.
- DCT_INTEGER: Integer value
- DCT_REAL: Floating point number
- DCT_MENU: A set of choice strings
- DCT_MENUFORM: A set of choice strings with associated form.
- DCT_INLINK: Input Link
- DCT_OUTLINK: Output Link
- DCT_FWDLINK: Forward Link
- DCT_NOACCESS: A private field for use by record access routines

A DCT_STRING field contains the address of a NULL terminated string. The field types DCT_INTEGER and DCT_REAL are used for numeric fields. A field that has any of these types can be accessed via the dbGetString, dbPutString, dbVerify, and dbGetRange routines.

The field type DCT_MENU has an associated set of strings defining the choices. Routines are available for accessing menu fields. A menu field can also be accessed via the dbGetString, dbPutString, dbVerify, and dbGetRange routines.

The field type DCT_MENUFORM is like DCT_MENU but in addition the field has an associated link field. The information for the link field can be entered via a set of form manipulation fields.

DCT_INLINK (input), DCT_OUTLINK (output), and DCT_FWDLINK (forward) specify that the field is a link, which has an associated set of static access routines described in the next subsection. A field that has any of these types can also be accessed via the dbGetString, dbPutString, dbVerify, and dbGetRange routines.

14.3 Allocating and Freeing DBBASE

14.3.1 dbAllocBase

```
DBBASE *dbAllocBase(void);
```

This routine allocates and initializes a DBBASE structure. It does not return if it is unable to allocate storage.

`dbAllocBase` allocates and initializes a DBBASE structure. Normally an application does not need to call `dbAllocBase` because a call to `dbReadDatabase` or `dbReadDatabaseFP` automatically calls this routine if `pdbname` is null. Thus the user only has to supply code like the following:

```
DBBASE *pdbname=0;
...
status = dbReadDatabase(&pdbname, "sample.db",
    "<path>", "<macro substitutions>");
```

The static database access library allows applications to work with multiple databases, each referenced via a different (DBBASE *) pointer. Such applications may find it necessary to call `dbAllocBase` directly.

`dbAllocBase` does not return if it is unable to allocate storage.

14.3.2 dbFreeBase

```
void dbFreeBase(DBBASE *pdbname);
```

`dbFreeBase` frees the entire database reference by `pdbname` including the DBBASE structure itself.

14.4 DBENTRY Routines

14.4.1 Alloc/Free DBENTRY

```
DBENTRY *dbAllocEntry(DBBASE *pdbname);
void dbFreeEntry(DBENTRY *pdbentry);
```

These routines allocate, initialize, and free DBENTRY structures. The user can allocate and free DBENTRY structures as necessary. Each DBENTRY is, however, tied to a particular database.

`dbAllocEntry` and `dbFreeEntry` act as a pair, i.e. the user calls `dbAllocEntry` to create a new DBENTRY and calls `dbFreeEntry` when done.

14.4.2 dbInitEntry dbFinishEntry

```
void dbInitEntry(DBBASE *pdbname, DBENTRY *pdbentry);
void dbFinishEntry(DBENTRY *pdbentry);
```

The routines `dbInitEntry` and `dbFinishEntry` are provided in case the user wants to allocate a DBENTRY structure on the stack. Note that the caller MUST call `dbFinishEntry` before returning from the routine that calls `dbInitEntry`. An example of how to use these routines is:

```
int xxx(DBBASE *pdbname)
{
    DBENTRY dbentry;
    DBENTRY *pdbentry = &dbentry;
    ...
}
```

```

    dbInitEntry (pdbname, pdbentry);
    ...
    dbFinishEntry (pdbentry);
}

```

14.4.3 dbCopyEntry

dbCopyEntry

Contents

```

DBENTRY *dbCopyEntry (DBENTRY *pdbentry);
void dbCopyEntryContents (DBENTRY *pfrom, DBENTRY *pto);

```

The routine `dbCopyEntry` allocates a new entry, via a call to `dbAllocEntry`, copies the information from the original entry, and returns the result. The caller must free the entry, via `dbFreeEntry` when finished with the `DBENTRY`.

The routine `dbCopyEntryContents` copies the contents of `pfrom` to `pto`. Code should never perform structure copies.

14.5 Read and Write Database

14.5.1 Read Database File

```

long dbReadDatabase (DBBASE **ppdbname, const char *filename,
    char *path, char *substitutions);
long dbReadDatabaseFP (DBBASE **ppdbname, FILE *fp,
    char *path, char *substitutions);
long dbPath (DBBASE *pdbname, const char *path);
long dbAddPath (DBBASE *pdbname, const char *path);

```

`dbReadDatabase` and `dbReadDatabaseFP` both read a file containing database definitions as described in chapter "Database Definitions". If `*ppdbname` is `NULL`, `dbAllocBase` is automatically invoked and the return address assigned to `*pdbname`. The only difference between the two routines is that one accepts a file name and the other a "FILE *". Any combination of these routines can be called multiple times. Each adds definitions with the rules described in chapter "Database Definitions".

The routines `dbPath` and `dbAddPath` specify paths for use by include statements in database definition files. These are not normally called by user code.

14.5.2 Write Database Definitions

```

long dbWriteMenu (DBBASE *pdbname, char *filename, char *menuName);
long dbWriteMenuFP (DBBASE *pdbname, FILE *fp, char *menuName);
long dbWriteRecordType (DBBASE *pdbname, char *filename, char *recordTypeName);
long dbWriteRecordTypeFP (DBBASE *pdbname, FILE *fp, char *recordTypeName);
long dbWriteDevice (DBBASE *pdbname, char *filename);
long dbWriteDeviceFP (DBBASE *pdbname, FILE *fp);
long dbWriteDriver (DBBASE *pdbname, char *filename);
long dbWriteDriverFP (DBBASE *pdbname, FILE *fp);
long dbWriteRegistrarFP (DBBASE *pdbname, FILE *fp);

```

```

long dbWriteFunctionFP (DBBASE *pdbname, FILE *fp);
long dbWriteVariableFP (DBBASE *pdbname, FILE *fp);
long dbWriteBreaktable (DBBASE *pdbname, const char *filename);
long dbWriteBreaktableFP (DBBASE *pdbname, FILE *fp);

```

Each of these routines writes files in the same format accepted by `dbReadDatabase` and `dbReadDatabaseFP`. Two versions of each type are provided. The only difference is that one accepts a filename and the other a "FILE *". Thus only one of each type has to be described.

`dbWriteMenu` writes the description of the specified menu or, if `menuName` is `NULL`, the descriptions of all menus.

`dbWriteRecordType` writes the description of the specified record type or, if `recordTypeName` is `NULL`, the descriptions of all record types to the named file.

`dbWriteDevice` writes the description of all devices to the named file.

`dbWriteDriver` writes the description of all drivers to the named file.

`dbWriteRegistrarFP` writes the list of all registrars to the given open file (no filename version is provided).

`dbWriteFunctionFP` writes the list of all functions to the given open file (no filename version is provided).

`dbWriteVariableFP` writes the list of all variables to the given open file (no filename version is provided).

`dbWriteBreaktable` writes the definitions of all breakpoint tables to the named file.

14.5.3 Write Record Instances

```

long dbWriteRecord (DBBASE *pdbname, char * file,
char *precordTypeName, int level);
long dbWriteRecordFP (DBBASE *pdbname, FILE *fp,
char *precordTypeName, int level);

```

Each of these routines writes files in the same format accepted by `dbReadDatabase` and `dbReadDatabaseFP`. Two versions of each type are provided. The only difference is that one accepts a filename and the other a "FILE *". Thus only one of each type has to be described.

`dbWriteRecord` writes record instances. If `precordTypeName` is `NULL`, then the record instances for all record types are written, otherwise only the records for the specified type are written. `level` has the following meaning:

- 0 - Write only prompt fields that are different than the default value.
- 1 - Write only the fields which are prompt fields.
- 2 - Write the values of all fields.

14.6 Manipulating Record Types

14.6.1 Get Number of Record Types

```

int dbGetNRecordTypes (DBENTRY *pdbentry);

```

This routine returns the number of record types in the database.

14.6.2 Locate Record Type

```
long dbFindRecordType (DBENTRY *pdbentry,
char *recordTypeName);
long dbFirstRecordType (DBENTRY *pdbentry);
long dbNextRecordType (DBENTRY *pdbentry);
```

`dbFindRecordType` locates a particular record type. `dbFirstRecordType` locates the first, in alphabetical order, record type. Given that `DBENTRY` points to a particular record type, `dbNextRecordType` locates the next record type. Each routine returns 0 for success and a non zero status value for failure. A typical code segment using these routines is:

```
status = dbFirstRecordType (pdbentry);
while (!status) {
    /*Do something*/
    status = dbNextRecordType (pdbentry)
}
```

14.6.3 Get Record Type Name

```
char *dbGetRecordTypeName (DBENTRY *pdbentry);
```

This routine returns the name of the record type that `DBENTRY` currently references. This routine should only be called after a successful call to `dbFindRecordType`, `dbFirstRecordType`, or `dbNextRecordType`. It returns NULL if `DBENTRY` does not point to a record description.

14.7 Manipulating Field Descriptions

The routines described in this section all assume that `DBENTRY` references a record type, i.e. that `dbFindRecordType`, `dbFirstRecordType`, or `dbNextRecordType` has returned success or that a record instance has been successfully located.

14.7.1 Get Number of Fields

```
int dbGetNFields (DBENTRY *pdbentry, int dctonly);
```

Returns the number of fields for the record instance that `DBENTRY` currently references.

14.7.2 Locate Field

```
long dbFirstField (DBENTRY *pdbentry, int dctonly);
long dbNextField (DBENTRY *pdbentry, int dctonly);
```

These routines are used to locate fields. If any of these routines returns success, then `DBENTRY` references that field description.

14.7.3 Get Field Type

```
int dbGetFieldType (DBENTRY *pdbentry);
```

This routine returns the integer value for a DCT field type, see Section 14.2.3 on page 198, for a description of the field types.

14.7.4 Get Field Name

```
char *dbGetFieldName(DBENTRY *pdbentry);
```

This routine returns the name of the field that DBENTRY currently references. It returns NULL if DBENTRY does not point to a field.

14.7.5 Get Default Value

```
char *dbGetDefault(DBENTRY *pdbentry);
```

This routine returns the default value for the field that DBENTRY currently references. It returns NULL if DBENTRY does not point to a field or if the default value is NULL.

14.7.6 Get Field Prompt

```
char *dbGetPrompt(DBENTRY *pdbentry);
int dbGetPromptGroup(DBENTRY *pdbentry);
```

The dbGetPrompt routine returns the character string prompt value, which describes the field. dbGetPromptGroup returns the field group as described in guigroup.h.

14.8 Manipulating Record Attributes

A record attribute is a psuedo-field definition attached to a record type. If a attribute value is assigned to a psuedo field name then all record instances of that record type appear to have that field with the defined value. All attribute fields are DCT_STRING fields.

Two field attributes are automatically created: RTYP and VERS. RTYP is set equal to ,the record type name. VERS is initialized to the value “none specified” but can be changed by record support.

14.8.1 dbPutRecord

Attribute

```
long dbPutRecordAttribute(DBENTRY *pdbentry, const char *name, const char *value)
```

This creates or modifies the attribute *name* on the record type referenced by *pdbentry* to *value*. Attribute names should be valid C identifiers, starting with a letter or underscore followed by any number of alphanumeric or underscore characters.

14.8.2 dbGetRecord

Attribute

```
long dbGetRecordAttribute(DBENTRY *pdbentry, const char *name);
```

Looks up the attribute *name* for the record type referenced by *pdbentry* and sets the the field pointer in *pdbentry* to refer to this string if it exists. The routine dbGetString can be used subsequently to read the current attribute value.

14.9 Manipulating Record Instances

With the exception of `dbFindRecord`, each of the routines described in this section require that `DBENTRY` references a valid record type, i.e. that `dbFindRecordType`, `dbFirstRecordType`, or `dbNextRecordType` has been called and returned success.

14.9.1 Get Number of Records

```
int dbGetNRecords(DBENTRY *pdbentry);
```

Returns the total number of record instances and aliases for the record type that `DBENTRY` currently references.

14.9.2 Get Number of Record Aliases

```
int dbGetNAliases(DBENTRY *pdbentry)
```

Returns the number of record aliases for the record type that `DBENTRY` currently references.

14.9.3 Locate Record

```
long dbFindRecord(DBENTRY *pdbentry, char *precordName);
long dbFirstRecord(DBENTRY *pdbentry);
long dbNextRecord(DBENTRY *pdbentry);
```

These routines are used to locate record instances and aliases. If any of these routines returns success, then `DBENTRY` references a record or a record alias (use `dbIsAlias` to distinguish the two). `dbFindRecord` may be called without `DBENTRY` referencing a valid record type. `dbFirstRecord` only works if `DBENTRY` references a record type. The `dbDumpRecords` example given at the end of this chapter shows how these routines can be used.

`dbFindRecord` also calls `dbFindField` if the record name includes a field name, i.e. it ends in “.XXX”. The routine `dbFoundField` returns (TRUE, FALSE) if the field (was, was not) found. If it was not found, then `dbFindField` must be called before individual fields can be accessed.

14.9.4 Get Record Name

```
char *dbGetRecordName(DBENTRY *pdbentry);
```

This routine only works properly if called after `dbFindRecord`, `dbFirstRecord`, or `dbNextRecord` has returned success. If `DBENTRY` refers to an alias, the name returned is that of the alias, not of the record it refers to.

14.9.5 Distinguishing Record Aliases

```
int dbIsAlias(DBENTRY *pdbentry)
```

This routine only works properly if called after `dbFindRecord`, `dbFirstRecord`, or `dbNextRecord` has returned success. If `DBENTRY` refers to an alias it returns a non-zero value, otherwise it returns zero.

14.9.6 Create/Delete/Free Records and Aliases

```
long dbCreateRecord(DBENTRY *pdbentry, char *precordName);
long dbCreateAlias(DBENTRY *pdbentry, const char *paliasName);
long dbDeleteRecord(DBENTRY *pdbentry);
long dbDeleteAliases(DBENTRY *pdbentry);
long dbFreeRecords(DBBASE *pdbname);
```

`dbCreateRecord`, which assumes that `DBENTRY` references a valid record type, creates a new record instance and initializes it as specified by the record description. If it returns success, then `DBENTRY` references the record just created. `dbCreateAlias` assumes that `DBENTRY` references a particular record instance and creates an alias for that record. If it returns success, then `DBENTRY` references the alias just created. `dbDeleteRecord` deletes either a single alias, or a single record instance and all the aliases that refer to it. `dbDeleteAliases` finds and deletes all aliases that refer to the current record. `dbFreeRecords` deletes all record instances.

14.9.7 Copy Record

```
long dbCopyRecord(DBENTRY *pdbentry, char *newRecordName
                 int overWriteOK)
```

This routine copies the record instance currently referenced by `DBENTRY` (it fails if `DBENTRY` references an alias). Thus it creates a new record with the name `newRecordName` that is of the same type as the original record and copies the original records field values to the new record. If `newRecordName` already exists and `overWriteOK` is true, then the original `newRecordName` is deleted and recreated. If `dbCopyRecord` completes successfully, `DBENTRY` references the new record.

14.9.8 Rename Record

```
long dbRenameRecord(DBENTRY *pdbentry, char *newname)
```

This routine renames the record instance currently referenced by `DBENTRY` (it fails if `DBENTRY` references an alias). If `dbRenameRecord` completes successfully, `DBENTRY` references the renamed record.

14.9.9 Record Visibility

These routines are for use by graphical configuration tools.

```
long dbVisibleRecord(DBENTRY *pdbentry);
long dbInvisibleRecord(DBENTRY *pdbentry);
int dbIsVisibleRecord(DBENTRY *pdbentry);
```

`dbVisibleRecord` sets a record to be visible. `dbInvisibleRecord` sets a record invisible. `dbIsVisibleRecord` returns `TRUE` if a record is visible and `FALSE` otherwise.

14.9.10 Find Field

```
long dbFindField(DBENTRY *pdbentry, char *pfieldName);
int dbFoundField(DBENTRY *pdbentry);
```

Given that a record instance has been located, `dbFindField` finds the specified field. If it returns success, then `DBENTRY` references that field. `dbFoundField` returns (`FALSE`, `TRUE`) if (no field instance is currently available, a field instance is available).

14.9.11 Get/Put Field Values

```
char *dbGetString(DBENTRY *pdbentry);
long dbPutString(DBENTRY *pdbentry, char *pstring);
char *dbVerify(DBENTRY *pdbentry, char *pstring);
char *dbGetRange(DBENTRY *pdbentry);
int dbIsDefaultValue(DBENTRY *pdbentry);
```

These routines are used to get or change field values. They work on all the database field types except `DCT_NOACCESS` but should NOT be used to prompt the user for information for `DCT_MENU`, `DCT_MENUFORM`, or `DCT_LINK_XXX` fields. `dbVerify` returns (`NULL`, a message) if the string is (valid, invalid). Please note that the strings returned are volatile, i.e. the next call to any routine that returns a string will overwrite the value returned by a previous call. Thus it is the caller's responsibility to copy the strings if the value must be kept.

`DCT_MENU`, `DCT_MENUFORM` and `DCT_LINK_XXX` fields can be manipulated via routines described in the following sections. If, however `dbGetString` and `dbPutString` are used, they do work correctly. For these field types `dbGetString` and `dbPutString` are intended to be used only for creating and restoring versions of a database.

14.10 Manipulating Menu Fields

These routines should only be used for `DCT_MENU` and `DCT_MENUFORM` fields. Thus they should only be called if `dbFindField`, `dbFirstField`, or `dbNextField` has returned success and the field type is `DCT_MENU` or `DCT_MENUFORM`.

14.10.1 Get Number of Menu Choices

```
int dbGetNMenuChoices(DBENTRY *pdbentry);
```

This routine returns the number of menu choices for menu.

14.10.2 Get Menu Choice

```
char **dbGetMenuChoices(DBENTRY *pdbentry);
```

This routine returns the address of an array of pointers to strings which contain the menu choices.

14.10.3 Get/Put Menu

```
int dbGetMenuIndex(DBENTRY *pdbentry);
long dbPutMenuIndex(DBENTRY *pdbentry, int index);
char *dbGetMenuStringFromIndex(DBENTRY *pdbentry, int index);
int dbGetMenuIndexFromString(DBENTRY *pdbentry,
char *choice);
```

NOTE: These routines do not work if the current field value contains a macro definition.

`dbGetMenuIndex` returns the index of the menu choice for the current field, i.e. it specifies which choice to which the field is currently set. `dbPutMenuIndex` sets the field to the choice specified by the index.

`dbGetMenuStringFromIndex` returns the string value for a menu index. If the index value is invalid `NULL` is returned. `dbGetMenuIndexFromString` returns the index for the given string. If the string is not a valid choice a -1 is returned.

14.10.4 Locate Menu

```
dbMenu *dbFindMenu(DBBASE *pdbname, char *name);
```

`dbFindMenu` is most useful for runtime use but is a static database access routine. This routine just finds a menu with the given name.

14.11 Manipulating Link Fields

14.11.1 Link Types

Links are the most complicated types of fields. A link can be a constant, reference a field in another record, or can refer to a hardware device. Two additional complications arise for hardware links. The first is that field `DTYP`, which is a menu field, determines if the `INP` or `OUT` field is a device link. The second is that the information that must be specified for a device link is bus dependent. In order to shelter database configuration tools from these complications the following is done for static database access.

- Static database access will treat `DTYP` as a `DCT_MENUFORM` field.
- The information for the link field related to the `DCT_MENUFORM` can be entered via a set of form manipulation routines associated with the `DCT_MENUFORM` field. Thus the link information can be entered via the `DTYP` field rather than the link field.
- The Form routines described in the next section can also be used with any link field.

Each link is one of the following types:

- `DCT_LINK_CONSTANT`: Constant value.
- `DCT_LINK_PV`: A process variable link.
- `DCT_LINK_FORM`: A link that can only be processed via the form routines described in the next chapter.

Database configuration tools can change any link between being a constant and a process variable link. Routines are provided to accomplish these tasks.

The routines `dbGetString`, `dbPutString`, and `dbVerify` can be used for link fields but the form routines can be used to provide a friendlier user interface.

14.11.2 All Link Fields

```
int dbGetNLinks(DBENTRY *pdbentry);
long dbGetLinkField(DBENTRY *pdbentry, int index)
int dbGetLinkType(DBENTRY *pdbentry);
```

These are routines for manipulating `DCT_XXXLINK` fields. `dbGetNLinks` and `dbGetLinkField` are used to walk through all the link fields of a record. `dbGetLinkType` returns one of the values: `DCT_LINK_CONSTANT`, `DCT_LINK_PV`, `DCT_LINK_FORM`, or the value -1 if it is called for an illegal field.

14.11.3 Constant and Process Variable Links

```
long dbCvtLinkToConstant(DBENTRY *pdbentry);
long dbCvtLinkToPvlink(DBENTRY *pdbentry);
```

These routines should be used for modifying `DCT_LINK_CONSTANT` or `DCT_LINK_PV` links. They should not be used for `DCT_LINK_FORM` links, which should be processed via the associated `DCT_MENUFORM` field described above.

14.12 Manipulating MenuForm Fields

These routines are used with a `DCT_MENUFORM` field (a `DTYP` field) to manipulate the associated `DCT_INLINK` or `DCT_OUTLINK` field. They can also be used on any `DCT_INLINK`, `DCT_OUTLINK`, or `DCT_FWDLINK` field.

14.12.1 Alloc/Free Form

```
int dbAllocForm(DBENTRY *pdbentry)
long dbFreeForm(DBENTRY *pdbentry)
```

`dbAllocForm` allocates storage needed to manipulate forms. The return value is the number of elements in the form. If the current field value contains a macro definition, the number of lines returned is 0.

14.12.2 Get/Put Form

```
char **dbGetFormPrompt(DBENTRY *pdbentry)
char **dbGetFormValue(DBENTRY *pdbentry)
long dbPutForm(DBENTRY *pdbentry, char **value)
```

`dbGetFormPrompt` returns a pointer to an array of pointers to character strings specifying the prompt string. `dbGetFormValue` returns the current values. `dbPutForm`, which can use the same array of values returned by `dbGetForm`, sets new values.

14.12.3 Verify Form

```
char **dbVerifyForm(DBENTRY *pdbentry, char **value)
```

`dbVerifyForm` can be called to verify user input. It returns `NULL` if no errors are present. If errors are present, it returns a pointer to an array of character strings containing error messages. Lines in error have a message and correct lines have a `NULL` string.

14.12.4 Get Related Field

```
char *dbGetRelatedField(DBENTRY *pdbentry)
```

This routine returns the field name of the related field for a `DCT_MENUFORM` field. If it is called for any other type of field it returns `NULL`.

14.12.5 Example

The following is code showing use of these routines:

```
char **value;
char **prompt;
char **error;
int n;
```

```

...
n = dbAllocForm(pdbentry);
if(n<=0) {<Error>}
prompt = dbGetFormPrompt(pdbentry);
value = dbGetFormValue(pdbentry);
for(i=0; i<n; i++) {
    printf("%s (%s) : \n",prompt[i],value[i]);
    /*The following accepts input from stdin*/
    scanf("%s",value[i]);
}
error = dbVerifyForm(pdbentry,value);
if(error) {
    for(i=0; i<n; i++) {
        if(error[i]) printf("Error: %s (%s) %s\n", prompt[i],
            value[i],error[i]);
    }
}else {
    dbPutForm(pdbentry,value)
}
dbFreeForm(pdbentry);

```

All value strings are MAX_STRING_SIZE in length.

A set of form calls for a particular DBENTRY, MUST begin with a call to dbAllocForm and end with a call to dbFreeForm. The values returned by dbGetFormPrompt, dbGetFormValue, and dbVerifyForm are valid only between the calls to dbAllocForm and dbFreeForm.

14.13 Manipulating Information Items

Information items are stored as a list attached to each individual record instance. All routines listed in this section require that the DBENTRY argument refer to a valid record instance.

14.13.1 Locate Item

```

long dbFirstInfo(DBENTRY *pdbentry);
long dbNextInfo(DBENTRY *pdbentry);
long dbFindInfo(DBENTRY *pdbentry, const char *name);

```

There are two ways to locate info items, by scanning through the list using first/next, or by asking for the item by name. These routines set pdbentry to refer to the info item and return 0, or return an error code if no info item is found.

14.13.2 Get Item Name

```

const char * dbGetInfoName(DBENTRY *pdbentry);

```

Returns the name of the info item referred to by pdbentry, or a NULL pointer if no item is referred to.

14.13.3 Get/Set Item String Value

```
const char * dbGetInfoString(DBENTRY *pdbentry);
long dbPutInfoString(DBENTRY *pdbentry, const char *string);
```

These routines provide access to the currently selected items' string value. When changing the string value using `dbPutInfoString`, the character string provided will be copied, with additional memory being allocated as necessary. Users are advised not to make continuously repeated calls to `dbPutInfoString` at IOC runtime as this could fragment the free memory heap. The Put routine returns 0 if Ok or an error code; the Get routine returns NULL on error.

14.13.4 Get/Set Item Pointer Value

```
void * dbGetInfoPointer(DBENTRY *pdbentry);
long dbPutInfoPointer(DBENTRY *pdbentry, void *pointer);
```

Each info item includes space to store a single `void*` pointer as well as the value string. Applications using the info item may set this as often as they wish. The Put routine returns 0 if Ok or an error code; the Get routine returns NULL on error.

14.13.5 Create/Delete Item

```
long dbPutInfo(DBENTRY *pdbentry, const char *name, const char *string);
long dbDeleteInfo(DBENTRY *pdbentry);
```

A new info item can be created by calling `dbPutInfo`. If an item by that name already exists its value will be replaced with the new string, otherwise storage is allocated and the name and value strings copied into it. The function returns 0 on success, or an error code.

When calling `dbDeleteInfo`, the `pdbentry` must refer to the item to be removed (using `dbFindFirst/dbFindNext` or `dbFindInfo`). The function returns 0 on success, or an error code.

14.13.6 Convenience Routine

```
const char * dbGetInfo(DBENTRY *pdbentry, const char *name);
```

It is common to want to look up the value of a named info item in one call, and `dbGetInfo` is provided for this purpose. It returns a NULL pointer if no info item exists with the given name.

14.14 Find Breakpoint Table

```
brkTable *dbFindBrkTable(DBBASE *pdbname, char *name)
```

This routine returns the address of the specified breakpoint table. It is normally used by the runtime breakpoint conversion routines so will not be discussed further.

14.15 Dump Routines

```
void dbDumpPath(DBBASE *pdbname)
void dbDumpRecord(DBBASE *pdbname, char *precordTypeName, int level);
void dbDumpMenu(DBBASE *pdbname, char *menuName);
```

```

void dbDumpRecordType (DBBASE *pdbbase, char *recordTypeName);
void dbDumpField (DBBASE *pdbbase, char *recordTypeName, char *fname);
void dbDumpDevice (DBBASE *pdbbase, char *recordTypeName);
void dbDumpDriver (DBBASE *pdbbase);
void dbDumpRegistrar (DBBASE *pdbbase);
void dbDumpFunction (DBBASE *pdbbase);
void dbDumpVariable (DBBASE *pdbbase);
void dbDumpBreaktable (DBBASE *pdbbase, char *name);
void dbPvdDump (DBBASE *pdbbase, int verbose);
void dbReportDeviceConfig (DBBASE *pdbbase, FILE *report);

```

These routines are used to dump information about the database. `dbDumpRecord`, `dbDumpMenu`, `dbDumpDriver`, `dbDumpRegistrar` and `dbDumpVariable` just call the corresponding `dbWritexxxFP` routine specifying stdout for the file. `dbDumpRecordType`, `dbDumpField`, and `dbDumpDevice` give internal information useful on an ioc. These commands can be executed via `iocsh`. Just specify `pdbbase` as the first argument.

14.16 Examples

14.16.1 Expand Include

This example is like the `dbExpand` utility, except that it doesn't allow path or macro substitution options, It reads a set of database definition files and writes all definitions to stdout. All include statements appearing in the input files are expanded.

```

/* dbExpand.c */
#include <stdlib.h>
#include <stddef.h>
#include <stdio.h>
#include <epicsPrint.h>
#include <dbStaticLib.h>

DBBASE *pdbbase = NULL;

int main(int argc, char **argv)
{
    long status;
    int i;
    int arg;

    if (argc < 2) {
        printf("usage: expandInclude file1.db file2.db...\n");
        exit 0;
    }

    for (i = 1; i < argc; i++) {
        status = dbReadDatabase(&pdbbase, argv[i], NULL, NULL);
        if (!status) continue;
        fprintf(stderr, "For input file %s", argv[i]);
        errMessage(status, "from dbReadDatabase");
    }
    dbWriteMenuFP (pdbbase, stdout, 0);
    dbWriteRecordTypeFP (pdbbase, stdout, 0);
    dbWriteDeviceFP (pdbbase, stdout);
}

```

```

    dbWriteDriverFP (pdbname.stdout);
    dbWriteRecordFP (pdbname, stdout, 0, 0);
    return (0);
}

```

14.16.2 dbDumpRecords

NOTE: This example is similar but not identical to the actual dbDumpRecords routine.

The following example demonstrates how to use the database access routines. The example shows how to iterate through the record types and instances and display field values.

```

void dbDumpRecords (DBBASE *pdbname)
{
    DBENTRY *pdbentry;
    long status;

    pdbentry = dbAllocEntry (pdbname);
    status = dbFirstRecordType (pdbentry);
    if (status) {
        printf ("No record types\n");
        return;
    }
    while (!status) {
        printf ("Record type: %s\n", dbGetRecordTypeName (pdbentry));
        status = dbFirstRecord (pdbentry);
        if (status)
            printf (" No record instances\n");
        else while (!status) {
            if (dbIsAlias (pdbentry))
                printf (" Alias: %s\n", dbGetRecordName (pdbentry));
            else {
                printf (" Record: %s\n", dbGetRecordName (pdbentry));
                status = dbFirstField (pdbentry, TRUE);
                if (status)
                    printf (" No fields\n");
                else while (!status) {
                    printf (" %s: %s\n", dbGetFieldName (pdbentry),
                        dbGetString (pdbentry));
                    status = dbNextField (pdbentry, TRUE);
                }
            }
            status = dbNextRecord (pdbentry);
        }
        status = dbNextRecordType (pdbentry);
    }
    printf ("End of record types\n");
    dbFreeEntry (pdbentry);
}

```

Chapter 15

Runtime Database Access

15.1 Overview

This chapter describes routines for manipulating and accessing an initialized IOC database.

This chapter is divided into the following sections:

- Database related include files. All of interest are listed and those of general interest are discussed briefly.
- Runtime database access overview.
- Description of each runtime database access routine.
- Runtime modification of link fields.
- Lock Set Routines
- Database to Channel Access Routines

15.2 Database Include Files

There are a number of database related include files. Of particular interest to this chapter are:

- dbDefs.h: Miscellaneous database related definitions
- dbFldTypes.h: Field type definitions
- dbAccess.h: Runtime database access definitions.
- link.h: Definitions for link fields.

15.2.1 dbDefs.h

This file contains a number of database related definitions. The most important are:

- PVNAME_STRINGSZ: The number of characters reserved for the record name, including the terminating zero byte.
- PVNAME_SZ: The maximum number of characters allowed in the record name.
- DB_MAX_CHOICES: The maximum number of choices for a choice field.

Note that `DB_MAX_CHOICES` applies for code using the runtime routines documented in this chapter, but for Channel Access clients the maximum number of choices and their choice string length are different, and are defined in the `db_access.h` file.

15.2.2 `dbFldTypes.h`

This file defines the possible field types. A field's type is perhaps its most important attribute. Changing the possible field types is a fundamental change to the IOC software, because many IOC software components are aware of the field types.

The field types are:

- `DBF_STRING`: 40 character ASCII string
- `DBF_CHAR`: Signed character
- `DBF_UCHAR`: Unsigned character
- `DBF_SHORT`: Short integer
- `DBF_USHORT`: Unsigned short integer
- `DBF_LONG`: Long integer
- `DBF_ULONG`: Unsigned long integer
- `DBF_FLOAT`: Floating point number
- `DBF_DOUBLE`: Double precision float
- `DBF_ENUM`: An enumerated field
- `DBF_MENU`: A menu choice field
- `DBF_DEVICE`: A device choice field
- `DBF_INLINK`: Input Link
- `DBF_OUTLINK`: Output Link
- `DBF_FWDLINK`: Forward Link
- `DBF_NOACCESS`: A private field for use by record access routines

A field of type `DBF_STRING`, ..., `DBF_DOUBLE` can be a scalar or an array. A `DBF_STRING` field contains a NULL terminated ascii string. The field types `DBF_CHAR`, ..., `DBF_DOUBLE` correspond to the standard C data types.

`DBF_ENUM` is used for enumerated items, which is analogous to the C language enumeration. An example of an enum field is field `VAL` of a multi bit binary record.

The field types `DBF_ENUM`, `DBF_MENU`, and `DBF_DEVICE` all have an associated set of ASCII strings defining the choices. For a `DBF_ENUM`, the record support module supplies values and thus are not available for static database access. The database access routines locate the choice strings for the other types.

`DBF_INLINK` and `DBF_OUTLINK` specify link fields. A link field can refer to a signal located in a hardware module, to a field located in a database record in the same IOC, or to a field located in a record in another IOC. A `DBF_FWDLINK` can only refer to a record in the same IOC. Link fields are described in a later chapter.

`DBF_INLINK` (input), `DBF_OUTLINK` (output), and `DBF_FWDLINK` (forward) specify that the field is a link structure as defined in `link.h`. There are three classes of links:

1. Constant - The value associated with the field is a floating point value initialized with a constant value. This is somewhat of a misnomer because constant link fields can be modified via `dbPutField` or `dbPutLink`.

2. Hardware links - The link contains a data structure which describes a signal connected to a particular hardware bus. See `link.h` for a description of the bus types currently supported.
3. Process Variable Links - This is one of three types:
 - (a) `PV_LINK`: The process variable name.
 - (b) `DB_LINK`: A reference to a process variable in the same IOC.
 - (c) `CA_LINK`: A reference to a variable located in another IOC.

When first loaded the field is always creates as a `PV_LINK`. When the IOC is initialized each `PV_LINK` is converted either to a `DB_LINK` or a `CA_LINK`.

`DBF_NOACCESS` fields are for private use by record processing routines.

15.2.3 dbAccess.h

This file is the interface definition for the run time database access library, i.e. for the routines described in this chapter.

An important structure defined in this header file is `DBADDR`

```
typedef struct dbAddr{
    struct dbCommon *precord; /* address of record*/
    void* pfield; /* address of field*/
    void* pfldDes; /* address of struct fldDes*/
    void* asPvt; /* Access Security Private*/
    long no_elements; /* number of elements (arrays)*/
    short field_type; /* type of database field*/
    short field_size; /* size (bytes) of the field*/
    short special; /* special processing*/
    short dbr_field_type; /*optimal database request type*/
}DBADDR;
```

- `precord`: Address of record. Note that its type is a pointer to a structure defining the fields common to all record types. The common fields appear at the beginning of each record. A record support module can cast `precord` to point to the specific record type.
- `pfield`: Address of the field within the record. Note that `pfield` provides direct access to the data value.
- `pfldDes`: This points to a structure containing all details concerning the field. See Chapter “Database Structures” for details.
- `asPvt`: A field used by access security.
- `no_elements`: A string or numeric field can be either a scalar or an array. For scalar fields `no_elements` has the value 1. For array fields it is the maximum number of elements that can be stored in the array.
- `field_type`: Field type.
- `field_size`: Size of one element of the field.
- `special`: Some fields require special processing. This specifies the type. Special processing is described later in this manual.
- `dbr_field_type`: This specifies the optimal database request type for this field, i.e. the request type that will require the least CPU overhead.

NOTE: `pfield`, `no_elements`, `field_type`, `field_size`, `special`, and `dbr_field_type` can all be set by record support (`cvt_dbaddr`). Thus `field_type`, `field_size`, and `special` can differ from that specified by `pfldDes`.

15.2.4 link.h

This header file describes the various types of link fields supported by EPICS.

15.3 Runtime Database Access Overview

With the exception of record and device support, all access to the database is via the channel or database access routines. Even record support routines access other records only via database or channel access. Channel Access, in turn, accesses the database via database access.

Perhaps the easiest way to describe the database access layer is to list and briefly describe the set of routines that constitute database access. This provides a good look at the facilities provided by the database.

Before describing database access, one caution must be mentioned. The only way to communicate with an IOC database from outside the IOC is via Channel Access. In addition, any special purpose software, i.e. any software not described in this document, should communicate with the database via Channel Access, not database access, even if it resides in the same IOC as the database. Since Channel Access provides network independent access to a database, it must ultimately call database access routines. The database access interface was changed in 1991, but Channel Access was never changed to use the new interface. Instead a module was written which translates old style database access calls to new. This interface between the old and new style database access calls is not discussed in this manual.

The database access routines are:

- dbNameToAddr: Locate a database variable.
- dbGetField: Get values associated with a database variable.
- dbGetLink: Get value of field referenced by database link (Macro)
- dbGetLinkValue: Get value of field referenced by database link (Subroutine)
- dbGet: Routine called by dbGetLinkValue and dbGetField
- dbPutField: Change the value of a database variable.
- dbPutLink: Change value referenced by database link (Macro)
- dbPutLinkValue: Change value referenced by database link (Subroutine)
- dbPut: Routine called by dbPutxxx functions.
- putNotifyInit: Initialize "struct putNotify"
- putNotifyCleanup: Cleanup "struct putNotify"
- dbPutNotify: A database put with notification on completion
- dbNotifyCancel: Cancel dbPutNotify
- dbNotifyAdd: Add a new record for to notify set.
- dbNotifyCompletion: Announce that put notify is complete.
- dbBufferSize: Determine number of bytes in request buffer.
- dbValueSize: Number of bytes for a value field.
- dbGetRset: Get pointer to Record Support Entry Table
- dbIsValueField: Is this field the VAL field.
- dbGetFieldIndex: Get field index. The first field in a record has index 0.
- dbGetNelement: Get number of elements in the field

- `dbIsLinkConnected`: Is the link field connected.
- `dbGetPdbAddrFromLink`: Get address of DBADDR.
- `dbGetLinkDBFtype`: Get field type of link.
- `dbGetControlLimits`: Get Control Limits.
- `dbGetGraphicLimits`: Get Graphic Limits.
- `dbGetAlarmLimits`: Get Alarm Limits
- `dbGetPrecision`: Get Precision
- `dbGetUnits`: Get Units
- `dbGetNelements`: Get Number of Elements
- `dbGetSevr`: Get Severity
- `dbGetTimeStamp`: Get Time Stamp
- `dbPutAttribute` Give a value to a record attribute.
- `dbScanPassive`: Process record if it is passive.
- `dbScanLink`: Process record referenced by link if it is passive.
- `dbProcess`: Process Record
- `dbScanFwdLink`: Scan a forward link.

15.3.1 Database Request Types and Options

Before describing database access structures, it is necessary to describe database request types and request options. When `dbPutField` or `dbGetField` are called one of the arguments is a database request type. This argument has one of the following values:

- `DBR_STRING`: Value is a NULL terminated string
- `DBR_CHAR`: Value is a signed char
- `DBR_UCHAR`: Value is an unsigned char
- `DBR_SHORT`: Value is a short integer
- `DBR_USHORT`: Value is an unsigned short integer
- `DBR_LONG`: Value is a long integer
- `DBR_ULONG`: Value is an unsigned long integer
- `DBR_FLOAT`: Value is an IEEE floating point value
- `DBR_DOUBLE`: Value is an IEEE double precision floating point value
- `DBR_ENUM`: Value is a short which is the enum item
- `DBR_PUT_ACKT`: Value is an unsigned short for setting the ACKT.
- `DBR_PUT_ACKS`: Value is an unsigned short for global alarm acknowledgment.

The request types `DBR_STRING`,..., `DBR_DOUBLE` correspond exactly to valid data types for database fields. `DBR_ENUM` corresponds to database fields that represent a set of choices or options. In particular it corresponds to the fields types `DBF_ENUM`, `DBF_DEVICE`, and `DBF_MENU`. The complete set of database field types are defined in `dbFldTypes.h`. `DBR_PUT_ACKT` and `DBR_PUT_ACKS` are used to perform global alarm acknowledgment.

`dbGetField` also accepts argument options which is a mask containing a bit for each additional type of information the caller desires. The complete set of options is:

- `DBR_STATUS`: returns the alarm status and severity
- `DBR_UNITS`: returns a string specifying the engineering units
- `DBR_PRECISION`: returns a long integer specifying floating point precision.
- `DBR_TIME`: returns the time
- `DBR_ENUM_STRS`: returns an array of strings
- `DBR_GR_LONG`: returns graphics info as long values
- `DBR_GR_DOUBLE`: returns graphics info as double values
- `DBR_CTRL_LONG`: returns control info as long values
- `DBR_CTRL_DOUBLE`: returns control info as double values
- `DBR_AL_LONG`: returns alarm info as long values
- `DBR_AL_DOUBLE`: returns alarm info as double values

15.3.2 Options

Example

The file `dbAccess.h` contains macros for using options. A brief example should show how they are used. The following example defines a buffer to accept an array of up to ten float values. In addition it contains fields for options `DBR_STATUS` and `DBR_TIME`.

```
struct buffer {
    DBRstatus
    DBRtime
    float value[10];
} buffer;
```

The associated `dbGetField` call is:

```
long options, number_elements, status;
...
options = DBR_STATUS | DBR_TIME;
number_elements = 10;
status = dbGetField(paddr, DBR_FLOAT, &buffer, &options, &number_elements);
```

Consult `dbAccess.h` for a complete list of macros.

Structure `dbAddr` contains a field `dbr_field_type`. This field is the database request type that most closely matches the database field type. Using this request type will put the smallest load on the IOC.

Channel Access provides routines similar to `dbGetField`, and `dbPutField`. It provides remote access to `dbGetField`, `dbPutField`, and to the database monitors described below.

15.3.3 ACKT and ACKS

The request types `DBR_PUT_ACKT` and `DBR_PUT_ACKS` are used for global alarm acknowledgment. The alarm handler uses these requests. For each of these types the user (normally channel access) passes an unsigned short value. This value represents:

DBR_PUT_ACKT - Do transient alarms have to be acknowledged? 0 means no, 1 means yes

DBR_PUT_ACKS - The highest alarm severity to acknowledge. If the current alarm severity is less than or equal to this value the alarm is acknowledged.

15.4 Database Access Routines

15.4.1 dbNameToAddr

Locate a process variable, format:

```
long dbNameToAddr(
    char *pname, /*ptr to process variable name */
    struct dbAddr *paddr);
```

The most important goal of database access can be stated simply: Provide quick access to database records and fields within records. The basic rules are:

1. Call `dbNameToAddr` once and only once for each field to be accessed.
2. Read field values via `dbGetField` and write values via `dbPutField`.

The routines described in this subsection are used by channel access, sequence programs, etc. Record processing routines, however, use the routines described in the next section rather than `dbGetField` and `dbPutField`.

Given a process variable name, this routine locates the process variable and fills in the fields of structure `dbAddr`. The format for an IOC process variable name is one of:

```
<record_name>
<record_name>.
<record_name>.<field_name>
<record_name>.<field_name><modifier>
<record_name>.<modifier>
```

For example the value field of a record with record named `sample_name` is:

“`sample_name.VAL`”.

The *record_name* is case sensitive. The *field_names* available depend on the record type of the record and usually consist of all upper-case letters. If omitted the field name `VAL` is used if it exists. Currently the only *modifier* supported is a single dollar sign `$` and is only valid on fields which are strings or links. Additional modifiers may be added in future releases.

`dbNameToAddr` locates a record via a process variable directory (PVD). It fills in a structure (`dbAddr`) describing the field. `dbAddr` contains the address of the record and also the field. Thus other routines can locate the record and field without a search. Although the PVD allows the record to be located via a hash algorithm and the field within a record via a binary search, it still takes about 80 microseconds (25MHz 68040) to locate a process variable. Once located the `dbAddr` structure allows the process variable to be accessed directly.

15.4.2 Get Routines

15.4.2.1 dbGetField

Get values associated with a process variable, format:

```
long dbGetField(
    struct dbAddr *paddr,
    short dbrType, /* DBR_xxx */
    void *pbuffer, /* ptr to returned data */
```

```

    long *options,    /* ptr to options */
    long *nRequest,  /* ptr to number of elements */
    void *pfl);     /* used by monitor routines */

```

This routine locks, calls dbGet, and unlocks.

15.4.2.2 dbGetLink and dbGetLinkValue

Get value from the field referenced by a database link, format:

```

long dbGetLink(
    struct db_link *plink, /* ptr to link field */
    short dbrType,       /* DBR_XXX */
    void *pbuffer,       /* ptr to returned data */
    long *options,       /* ptr to options */
    long *nRequest);    /* ptr to number of elements desired */

```

NOTES:

- options can be NULL if no options are desired.
- nRequest can be NULL for a scalar.

dbGetLink is implemented as a macro that calls dbGetLinkValue and can reference its arguments more than once. The macro skips the call for constant links. User code should never call dbGetLinkValue.

This routine is called by database access itself and by record support and/or device support routines in order to get values for input links. The value can be obtained directly from other records or via a channel access client. This routine honors the link options (process and maximize severity). In addition it has code that optimizes the case of no options and scalar.

15.4.2.3 dbGet

Get values associated with a process variable, format:

```

long dbGet(
    struct dbAddr *paddr,
    short dbrType, /* DBR_XXX */
    void *pbuffer, /* ptr to returned data */
    long *options, /* ptr to options */
    long *nRequest, /* ptr to number of elements */
    void *pfl);    /* used by monitor routines */

```

This routine retrieves the data referenced by paddr and converts it to the format specified by dbrType.

- options is a read/write field. Upon entry to dbGet, options specifies the desired options. When dbGetField returns, options specifies the options actually honored. If an option is not honored, the corresponding fields in buffer are filled with zeros.
- nRequest is also a read/write field. Upon entry to dbGet it specifies the maximum number of data elements the caller is willing to receive. When dbGet returns it has been set to the actual number of elements returned. It is permissible to request zero elements. This is useful when only option data is desired.
- The pfl argument is for use by the Channel Access monitor routines. All other users must give pfl=NULL.

dbGet calls one of a number of conversion routines in order to convert data from the DBF types to the DBR types. It calls record support routines for special cases such as arrays. For example, if the number of field elements is greater than 1 and record support routine get_array_info exists, then it is called. It returns two values: the current number of valid field elements and an offset. The number of valid elements may not match dbAddr.no_elements, which

is really the maximum number of elements allowed. The offset is for use by records which implement circular buffers, and provides the offset to the current beginning of the array data.

15.4.3 Put Routines

15.4.3.1 dbPutField

Change the value of a process variable, format:

```
long dbPutField(
    struct dbAddr *paddr,
    short dbrType,    /* DBR_xxx */
    void *pbuffer,    /* ptr to data */
    long nRequest); /* number of elements to write */
```

This routine is responsible for accepting data in one of the DBR_xxx formats, converting it as necessary, and modifying the database. Similar to dbGetField, this routine calls one of a number of conversion routines to do the actual conversion and relies on record support routines to handle arrays and other special cases.

It should be noted that routine dbPut does most of the work. The actual algorithm for dbPutField is:

1. If the DISP field is TRUE then, unless it is the DISP field itself which is being modified, the field is not written.
2. The record is locked.
3. dbPut is called.
4. If the dbPut is successful then:

If this is the PROC field or if both of the following are TRUE: 1) the field is a process passive field, 2) the record is passive.

- (a) If the record is already active, ask for the record to be reprocessed when it completes.
- (b) Call dbScanPassive after setting putf TRUE to show the process request came from dbPutField.
5. The record is unlocked.

15.4.3.2 dbPutLink and dbPutLinkValue

Change the value referenced by a database link, format:

```
long dbPutLink(
    struct db_link *plink, /* ptr to link field */
    short dbrType,        /* DBR_xxx */
    void *pbuffer,        /* ptr to data to write */
    long nRequest);      /* number of elements to write */
```

dbPutLink is actually a macro that calls dbPutLinkValue and can reference its arguments more than once. The macro skips the call for constant links. User code should never call dbPutLinkValue.

This routine is called by database access itself and by record support and/or device support routines in order to put values into other database records via output links.

For Channel Access links it calls dbCaPutLink.

For database links it performs the following functions:

1. Calls dbPut.
2. Implements maximize severity.

3. If the field being referenced is PROC or if both of the following are true: 1) `process_passive` is TRUE and 2) the record is passive then:
 - (a) If the record is already active because of a `dbPutField` request then ask for the record to be reprocessed when it completes.
 - (b) otherwise call `dbScanPassive`.

15.4.3.3 dbPut

Put a value to a database field, format:

```
long dbPut (
    struct dbAddr *paddr,
    short dbrType, /* DBR_*** */
    void *pbuffer, /* addr of data */
    long nRequest); /* number of elements to write */
```

This routine is responsible for accepting data in one of the DBR_*** formats, converting it as necessary, and modifying the database. Similar to `dbGet`, this routine calls one of a number of conversion routines to do the actual conversion and relies on record support routines to handle arrays and other special cases.

15.4.4 Put Notify Routines

`dbPutNotify` is a request to notify the caller when all records that are processed as a result of a put complete processing. Complications occur because of record linking and asynchronous records. A put can cause an entire chain of records to process. If any record is an asynchronous record then record completion means asynchronous completion.

The following rules are implemented:

1. The user code must allocate a `putNotify` control block. Before calling `dbPutNotify` the user must set fields `paddr`, `pbuffer`, `nRequest`, and `dbrType`. `paddr` is the value returned by `dbNameToAddr`. The field `userPvt` is never accessed by the `dbNotify` code and is for the user. If a `putNotify` is already in use, i.e. a `dbPutNotify` has been issued and the `userCallback` has not been called, it is illegal to issue a new `dbPutNotify` with the same `putNotify` control block. Any such attempt will cause an assert failure.
2. The `userCallback` routine will be always be called unless `dbNotifyCancel` is called. The `userCallback` is called when the `dbPutNotify` is complete. The user is then free to reuse or delete the `putNotify` control block. The user supplied callback is called when all processing is complete or when an error is detected. If everything completes synchronously the callback routine will be called BEFORE `dbPutNotify` returns. The `userCallback` is called without anything locked.
3. If the user calls `dbNotifyCancel` then the `userCallback` will NOT be called after `dbNotifyCancel` returns. It may get called while `dbNotifyCancel` is active. If it is active `dbNotifyCancel` will not return until `userCallback` completes. Thus after `dbNotifyCancel` returns the user may reuse or delete the `putNotify` control block. The `putNotify` control block MUST NOT be deleted by `userCallback`.
4. If another `putNotify` is already active on the record associated with the `putNotify`, the new `putNotify` is put on a restart list and automatically restarted.
5. If the record associated with the `putNotify` is already active for some other reason, the `putNotify` takes ownership of the record and starts the put request when the record completes processing.
6. In general a set of records may be processed as a result of a single `dbPutNotify`. If a record in the set is found to be active, either because PACT is true or because a `putNotify` already owns the record, then that record is not made part of the set of records that must complete before the `putNotify` request completes.

15.4.4.1 putNotifyInit

This is called by iocInit.

Format:

```
void putNotifyInit(void);
```

15.4.4.2 dbPutNotify

Perform a database put and notify when record processing is complete.

Format:

```
void dbPutNotify(putNotify *pputnotify);
```

where

```
typedef enum {
    putNotifyOK,
    putNotifyCanceled,
    putNotifyError
    putNotifyPutDisabled
}putNotifyStatus;

typedef struct putNotify{
    /*The following must be set by the user*/
    void(*userCallback)(struct putNotify *);
    struct dbAddr*paddr;/*dbAddr set by dbNameToAddr*/
    void*pbuffer;/*address of data*/
    longnRequest;/*number of elements to be written*/
    shortdbrType;/*database request type*/
    void*usrPvt;/*for private use of user*/
    /*The following is status of request.Set by dbPutNotify*/
    putNotifyStatus status;
    /*fields private to database access*/
    ...
}putNotify;
```

The caller must allocate a putNotify structure and set the fields:

```
userCallback - Routine that is called upon completion
paddr - address of a dbAddr. Returned by dbNameToAddr.
pbuffer - address of data
nRequest - number of data elements
dbrType - database request type
usrPvt - a void * field that caller can use as needed.
```

The status value in putNotify.status is one of

- putNotifyOK Success
- putNotifyCanceled User issued a dbNotifyCancel.
- putNotifyError Either dbPut or dbProcess returned an error.
- putNotifyPutDisabled Puts have been disabled for the record.

The user callback is always called unless `dbNotifyCancel` is called before the put notify competes. It may be called while `dbPutNotify` or `dbNotifyCancel` is active.

15.4.4.3 `dbNotifyCancel`

Cancel an outstanding `dbPutNotify`.

Format:

```
void dbNotifyCancel(putNotify *pputnotify);
```

This cancels an active `dbPutNotify`.

15.4.4.4 `dbNotifyAdd`

This routine is called by database access itself. It should never be called by user code.

15.4.4.5 `dbNotifyCompletion`

This routine is called by database access itself. It should never be called by user code.

15.4.5 Utility Routines

15.4.5.1 `dbBufferSize`

Determine the buffer size for a `dbGetField` request, format:

```
long dbBufferSize(
    short dbrType, /* DBR_xxx*/
    long options, /* options mask*/
    long nRequest); /* number of elements*/
```

This routine returns the number of bytes that will be returned to `dbGetField` if the request type, options, and number of elements are specified as given to `dbBufferSize`. Thus it can be used to allocate storage for buffers.

NOTE: This should become a Channel Access routine

15.4.5.2 `dbValueSize`

Determine the size a value field, format:

```
dbValueSize(short dbrType); /* DBR_xxx*/
```

This routine returns the number of bytes for each element of type `dbrType`.

NOTE: This should become a Channel Access routine

15.4.5.3 `dbGetRset`

Get address of a record support entry table.

Format:

```
struct rset *dbGetRset(DBADDR *paddr);
```

This routine returns the address of the record support entry table for the record referenced by the DBADDR.

15.4.5.4 dbIsValueField

Is this field the VAL field of the record?

Format:

```
int dbIsValueField(struct dbFldDes *pdbFldDes);
```

This is the routine that makes the `get_value` record support routine obsolete.

15.4.5.5 dbGetFieldIndex

Get field index.

Format:

```
int dbGetFieldIndex(DBADDR *paddr);
```

Record support routines such as `special` and `cvt_dbaddr` need to know which field the DBADDR references. The include file describing the record contains define statements for each field. `dbGetFieldIndex` returns the index that can be matched against the define statements (normally via a switch statement).

15.4.5.6 dbGetNelements

Get number of elements in a field.

Format:

```
long dbGetNelements(struct link *plink, long *nelements);
```

This sets `*nelements` to the number of elements in the field referenced by `plink`.

15.4.5.7 dbIsLinkConnected

Is the link connected.

Format:

```
int dbIsLinkConnected(struct link *plink);
```

This routine returns (TRUE, FALSE) if the link (is, is not) connected.

15.4.5.8 dbGetPdbAddrFromLink

Get address of DBADDR from link.

Format:

```
DBADDR *dbGetPdbAddrFromLink(struct link *plink);
```

This macro returns the address of the DBADDR for a database link and NULL for all other link types.

15.4.5.9 dbGetLinkDBFtype

Get field type of a link.

Format:

```
int dbGetLinkDBFtype(struct link *plink);
```

15.4.5.10 dbGetControlLimits

Get Control Limits for link.

Format:

```
long dbGetControlLimits(struct link *plink, double *low, double *high);
```

15.4.5.11 dbGetGraphicLimits

Get Graphic Limits for link.

Format:

```
long dbGetGraphicLimits(struct link *plink, double *low, double *high);
```

15.4.5.12 dbGetAlarmLimits

Get Alarm Limits for link.

Format:

```
long dbGetAlarmLimits(struct link *plink,  
double lolo, double *low, double *high, double hihi);
```

15.4.5.13 dbGetPrecision

Get Precision for link.

Format:

```
long dbGetPrecision(struct link *plink, short *precision);
```

15.4.5.14 dbGetUnits

Get Units for link.

Format:

```
long dbGetUnits(struct link *plink, char *units, int unitsSize);
```

15.4.5.15 dbGetSevr

Get Severity for link.

Format:

```
long dbGetSevr(struct link *plink, short *sevr);
```

15.4.5.16 dbGetTimeStamp

Get Time Stamp for record containing link.

Format:

```
long dbGetTimeStamp(struct link *plink, TS_STAMP *pstamp);
```

15.4.6 Attribute Routine**15.4.6.1 dbPutAttribute**

Give a value to a record attribute.

```
long dbPutAttribute(char *recordTypename,
char *name, char*value);
```

This sets the record attribute name for record type recordTypename to value. For example the following would set the version for the ai record.

```
dbPutAttribute("ai", "VERS", "V800.6.95")
```

15.4.7 Process Routines**15.4.7.1 dbScanPassive**

dbScanLink

dbScanFwdLink

Process record if it is passive, format:

```
long dbScanPassive(
    struct dbCommon *pfrom,
    struct dbCommon *pto); /* addr of record*/
long dbScanLink(
    struct dbCommon *pfrom, struct dbCommon *pto);
void dbScanFwdLink(struct link *plink);
```

dbScanPassive and dbScanLink are given the record requesting the scan, which may be NULL, and the record to be processed. If the record is passive and pact is FALSE then dbProcess is called. Note that these routine are called by dbGetLink, dbPutField, and by recGblFwdLink.

dbScanFwdLink is given a link that must be a forward link field. It follows the rules for scanning a forward link. That is for DB_LINKS it calls dbScanPassive and for CA_LINKS it does a dbCaPutLink if the PROC field of record is being addressed.

15.4.7.2 dbProcess

Request that a database record be processed, format:

```
long dbProcess(struct dbCommon *precord);
```

Request that record be processed. Record processing is described in detail below.

15.5 Runtime Link Modification

Database links can be changed at run time but only via a channel access client, i.e. via calls to `dbPutField` but not to `dbPutLink`. The following restrictions apply:

- The date type may be `DBR_STRING` or a nil-terminated array of `DBR_CHAR` or `DBR_UCHAR` characters
- If a link is being changed to a different hardware link type then the `DTYP` field must be set before the link field.
- The syntax for the string is exactly the same as described for link fields in chapter “Database Definition”

15.6 Channel Access Monitors

There are facilities within the Channel Access communication infrastructure which allow the value of a process variable to be monitored by a channel access client. It is a responsibility of record support (and db common) to notify the channel access server when the internal state of a process variable has been modified. State changes can include changes in the value of a process variable and also changes in the alarm state of a process variable. The routine `db_post_events` is called to inform the channel access server that a process variable state change event has occurred.

```
#include <caeventmask.h>

int db_post_events(void *precord, void *pfield,
unsigned intselect);
```

The first argument, “precord”, should be passed a pointer to the record which is posting the event(s). The second argument, “pfield”, should be passed a pointer to the field in the record that contains the process variable that has been modified. The third argument, “select”, should be passed an event select mask. This mask can be any logical or combination of `{DBE_VALUE, DBE_LOG, DBE_ALARM}`. A description of the purpose of each flag in the event select mask follows.

- `DBE_VALUE` This indicates that a significant change in the process variable’s value has occurred. A significant change is often determined by the magnitude of the monitor “dead band” field in the record.
- `DBE_LOG` This indicates that a change in the process variable’s value significant to archival clients has occurred. A significant change to archival clients is often determined by the magnitude of the archive “dead band” field in the record.
- `DBE_ALARM` This indicates that a change in the process variable’s alarm state has occurred.

The function `db_post_events` returns 0 if it is successful and -1 if it fails. It appears to be common practice within EPICS record support to ignore the status from `db_post_events`. At this time `db_post_events` always returns 0 (success). Because so many records at this time depend on this behavior it is unlikely that it will be changed in the future.

The function `db_post_events` is written so that record support will never be blocked attempting to post an event because a slow client is not able to process events fast enough. Each call to `db_post_events` causes the current value, alarm status, and time stamp for the field to be copied into a ring buffer. The thread calling `db_post_events`

will not be delayed by any network or memory allocation overhead. A lower priority thread in the server is responsible for transferring the events in the event queue to the channel access clients that may be monitoring the process variable.

Currently, when an event is posted for a DBF_STRING field or a field containing array data the value is NOT saved in the ring buffer and the client will receive whatever value happens to be in the field when the lower priority thread transfers the event to the client. This behavior may be improved in the future.

15.7 Lock Set Routines

User code only calls `dbScanLock` and `dbScanUnlock`. All other routines are called by `iocCore`.

15.7.0.3 dbScanLock

Lock a lock set:

```
long void dbScanLock(struct dbCommon *precord);
```

Lock the lock set to which the specified record belongs.

15.7.0.4 dbScanUnlock

Unlock a lock set:

```
long void dbScanUnlock(struct dbCommon *precord);
```

Lock the lock set to which the specified record belongs

15.7.0.5 dbLockGetLockId

Get lock set id:

```
long dbLockGetLockId(struct dbCommon *precord);
```

Each lock set is assigned a unique ID. This routine retrieves it. This is most useful to determine if two records are in the same lock set.

15.7.0.6 dbLockInitRecords

Determine lock sets for each record in database.

```
void dbLockInitRecords(dbBase *pdbname);
```

Called by `iocInit`.

15.7.0.7 dbLockSetMerge

Merge records into same lock set.

```
void dbLockSetMerge(struct dbCommon *pfirst,
                    struct dbCommon *psecond);
```

If specified records are not in same lock set the lock sets are merged. Called by `dbLockInitRecords` and also when links are modified by `dbPutField`.

15.7.0.8 dbLockSetSplitSI

Recompute lock sets for given lock set

```
void dbLockSetSplit(struct dbCommon *psource);
```

This is called when dbPutField modifies links.

15.7.0.9 dbLockSetGblLock

Global lock for modifying links.

```
void dbLockSetGblLock(void);
```

Only one task at a time can modify link fields. This routine provides a global lock to prevent conflicts.

15.7.0.10 dbLockSetGblUnlock

Unlock the global lock.

```
void dbLockSetGblUnlock(void);
```

15.7.0.11 dbLockSetRecordLock

If record is not already scan locked lock it.

```
void dbLockSetRecordLock(struct dbCommon *precord);
```

15.8 Channel Access Database Links

The routines described here are used to create and manipulate Channel Access connections from database input or output links. At IOC initialization an attempt is made to convert all process variable links to database links. For any link that fails, it is assumed that the link is a Channel Access link, i.e. a link to a process variable defined in another IOC. The routines described here are used to manage these links. User code never needs to call these routines. They are automatically called by iocInit and database access.

At iocInit time a task dbCaLink is spawned. This task is a channel access client that issues channel access requests for all channel access links in the database. For each link a channel access search request is issued. When the search succeeds a channel access monitor is established. The monitor is issued specifying `ca_field_type` and `ca_element_count`. A buffer is also allocated to hold monitor return data as well as severity. When dbCaGetLink is called data is taken from the buffer, converted if necessary, and placed in the location specified by the `pbuffer` argument.

When the first dbCaPutLink is called for a link an output buffer is allocated, again using `ca_field_type` and `ca_element_count`. The data specified by the `pbuffer` argument is converted and stored in the buffer. A request is then made to dbCaLink task to issue a `ca_put`. Subsequent calls to dbCaPutLink reuse the same buffer.

15.8.1 Basic Routines

Except for dbCaPutLinkCallback, these routines are normally only called by database access, i.e. they are not called by record support modules.

15.8.1.1 dbCaLinkInit

Called by `iocInit` to initialize the `dbCa` library

```
void dbCaLinkInit(void);
```

15.8.1.2 dbCaAddLink

Add a new channel access link

```
void dbCaAddLink(struct link *plink);
```

15.8.1.3 dbCaAddLinkCallback

```
void dbCaAddLinkCallback(struct link *plink,
    dbCaCallback connect, dbCaCallback monitor, void *userPvt);
```

`connect` will be called whenever the link connects or disconnects. `monitor` will be called whenever a monitor event occurs. `connect` and or `monitor` may be null.

15.8.1.4 dbCaRemoveLink

Remove channel access link.

```
void dbCaRemoveLink(struct link *plink);
```

15.8.1.5 dbCaGetLink

Get link value

```
long dbCaGetLink(struct link *plink, short dbrType,
    void *pbuffer, unsigned short *psevr, long *nRequest);
```

15.8.1.6 dbCaPutLink

Put link value

```
long dbCaPutLink(struct link *plink, short dbrType,
    void *buffering nRequest);
```

15.8.1.7 dbCaPutLinkCallback

This is meant for use by device or record support that wants a put to complete before completing record processing.

```
long dbCaPutLinkCallback(struct link *plink, short dbrType,
    const void *pbuffer, long nRequest, dbCaPutCallback callback);
```

<base>/src/dev/devSoft/devAoSoftCallback.c provides an example of how to use this function. It contains:

```
static long write_ao(aoRecord *pao)
{
    struct link *plink = &pao->out;
    long status;
```

```

    if(pao->pact) return(0);
    if(plink->type!=CA_LINK) {
        status = dbPutLink(&pao->out,DBR_DOUBLE,&pao->oval,1);
        return(status);
    }
    status = dbCaPutLinkCallback(plink,DBR_DOUBLE,&pao->oval,1,
        (dbCaCallback)dbCaCallbackProcess,plink);
    if(status) {
        recGblSetSevr(pao, LINK_ALARM, INVALID_ALARM);
        return(status);
    }
    pao->pact = TRUE;
    return(0);
}

```

What happens is the following:

When the record is processed write_ao is called with pact=0.

If the link is not a CA_LINK it just calls dbPutLink. It leaves pact 0. Thus record support completes.

If it is a CA_LINK it calls dbCaPutLinkCallback and sets pact true. Thus record is asynchronous.

If the record is asynchronous then sometime later dbCaCallbackProcess is called. It calls the process routine of record support, which calls write_ao with pact true. write_ao just returns success. Record support then completes the second phase of record processing.

There is a possibility that the link is changed between the two phases of record processing. If this happens the user supplied callback will still get called exactly once but the link may have been modified.

15.8.2 Attributes of Link

The routines in this section are meant for use by device support to find out information about link fields. They must be called with dbScanLock held, i.e. normally they are called by the read or write method provided by device support.

15.8.2.1 dbCaIsLinkConnected

Is Channel Connected

```
int dbCaIsLinkConnected(struct link *plink)
```

This routine returns (TRUE, FALSE) if the link (is, is not) connected.

15.8.2.2 dbCaGetNelements

Get Number of Elements

```
long dbCaGetNelements(struct link *plink,long *nelements);
```

This call, which returns an error if the link is not connected, sets the native number of elements.

15.8.2.3 dbCaGetSevr

Get Alarm Severity

```
long dbCaGetSevr(struct link *plink,short *severity);
```

This call, which returns an error if the link is not connected, sets the alarm severity.

15.8.2.4 dbCaGetTimeStamp

Get Time Stamp

```
long dbCaGetTimeStamp(struct link *plink, TS_STAMP *pstamp);
```

This call, which returns an error if the link is not connected, sets pstamp to the time obtained by the last CA monitor.

15.8.2.5 dbCaGetLinkDBFtype

Get link type

```
int dbCaGetLinkDBFtype(struct link *plink);
```

This call, which returns an error if the link is not connected, returns the field type.

15.8.2.6 dbCaGetAttributes

Get Attributes

```
long dbCaGetAttributes(struct link *plink,
    void (*callback)(void *usrPvt), void *usrPvt);
```

When ever dbCa receives a connection it issues a CA get request to obtain the control, graphic, and alarm limits and to obtain the precision and units. By calling dbCaGetAttributes the caller can be notified when this get completes.

15.8.2.7 dbCaGetControlLimits

Get Control Limits

```
long dbCaGetControlLimits(struct link *plink, double *low, double *high);
```

This call returns an error if the link is not connected or if the CA get request for limits, etc. has not completed. If it returns success it has set the control limits.

15.8.2.8 dbCaGetGraphicLimits

Get graphic Limits

```
long dbCaGetGraphicLimits(struct link *plink, double *low, double *high);
```

This call returns an error if the link is not connected or if the CA get request for limits, etc. has not completed. If it returns success it has set the graphic limits.

15.8.2.9 dbCaGetAlarmLimits

Get Alarm Limits

```
long dbCaGetAlarmLimits(struct link *plink,
    double *lolo, double *low, double *high, double *hihi);
```

This call returns an error if the link is not connected or if the CA get request for limits, etc. has not completed. If it returns success it has set the alarm limits.

15.8.2.10 dbCaGetPrecision

Get Precision

```
long dbCaGetPrecision(struct link *plink, short *precision);
```

This call returns an error if the link is not connected or if the CA get request for limits, etc. has not completed. If it returns success it has set the precision.

15.8.2.11 dbCaGetUnits

Get Units

```
long dbCaGetUnits(struct link *plink, char *units, int unitsSize);
```

This call returns an error if the link is not connected or if the CA get request for limits, etc. has not completed. If it returns success it has set the units.

Chapter 16

EPICS General Purpose Tasks

16.1 Overview

This chapter describes two sets of EPICS supplied general purpose tasks: 1) Callback, and 2) Task Watchdog.

Often when writing code for an IOC there is no obvious task under which to execute. A good example is completion code for an asynchronous device support module. EPICS supplies the callback tasks for such code.

If an IOC tasks “crashes” there is normally no one monitoring the vxWorks shell to detect the problem. EPICS provides a task watchdog task which periodically checks the state of other tasks. If it finds that a monitored task has terminated or suspended it issues an error message and can also call other routines which can take additional actions. For example a subroutine record can arrange to be put into alarm if a monitored task crashes.

Since IOCs normally run autonomously, i.e. no one is monitoring the vxWorks shell, IOC code that issues `printf` calls generates errors messages that are never seen. In addition the vxWorks implementation of `fprintf` requires much more stack space than `printf` calls. Another problem with vxWorks is the `logMsg` facility. `logMsg` generates messages at higher priority than all other tasks except the shell. EPICS solves all of these problems via an error message handling facility. Code can call any of the routines `errMessage`, `errPrintf`, or `errlogPrintf`. Any of these result in the error message being generated by a separate low priority task. The calling task has to wait until the message is handled but other tasks are not delayed. In addition the message can be sent to a system wide error message file.

16.2 General Purpose Callback Tasks

16.2.1 Overview

EPICS provides three general purpose IOC callback tasks. The only difference between the tasks is their scheduling priority; low, medium or high. The low priority task runs at a priority just higher than Channel Access, the medium at a priority about equal to the median of the periodic scan tasks, and the high at a priority higher than the event scan task. The callback tasks are available for any software component that needs a task under which to run some job either immediately or after some delay. Jobs can also be cancelled during their delay period. The three callback tasks register themselves with the task watchdog (described below). They are created with a generous amount of stack space and can thus be used for invoking record processing. For example the I/O event scanner uses the general purpose callback tasks.

The following steps must be taken in order to use the general purpose callback tasks:

1. Include callback definitions:

```
#include <callback.h>
```

2. Provide storage for a structure that is a private structure for the callback tasks:

```
CALLBACK mycallback;
```

It is permissible for this to be part of a larger structure, e.g.

```
struct {
    ...
    CALLBACK mycallback;
    ...
} ...
```

3. Make calls (in most cases these are actually macros) to initialize the fields in the CALLBACK:

```
callbackSetCallback(CALLBACKFUNC func, CALLBACK *pcb);
```

This defines the callback routine to be executed. The first argument is the address of a function that will be given the address of the CALLBACK and returns void. The second argument is the address of the CALLBACK structure.

```
callbackSetPriority(int, CALLBACK *pcb);
```

The first argument is the priority, which can have one of the values: `priorityLow`, `priorityMedium`, or `priorityHigh`. These values are defined in `callback.h`. The second argument is again the address of the CALLBACK structure.

```
callbackSetUser(void *, CALLBACK *pcb);
```

This call is used to save a pointer value that can be retrieved again using the macro:

```
callbackGetUser(void *, CALLBACK *pcb);
```

If your callback function exists to process a single record inside calls to `dbScanLock/dbScanUnlock`, you can use this shortcut which provides the callback routine for you and sets the other two parameters at the same time (the user parameter here is a pointer to the record instance):

```
callbackSetProcess(CALLBACK *pcb, int prio, void *prec);
```

4. Whenever a callback request is desired just call one of the following:

```
callbackRequest(CALLBACK *pcb);
callbackRequestProcessCallback(CALLBACK *pcb, int prio, void *prec);
```

Both can be called from interrupt level code. The callback routine is passed a single argument, which is the same argument that was passed to `callbackRequest`, i.e., the address of the CALLBACK structure. The second routine is a shortcut for calling both `callbackSetProcess` and `callbackRequest`. The following delayed versions wait for the given time before queueing the callback routine for the relevant task to execute.

```
callbackRequestDelayed(CALLBACK *pCallback, double seconds);
callbackRequestProcessCallbackDelayed(CALLBACK *pCallback,
int Priority, void *pRec, double seconds);
```

These routines cannot be called from interrupt level code.

16.2.2 Syntax

The following calls are provided:

Notes:

```

void callbackInit(void);

void callbackSetCallback(void *pcallbackFunction,
    CALLBACK *pcallback);
void callbackSetPriority(int priority, CALLBACK *pcallback);
void callbackSetUser(void *user, CALLBACK *pcallback);
void callbackGetUser(void *user, CALLBACK *pcallback);
void callbackSetProcess(CALLBACK *pcallback, int Priority, void *prec);

void callbackRequest(CALLBACK *);
void callbackRequestProcessCallback(CALLBACK *pCallback,
    int Priority, void *prec);
void callbackRequestDelayed(CALLBACK *pCallback, double seconds);
void callbackRequestProcessCallbackDelayed(
    CALLBACK *pCallback, int Priority, void *prec, double seconds);
void callbackCancelDelayed(CALLBACK *pcallback);
int callbackSetQueueSize(int size);

```

- `callbackInit` is performed automatically at IOC initialization, thus user code never calls this function.
- `callbackSetCallback`, `callbackSetPriority`, `callbackSetUser`, and `callbackGetUser` are actually macros.
- Both `callbackRequest` and `callbackRequestProcessCallback` may be called from interrupt context.
- The delayed versions of the `callbackRequest` routines wait the given time before queueing the callback.
- `callbackCancelDelayed` can be used to cancel a delayed callback.
- `callbackRequestProcessCallback` issues the calls:

```

callbackSetCallback(ProcessCallback, pCallback);
callbackSetPriority(Priority, pCallback);
callbackSetUser(pRec, pCallback);
callbackRequest(pCallback);

```

The routine `ProcessCallback` was designed for asynchronous device completion and is defined as:

```

static void ProcessCallback(CALLBACK *pCallback)
{
    dbCommon    *pRec;
    struct rset *prset;

    callbackGetUser(pRec, pCallback);
    prset = (struct rset *)pRec->rset;
    dbScanLock(pRec);
    (*prset->process)(pRec);
    dbScanUnlock(pRec);
}

```

16.2.3 Example

An example use of the callback tasks.

```

#include <callback.h>

static structure {

```

```

    char    begid[80];
    CALLBACK callback;
    char    endid[80];
}myStruct;

void myCallback(CALLBACK *pcallback)
{
    struct myStruct *pmyStruct;
    callbackGetUser(pmyStruct,pcallback)
    printf("begid=%s endid=%s\n",&pmyStruct->begid[0],
        &pmyStruct->endid[0]);
}
example(char *pbegid, char*pendid)
{
    strcpy(&myStruct.begid[0],pbegid);
    strcpy(&myStruct.endid[0],pendid);
    callbackSetCallback(myCallback,&myStruct.callback);
    callbackSetPriority(priorityLow,&myStruct.callback);
    callbackSetUser(&myStruct,&myStruct.callback);
    callbackRequest(&myStruct.callback);
}

```

The example can be tested by issuing the following command to the vxWorks shell:

```
example("begin", "end")
```

This simple example shows how to use the callback tasks with your own structure that contains the CALLBACK structure at an arbitrary location.

16.2.4 Callback Queue

The callback requests put the requests for each callback priority into a separate ring buffer. These buffers can by default hold up to 2000 requests. This limit can be changed by calling `callbackSetQueueSize` before `iocInit` in the startup file. The syntax is:

```
int callbackSetQueueSize(int size)
```

16.3 Task Watchdog

EPICS provides a task that acts as a watchdog for other tasks. Any task can request to be watched, and most of the IOC tasks do this. A status monitoring subsystem in the IOC can register to be notified about any changes that occur. The watchdog task runs periodically and checks each task in its task list. If any task is suspended, an error message is displayed and any notifications made. The task watchdog provides the following features:

1. Include module:

```
#include <taskwd.h>
```

2. Request by a task to be monitored:

```
taskwdInsert (epicsThreadId tid, TASKWDFUNC callback, VOID *usr);
```

This adds the task with the specified `tid` to the list of tasks to be watched, and makes any requested notifications that a new task has been registered. If `tid` is given as zero, the `epicsThreadId` of the calling thread is used

instead. If `callback` is not `NULL` and the task later becomes suspended, the callback routine will be called with the single argument `usr`.

3. Remove task from list:

```
taskwdRemove(epicsThreadId tid);
```

This routine must be called before the monitored task exits. It makes any requested notifications and removes the task from the list of tasks being watched. If `tid` is given as zero, the `epicsThreadId` of the calling thread is used instead.

4. Request to be notified of changes:

```
typedef struct {
    void (*insert)(void *usr, epicsThreadId tid);
    void (*notify)(void *usr, epicsThreadId tid, int suspended);
    void (*remove)(void *usr, epicsThreadId tid);
} taskwdMonitor;
```

```
taskwdMonitorAdd(const taskwdMonitor *funcs, void *usr);
```

This call provides a set of callbacks for the task watchdog to call when a task is registered or removed or when any task gets suspended. The `usr` pointer given at registration is passed to the callback routine along with the `tid` of the thread the notification is about. In many cases the `insert` and `remove` callbacks will be called from the context of the thread itself, although this is not guaranteed (the registration could be made by a parent thread for instance). The `notify` callback also indicates whether the task went into or out of suspension; it is called in both cases, unlike the callbacks registered with `taskwdInsert` and `taskwdAnyInsert`.

5. Rescind notification request:

```
taskwdMonitorDel(const taskwdMonitor *funcs, void *usr);
```

This call removes a previously registered notification. Both `funcs` and `usr` must match the values given to `taskwdMonitorAdd` when originally registered.

6. Print a report:

```
taskwdShow(int level);
```

If `level` is zero, the number of tasks and monitors registered is displayed. For higher values the registered task names and their current states are also shown in tabular form.

7. The following routines are provided for backwards compatibility purposes, but are now deprecated:

```
taskwdAnyInsert(void *key, TASKWDANYFUNC callback, VOID *usr);
```

The callback routine will be called whenever any of the tasks being monitored by the task watchdog become suspended. `key` must have a unique value because the task watchdog system uses this value to determine which entry to remove when `taskwdAnyRemove` is called.

```
taskwdAnyRemove(void *key);
```

`key` is the same value that was passed to `taskwdAnyInsert`.

Chapter 17

Database Scanning

17.1 Overview

Database scanning is the mechanism for deciding when to process a record. Five types of scanning are possible:

- **Periodic:** A record can be processed periodically. A number of time intervals are supported.
- **Event:** Event scanning is based on the posting of an event by another component of the software via a call to the routine `post_event`.
- **I/O Event:** The original meaning of this scan type is a request for record processing as a result of a hardware interrupt. The mechanism supports hardware interrupts as well as software generated events.
- **Passive:** Passive records are processed only via requests to `dbScanPassive`. This happens when database links (Forward, Input, or Output), which have been declared “Process Passive” are accessed during record processing. It can also happen as a result of `dbPutField` being called (which normally results from a Channel Access put request).
- **Scan Once:** In order to provide for caching puts, the scanning system provides a routine `scanOnce` which arranges for a record to be processed one time.

This chapter explains database scanning in increasing order of detail. It first explains database fields involved with scanning. It next discusses the interface to the scanning system. The last section gives a brief overview of how the scanners are implemented.

17.2 Scan Related Database Fields

The following fields are normally defined via DCT. It should be noted, however, that it is quite permissible to change any of the scan related fields of a record dynamically. For example, a display manager screen could tie a menu control to the `SCAN` field of a record and allow the operator to dynamically change the scan mechanism.

17.2.1 SCAN

This field, which specifies the scan mechanism, has an associated menu of the following form:

Passive: Passively scanned.

Event: Event Scanned. The field `EVNT` specifies event number.

I/O Intr: I/O Event scanned.

10 Second: Periodically scanned every 10 seconds

...

.1 Second: Periodically scanned every .1 seconds

17.2.2 PHAS

This field determines processing order for records that are in the same scan set. For example all records periodically scanned at a 2 second rate are in the same scan set. All Event scanned records with the same EVNT are in the same scan set, etc. For records in the same scan set, all records with PHAS=0 are processed before records with PHAS=1, which are processed before all records with PHAS=2, etc.

In general it is not a good idea to rely on PHAS to enforce processing order. It is better to use database links.

17.2.3 EVNT - Event Number

This field only has meaning when SCAN is set to Event scanning, in which case it specifies the event number. In order for a record to be event scanned, EVNT must be in the range 0,...255. It should also be noted that some EPICS software components will not request event scanning for event 0. One example is the eventRecord record support module. Thus the application developer will normally want to define events in the range 1,...,255.

17.2.4 PRIO - Scheduling Priority

This field can be used by any software component that needs to specify scheduling priority, e.g. the event and I/O event scan facility uses this field.

17.3 Scan Related Software Components

17.3.1 menuScan.dbd

This file contains definitions for a menu related to field SCAN. The definitions are of the form:

```
menu(menuScan) {
    choice(menuScanPassive, "Passive")
    choice(menuScanEvent, "Event")
    choice(menuScanI_O_Intr, "I/O Intr")
    choice(menuScan10_second, "10 second")
    choice(menuScan5_second, "5 second")
    choice(menuScan2_second, "2 second")
    choice(menuScan1_second, "1 second")
    choice(menuScan_5_second, ".5 second")
    choice(menuScan_2_second, ".2 second")
    choice(menuScan_1_second, ".1 second")
}
```

The first three choices must appear in the order and location shown. The remaining definitions are for the periodic scan rates, which must appear in the order slowest to fastest (the order directly controls the thread priority assigned to the particular scan rate, and faster scan rates should be assigned higher thread priorities). At IOC initialization, the menu choice strings are read at scan initialization. The number of periodic scan rates and the period of each rate is determined from the menu choice strings. Thus periodic scan rates can be changed by changing menuScan.dbd

and loading this version via `dbLoadDatabase`. The only requirement is that each periodic choice string must begin with a numeric value specified in units of seconds.

17.3.2 `dbScan.h`

All software components that interact with the scanning system must include this file.

The most important definitions in this file are:

```
#define SCAN_PASSIVE      menuScanPassive
#define SCAN_EVENT       menuScanEvent
#define SCAN_IO_EVENT    menuScanI_O_Intr
#define SCAN_1ST_PERIODIC (menuScanI_O_Intr + 1)

/*definitions for I/O Interrupt Scanning */
typedef struct io_scan_list *IOSCANPVT;

long scanInit(void);
void scanRun(void);
void scanPause(void);

void post_event(int event);
void scanAdd(struct dbCommon *);
void scanDelete(struct dbCommon *);
double scanPeriod(int scan);
void scanOnce(struct dbCommon *precord);
int scanOnceSetQueueSize(int size);

int scanppl(void); /* print periodic lists*/
int scanpel(void); /* print event lists*/
int scanpiol(void); /* print io_event list*/

void scanIoInit(IOSCANPVT *);
void scanIoRequest(IOSCANPVT);
```

The first set of definitions defines the various scan types. The next definition `IOSCANPVT` is used when interfacing with the I/O interrupt scanner. The remaining definitions define the public scan access routines. These are described in the following subsections.

17.3.3 Initializing And Controlling Database Scanning

```
scanInit(void);
```

The routine `scanInit` is called by `iocInit`. It initializes the scanning system.

```
scanRun(void);
scanPause(void);
```

These routines start and stop all the scan tasks respectively. They are used by the `iocInit`, `iocRun` and `iocPause` commands.

17.3.4 Adding And Deleting Records From Scan List

The following routines are called each time a record is added or deleted from a scan list.

```
scanAdd(struct dbCommon *);
scanDelete(struct dbCommon *);
```

These routines are called by `scanInit` at IOC initialization time in order to enter all records created via DCT into the correct scan list. The routine `dbPut` calls `scanDelete` and `scanAdd` each time a scan related field is changed (each scan related field is declared to be `SPC_SCAN` in `dbCommon.dbd`). `scanDelete` is called before the field is modified and `scanAdd` after the field is modified.

17.3.5 Obtaining the scan period from the SCAN field

```
double scanPeriod(int scan);
```

The argument is an offset into the set of enum choices for `menuScan.h`. Most users will just use the `SCAN` field of a database record. It returns the scan period in seconds. The result will be 0.0 if scan doesn't refer to a periodic rate.

17.3.6 Declaring Database Event

Whenever any software component wants to declare a database event, it just calls:

```
post_event(event)
```

This can be called by virtually any IOC software component. For example sequence programs can call it. The record support module for `eventRecord` calls it.

17.3.7 Interfacing to I/O Event Scanning

Interfacing to the I/O event scanner is done via some combination of device and driver support.

1. Include `<dbScan.h>`
2. For each separate event source the following must be done:

- (a) Declare an `IOSCANPVT` variable, e.g.

```
static IOSCANPVT ioscanpvt;
```

- (b) Call `scanIoInit`, e.g.

```
scanIoInit(&ioscanpvt);
```

3. Provide the device support `get_ioint_info` routine. This routine has the format:

```
long get_ioint_info(
    int cmd,
    struct dbCommon *precord,
    IOSCANPVT *ppvt);
```

This routine is called each time the record pointed to by `precord` is added or deleted from an I/O event scan list. `cmd` has the value (0,1) if the record is being (added to, deleted from) an I/O event list. This routine must give a value to `*ppvt`.

4. Whenever an I/O event is detected call `scanIoRequest`, e.g.

```
scanIoRequest(ioscanpvt)
```

This routine can be called from interrupt level. The request is actually directed to one of the standard callback tasks. The actual one is determined by the `PRIO` field of `dbCommon`.

The following code fragment shows an event record device support module that supports I/O event scanning:

```

#include <vxWorks.h>
#include <types.h>
#include <stdioLib.h>
#include <intLib.h>
#include <dbDefs.h>
#include <dbAccess.h>
#include <dbScan.h>
#include <recSup.h>
#include <devSup.h>
#include <eventRecord.h>
/* Create the dset for devEventXXX */
long init();
long get_ioint_info();
struct {
    long number;
    DEVSUPFUN report;
    DEVSUPFUN init;
    DEVSUPFUN init_record;
    DEVSUPFUN get_ioint_info;
    DEVSUPFUN read_event;
} devEventTestIoEvent={
    5,
    NULL,
    init,
    NULL,
    get_ioint_info,
    NULL};
static IOSCANPVT ioscanpvt;
static void int_service(IOSCANPVT ioscanpvt)
{
    scanIoRequest(ioscanpvt);
}

static long init()
{
    scanIoInit(&ioscanpvt);
    intConnect(<vector>, (FUNCPTR)int_service, ioscanpvt);
    return(0);
}
static long get_ioint_info(
int cmd,
struct eventRecord *pr,
IOSCANPVT *ppvt)
{
    *ppvt = ioscanpvt;
    return(0);
}

```

17.4 Implementation Overview

The code for the entire scanning system resides in `dbScan.c`, i.e. periodic, event, and I/O event. This section gives an overview of how the code in `dbScan.c` is organized. The listing of `dbScan.c` must be studied for a complete

understanding of how the scanning system works.

17.4.1 Definitions And Routines Common To All Scan Types

Everything is built around two basic structures:

```
typedef struct scan_list {
    epicsMutexId lock;
    ELLLIST      list;
    short        modified;
};

typedef struct scan_element{
    ELLNODE      node;
    scan_list    *pscan_list;
    struct dbCommon *precord;
}
```

Later we will see how `scan_lists` are determined. For now just realize that `scan_list.list` is the head of a list of records that belong to the same scan set (for example, all records that are periodically scanned at a 1 second rate are in the same scan set). The `node` field in `scan_element` contain the list links. The normal `libCom ellLib` routines are used to access the list. Each record that appears in some scan list has an associated `scan_element`. The `SPVT` field which appears in `dbCommon` holds the address of the associated `scan_element`.

The `lock`, `modified`, and `pscan_list` fields allow `scan_elements`, i.e. records, to be dynamically removed and added to scan lists. If `scanList`, the routine which actually processes a scan list, is studied it can be seen that these fields allow the list to be scanned very efficiently if no modifications are made to the list while it is being scanned. This is, of course, the normal case.

The `dbScan.c` module contains several private routines. The following access a single scan set:

- `printList`: Prints the names of all records in a scan set.
- `scanList`: This routine is the heart of the scanning system. For each record in a scan set it does the following:

```
dbScanLock (precord);
dbProcess (precord);
dbScanUnlock (precord);
```

It also has code to recognize when a scan list is modified while the scan set is being processed.


- `addToList`: This routine adds a new element to a scan list.
- `deleteFromList`: This routine deletes an element from a scan list.

17.4.2 Event Scanning

Event scanning is built around the following definitions:

```
#define MAX_EVENTS 256
typedef struct event_scan_list {
    CALLBACK      callback;
    scan_list      scan_list;
} event_scan_list;
static event_scan_list
*pevent_list [NUM_CALLBACK_PRIORITIES] [MAX_EVENTS];
```


`pevent_list` is a 2d array of pointers to `scan_lists`. Note that the array allows for 256 events, i.e. one for each possible event number. In other words, each event number and priority has its own scan list. No `scan_list` is actually created until the first request to add an element for that event number. The event scan lists have the memory layout illustrated below:



17.4.2.1 post_event

```
post_event(int event)
```

This routine is called to request event scanning. It can be called from interrupt level. It looks at each `event_scan_list` referenced by `pevent_list[*][event]` (one for each callback priority) and if any elements are present in the `scan_list` a `callbackRequest` is issued. The appropriate callback task calls routine `eventCallback`, which just calls `scanList`.

17.4.3 I/O Event Scanning

I/O event scanning is built around the following definitions:

```
typedef struct io_scan_list {
    CALLBACK          callback;
    scan_list         scan_list;
    struct io_scan_list *next;
}
static io_scan_list *iosl_head[NUM_CALLBACK_PRIORITIES] = {
    NULL, NULL, NULL};
```


The array `iosl_head` and the field `next` are only kept so that `scanpio1` can be implemented and will not be discussed further. I/O event scanning uses the general purpose callback tasks to perform record processing, i.e. no task is spawned for I/O event. The `callback` field of `io_scan_list` is used to communicate with the callback tasks.

The following routines implement I/O event scanning:

17.4.3.1 scanIoInit

```
scanIoInit (IOSCANPVT *ppioscanpvt)
```

This routine is called by device or driver support. It is called once for each interrupt source. `scanIoInit` allocates and initializes an array of `io_scan_list` structures; one for each callback priority and puts the address in `pioscanpvt`. Three callback priorities are supported; low, medium, and high. Thus for each interrupt source the structures are as illustrated below:



When `scanAdd` or `scanDelete` are called, they call the device support routine `get_ioint_info` which returns `pioscanpvt`. The `scan_element` is added or deleted from the correct scan list.

17.4.3.2 scanIoRequest

```
scanIoRequest (IOSCANPVT pioscanpvt)
```


This routine is called to request I/O event scanning. It can be called from interrupt level. It looks at each `io_scan_list` referenced by `pioscanpvt` (one for each callback priority) and if any elements are present in the `scan_list` a `callbackRequest` is issued. The appropriate callback task calls routine `ioeventCallback`, which just calls `scanList`.

17.4.4 Periodic Scanning

Periodic scanning is built around the following definitions:

```
typedef struct periodic_scan_list {
    scan_list      scan_list;
    double         period;
    volatile enum ctl scanCtl;
    epicsEventId  loopEvent;
} periodic_scan_list;
static int nPeriodic;
static periodic_scan_list **papPeriodic;
static epicsThreadId *periodicTaskId;
```

`nPeriodic`, which is determined at `iocInit` time, is the number of periodic rates. `papPeriodic` is a pointer to an array of pointers to `scan_lists`. There is an array element for each scan rate. Thus the structure illustrated in the figure below exists after `iocInit`.



A periodic scan task is created for each scan rate. The following routines implement periodic scanning:

17.4.4.1 initPeriodic

```
initPeriodic()
```

This routine first determines the scan rates. It does this by accessing the `SCAN` field of the first record it finds. It issues a call to `dbGetField` with a `DBR_ENUM` request. This returns the menu choices for `SCAN`. From this the periodic rates are determined. The array of pointers referenced by `papPeriodic` is allocated. For each scan rate a `scan_list` is allocated and a `periodicTask` is spawned.

17.4.4.2 periodicTask

```
periodicTask (struct scan_list *psl)
```

This task just performs an infinite loop, calling `scanList` and then waiting until the start of the next scan interval, allowing for the time it took to scan the list. If a periodic scan list takes longer to process than its defined scan period, the next scan will be delayed by half a scan period, with a maximum of 1 second delay. This does not limit what scan rates can actually be implemented, as long as all the records in the list can be processed within the requested period. Persistent over-runs (more than 10 times in a row) will result in a warning message being logged. The total number of over-runs is counted by each scan thread and can be displayed using the `scanppl` command.

17.4.5 Scan Once

17.4.5.1 scanOnce

```
void scanOnce (dbCommon *precord)
```

A task `onceTask` waits for requests to issue a `dbProcess` request. The routine `scanOnce` puts the address of the record to be processed in a ring buffer and wakes up `onceTask`.

This routine can be called from interrupt level.

17.4.5.2 SetQueueSize

`scanOnce` places its request on a ring buffer. This is set by default to 1000 entries. It can be changed by executing the following command in the startup script before `iocInit`:

```
int scanOnceSetQueueSize(int size);
```


Chapter 18

IOC Shell

18.1 Introduction

The EPICS IOC shell is a simple command interpreter which provides a subset of the capabilities of the vxWorks shell. It is used to interpret startup scripts (st.cmd) and to execute commands entered at the console terminal. In most cases vxWorks startup scripts can be interpreted by the IOC shell without modification. The following sections of this chapter describe the operation of the IOC shell from the user's and programmer's points of view.

18.2 IOC Shell Operation

The IOC shell reads lines of input, expands environment variable parameters, breaks the line into commands and arguments then calls functions corresponding to the decoded command. Commands and arguments are separated by one or more 'space' characters. Characters interpreted as spaces include the actual space character and the tab character as well as commas and open and close parentheses. Thus, the command line

```
dbLoadRecords ("db/dbExample1.db", "user=mrk")
```

would be interpreted by the IOC shell as the `dbLoadRecords` command with arguments `db/dbExample1.db` and `user=mrk`.

Unrecognized commands result in a diagnostic message but are otherwise ignored. Missing arguments are given a default value (0 for numeric arguments, NULL for string arguments). Extra arguments are ignored.

Unlike the vxWorks shell, string arguments do not have to be enclosed in quotes unless they contain one or more of the space characters, in which case one of the quoting mechanisms described in the following section must be used.

18.2.1 Environment variable expansion

Lines of input not beginning with a comment character (#) are searched for macro references in the form `${name}` or `$(name)`. The documentation for the `macLib` facility (chapter 19) describes some possible syntax variations for macro references. Such references are replaced with the value of the environment variable they name before any other processing takes place. Macro expansion is recursive so, for example,

```
epics> epicsEnvSet v1 \${v2}
epics> epicsEnvSet v2 \${v3}
epics> epicsEnvSet v3 somePV
epics> dbpr ${v1}
```

will print information about the `somePV` process variable - the `${v1}` argument to the `dbpr` command expands to `${v2}` which expands to `${v3}` which expands to `somePV`. The backslashes in the definitions are needed to postpone the substitution of the following variables, which would otherwise be performed before the `epicsEnvSet` command was run.

Macro references that appear inside single-quotes are not expanded.

18.2.2 Quoting

Quoting is used to remove the special meaning normally assigned to certain characters and can be used to include space or quote characters in arguments. Quoting takes place after the macro expansion described above has been performed, and cannot be used to extend a command over more than one input line.

There are three quoting mechanisms: the backslash character, single quotes, and double quotes. A backslash (`\`) preserves the literal value of the following character. Enclosing characters in single or double quotes preserves the literal value of each character (including backslashes) within the quotes. A single quote may occur between double quotes and a double quote may occur between single quotes. Note that commands called from the shell may perform additional unescaping and macro expansion on their argument strings.

18.2.3 Command-line editing and history

The IOC shell can use the `readline` or `tecla` library to obtain input from the console terminal. This provides full command-line editing as well as easy access to previous commands through the command-line history capabilities provided by these libraries. For full details, refer to the `readline` or `tecla` library documentation. Command and argument completion is not supported.

If neither the `readline` nor `tecla` library is used the only command-line editing and history capabilities will be those supplied by the underlying operating system. The console keyboard driver in Windows, for example, provides its own command-line editing and history commands. On `vxWorks` the `ledLib` command-line input library routines are used.

18.2.4 Redirection

The IOC shell recognizes a subset of UNIX shell I/O redirection operators. The redirection operators may precede, or appear anywhere within, or follow a command. Redirections are processed in the order they appear, from left to right. Failure to open or create a file causes the redirection to fail and the command to be ignored.

Redirection of input causes the file whose name results from the expansion of `filename` to be opened for reading on file descriptor `n`, or the standard input (file descriptor 0) if `n` is not specified. The general format for redirecting input is:

```
[n]<filename
```

As a special case, the IOC shell recognizes a standard input redirection appearing by itself (i.e. with no command) as a request to read commands from `filename` until an exit command or EOF is encountered. The IOC shell then resumes reading commands from the current source. Commands read from `filename` are not added to the `readline` command history. The level of nesting is limited only by the maximum number of files that can be open simultaneously.

Redirection of output causes the file whose name results from the expansion of `filename` to be opened for writing on file descriptor `n`, or the standard output (file descriptor 1) if `n` is not specified. If the file does not exist it is created; if it does exist it is truncated to zero size. The general format for redirecting output is:

```
[n]>filename
```

The general format for appending output is:

```
[n]>>filename
```

Redirection of output in this fashion causes the `filename` to be opened for appending on file descriptor `n`, or the standard output (file descriptor 1) if `n` is not specified. If the file does not exist it is created.

18.2.5 Utility Commands

The IOC shell recognizes the following commands as well as the commands described in chapter 6 (Database Definition) and chapter 9 (IOC Test Facilities) among others. The commands described in the sequencer documentation will also be recognized if the sequencer is included.

Command	Description
<code>help [command ...]</code>	Display synopsis of specified commands. Wild-card matching is applied so <code>'help db*'</code> displays a synopsis of all commands beginning with the letters <code>'db'</code> . With no arguments this displays a list of all commands.
<code>#</code>	A <code>'#'</code> as the first non-whitespace character on a line marks the beginning of a comment, which continues to the end of the line (however some versions of Base may require a space after the <code>'#'</code> character to properly recognize it as a comment)
<code>exit</code>	Stop reading commands. When the top-level command interpreter encounters an exit command or end-of-file (EOF) it returns to its caller.
<code>cd directory</code>	Change working directory to directory.
<code>pwd</code>	Print the name of the working directory.
<code>var [name [value]]</code>	If both arguments are present, assign the value to the named variable. If only the name argument is present, print the current value of that variable. If neither argument is present, print the value of all variables registered with the shell. Variables are registered in application database definitions using the variable keyword as described in Section 6.9 on page 104.
<code>show [-level] [task ...]</code>	Show information about specified tasks. If no task arguments are present, show information on all tasks. The level argument controls the amount of information printed. The default level is 0. The task arguments can be task names or task i.d. numbers.
<code>system command_string</code>	Send <code>command_string</code> to the system command interpreter for execution. This command is present only if some application database definition file contains <code>registrar(iocshSystemCommand)</code> and if the system provides a suitable command interpreter (vx-Works does not).
<code>epicsEnvSet name value</code>	Set environment variable name to the specified value.
<code>epicsEnvShow [name]</code>	If no name is specified the names and values of all environment variables will be shown. If a name is specified the value of that environment variable will be shown.
<code>epicsParamShow</code>	Show names and values of all EPICS configuration parameters.
<code>iocLogInit</code>	Initialize IOC logging.
<code>epicsThreadSleep sec</code>	Pause execution of IOC shell for <code>sec</code> seconds.

The `var` command is intended for simple applications such as setting the value of debugging flags. Applications which require more complex expression handling should use the `cexp` package.

A `spy` command to show periodic activity reports is available on RTEMS as part of the `RTEMS_UTILS` support module. The following changes must be made to add this command to an application.

- Add an `RTEMS_UTILS` entry to the application `configure/RELEASE` file.
- Add `spy.dbd` to the list of application `dbd` files and `rtemsutils` to the list of application libraries in the application `Makefile`.

18.2.6 Environment Variables

The IOC shell uses the following environment variables to control its operation.

Variable	Description
IOCSH_PS1	Prompt string. Default is "epics>".
IOCSH_HISTSZIE	Number of previous command lines to remember. If the IOCSH_HISTSZIE environment variable is not present the value of the HISTSIZE environment variable is used. In the absence of both environment variables, 10 command lines will be remembered.
TERM, INPUTRC	These and other environment variables are used by the readline and termcap libraries and are described in the documentation for those libraries.

18.2.7 Conditionals

The IOC shell does not provide operators for conditionally executing commands but the effect can be simulated using macro expansion. The simplest technique is to precede a command with a macro that expands to either '#' or '' (or ' '). The following startup script line shows how this can be done:

```
...
$(LOAD_DEBUG=#) $(DEBUG) dbLoadRecords("db/debugRec.db", "P=$(P),R=debug")
...
```

Starting the IOC in the normal fashion will result in the above line being commented out and the debugRec.db file being omitted:

```
./st.cmd
```

Setting the LOAD_DEBUG environment variable to an empty string before starting the IOC will result in the debugRec.db file being loaded:

```
LOAD_DEBUG="" ./st.cmd
```

A similar technique can be used to execute external scripts conditionally. The startup command file contains code like:

```
epicsEnvSet PILATUS_ENABLED "$ (PILATUS_ENABLED=NO) "
...
< pilatus-$(PILATUS_ENABLED).cmd
```

with one set of conditional code in a file named pilatus-YES.cmd and the other set of conditional code in a file named pilatus-NO.cmd This technique can be expanded to a form similar to a C 'switch' statement for the example above by providing additional pilatus-XXX.cmd scripts.

18.3 IOC Shell Programming

The declarations described in this section are included in the `iocsh.h` header file.

18.3.1 Invoking the IOC shell

The prototypes for calling the IOC shell command interpreter are:

```
int iocsh(const char *pathname);
int iocshCmd(const char *cmd);
```

The pathname argument to the `iocsh` function is the name of the file from which commands are to be read. If the pathname argument is NULL, commands are read from the standard input and prompts are issued to the standard output. Commands are read until an `exit` command is encountered or until end-of-file is reached, at which point `iocsh` returns a value of 0. If the specified file can not be opened `iocsh` returns -1.

The IOC shell can be invoked from the vxWorks shell, either from within a vxWorks startup script or from vxWorks command-line interpreter, using

```
iocsh "script"
```

to read from an IOC shell script. It can also be invoked from the vxWorks command-line interpreter with no argument, in which case the IOC shell takes over the duties of command-line interaction.

The `iocshCmd` function takes a single IOC shell command and executes it. The function may be called from any thread, but many of the commands are not necessarily thread-safe so this should only be used with care. The function is most useful to execute a single IOC shell command from a vxWorks startup script or command line, like this:

```
iocshCmd "iocsh command string"
```

The `stdio` stream redirection and environment variable expansion processes described above are performed on the string as part of the execution process.

18.3.2 Registering Commands

Commands must be registered before they can be recognized by the IOC shell. Registration is achieved by calling the registration function:

```
void iocshRegister(const iocshFuncDef *piocshFuncDef, iocshCallFunc func);
```

The first argument is a pointer to a data structure which describes the command and any arguments it may take. The second argument is a pointer to a function which will be called by `iocsh` when the corresponding command is encountered.

The command is described by the `iocshFuncDef` structure:

```
struct iocshFuncDef {
    const char *name;
    int nargs;
    const iocshArg * const *arg;
};
```

The `name` element is the name of the command. The `arg` element is a pointer to an array of pointers to structures each of which defines a single argument. The `nargs` element declares the number of entries in the array of pointers to the argument descriptions. If `nargs` is zero, `arg` can be `NULL`. The structures which define each of the arguments is:

```
struct iocshArg {
    const char *name;
    iocshArgType type;
}iocshArg;
```

The `name` element is used by the `help` command to print a synopsis for the command. The `type` element describes the type of the argument and takes one of the following values:

Type Specifier	Description
<code>iocshArgInt</code>	The argument will be converted to an integer value.
<code>iocshArgDouble</code>	The argument will be converted to a double-precision floating point value.
<code>iocshArgString</code>	The argument will be left as a string. The memory used to hold the string is 'owned' by <code>iocsh</code> and will be reused once the handler function returns.
<code>iocshArgPersistentString</code>	A copy of the argument will be made and a pointer to the copy will be passed to the handler. The called function can release this copy by using the pointer as an argument to <code>free()</code> .
<code>iocshArgPdbbase</code>	The argument must be <code>pdbbase</code> .

`iocshArgArgv` An arbitrary number of arguments is expected. Subsequent `iocshArg` structures will be ignored.

The ‘handler’ function which is called when its corresponding command is recognized should be of the form:

```
void showCallFunc(const iocshArgBuf *args);
```

The argument to the handler function is a pointer to an array of unions. The number of elements in this array is equal to the number of arguments specified in the structure describing the command. The type and name of the union element which contains the argument value depends on the ‘type’ element of the corresponding argument descriptor:

Type Specifier	Type	Union element
<code>iocshArgInt</code>	<code>int</code>	<code>args[i].ival</code>
<code>iocshArgDouble</code>	<code>double</code>	<code>args[i].dval</code>
<code>iocshArgString</code>	<code>char *</code>	<code>args[i].sval</code>
<code>iocshArgPersistentString</code>	<code>char *</code>	<code>args[i].sval</code>
<code>iocshArgPdbbase</code>	<code>void *</code>	<code>args[i].vval</code>
<code>iocshArgArgv</code>	<code>int</code>	<code>args[i].aval.ac</code>
	<code>char **</code>	<code>args[i].aval.av</code>

If an `iocshArgArgv` argument type is present it is often the first and only argument specified for the command. In this case, `args[0].aval.av[0]` will be the name of the command, `args[0].aval.av[1]` will be the first argument, and so on.

18.3.3 Registrar Command Registration

Commands are normally registered with the IOC shell in a registrar function. The application’s database description file uses the `registrar` keyword to specify a function which will be called from the EPICS initialization code during the application startup process. This function then calls `iocshRegister` to register its commands with the `iocsh`.

The following code fragments shows how this can be performed for an example driver.

```
#include <iocsh.h>
#include <epicsExport.h>

/* drvXxx code, FuncDef and CallFunc definitions ... */

static void drvXxxRegistrar(void)
{
    iocshRegister(&drvXxxConfigureFuncDef, drvXxxConfigureCallFunc);
}
epicsExportRegistrar(drvXxxRegistrar);
```

To include this driver in an application a developer would then add

```
registrar(drvXxxRegistrar)
```

to an application database description file.

18.3.4 Automatic Command Registration

A C++ static constructor can also be used to register IOC shell commands before the EPICS application begins. The following example shows how the `epicsThreadSleep` command could be described and registered.

```
#include <iocsh.h>

static const iocshArg epicsThreadSleepArg0 = { "seconds",iocshArgDouble};
static const iocshArg *const epicsThreadSleepArgs[1] =
    {&epicsThreadSleepArg0};
static const iocshFuncDef epicsThreadSleepFuncDef =
    {"epicsThreadSleep",1,epicsThreadSleepArgs};
static void epicsThreadSleepCallFunc(const iocshArgBuf *args)
{
    epicsThreadSleep(args[0].dval);
}

static int doRegister(void)
{
    iocshRegister(epicsThreadSleepFuncDef, epicsThreadSleepCallFunc);
    return 1;
}
static int done = doRegister();
```


Chapter 19

libCom

This chapter and the next describe the facilities provided in `<base>/src/libCom`. This chapter describes facilities which are platform independent. The next chapter describes facilities which have different implementations on different platforms.

19.1 bucketLib

`bucketLib.h` describes a hash facility for integers, pointers, and strings. It is used by the Channel Access Server. It is currently undocumented.

19.2 calc

`postfix.h` defines several macros and the routines used by the calculation record type `calcRecord`, access security, and other code, to compile and evaluate mathematical expressions. The syntax of the infix expressions accepted is described below.

```
long postfix(const char *psrc, char *ppostfix, short *perror);
long calcArgUsage(const char *ppostfix, unsigned long *pinputs, unsigned long *pstores);
const char * calcErrorStr(short error);
long calcPerform(double *parg, double *presult, const char *ppostfix);
```

The `postfix()` routine converts an expression from infix to postfix notation. It is the callers's responsibility to ensure that `ppostfix` points to sufficient storage to hold the postfix expression. The macro `INFIX_TO_POSTFIX_SIZE(n)` can be used to calculate an appropriate postfix buffer size from the length of the infix buffer (note that `n` must count the terminating nil byte too).

There is no maximum length to the input expression that can be accepted, although there are internal limits to the complexity of the expressions that can be converted and evaluated. If `postfix()` returns a non-zero value it will have placed an error code at the location pointed to by `perror`.

The error codes used are defined in `postfix.h` as a series of macros with names starting `CALC_ERR_`, but a string representation of the error code is more useful and can be obtained by passing the value to the `calcErrorStr()` routine, which returns a static error message string explaining the error.

Software using the `calc` subsystem may need to know what expression arguments are used and/or modified by a particular expression. It can discover this from the postfix string by calling `calcArgUsage()`, which takes two pointers `pinputs` and `pstores` to a pair of unsigned long bitmaps which return that information to the caller. Passing a NULL value for either of these pointers is legal if only the other is needed.

The least significant bit (bit 0) of the bitmap at **pinputs* will be set if the expression depends on the argument A, and so on through bit 11 for the argument L. Similarly, bit 0 of the bitmap at **pstores* will be set if the expression assigns a value to the argument A. An argument that is not used until after a value has been assigned to it will not be set in the *pinputs* bitmap, thus the bits can be used to determine whether a value needs to be supplied for their associated argument or not for the purposes of evaluating the expression. The return value from `calcArgUsage()` will be non-zero if the *ppostfix* expression was illegal, otherwise 0.

The postfix expression is evaluated by calling the `calcPerform()` routine, which returns the status values 0 for OK, or non-zero if an error is discovered during the evaluation process.

The arguments to `calcPerform()` are:

parg - Pointer to an array of double values for the arguments A-L that can appear in the expression. Note that the argument values may be modified if the expression uses the assignment operator.

presult - Where to put the calculated result, which may be a NaN or Infinity.

ppostfix - The postfix expression created by `postfix()`.

19.2.1 Infix Expression Syntax

The infix expressions that can be used are very similar to the C expression syntax, but with some additions and subtle differences in operator meaning and precedence. The string may contain a series of expressions separated by a semi-colon character ‘;’ any one of which may actually provide the calculation result; however all of the other expressions included must assign their result to a variable. All alphabetic elements described below are case independent, so upper and lower case letters may be used and mixed in the variable and function names as desired. Spaces may be used anywhere within an expression except between the characters that make up a single expression element.

19.2.1.1 Numeric Literals

The simplest expression element is a numeric literal, any (positive) number expressed using the standard floating point syntax that can be stored as a double precision value. This now includes the values `Infinity` and `NaN` (not a number). Note that negative numbers will be encoded as a positive literal to which the unary negate operator is applied.

Examples:

```
1
2.718281828459
Inf
```

19.2.1.2 Constants

There are three trigonometric constants available to any expression which return a value:

- `pi` returns the value of the mathematical constant π .
- `D2R` evaluates to $\pi/180$ which, when used as a multiplier, converts an angle from degrees to radians.
- `R2D` evaluates to $180/\pi$ which as a multiplier converts an angle from radians to degrees.

19.2.1.3 Variables

Variables are used to provide inputs to an expression, and are named using the single letters A through L inclusive or the keyword `VAL` which refers to the previous result of this calculation. The software that makes use of the expression evaluation code should document how the individual variables are given values; for the `calc` record type the input links

INPA through INPL can be used to obtain these from other record fields, and VAL refers to the the VAL field (which can be overwritten from outside the record via Channel Access or a database link).

19.2.1.4 Variable Assignment Operator

Recently added is the ability to assign the result of a sub-expression to any of the single letter variables, which can then be used in another sub-expression. The variable assignment operator is the character pair := and must immediately follow the name of the variable to receive the expression value. Since the infix string must return exactly one value, every expression string must have exactly one sub-expression that is not an assignment, which can appear anywhere in the string. Sub-expressions within the string are separated by a semi-colon character.

Examples:

```
B; B:=A
i:=i+1; a*sin(i*D2R)
```

19.2.1.5 Arithmetic Operators

The usual binary arithmetic operators are provided: + - * and / with their usual relative precedence and left-to-right associativity, and - may also be used as a unary negate operator where it has a higher precedence and associates from right to left. There is no unary plus operator, so numeric literals cannot begin with a + sign.

Examples:

```
a*b + c
a/-4 - b
```

Three other binary operators are also provided: % is the integer modulo operator, while the synonymous operators ** and ^ raise their left operand to the power of the right operand. % has the same precedence and associativity as * and /, while the power operators associate left-to-right and have a precedence in between * and unary minus.

Examples:

```
e:=a%10; d:=a/10%10; c:=a/100%10; b:=a/1000%10; b*4096+c*256+d*16+e
sqrt(a**2 + b**2)
```

19.2.1.6 Algebraic Functions

Various algebraic functions are available which take parameters inside parentheses. The parameter separator is a comma.

- Absolute value: `abs(a)`
- Exponential e^a : `exp(a)`
- Logarithm, base 10: `log(a)`
- Natural logarithm (base e): `ln(a)` *or* `loge(a)`
- n parameter maximum value: `max(a, b, ...)`
- n parameter minimum value: `min(a, b, ...)`
- Square root: `sqr(a)` *or* `sqrt(a)`

19.2.1.7 Trigonometric Functions

Standard circular trigonometric functions, with angles expressed in radians:

- Sine: `sin(a)`
- Cosine: `cos(a)`
- Tangent: `tan(a)`
- Arcsine: `asin(a)`
- Arccosine: `acos(a)`
- Arctangent: `atan(a)`
- 2 parameter arctangent: `atan2(a, b)` - *Note that these arguments are the reverse of the ANSI C function, so while C would return $\arctan(a/b)$ the calc expression engine returns $\arctan(b/a)$*

19.2.1.8 Hyperbolic Trigonometry

The basic hyperbolic functions are provided, but no inverse functions (which are not provided by the ANSI C math library either).

- Hyperbolic sine: `sinh(a)`
- Hyperbolic cosine: `cosh(a)`
- Hyperbolic tangent: `tanh(a)`

19.2.1.9 Numeric Functions

The numeric functions perform operations related to the floating point numeric representation and truncation or rounding.

- Round up to next integer: `ceil(a)`
- Round down to next integer: `floor(a)`
- Round to nearest integer: `nint(a)`
- Test for infinite result: `isinf(a)`
- Test for any non-numeric values: `isnan(a, ...)`
- Test for all finite, numeric values: `finite(a, ...)`
- Random number between 0 and 1: `rndm`

19.2.1.10 Boolean Operators

These operators regard their arguments as true or false, where 0.0 is false and any other value is true.

- Boolean and: `a && b`
- Boolean or: `a || b`
- Boolean not: `!a`

19.2.1.11 Bitwise Operators

The bitwise operators convert their arguments to an integer (by truncation), perform the appropriate bitwise operation and convert back to a floating point value. Unlike in C though, `^` is *not* a bitwise exclusive-or operator.

- Bitwise and: `a & b` or `a and b`
- Bitwise or: `a | b` or `a or b`
- Bitwise exclusive or: `a xor b`
- Bitwise not (ones complement): `~a` or `not a`
- Bitwise left shift: `a << b`
- Bitwise right shift: `a >> b`

19.2.1.12 Relational Operators

Standard numeric comparisons between two values:

- Less than: `a < b`
- Less than or equal to: `a <= b`
- Equal to: `a = b` or `a == b`
- Greater than or equal to: `a >= b`
- Greater than: `a > b`
- Not equal to: `a != b` or `a # b`

19.2.1.13 Conditional Operator

Expressions can use the C conditional operator, which has a lower precedence than all of the other operators except for the assignment operator.

- *condition ? true result : false result*

Example:

```
a < 360 ? a+1 : 0
```

19.2.1.14 Parentheses

Sub-expressions can be placed within parentheses to override operator precedence rules. Parentheses can be nested to any depth, but the intermediate value stack used by the expression evaluation engine is limited to 80 results (which requires an expression at least 321 characters long to reach).

19.3 cppStd

This subdirectory of libCom is intended for facilities such as class and function templates that implement parts of the ISO standard C++ library where such facilities are not available or not efficient on all the target platforms on which EPICS is supported. EPICS does not make use of the C++ container templates because the large number of memory allocation and deletion operations that these use causes memory pool fragmentation on some platforms, threatening the lifetime of an individual IOC.

19.3.1 epicsAlgorithm

`epicsAlgorithm.h` contains a few templates that are also available in the C++ standard header `algorithm`, but are provided here in a much smaller file. `algorithm` contains many templates for sorting and searching through C++ template containers which are not used in EPICS. If all you need from there is `std::min()`, `std::max()` and/or `std::swap()` your code may compile faster if you include `epicsAlgorithm.h` and use `epicsMin()`, `epicsMax()` and `epicsSwap()` instead.

template <class T>	Meaning
<code>const T& epicsMin(const T& a, const T& b)</code>	Returns the smaller of a or b compared using <code>a<b</code> . Handles NaNs correctly.
<code>const T& epicsMax(const T& a, const T& b)</code>	Returns the larger of a or b compared using <code>a<b</code> . Handles NaNs correctly.
<code>void epicsSwap(T& a, T& b)</code>	Swaps the values of a and b; T must have a copy-constructor and operator=.

19.4 epicsExit

```
void epicsExit(int status);
void epicsExitCallAtExits(void);
void epicsAtExit(void (*epicsExitFunc)(void *arg), void *arg);
void epicsExitCallAtThreadExits(void);
int epicsAtThreadExit(void (*epicsExitFunc)(void *arg), void *arg);
```

This is an extended replacement for the Posix `exit` and `atexit` routines, which also provides a pointer argument to pass to the exit handlers. This facility was created because of problems on vxWorks and windows with the implementation of `atexit`, i.e. neither of these systems implement `exit` and `atexit` according to the POSIX standard.

Method	Meaning
<code>epicsExit</code>	This calls <code>epicsExitCallAtExits</code> and then passes <code>status</code> on to <code>exit</code> .
<code>epicsExitCallAtExits</code>	This calls each of the functions registered by prior calls to <code>epicsAtExit</code> , in reverse order of their registration. Most applications will not call this routine directly.
<code>epicsAtExit</code>	Register a function and an associated context parameter, to be called with the given parameter when <code>epicsExitCallAtExits</code> is invoked.
<code>epicsExitCallAtThreadExits</code>	This calls each of the functions that were registered by the current thread calling <code>epicsAtThreadExit</code> , in reverse order of the function registration. This routine is called automatically when an <code>epicsThread</code> 's main entry method returns, but will not be run if the thread is stopped by other means.
<code>epicsAtThreadExit</code>	Register a function and an associated context parameter. The function will be called with the given parameter when <code>epicsExitCallAtThreadExits</code> is invoked by the current thread ending normally, i.e. when the thread function returns.

19.5 cvtFast

`cvtFast.h` provides routines for converting various numeric types into an ascii string. They offer a combination of speed and convenience not available with `sprintf()`.

```
/* These functions return the number of ASCII characters generated */
int cvtFloatToString(float value, char *pstr, unsigned short precision);
int cvtDoubleToString(double value, char *pstr, unsigned short prec);
```

```

int cvtFloatToExpString(float value, char *pstr, unsigned short prec);
int cvtDoubleToExpString(double value, char *pstr, unsigned short prec);
int cvtFloatToCompactString(float value, char *pstr, unsigned short prec);
int cvtDoubleToCompactString(double value, char *pstr, unsigned short prec);
int cvtCharToString(char value, char *pstring);
int cvtUcharToString(unsigned char value, char *pstr);
int cvtShortToString(short value, char *pstr);
int cvtUshortToString(unsigned short value, char *pstr);
int cvtLongToString(epicsInt32 value, char *pstr);
int cvtUlongToString(epicsUInt32 value, char *pstr);
int cvtLongToHexString(epicsInt32 value, char *pstr);
int cvtLongToOctalString(epicsInt32 value, char *pstr);
unsigned long cvtBitsToUlong(
    epicsUInt32 src,
    unsigned bitFieldOffset,
    unsigned bitFieldLength);
unsigned long cvtUlongToBits(
    epicsUInt32 src,
    epicsUInt32 dest,
    unsigned bitFieldOffset,
    unsigned bitFieldLength);

```

19.6 cxxTemplates

This directory contains the following C++ template headers:

- resourceLib.h - A C++ hash facility that implements the same functionality as bucketLib
- tsBTree.h - Binary tree.
- tsDLLList.h - Double Linked List
- tsFreeList.h - Free List for efficient new/delete
- tsMinMax.h - min and max.
- tsSLLList.h - Single Linked List

Currently these are only being used by Channel Access Clients and the portable Channel Access Server. It has not been decided if any of these will remain in libCom.

19.7 dbmf

dbmf.h (Database Macro/Free) describes a facility that prevents memory fragmentation when memory is allocated and then freed a short time later.

Routines within iocCore like dbLoadDatabase() have the following attributes:

- They repeatedly call malloc() followed soon afterwards by a call to free() the temporarily allocated storage.
- Between those calls to malloc() and free(), an additional call to malloc() is made that does NOT have an associated free().

In some environments, e.g. vxWorks 5.x, this behavior causes severe memory fragmentation.

The dbmf facility stops the memory fragmentation. It should NOT be used by code that allocates storage and then keeps it for a considerable period of time before releasing. Such code can use the freeList library described below.

```

int dbmfInit(size_t size, int chunkItems);
void *dbmfMalloc(size_t bytes);
void dbmfFree(void* bytes);
void dbmfFreeChunks(void);
int dbmfShow(int level);

```

Routine	Meaning
dbmfInit()	Initialize the facility. Each time malloc() must be called size*chunkItems bytes are allocated. size is the maximum size request from dbmfMalloc() that will be allocated from the dbmf pool. If dbmfInit() was not called before one of the other routines then it is automatically called with size=64 and chunkItems=10.
dbmfMalloc()	Allocate memory. If bytes is > size then malloc() is used to allocate the memory.
dbmfFree()	Free the memory allocated by dbmfMalloc().
dbmfFreeChunks()	Free all chunks that have contain only free items.
dbmfShow()	Show the status of the dbmf memory pool.

19.8 ellLib

ellLib.h describes a double linked list library. It provides functionality similar to the vxWorks lstLib library. See the vxWorks documentation for details. There is an ellXXX() routine to replace most vxWorks lstXXX() routines.

```

typedef struct ELLNODE {
    struct ELLNODE *next;
    struct ELLNODE *previous;
} ELLNODE;

typedef void (*FREEFUNC)(void *);

typedef struct ELLLIST {
    ELLNODE node;
    int count;
} ELLLIST;

void ellInit (ELLLIST *pList);
int ellCount (ELLLIST *pList);
ELLNODE *ellFirst (ELLLIST *pList);
ELLNODE *ellLast (ELLLIST *pList);
ELLNODE *ellNext (ELLNODE *pNode);
ELLNODE *ellPrevious (ELLNODE *pNode);
void ellAdd (ELLLIST *pList, ELLNODE *pNode);
void ellConcat (ELLLIST *pDstList, ELLLIST *pAddList);
void ellDelete (ELLLIST *pList, ELLNODE *pNode);
void ellExtract (ELLLIST *pSrcList, ELLNODE *pStartNode,
    ELLNODE *pEndNode, ELLLIST *pDstList);
ELLNODE *ellGet (ELLLIST *pList);
void ellInsert (ELLLIST *plist, ELLNODE *pPrev, ELLNODE *pNode);
ELLNODE *ellNth (ELLLIST *pList, int nodeNum);
ELLNODE *ellNStep (ELLNODE *pNode, int nStep);
int ellFind (ELLLIST *pList, ELLNODE *pNode);
void ellFree2 (ELLLIST *pList, FREEFUNC freeFunc);
void ellFree (ELLLIST *pList); // Use only if freeFunc is free()
void ellVerify (ELLLIST *pList);

```

19.9 epicsRingBytes

`epicsRingBytes.h` contains

```
epicsRingBytesId epicsRingBytesCreate(int nbytes);
void epicsRingBytesDelete(epicsRingBytesId id);
int epicsRingBytesGet(epicsRingBytesId id, char *value, int nbytes);
int epicsRingBytesPut(epicsRingBytesId id, char *value, int nbytes);
void epicsRingBytesFlush(epicsRingBytesId id);
int epicsRingBytesFreeBytes(epicsRingBytesId id);
int epicsRingBytesUsedBytes(epicsRingBytesId id);
int epicsRingBytesSize(epicsRingBytesId id);
int epicsRingBytesIsEmpty(epicsRingBytesId id);
int epicsRingBytesIsFull(epicsRingBytesId id);
```

Method	Meaning
<code>epicsRingBytesCreate()</code>	Create a new ring buffer of size <code>nbytes</code> . The returned <code>epicsRingBytesId</code> is passed to the other ring methods.
<code>epicsRingBytesDelete()</code>	Delete the ring buffer and free any associated memory.
<code>epicsRingBytesGet()</code>	Move up to <code>nbytes</code> from the ring buffer to <code>value</code> . The number of bytes actually moved is returned.
<code>epicsRingBytesPut()</code>	Move <code>nbytes</code> from <code>value</code> to the ring buffer if there is enough free space available to hold them. The number of bytes actually moved is returned, which will be zero if insufficient space exists.
<code>epicsRingBytesFlush()</code>	Make the ring buffer empty.
<code>epicsRingBytesFreeBytes()</code>	Return the number of free bytes in the ring buffer.
<code>epicsRingBytesUsedBytes()</code>	Return the number of bytes currently stored in the ring buffer.
<code>epicsRingBytesSize()</code>	Return the size of the ring buffer, i.e., <code>nbytes</code> specified in the call to <code>epicsRingBytesCreate()</code> .
<code>epicsRingBytesIsEmpty()</code>	Return (true, false) if the ring buffer is currently empty.
<code>epicsRingBytesIsFull()</code>	Return (true, false) if the ring buffer is currently empty.

`epicsRingBytes` has the following properties:

- For a ring buffer with a single writer it is not necessary to lock `epicsRingBytesPut()` calls.
- For a ring buffer with a single reader it is not necessary to lock `epicsRingBytesGet()` calls.
- `epicsRingBytesFlush()` should only be used when both gets and puts are locked out.

19.10 epicsRingPointer

`epicsRingPointer.h` describes a C++ and a C facility for a commonly used type of ring buffer.

19.10.1 C++ Interface

`EpicsRingPointer` provides methods for creating and using ring buffers (first in first out circular buffers) that store pointers. It is designed so that a writer thread and reader thread can access the ring simultaneously without requiring mutual exclusion.

```
template <class T>
class epicsRingPointer {
public:
```

```

    epicsRingPointer(int size);
    ~epicsRingPointer();
    bool push(T *p);
    T* pop();
    void flush();
    int getFree() const;
    int getUsed() const;
    int getSize() const;
    bool isEmpty() const;
    bool isFull() const;

private: // Prevent compiler-generated member functions
    // default constructor, copy constructor, assignment operator
    epicsRingPointer();
    epicsRingPointer(const epicsRingPointer &);
    epicsRingPointer& operator=(const epicsRingPointer &);

private: // Data
    ...
};

```

An `epicsRingPointer` cannot be assigned to, copy-constructed, or constructed without giving the *size* argument. The C++ compiler will object to some of the statements below:

```

    epicsRingPointer rp0(); // Error: default constructor is private
    epicsRingPointer rp1(10); // OK
    epicsRingPointer rp2(t1); // Error: copy constructor is private
    epicsRingPointer *prp; // OK, pointer
    *prp = rp1; // Error: assignment operator is private
    prp = &rp1; // OK, pointer assignment and address-of

```

Method	Meaning
<code>epicsRingPointer()</code>	Constructor. The size is the maximum number of elements (pointers) that can be stored in the ring.
<code>~epicsRingPointer()</code>	Destructor.
<code>push()</code>	Push a new entry on the ring. It returns (false,true) is (failure, success). Failure means the ring was full. If a single writer is present it does not have to use a lock while performing the push. If multiple writers are present they must use a common lock while issuing the push.
<code>pop()</code>	Take a element off the ring. It returns 0 (null) if the ring was empty. If a single reader is present it does not have to lock while issuing the pop. If multiple readers are present they must use a common lock while issuing the pop.
<code>flush()</code>	Remove all elements from the ring. If this operation is performed then all access to the ring should be locked.
<code>getFree()</code>	Return the amount of empty space in the ring, i.e. how many additional elements it can hold.
<code>getUsed()</code>	Return the number of elements stored on the ring
<code>getSize()</code>	Return the size of the ring, i.e. the value of size specified when the ring was created.
<code>isEmpty()</code>	Returns true if the ring is empty, else false.
<code>isFull()</code>	Returns true if the ring is full, else false.

19.10.2 C interface

```
typedef void *epicsRingPointerId;
epicsRingPointerId epicsRingPointerCreate(int size);
void epicsRingPointerDelete(epicsRingPointerId id);
/*epicsRingPointerPop returns 0 if the ring was empty */
void * epicsRingPointerPop(epicsRingPointerId id) ;
/*epicsRingPointerPush returns (0,1) if p (was not, was) put on ring*/
int epicsRingPointerPush(epicsRingPointerId id,void *p);
void epicsRingPointerFlush(epicsRingPointerId id);
int epicsRingPointerGetFree(epicsRingPointerId id);
int epicsRingPointerGetUsed(epicsRingPointerId id);
int epicsRingPointerGetSize(epicsRingPointerId id);
int epicsRingPointerIsEmpty(epicsRingPointerId id);
int epicsRingPointerIsFull(epicsRingPointerId id);
```

Each C function corresponds to one of the C++ methods.

19.11 epicsTimer

epicsTimer.h describes a C++ and a C timer facility.

19.11.1 C++ Interface

19.11.1.1 epicsTimerNotify and epicsTimer

```
class epicsTimerNotify {
public:
enum restart_t { noRestart, restart };
class expireStatus {
public:
expireStatus ( restart_t );
expireStatus ( restart_t, const double &expireDelaySec );
bool restart () const;
double expirationDelay () const;
private:
double delay;
};
virtual ~epicsTimerNotify ();
// return noRestart OR return expireStatus ( restart, 30.0 /* sec */ );
virtual expireStatus expire ( const epicsTime & currentTime ) = 0;
virtual void show ( unsigned int level ) const;
};

class epicsTimer {
public:
virtual void destroy () = 0; // requires existence of timer queue
virtual void start ( epicsTimerNotify &, const epicsTime & ) = 0;
virtual void start ( epicsTimerNotify &, double delaySeconds ) = 0;
virtual void cancel () = 0;
struct expireInfo {
expireInfo ( bool active, const epicsTime & expireTime );
```

```

        bool active;
        epicsTime expireTime;
    };
    virtual expireInfo getExpireInfo () const = 0;
    double getExpireDelay ();
    virtual void show ( unsigned int level ) const = 0;
protected:
    virtual ~epicsTimer () = 0; // use destroy
};

```

Method	Meaning
epicsTimerNotify-expire()	Code using an epicsTimer must include a class that inherits from epicsTimerNotify. The derived class must implement the method expire(), which is called by the epicsTimer when the associated timer expires. epicsTimerNotify defines a class expireStatus which makes it easy to implement both one shot and periodic timers. A one-shot expire() returns with the statement: return(noRestart);A periodic timer returns with a statement like: return(restart,10.0);where is second argument is the delay until the next callback.
epicsTimer	epicsTimer is an abstract base class. An epics timer can only be created by calling createTimer, which is a method of epicsTimerQueue.
destroy	This is provided instead of a destructor. This will automatically call cancel before freeing all resources used by the timer.
start()	Starts the timer to expire either at the specified time or the specified number of seconds in the future. If the timer is already active when start is called, it is first canceled.
cancel()	If the timer is scheduled, cancel it. If it is not scheduled do nothing. Note that if the expire() method is already running, this call delays until the expire() completes.
getExpireInfo	Get expireInfo, which says if timer is active and if so when it expires.
getExpireDelay()	Return the number of seconds until the timer will expire. If the timer is not active it returns DBL_MAX
show()	Display info about object.

19.11.1.2 epicsTimerQueue

```

class epicsTimerQueue {
public:
    virtual epicsTimer & createTimer () = 0;
    virtual void show ( unsigned int level ) const = 0;
protected:
    virtual ~epicsTimerQueue () = 0;
};

```

Method	Meaning
createTimer()	This is a “factory” method to create timers which use this queue.
show()	Display info about object

19.11.1.3 epicsTimerQueueActive

```

class epicsTimerQueueActive : public epicsTimerQueue {
public:

```



```

    static epicsTimerQueueActive & allocate (
        bool okToShare, unsigned threadPriority = epicsThreadPriorityMin + 10 );
    virtual void release () = 0;
protected:
    virtual ~epicsTimerQueueActive () = 0;
};

```

Method	Meaning
allocate()	This is a “factory” method to create a timer queue. If okToShare is (true,false) then a (shared, separate) thread will manage the timer requests.If the okToShare constructor parameter is true then if a timer queue is already running at the specified priority then it will be referenced for shared use by the application, and an independent timer queue will not be created.This method should not be called from within a C++ static constructor, since the queue thread requires that a current time provider be available and the last-resort time provider is not guaranteed to have been registered until all constructors have run.Editorial note: It is useful for two independent timer queues to run at the same priority if there are multiple processors, or if there is an application with well behaved timer expire functions that needs to be independent of applications with computationally intensive, mutex locking, or IO blocking timer expire functions.
release()	Release the queue, i.e. the calling facility will no longer use the queue. The caller MUST ensure that it does not own any active timers. When the last facility using the queue calls release, all resources used by the queue are freed.

19.11.1.4 epicsTimerQueueNotify and epicsTimerQueuePassive

These two classes manage a timer queue for single threaded applications. Since it is single threaded, the application is responsible for requesting that the queue be processed.

```

class epicsTimerQueueNotify {
public:
    // called when a new timer is inserted into the queue and the
    // delay to the next expire has changed
    virtual void reschedule () = 0;
    // if there is a quantum in the scheduling of timer intervals
    // return this quantum in seconds. If unknown then return zero.
    virtual double quantum () = 0;
protected:
    virtual ~epicsTimerQueueNotify () = 0;
};

class epicsTimerQueuePassive {
public:
    static epicsTimerQueuePassive & create ( epicsTimerQueueNotify & );
    virtual ~epicsTimerQueuePassive () = 0;
    // process returns the delay to the next expire
    virtual double process (const epicsTime & currentTime) = 0;
};

```

Method	Meaning
epicsTimerQueueNotify-reschedule()	The virtual function epicsTimerQueueNotify::reschedule() is called when the delay to the next timer to expire on the timer queue changes.

<code>epicsTimerQueueNotify-quantum</code>	The virtual function <code>epicsTimerQueueNotify::quantum()</code> returns the timer expire interval scheduling quantum in seconds. This allows different types of timer queues to use application specific timer expire delay scheduling policies. The implementation of <code>epicsTimerQueueActive</code> employs <code>epicsThreadSleep()</code> for this purpose, and therefore <code>epicsTimerQueueActive::quantum()</code> returns the returned value from <code>epicsThreadSleepQuantum()</code> . Other types of timer queues might choose to schedule timer expiration using specialized hardware interrupts. In this case <code>epicsTimerQueueNotify::quantum()</code> might return a value reflecting the precision of a hardware timer. If unknown, then <code>epicsTimerQueueNotify::quantum()</code> should return zero.
<code>epicsTimerQueuePassive</code>	<code>epicsTimerQueuePassive</code> is an abstract base class so cannot be instantiated directly, but contains a static member function to create a concrete passive timer queue object of a (hidden) derived class.
<code>create()</code>	A “factory” method to create a non-threaded timer queue. The calling software also passes an object derived from <code>epicsTimerQueueNotify</code> to receive <code>reschedule()</code> callbacks.
<code>~epicsTimerQueuePassive()</code>	Destructor. The caller MUST ensure that it does not own any active timers, i.e. it must cancel any active timers before deleting the <code>epicsTimerQueuePassive</code> object.
<code>process()</code>	This calls <code>expire()</code> for all timers that have expired. The facility that creates the queue MUST call this. It returns the delay until the next timer will expire.

19.11.2 C Interface

```

typedef struct epicsTimerForC * epicsTimerId;
typedef void ( *epicsTimerCallback ) ( void *pPrivate );

/* thread managed timer queue */
typedef struct epicsTimerQueueActiveForC * epicsTimerQueueId;
epicsTimerQueueId epicsTimerQueueAllocate(
    int okToShare, unsigned int threadPriority );
void epicsTimerQueueRelease ( epicsTimerQueueId );
epicsTimerId epicsTimerQueueCreateTimer ( epicsTimerQueueId queueid,
    epicsTimerCallback callback, void *arg );
void epicsTimerQueueDestroyTimer ( epicsTimerQueueId queueid, epicsTimerId id );
void epicsTimerQueueShow ( epicsTimerQueueId id, unsigned int level );

/* passive timer queue */
typedef struct epicsTimerQueuePassiveForC * epicsTimerQueuePassiveId;
typedef void ( *epicsTimerQueueNotifyReschedule ) ( void *pPrivate );
typedef double ( * epicsTimerQueueNotifyQuantum ) ( void * pPrivate );
epicsTimerQueuePassiveId epicsTimerQueuePassiveCreate(
    epicsTimerQueueNotifyReschedule, epicsTimerQueueNotifyQuantum,
    void *pPrivate );
void epicsTimerQueuePassiveDestroy ( epicsTimerQueuePassiveId );
epicsTimerId epicsTimerQueuePassiveCreateTimer ( epicsTimerQueuePassiveId queueid,
    epicsTimerCallback pCallback, void *pArg );
void epicsTimerQueuePassiveDestroyTimer (
    epicsTimerQueuePassiveId queueid, epicsTimerId id );
double epicsTimerQueuePassiveProcess ( epicsTimerQueuePassiveId );
void epicsTimerQueuePassiveShow( epicsTimerQueuePassiveId id, unsigned int level);
/* timer */
void epicsTimerStartTime( epicsTimerId id, const epicsTimeStamp *pTime);
void epicsTimerStartDelay( epicsTimerId id, double delaySeconds);

```

```
void epicsTimerCancel ( epicsTimerId id );
double epicsTimerGetExpireDelay ( epicsTimerId id );
void epicsTimerShow ( epicsTimerId id, unsigned int level );
```

The C interface provides most of the facilities as the C++ interface. It does not support the periodic timer features. The typedefs `epicsTimerQueueNotifyReschedule` and `epicsTimerQueueNotifyQuantum` are the “C” interface equivalents to `epicsTimerQueueNotify::reschedule()` and `epicsTimerQueueNotify::quantum()`.

19.11.3 Example

This example allocates a timer queue and two objects which have a timer that uses the queue. Each object is requested to schedule itself. The `expire()` callback just prints the name of the object. After scheduling each object the main thread just sleeps long enough for each expire to occur and then just returns after releasing the queue.

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

class something : public epicsTimerNotify {
public:
    something(const char* nm,epicsTimerQueueActive &queue)
        : name(nm), timer(queue.createTimer()) {}
    virtual ~something() { timer.destroy();}
    void start(double delay) {timer.start(*this,delay);}
    virtual expireStatus expire(const epicsTime & currentTime) {
        printf("%s\n",name);
        currentTime.show(1);
        return(noRestart);
    }
private:
    const char* name;
    epicsTimer &timer;
};

void epicsTimerExample()
{
    epicsTimerQueueActive &queue = epicsTimerQueueActive::allocate(true);
    {
        something first("first",queue);
        something second("second",queue);

        first.start(1.0);
        second.start(1.5);
        epicsThreadSleep(2.0);
    }
    queue.release();
}
```

19.11.4 C Example

This example shows how C programs can use EPICS timers.

```
#include <stdio.h>
#include <epicsTimer.h>
#include <epicsThread.h>
```

```

static void
handler (void *arg)
{
    printf ("%s timer tripped.\n", (char *)arg);
}

int
main(int argc, char **argv)
{
    epicsTimerQueueId timerQueue;
    epicsTimerId first, second;

    /*
     * Create the queue of timer requests
     */
    timerQueue = epicsTimerQueueAllocate(1, epicsThreadPriorityScanHigh);

    /*
     * Create the timers
     */
    first = epicsTimerQueueCreateTimer(timerQueue, handler, "First");
    second = epicsTimerQueueCreateTimer(timerQueue, handler, "Second");

    /*
     * Start a timer
     */
    printf("First timer should trip in 3 seconds.\n");
    epicsTimerStartDelay(first, 3.0);
    epicsThreadSleep(5.0);
    printf("First timer should have tripped by now.\n");

    /*
     * Try starting and then cancelling a request
     */
    printf("Second timer should trip in 3 seconds.\n");
    epicsTimerStartDelay(first, 3.0);
    epicsTimerStartDelay(second, 3.0);
    epicsThreadSleep(1.0);
    epicsTimerCancel(first);
    epicsThreadSleep(5.0);
    printf("Second timer should have tripped, first timer should not have
tripped.\n");

    /*
     * Clean up a single timer
     */
    epicsTimerQueueDestroyTimer(timerQueue, first);

    /*
     * Clean up an entire queue of timers
     */
    epicsTimerQueueRelease(timerQueue);
    return 0;
}

```

19.12 fdmgr

File Descriptor Manager. `fdManager.h` describes a C++ implementation. `fdmgr.h` describes a C implementation. Neither is currently documented.

19.13 freeList

`freeList.h` describes routines to allocate and free fixed size memory elements. Free elements are maintained on a free list rather than being returned to the heap via calls to `free`. When it is necessary to call `malloc()`, memory is allocated in multiples of the element size.

```
void freeListInitPvt(void **ppvt, int size, int nmalloc);
void *freeListCalloc(void *pvt);
void *freeListMalloc(void *pvt);
void freeListFree(void *pvt, void *pmem);
void freeListCleanup(void *pvt);
size_t freeListItemsAvail(void *pvt);
```

where

pvt - For internal use by the freelist library. Caller must provide storage for a “void *pvt”

size - Size in bytes of each element. Note that all elements must be same size

nmalloc - Number of elements to allocate when regular `malloc()` must be called.

19.14 gpHash

`gpHash.h` describes a general purpose hash table for character strings. The hash table contains *tableSize* entries. Each entry is a list of members that hash to the same value. The user can maintain separate directories which share the same table by having a different *pvt* value for each directory.

```
typedef struct{
    ELLNODE    node;
    const char *name;           /*address of name placed in directory*/
    void       *pvtid;         /*private name for subsystem user*/
    void       *userPvt;      /*private for user*/
} GPENTRY;

struct gphPvt;

/*tableSize must be power of 2 in range 256 to 65536*/
void gphInitPvt(struct gphPvt **ppvt, int tableSize);
GPENTRY *gphFind(struct gphPvt *pvt, const char *name, void *pvtid);
GPENTRY *gphAdd(struct gphPvt *pvt, const char *name, void *pvtid);
void gphDelete(struct gphPvt *pvt, const char *name, void *pvtid);
void gphFreeMem(struct gphPvt *pvt);
void gphDump(struct gphPvt *pvt);
void gphDumpFP(FILE *fp, struct gphPvt *pvt);
```

where

pvt - For internal use by the gpHash library. Caller must provide storage for a `struct gphPvt *pvt`

name - The character string that will be hashed and added to table.

pvtid - The name plus the value of this pointer constitute a unique entry.

19.15 logClient

Together with the program `iocLogServer` this provides generic support for logging text messages from an IOC or other program to a file on the log server host machine.

A log client runs on the IOC. It accepts string messages and forwards them over a TCP connection to its designated log server (normally running on a host machine).

A log server accepts connections from multiple clients and writes the messages it receives into a rotating file. A log server program (`'iocLogServer'`) is also part of EPICS base.

Configuration of the `iocLogServer`, as well as the standard `iocLogClient` that internally uses this library, are described in Section 10.7 on page 167.

The header file `logClient.h` exports the following types and routines:

```
typedef void *logClientId;
```

An abstract data type, representing a log client.

```
logClientId logClientCreate (
    struct in_addr server_addr, unsigned short server_port);
```

Create a new log client. Will block the calling task for a maximum of 2 seconds trying to connect to a server with the given ip address and port. If a connection cannot be established, an error message is printed on the console, but the log client will keep trying to connect in the background. This is done by a background task, that will also periodically (every 5 seconds) flush pending messages out to the server.

```
void logClientSend (logClientId id, const char *message);
```

Send the given message to the given log client. Messages are not immediately sent to the log server. Instead they are sent whenever the cache overflows, or `logClientFlush()` is called.

```
void logClientFlush (logClientId id);
```

Immediately send all outstanding messages to the server.

```
void logClientShow (logClientId id, unsigned level);
```

Print information about the log clients internal state to stdout.

For backward compatibility with older versions of the `logClient` library, the header file also includes `iocLog.h`, which exports the definitions for the standard `iocLogClient` for error logging. See Chapter 10.7.2.

Also for backward compatibility, the following deprecated routines are exported.

```
logClientId logClientInit (void);
```

Create a log client that uses the standard ioc log client environment variables (`EPICS_IOC_LOG_INET` and `EPICS_IOC_LOG_PORT`) as input to `logClientCreate` and also registers the log client with the `errlog` task using `errlogAddListener`.

```
void logClientSendMessage (logClientId id, const char *message);
```

Check the global variable `iocLogDisable` before calling `logClientSend`.

19.16 macLib

`macLib.h` describes a general purpose macro substitution library. It is used for all macro substitution in base.

```

long macCreateHandle(
    MAC_HANDLE **handle,      /* address of variable to receive pointer */
                                /* to new macro substitution context */
    char *pairs[]            /* pointer to NULL-terminated array of */
                                /* {name,value} pair strings; a NULL */
                                /* value implies undefined; a NULL */
                                /* argument implies no macros */
);

void macSuppressWarning(
    MAC_HANDLE *handle,      /* opaque handle */
    int falseTrue          /* *0 means ussue, 1 means suppress*/
);

/*following returns #chars copied, <0 if any macros are undefined*/
long macExpandString(
    MAC_HANDLE *handle,      /* opaque handle */
    char *src,              /* source string */
    char *dest,             /* destination string */
    long maxlen             /* maximum number of characters to copy */
                                /* to destination string */
);

/*following returns length of value */
long macPutValue(
    MAC_HANDLE *handle,      /* opaque handle */
    char *name,             /* macro name */
    char *value             /* macro value */
);

/*following returns #chars copied (<0 if undefined) */
long macGetValue(
    MAC_HANDLE *handle,      /* opaque handle */
    char *name,             /* macro name or reference */
    char *value,            /* string to receive macro value or name */
                                /* argument if macro is undefined */
    long maxlen             /* maximum number of characters to copy */
                                /* to value */
);

long macDeleteHandle(MAC_HANDLE *handle);
long macPushScope(MAC_HANDLE *handle);
long macPopScope(MAC_HANDLE *handle);
long macReportMacros(MAC_HANDLE *handle);

/* Function prototypes (utility library) */

/*following returns #defs encountered; <0 = ERROR */
long macParseDefs(
    MAC_HANDLE *handle,      /* opaque handle; can be NULL if default */
                                /* special characters are to be used */
    char *defs,              /* macro definitions in "a=xxx,b=yyy" */
                                /* format */
);

```

```

    char        **pairs[]        /* address of variable to receive pointer */
                                /* to NULL-terminated array of {name, */
                                /* value} pair strings; all storage is */
                                /* allocated contiguously */

);

/*following returns #macros defined; <0 = ERROR */
long macInstallMacros(MAC_HANDLE *handle,
    char        *pairs[]        /* pointer to NULL-terminated array of */
                                /* {name,value} pair strings; a NULL */
                                /* value implies undefined; a NULL */
                                /* argument implies no macros */

);

/*Expand string using environment variables as macro definitions */
epicsShareFunc char *          /* expanded string; NULL if any undefined macros */
epicsShareAPI macEnvExpand(
    char *str                  /* string to be expanded */

);

```

NOTE: The directory <base>/src/libCom/macLib contains two files macLibNOTES and macLibREADME that explain this library.

19.17 misc

19.17.1 aToIPAddr

The function prototype for this routine appears in `osiSock.h`

```

int aToIPAddr(const char *pAddrString, unsigned short defaultPort,
              struct sockaddr_in *pIP);

```

`aToIPAddr()` fills in the structure pointed to by the *pIP* argument with the Internet address and port number specified by the *pAddrString* argument.

Three forms of *pAddrString* are accepted:

1. n.n.n.n:p

The Internet address of the host, specified as four (usually decimal) numbers separated by periods.

2. xxxxxxxx:p

The Internet address number of the host, specified as a single (usually hexadecimal) number.

3. hostname:p

The Internet host name of the host.

In all cases the ‘:p’ may be omitted in which case the port number is set to the value of the *defaultPort* argument. The numbers are normally interpreted in base 16 if they begin with ‘0x’ or ‘0X’, in base 8 if they begin with ‘0’, and in base 10 otherwise. However the numeric forms are interpreted by the operating system’s `gethostbyname()` function, thus the acceptable bases may be OS-specific.

19.17.2 adjustment

`adjustment.h` describes a single function:


```
size_t adjustToWorstCaseAlignment(size_t size);
```

`adjustToWorstCaseAlignment()` returns a value \geq *size* that an exact multiple of the worst case alignment for the architecture on which the routine is executed.

19.17.3 cantProceed

`cantProceed.h` declares routines that are provided for code that can't proceed when an error occurs.

```
void cantProceed(const char *errorMessage, ...);
void *callocMustSucceed(size_t count, size_t size, const char *errorMessage);
void *mallocMustSucceed(size_t size, const char *errorMessage);
```

`cantProceed()` accepts a `printf` format string and variable number of arguments; it displays the error message and suspends the current task. It will never return. `callocMustSucceed()` and `mallocMustSucceed()` can be used in place of `calloc()` and `malloc()`. If `size` or `count` are zero, or the memory allocation fails, they output a message and call `cantProceed()`.

19.17.4 dbDefs

`dbDefs.h` includes the C header `stddef.h` and then defines several generally-useful macros if they have not already been defined:

- `TRUE` - 1
- `FALSE` - 0
- `NELEMENTS(array)` - number of elements in array.
- `CONTAINER(pointer, structure, member)` - returns a pointer to the parent structure given a pointer to a member. The `structure` argument is a type name, `member` is the name of the member in that structure that `pointer` refers to.
- `LOCAL` - synonym for `static`, deprecated.
- `OFFSET(structure, member)` - synonym for `offsetof`, deprecated.

19.17.5 epicsConvert

`epicsConvert.h` currently describes:

```
float epicsConvertDoubleToFloat(double value);
```

`epicsConvertDoubleToFloat` converts a double to a float. If the double value is outside the range that can be represented as a float the value assigned will be `FLT_MIN` or `FLT_MAX` with the appropriate matching sign. A floating exception is never raised.

19.17.6 epicsString

`epicsString.h` currently describes:

```
int epicsStrnRawFromEscaped(char *outbuf, size_t outsize, const char *inbuf,
    size_t inlen);
int epicsStrnEscapedFromRaw(char *outbuf, size_t outsize, const char *inbuf,
    size_t inlen);
size_t epicsStrnEscapedFromRawSize(const char *inbuf, size_t inlen);
```

```

int epicsStrCaseCmp(const char *s1, const char *s2);
int epicsStrnCaseCmp(const char *s1, const char *s2, int n);
char *epicsStrDup(const char *s);
int epicsStrPrintEscaped(FILE *fp, const char *s, int n);
int epicsStrGlobMatch(const char *str, const char *pattern);
char *epicsStrtok_r(char *s, const char *delim, char **lasts);
unsigned int epicsStrHash(const char *str, unsigned int seed);
unsigned int epicsMemHash(const char *str, size_t length,
    unsigned int seed);

```

`epicsStrnRawFromEscaped` copies the string *inbuf* to *outbuf* while substituting C-style escape sequences. It returns the length of the resultant string (which may contain null bytes). It will never write more than *outsize* characters into *outbuf* or read more than *inlen* characters from *inbuf*. Since the destination string can never be longer than the source, it is legal for *inbuf* and *outbuf* to point to the same location and *inlen* and *outsize* to have the same value, thus performing the translation in-place.

`epicsStrnEscapedFromRaw` does the opposite of `epicsStrnRawFromEscaped`: non-printable characters are converted into escape sequences. It returns the number of characters stored in the output buffer, or a value greater than or equal to *outsize* if the output has been truncated to fit in the specified size. As the destination string might be larger than the source, in-place translations are not allowed.

The following C-style escaped character constants are output:

```
\a \b \f \n \r \t \v \\ \' \"
```

All other non-printable characters appear as:

```
\ooo
```

where `ooo` comprises three octal digits (0-7).

`epicsStrnEscapedFromRawSize` scans a string that may contain non-printable characters and returns the number of characters needed for the escaped destination buffer.

`epicsStrPrintEscaped` prints the contents of its input buffer, substituting escape sequences for non-printable characters.

`epicsStrCaseCmp` and `epicsStrnCaseCmp` implement `strcasecmp` and `strncasecmp`, respectively, which are not available on all operating systems. They operate like `strcmp` and `strncmp`, but are case insensitive.

`epicsStrDup` implements `strdup`, which is not available on all operating systems. It allocates sufficient memory for the string, copies it and returns a pointer to the new copy. The pointer should eventually be passed to the function `free()`. If insufficient memory is available `cantProceed()` is called.

`epicsStrGlobMatch` returns non-zero if the `str` matches the shell wild-card pattern.

`epicsStrtok_r` implements `strtok_r`, which is not available on all operating systems.

`epicsStrHash` calculates a hash of a zero-terminated string *str*, while `epicsMemHash` uses the same algorithm on a fixed-length memory buffer that may contain zero bytes. In both cases an initial *seed* value may be provided which permits multiple strings or buffers to be combined into a single hash result. The final result should be masked to achieve the desired number of bits in the hash value.

19.17.7 epicsTypes

`epicsTypes.h` provides typedefs for architecture independent data types.

```

typedef char          epicsInt8;
typedef unsigned char epicsUInt8;
typedef short        epicsInt16;

```

```

typedef unsigned short  epicsUInt16;
typedef epicsUInt16    epicsEnum16;
typedef int             epicsInt32;
typedef unsigned       epicsUInt32;
typedef float          epicsFloat32;
typedef double         epicsFloat64;
typedef unsigned long  epicsIndex;
typedef epicsInt32     epicsStatus;

```

So far the definitions provided in this header file have worked on all architectures. In addition to the above definitions `epicsTypes.h` has a number of definitions for displaying the types and other useful definitions. See the header file for details.

19.17.8 locationException

A C++ template for use as an exception object, used inside Channel Access. Not documented here.

19.17.9 shareLib.h

This is the header file for the “decorated names” that appear in header files, e.g.

```

#define epicsExportSharedSymbols
epicsShareFunc int epicsShareAPI a_func(int arg);

```

These are needed to properly create DLLs on windows. Read the comments in the `shareLib.h` file for a detailed description of where they should be used. Note that the `epicsShareAPI` decorator is deprecated for all new EPICS APIs and is being removed from APIs that are only used within the IOC.

19.17.10 truncateFile.h

```

enum TF_RETURN {TF_OK=0, TF_ERROR=1};
TF_RETURN truncateFile (const char *pFileName, unsigned size);

```

where

pFileName - name (and optionally path) of file

`truncateFile()` truncates the file to the specified size. `truncate()` is not used because it is not portable. It returns `TF_OK` if the file is less than size bytes or if it was successfully truncated. It returns `TF_ERROR` if the file could not be truncated.

19.17.11 unixFileName.h

Defines macros `OSI_PATH_LIST_SEPARATOR` and `OSI_PATH_SEPARATOR`

19.17.12 epicsUnitTest.h

The unit test routines make it easy for a test program to generate output that is compatible with the Test Anything Protocol and can thus be used with Perl’s automated `Test::Harness` as well as generating human-readable output. The routines detect whether they are being run automatically and print a summary of the results at the end if not.

```

void testPlan(int tests);
int testOk(int pass, const char *fmt, ...);
#define testOk1(cond) testOk(cond, "%s", #cond)
void testPass(const char *fmt, ...);
void testFail(const char *fmt, ...);
int testOkV(int pass, const char *fmt, va_list pvar);
void testSkip(int skip, const char *why)
void testTodoBegin(const char *why);
void testTodoEnd();
int testDiag(const char *fmt, ...);
void testAbort(const char *fmt, ...);
int testDone(void);

typedef int (*TESTFUNC)(void);
epicsShareFunc void testHarness(void);
epicsShareFunc void runTestFunc(const char *name, TESTFUNC func);

#define runTest(func) runTestFunc(#func, func)

```

A test program starts with a call to `testPlan()`, announcing how many tests are to be conducted. If this number is not known a value of zero can be used during development, but it is recommended that the correct value be substituted after the test program has been completed.

Individual test results are reported using any of `testOk()`, `testOk1()`, `testOkV()`, `testPass()` or `testFail()`. The `testOk()` call takes and also returns a logical pass/fail result (zero means failure, any other value is success) and a printf-like format string and arguments which describe the test. The convenience macro `testOk1()` is provided which stringifies its single condition argument, reducing the effort needed when writing test programs. The individual `testPass()` and `testFail()` routines can be used when the test program takes a different path on success than on failure, but one or other must always be called for any particular test. The `testOkV()` routine is a varargs form of `testOk()` included for internal purposes which may prove useful in some cases.

If some program condition or failure makes it impossible to run some tests, the `testSkip()` routine can be used to indicate how many tests are being omitted from the run, thus keeping the test counts correct; the constant string why is displayed as an explanation to the user (this string is not printf-like).

If some tests are expected to fail because functionality in the module under test has not yet been fully implemented, these tests may still be executed, wrapped between calls to `testTodoBegin()` and `testTodoEnd()`. `testTodoBegin()` takes a constant string indicating why these tests are not expected to succeed. This modifies the counting of the results so the wrapped tests will not be recorded as failures.

Additional information can be supplied using the `testDiag()` routine, which displays the relevant information as a comment in the result output. None of the printable strings passed to any `testXxx()` routine should contain a newline ‘\n’ character, newlines will be added by the test routines as part of the Test Anything Protocol. For multiple lines of diagnostic output, call `testDiag()` as many times as necessary.

If at any time the test program is unable to continue for some catastrophic reason, calling `testAbort()` with an appropriate message will ensure that the test harness understands this. `testAbort()` does not return, but calls the ANSI C routine `abort()` to cause the program to stop immediately.

After all of the tests have been completed, the return value from `testDone()` can be used as the return status code from the program’s `main()` routine.

On vxWorks and RTEMS, an alternative test harness can be used to run a series of tests in order and summarize the results from them all at the end just like the Perl harness does. The routine `testHarness()` is called once at the beginning of the test harness program. Each test program is run by passing its main routine name to the `runTest()` macro which expands into a call to the `runTestFunc()` routine. The last test program or the harness program itself must finish by calling `epicsExit()` which triggers the test summary mechanism to generate its result outputs (from an `epicsAtExit` callback routine).

To make it easier to create a single test program that can be built for both the embedded and workstation operating system harnesses, the header file `testMain.h` provides a convenience macro `MAIN()` that adjusts the name of the test program according to the platform it is running on: `main()` on workstations and a regular function name on embedded systems.

The following is a simple example of a test program using the `epicsUnitTest` routines:

```
#include <math.h>
#include "epicsUnitTest.h"
#include "testMain.h"

MAIN(mathTest)
{
    testPlan(3);
    testOk(sin(0.0) == 0.0, "Sine starts");
    testOk(cos(0.0) == 1.0, "Cosine continues");
    if (!testOk1(M_PI == 4.0*atan(1.0)))
        testDiag("4*atan(1) = %g", 4.0*atan(1.0));
    return testDone();
}
```

The output from running the above program looks like this:

```
1..3
ok 1 - Sine starts
ok 2 - Cosine continues
ok 3 - M_PI == 4.0*atan(1.0)
```

```
Results
=====
    Tests: 3
    Passed: 3 = 100%
```


Chapter 20

libCom OSI libraries

20.1 Overview

Most code in base is operating system independent, i.e. the code is exactly the same for all supported operating systems. This is accomplished by providing epics defined libraries for facilities that are different on the various systems. The code is called Operating System Independent or OSI. OSI libraries have multiple implementations, which are Operating System Dependent or OSD.

20.1.1 OSI source directory

Directory `<base>/src/libCom/osi` contains the code for the operating system independent libraries. The structure of this directory is:

```
osi/  
  epics*.h  
  *.cpp - A few generic c++ implementations  
  os/  
    Linux/  
    Darwin/  
    RTEMS/  
    WIN32/  
    default/  
    posix/  
    solaris/  
    vxWorks/
```

Code for additional operating systems may also be present.

20.1.2 Rules for building OSI code

The `osi` directory contains header files with names starting with `epics`. These headers provide definitions for user code. Each of the directories under `osi/<arch>` contain architecture-specific code in filenames starting with `osd`. In most cases both a header and source file are present.

Installing header files residing under `src/libCom/osi` into `<base>/include`

- Header files in `osi` are installed into `<base>/include`

- Header files in a subdirectory below `osi/os` are installed into `<base>/include/os/<arch>`. The search order for locating the specific file to be installed is:
 - `libCom/osi/os/<arch>`
 - `libCom/osi/os/posix`
 - `libCom/osi/os/default`

The search order for locating `osd` source files is:

- `libCom/osi/os/<arch>`
- `libCom/osi/os/posix`
- `libCom/osi/os/default`

20.1.3 Locating OSI header files.

When code is compiled the search order for locating header files in `base/include` is:

- `<base>/include/os/<arch>`
- `<base>/include`

20.2 epicsAssert

This is a replacement for ANSI C's `assert`. To use this version just include:

```
#include "epicsAssert.h"
```

instead of

```
#include <assert.h>
```

If an `assert` fails, it calls `errlog` indicating the program's author, file name, and line number. Under each OS there are specialized instructions assisting the user to diagnose the problem and generate a good bug report. For instance, under `vxWorks`, there are instructions on how to generate a stack trace, and on `posix` there are instructions about saving the core file. After printing the message the calling thread is suspended.

An author may, before the above include line, optionally define a preprocessor macro named `epicsAssertAuthor` as a string that provides their name and email address if they wish to be contacted when the assertion fires.

20.3 epicsEndian

`epicsEndian.h` provides an operating-system independent means of discovering the native byte order of the CPU which the compiler is targeting, and works for both C and C++ code. It defines the following preprocessor macros, the values of which are integers:

- `EPICS_ENDIAN_LITTLE`
- `EPICS_ENDIAN_BIG`
- `EPICS_BYTE_ORDER`
- `EPICS_FLOAT_WORD_ORDER`

The latter two macros are defined to be one or other of the first two and may be compared with them to determine conditional compilation or execution of code that performs byte or word swapping as necessary.

20.4 epicsEvent

epicsEvent.h defines the C++ and a C APIs for a simple binary semaphore. If multiple threads are waiting on the same event, only one of them will be woken when the event is signalled.

20.4.1 C++ Interface

```
typedef enum {
    epicsEventWaitOK, epicsEventWaitTimeout, epicsEventWaitError
} epicsEventWaitStatus;

typedef enum {epicsEventEmpty, epicsEventFull} epicsEventInitialState;

class epicsEvent{
public:
    epicsEvent(epicsEventInitialState initial=epicsEventEmpty);
    ~epicsEvent();
    void signal();
    void wait(); /*blocks until full*/
    bool wait( double timeOut ); /* false if empty at time out */
    bool tryWait(); /* false if empty */
    void show( unsigned level ) const;

    class invalidSemaphore {}; /* exception */
private:
    ...
};
```

Method	Meaning
epicsEvent	An epicsEvent can be created empty or full. If it is created empty then a wait issued before a signal will block. If created full then the first wait will always succeed. Multiple signals may be issued between waits but have the same effect as a single signal.
~epicsEvent	Remove the event and any resources it uses. Any further use of the semaphore result in unknown (most certainly bad) behavior. No outstanding take can be active when this call is made.
signal	Signal the event i.e. ensures that the next or current call to wait completes. This method may be called from a vxWorks or RTEMS interrupt handler.
wait()	Wait for the event.
wait(double timeOut)	Similar to wait except that if event does not happen the call completes after the specified time out. The return value is true if the event was signalled, false if it timed out.
tryWait()	Similar to wait except that if event does not happen the call completes immediately. The return value is true if an unused event has already been signalled, false if not.
show	Display information about the semaphore. The information displayed is architecture dependent.

The primary use of an event semaphore is for thread synchronization. An example of using an event semaphore is a consumer thread that processes requests from one or more producer threads. For example:

- When creating the consumer thread also create an epicsEvent.

```
epicsEvent *pevent = new epicsEvent;
```

- The consumer thread has code containing:

```
while(1) {
    pevent->wait();
    while(/*more work*/) {
        /*process work*/
    }
}
```

- Producers create requests and issue the statement:

```
pevent->signal();
```

20.4.2 C Interface

```
typedef struct epicsEventOSD *epicsEventId;

epicsEventId epicsEventCreate(epicsEventInitialState initialState);
epicsEventId epicsEventMustCreate (epicsEventInitialState initialState);
void epicsEventDestroy(epicsEventId id);
void epicsEventSignal(epicsEventId id);
epicsEventWaitStatus epicsEventWait(epicsEventId id);

void epicsEventMustWait(epicsEventId id);
epicsEventWaitStatus epicsEventWaitWithTimeout(epicsEventId id, double timeout);
epicsEventWaitStatus epicsEventTryWait(epicsEventId id);
void epicsEventShow(epicsEventId id, unsigned int level);
```

Each C routine corresponds to one of the C++ methods. `epicsEventMustCreate` and `epicsEventMustWait` do not return if they fail.

20.5 epicsFindSymbol

`epicsFindSymbol.h` contains the following definition:

```
void * epicsFindSymbol(const char *name);
```

Method	Meaning
<code>epicsFindSymbol</code>	Return the address of the variable name

`vxWorks` provides a function `symFindByName`, which finds and returns the address of global variables. The registry, described in the next chapter, provides an alternative but also requires extra work by `iocCore` and/or user code. If the registry is asked for a name that has not been registered, it calls `epicsFindSymbol`. If `epicsFindSymbol` can locate the global symbol it returns the address, otherwise it returns null.

On `vxWorks` `epicsFindSymbol` calls `symFindByName`.

A default version just returns null, i.e. it always fails.

20.6 epicsGeneralTime

The generalTime framework provides a mechanism for several time providers to be present within the system. There are two types of provider, one type for the current time and one type for providing Time Event times. Each time provider has a priority, and installed providers are queried in priority order whenever a time is requested, until one returns successfully. Thus there is a fallback from higher priority providers (smaller value of priority) to lower priority providers (larger value of priority) if the higher priority ones fail. Each architecture has a “last resort” provider, installed at priority 999, usually based on the system clock, which is used in the absence of any other provider.

Targets running vxWorks and RTEMS have an NTP provider installed at priority 100.

Registered providers may also add an interrupt-safe routine that will be called from the epicsTimeGetCurrentInt() or epicsTimeGetEventInt() API routines. These interfaces do not check the priority queue, and will only succeed if the last- used provider has registered a suitable routine.

There are two interfaces to this framework, epicsGeneralTime.h for consumers that wish to get a time and query the framework, and generalTimeSup.h for providers that supply timestamps.

20.6.1 Consumer interface

The epicsGeneralTime.h header contains the following:

```
void generalTime_Init(void);
int installLastResortEventProvider(void);
void generalTimeResetErrorCounts();
int generalTimeGetErrorCounts();

const char * generalTimeHighestCurrentName(void);
const char * generalTimeCurrentProviderName(void);
const char * generalTimeEventProviderName(void);

/* Original names, for compatibility */
#define generalTimeCurrentTpName generalTimeCurrentProviderName
#define generalTimeEventTpName generalTimeEventProviderName

long generalTimeReport(int interest);
```

Method	Meaning
generalTime_Init	Initialise the framework. This is called automatically by any function that requires the framework. It does not need to be called explicitly.
installLastResortEventProvider	Install a Time Event time provider that returns the current time for any Time Event number. This is optional as it is site policy whether the last resort for a Time Event time in the absence of any working provider should be a failure, or the current time.
generalTimeResetErrorCounts	Reset the internal counter of the number of times the time returned was earlier than when previously requested. Used by device support for bo record with DTYP = “General Time” OUT = “@RSTERRCNT”
generalTimeGetErrorCounts	Return the internal counter of the number of times the time returned was earlier than when previously requested. Used by device support for longin record with DTYP = “General Time” INP = “@GETERRCNT”
generalTimeCurrentProviderName	Return the name of the provider that last returned a valid current time, or NULL if none. Used by stringin device support with DTYP = “General Time” INP = “@BESTTCP”
generalTimeEventProviderName	Return the name of the provider that last returned a valid Time Event time, or NULL if none. Used by stringin device support with DTYP = “General Time” INP = “@BESTTEP”

<code>generalTimeHighestCurrentName</code>	Return the name of the registered current time provider that has the highest priority. Used by stringin device support with DTYP = "General Time" INP = "@TOPTCP"
<code>generalTimeReport</code>	Provide information about the installed providers and their current best times.

20.6.2 Time Provider Interface

The `generalTimeSup.h` header for time providers contains the following:

```
typedef int (*TIMECURRENTFUN)(epicsTimeStamp *pDest);
typedef int (*TIMEEVENTFUN)(epicsTimeStamp *pDest, int event);

int generalTimeRegisterCurrentProvider(const char *name,
    int priority, TIMECURRENTFUN getTime);
int generalTimeRegisterEventProvider(const char *name,
    int priority, TIMEEVENTFUN getEvent);

/* Original names, for compatibility */
#define generalTimeCurrentTpRegister generalTimeRegisterCurrentProvider
#define generalTimeEventTpRegister generalTimeRegisterEventProvider

int generalTimeAddIntCurrentProvider(const char *name,
    int priority, TIMECURRENTFUN getTime);
int generalTimeAddIntEventProvider(const char *name,
    int priority, TIMEEVENTFUN getEvent);

int generalTimeGetExceptPriority(epicsTimeStamp *pDest, int *pPrio,
    int ignorePrio);
```

Method	Meaning
<code>generalTimeRegisterCurrentProvider</code>	Register a current time provider with the framework. The <code>getTime</code> routine must return <code>epicsTimeOK</code> if it provided a valid time, or <code>epicsTimeERROR</code> if it could not.
<code>generalTimeRegisterEventProvider</code>	Register a provider of Time Event times with the framework. The <code>getEvent</code> routine must return <code>epicsTimeOK</code> if it provided a valid time for the requested Time Event, or <code>epicsTimeERROR</code> if it could not. It is an implementation decision by the provider whether a request for an Event that has never happened should return an error and/or a valid timestamp.
<code>generalTimeAddIntCurrentProvider</code>	Add or replace an interrupt-safe provider routine for an already-registered current time provider with the given name and priority.
<code>generalTimeAddIntEventProvider</code>	Add or replace an interrupt-safe provider routine for an already-registered event time provider with the given name and priority.
<code>generalTimeGetExceptPriority</code>	Request the current time from the framework, but exclude providers with priority <code>ignorePrio</code> . This allows a provider without an absolute time source to synchronise itself with other providers that do provide an absolute time. <code>pPrio</code> returns the priority of the provider that supplied the result, which may be higher or lower than <code>ignorePrio</code> .

If multiple providers are registered at the same priority, they will be queried in the order in which they were registered until one is able to provide the time when requested.

Some providers may start a task that periodically synchronizes themselves with a higher priority provider, using `generalTimeGetExceptPriority` to ensure that they are themselves excluded from this time request.

Interrupt-safe providers are optional, but an IOC that needs to request the time from interrupt context must be using a current or event time source that supports the appropriate request because only the most recently successful provider will be used (the priority list is not traversed for these requests). The result returned is not protected against backwards movement.

20.6.3 Internal Interface

The generalTime framework also now provides the implementations of `epicsTimeGetCurrent()` and `epicsTimeGetEvent()`. If `epicsTimeGetEvent()` is called with an event number of 0 (`epicsTimeEventCurrentTime`) then it will get the time from the best available current time provider. Thus providers do not need to provide event times if they do not implement an event system.

20.6.4 Example

Soft device support is provided for `ai`, `bo`, `longin` and `stringin` records. A typical example is:

```
record(ai, "$(IOC):GTIM_CURTIME") {
    field(DESC, "Get Time")
    field(DTYP, "General Time")
    field(INP, "@TIME")
}

record(bo, "$(IOC):GTIM_RSTERR") {
    field(DESC, "Reset ErrorCounts")
    field(DTYP, "General Time")
    field(OUT, "@RSTERRCNT")
}

record(longin, "$(IOC):GTIM_ERRCNT") {
    field(DESC, "Get ErrorCounts")
    field(DTYP, "General Time")
    field(INP, "@GETERRCNT")
}

record(stringin, "$(IOC):GTIM_BESTTCP") {
    field(DESC, "Best Time-Current-Provider")
    field(DTYP, "General Time")
    field(INP, "@BESTTCP")
}

record(stringin, "$(IOC):GTIM_BESTTEP") {
    field(DESC, "Best Time-Event-Provider")
    field(DTYP, "General Time")
    field(INP, "@BESTTEP")
}
```

20.7 epicsInterrupt

`epicsInterrupt.h` contains the following:

20.7.1 C Interface

```
int epicsInterruptLock();

void epicsInterruptUnlock(int key);

int epicsInterruptIsInterruptContext();
void epicsInterruptContextMessage(const char *message);
```

Method	Meaning
epicsInterruptLock	Lock interrupts and return a key to be passed to epicsInterruptUnlock To lock the following is done. int key; ... key = epicsInterruptLock(); ... epicsInterruptUnlock(key);
epicsInterruptUnlock	Unlock interrupts.
epicsInterruptIsInterruptContext	Return (true, false) if current context is interrupt context.
epicsInterruptContextMessage	Generate a message while interrupt context is true.

20.7.2 Implementation notes

A vxWorks specific version is provided. It maps directly to intLib calls.

An RTEMS version is provided that maps to rtems_ calls.

A default version is provided that uses a global semaphore to lock. This version is intended for operating systems in which iocCore will run as a multi threaded process. The global semaphore is thus only global within the process. This version is intended for use on all except real time operating systems.

The vxWorks implementation will most likely not work on symmetric multiprocessing systems.

The reason epicsInterrupt is needed is:

- callbackRequest and scanOnce can be issued from interrupt level.
- The errlog routines can be called while at interrupt level.

20.8 epicsMath

epicsMath.h includes math.h and also ensures that isnan and isinf are defined.

20.9 epicsMessageQueue

epicsMessageQueue.h describes a C++ and a C facility for interlocked communication between threads.

20.9.1 C++ Interface

EpicsMessageQueue provides methods for sending messages between threads on a first in, first out basis. It is designed so that a single message queue can be used with multiple writer and reader threads.

```

class epicsMessageQueue {
public:
    epicsMessageQueue(unsigned int capacity, unsigned int maximumMessageSize);
    ~epicsMessageQueue();
    int trySend(void *message, unsigned int messageSize);
    int send(void *message, unsigned int messageSize);
    int send(void *message, unsigned int messageSize, double timeout);
    int tryReceive(void *message, unsigned int messageBufferSize);
    int receive(void *message, unsigned int messageBufferSize);
    int receive(void *message, unsigned int messageBufferSize, double timeout);
    void show(int level) const;
    int pending() const;

private: // Prevent compiler-generated member functions
    // default constructor, copy constructor, assignment operator
    epicsMessageQueue();
    epicsMessageQueue(const epicsMessageQueue &);
    epicsMessageQueue& operator=(const epicsMessageQueue &);

private: // Data
    ...
};

```

An `epicsMessageQueue` cannot be assigned to, copy-constructed, or constructed without giving the *capacity* and *maximumMessageSize* arguments. The C++ compiler will object to some of the statements below:

```

epicsMessageQueue mq0(); // Error: default constructor is private
epicsMessageQueue mq1(10, 20); // OK
epicsMessageQueue mq2(t1); // Error: copy constructor is private
epicsMessageQueue *pmq; // OK, pointer
*pmq = mq1; // Error: assignment operator is private
pmq = &mq1; // OK, pointer assignment and address-of

```

Method	Meaning
<code>epicsMessageQueue()</code>	Constructor. The capacity is the maximum number of messages, each containing 0 to <code>maximumMessageSize</code> bytes, that can be stored in the message queue.
<code>~epicsMessageQueue()</code>	Destructor.
<code>trySend()</code>	Try to send a message. Return 0 if the message was sent to a receiver or queued for future delivery. Return -1 if no more messages can be queued or if the message is larger than the queue's maximum message size. This method may be called from a vxWorks or RTEMS interrupt handler.
<code>send()</code>	Send a message. Return 0 if the message was sent to a receiver or queued for future delivery. Return -1 if the timeout, if any, was reached before the message could be sent or queued, or if the message is larger than the queue's maximum message size.
<code>tryReceive()</code>	Try to receive a message. If the message queue is non-empty, the first message on the queue is copied to the specified location and the length of the message is returned. Returns -1 if the message queue is empty. If the pending message is larger than the specified <code>messageBufferSize</code> it may either return -1, or truncate the message. It is most efficient if the <code>messageBufferSize</code> is equal to the <code>maximumMessageSize</code> with which the message queue was created.
<code>receive()</code>	Wait until a message is sent and store it in the specified location. The number of bytes in the message is returned. Returns -1 if a message is not received within the timeout interval. If the received message is larger than the specified <code>messageBufferSize</code> it may either return -1, or truncate the message. It is most efficient if the <code>messageBufferSize</code> is equal to the <code>maximumMessageSize</code> with which the message queue was created.

show()	Displays some information about the message queue. The level argument controls the amount of information displayed.
pending()	Returns the number of messages presently in the queue.

20.9.2 C interface

```
typedef struct epicsMessageQueueOSD *epicsMessageQueueId;

epicsMessageQueueId epicsMessageQueueCreate(unsigned int capacity,
      unsigned int maximumMessageSize);
void epicsMessageQueueDestroy(epicsMessageQueueId id);
int epicsMessageQueueTrySend(epicsMessageQueueId id,
      void *message, unsigned int size);
int epicsMessageQueueSend(epicsMessageQueueId id,
      void *message, unsigned int size);
int epicsMessageQueueSendWithTimeout(epicsMessageQueueId id,
      void *message, unsigned int size, double timeout);
int epicsMessageQueueTryReceive(epicsMessageQueueId id,
      void *message, unsigned int size);
int epicsMessageQueueReceive(epicsMessageQueueId id,
      void *message, unsigned int size);
int epicsMessageQueueReceiveWithTimeout(epicsMessageQueueId id,
      void *message, unsigned int size, double timeout);
void epicsMessageQueueShow(epicsMessageQueueId id);
int epicsMessageQueuePending(epicsMessageQueueId id);
```

Each C function corresponds to one of the C++ methods.

20.10 epicsMutex

epicsMutex.h contains both C++ and C descriptions for a mutual exclusion semaphore.

20.10.1 C++ Interface

```
typedef enum {
    epicsMutexLockOK, epicsMutexLockTimeout, epicsMutexLockError
} epicsMutexLockStatus;
class epicsMutex {
public:
    epicsMutex ();
    ~epicsMutex ();
    void lock (); /* blocks until success */
    bool tryLock (); /* true if successful */
    void unlock ();
    void show ( unsigned level ) const;

    class invalidSemaphore {}; /* exception */
private:
};
```


Method	Meaning
epicsMutex	Create a mutual exclusion semaphore.
~epicsMutex	Remove the semaphore and any resources it uses. Any further use of the semaphore result in unknown (most certainly bad) results.
lock()	Wait until the resource is free. After a successful lock additional, i.e. recursive, locks of any type can be issued but each must have an associated unlock.
tryLock()	Similar to lock except that, if the resource is owned by another thread, the call completes immediately. The return value is (false,true) if the resource (is not, is) owned by the caller.
unlock	Release the resource. If a thread issues recursive locks, there must be an unlock for each lock
show	Display information about the semaphore. The results are architecture dependent.

Mutual exclusion semaphores are for situations requiring mutually exclusive access to resources. A mutual exclusion semaphore may be taken recursively, i.e. can be taken more than once by the owner thread before releasing it. Recursive takes are useful for a set of routines that call each other while working on a mutually exclusive resource.

The typical use of a mutual exclusion semaphore is:

```
epicsMutex *plock = new epicsMutex;
...
...
plock->lock();
/* process resource */
plock->unlock();
```

20.10.2 C Interface

```
typedef struct epicsMutexOSD* epicsMutexId;

epicsMutexId epicsMutexCreate(void);
epicsMutexId epicsMutexMustCreate (void);
void epicsMutexDestroy(epicsMutexId id);
void epicsMutexUnlock(epicsMutexId id);
epicsMutexLockStatus epicsMutexLock(epicsMutexId id);

void epicsMutexMustLock(epicsMutexId id);
epicsMutexLockStatus epicsMutexTryLock(epicsMutexId id);
void epicsMutexShow(epicsMutexId id,unsigned int level);
void epicsMutexShowAll(int onlyLocked,unsigned int level);
```

Each C routine corresponds to one of the C++ methods. `epicsMutexMustCreate` and `epicsMutexMustLock` do not return if they fail.

20.10.3 Implementation Notes

The implementation:

- Must implement recursive locking
- May implement priority inheritance and be deletion safe

A posix version is implemented via pthreads.

20.11 epicsStdlib

epicsStdlib.h includes stdlib.h and contains definitions for the following functions/macros.

```
double epicsStrtod(const char *str, char **endp);
int epicsScanFloat(const char *str, float *dest);
int epicsScanDouble(const char *str, double *dest);
```

epicsStrtod has the same semantics as the C99 function `strtod`. It is provided because some architectures have implementations which do not accept NAN or INFINITY. On architectures which provide an acceptable version of `strtod` this is implemented as a simple macro expansion.

epicsScanFloat and epicsScanDouble behave like `sscanf` with a "%f" and "%lf" format string, respectively. They are provided because some architectures have implementations of `scanf` which do not accept NAN or INFINITY.

20.12 epicsStdio

epicsStdio.h contains definitions for the following functions.

```
int epicsSnprintf(char *str, size_t size,
    const char *format, ...);
int epicsVsnprintf(char *str, size_t size,
    const char *format, va_list ap);
void epicsTempName(char * pNameBuf, size_t nameBufLength );
FILE * epicsShareAPI epicsTempFile ();
enum TF_RETURN {TF_OK=0, TF_ERROR=1};
enum TF_RETURN truncateFile(const char *pFileName, unsigned size );
FILE * epicsGetStdin(void);
FILE * epicsGetStdout(void);
FILE * epicsGetStderr(void);
FILE * epicsGetThreadStdin(void);
FILE * epicsGetThreadStdout(void);
FILE * epicsGetThreadStderr(void);
void epicsSetThreadStdin(FILE *);
void epicsSetThreadStdout(FILE *);
void epicsSetThreadStderr(FILE *);
int epicsStdoutPrintf(const char *pformat, ...);
```

epicsSnprintf and epicsVsnprintf are meant to have the same semantics as the C99 functions `snprintf` and `vsnprintf`. They are provided because some architectures do not implement these functions, while others do not implement the correct semantics. If you haven't heard of these C99 functions, `snprintf` is like `sprintf` except that the `size` argument gives the maximum number of characters (including the trailing zero byte) that may be placed in `str`. Similarly `vsnprintf` is a size-limited version of `vsprintf`. In both cases the return value is supposed to be the number characters (not counting the terminating zero byte) that would be written to `str` if it was large enough to hold them all; the output has been truncated if the return value is `size` or more.

On some operating systems the implementations of these functions may not always return the correct value. If the OS implementation does not correctly return the number of characters that would have been written when the output gets truncated, it is not worth trying to fix this as long as both return `size-1` instead; the resulting string must always be correctly terminated with a zero byte.

Operating systems such as Solaris which follow the Single Unix Specification V2, `epicsSnprintf` and `epicsVsnprintf` do not provide correct C99 semantics for the return value when `size`

is given as zero. On these systems `epicsSnprintf` and `epicsVsnprintf` may return an error (a value less than zero) if a buffer length of zero is passed in, so callers should not use that technique.

`epicsTempName` and `epicsTempFile` can be called to get unique filenames and files.

`truncateFile` returns `TF_OK` if the file is less than size bytes or if it was successfully truncated. Returns `TF_ERROR` if the file could not be truncated.

The `Stdin/Stdout/Stderr` routines allow these file streams to be redirected on a per thread basis, e.g. calling `epicsSetThreadStdout` will affect only the thread which calls it.

`epicsGetStdin`, ..., `epicsStdoutPrintf` are not normally called by user code. Instead code that wants to allow redirection needs only to include `epicsStdioRedirect.h`

20.13 epicsStdioRedirect

Including this file cause the following names to be changed:

- `stdin` becomes `epicsGetStdin()`
- `stdout` becomes `epicsGetStdout()`
- `stderr` becomes `epicsGetStderr()`
- `printf` becomes `epicsStdoutPrintf`

This is done so that redirection can occur. A primary use of this facility is `iocsh`. It allows redirection of input and output. In order for it to work, all modules involved in I/O can just include this header.

20.14 epicsThread

`epicsThread.h` contains C++ and C descriptions for a thread.

20.14.1 C Interface

```
typedef void (*EPICSTHREADFUNC)(void *parm);

static const unsigned epicsThreadPriorityMax = 99;
static const unsigned epicsThreadPriorityMin = 0;
/* some generic values */
static const unsigned epicsThreadPriorityLow = 10;
static const unsigned epicsThreadPriorityMedium = 50;
static const unsigned epicsThreadPriorityHigh = 90;
/* some iocCore specific values */
static const unsigned epicsThreadPriorityCAServerLow = 20;
static const unsigned epicsThreadPriorityCAServerHigh = 40;
static const unsigned epicsThreadPriorityScanLow = 60;
static const unsigned epicsThreadPriorityScanHigh = 70;
static const unsigned epicsThreadPriorityIocsh = 91;
static const unsigned epicsThreadPriorityBaseMax = 91;

/* stack sizes for each stackSizeClass are implementation and CPU dependent */
typedef enum {
    epicsThreadStackSmall, epicsThreadStackMedium, epicsThreadStackBig
```

```

} epicsThreadStackSizeClass;

typedef enum {
    epicsThreadBooleanStatusFail, epicsThreadBooleanStatusSuccess
} epicsThreadBooleanStatus;

unsigned int epicsThreadGetStackSize(epicsThreadStackSizeClass size);

typedef int epicsThreadOnceId;
#define EPICS_THREAD_ONCE_INIT 0

void epicsThreadOnce(epicsThreadOnceId *id, EPICSTHREADFUNC, void *arg);

void epicsThreadExitMain(void);

/* (epicsThreadId)0 is guaranteed to be an invalid thread id */
typedef struct epicsThreadOSD *epicsThreadId;

epicsThreadId epicsThreadCreate(const char *name,
    unsigned int priority, unsigned int stackSize,
    EPICSTHREADFUNC funptr, void *parm);
void epicsThreadSuspendSelf(void);
void epicsThreadResume(epicsThreadId id);
unsigned int epicsThreadGetPriority(epicsThreadId id);
unsigned int epicsThreadGetPrioritySelf();
void epicsThreadSetPriority(epicsThreadId id, unsigned int priority);
epicsThreadBooleanStatus epicsThreadHighestPriorityLevelBelow (
    unsigned int priority, unsigned *pPriorityJustBelow);
epicsThreadBooleanStatus epicsThreadLowestPriorityLevelAbove (
    unsigned int priority, unsigned *pPriorityJustAbove);
int epicsThreadIsEqual(epicsThreadId id1, epicsThreadId id2);
int epicsThreadIsSuspended(epicsThreadId id);
void epicsThreadSleep(double seconds);
double epicsThreadSleepQuantum(void);
epicsThreadId epicsThreadGetIdSelf(void);
epicsThreadId epicsThreadGetId(const char *name);

const char * epicsThreadGetNameSelf(void);
void epicsThreadGetName(epicsThreadId id, char *name, size_t size);
int epicsThreadIsOkToBlock(void);
void epicsThreadSetOkToBlock(int isOkToBlock);

void epicsThreadShowAll(unsigned int level);
void epicsThreadShow(epicsThreadId id, unsigned int level);

typedef void * epicsThreadPrivateId;
epicsThreadPrivateId epicsThreadPrivateCreate(void);
void epicsThreadPrivateDelete(epicsThreadPrivateId id);
void epicsThreadPrivateSet(epicsThreadPrivateId, void *);
void * epicsThreadPrivateGet(epicsThreadPrivateId);

```

Method**Meaning**

epicsThreadGetStackSize	Get a stack size value that can be given to epicsThreadCreate. The size argument should be one of the values epicsThreadStackSmall, epicsThreadStackMedium or epicsThreadStackBig.
epicsThreadOnce	This is used as follows: void myInitFunc(void * arg) { ... } epicsThreadOnceId onceFlag = EPICS_THREAD_ONCE_INIT; ... epicsThreadOnce(&onceFlag,myInitFunc,(void *)myParm);For each unique epicsThreadOnceId, epicsThreadOnce guarantees 1) myInitFunc is called only once. 2) myInitFunc completes before any epicsThreadOnce call completes.Note that myInitFunc must not call epicsThreadOnce with the same onceId.
epicsThreadExitMain	If the main routine is done but wants to let other threads run it can call this routine. This should be the last call in main, except the final return. On most systems epicsThreadExitMain never returns. This must only be called by the main thread.
epicsThreadCreate	Create a new thread. The use made of the priority, and stackSize arguments is implementation dependent. Some implementations may ignore one or other of these, but for portability appropriate values should be given for both. The value passed as the stackSize parameter should be obtained by calling epicsThreadGetStackSize. The funptr argument specifies a function that implements the thread, and parm is the single argument passed to funptr. A thread terminates when funptr returns.
epicsThreadSuspendSelf	This causes the calling thread to suspend. The only way it can resume is for another thread to call epicsThreadResume.
epicsThreadResume	Resume a suspended thread. Only do this if you know that it is safe to resume a suspended thread.
epicsThreadGetPriority	Get the priority of the specified thread.
epicsThreadGetPrioritySelf	Get the priority of this thread.
epicsThreadSetPriority	Set a new priority for the specified thread. The result is implementation dependent.
epicsThreadHighestPriorityLevelBelow	Get a priority that is just lower than the specified priority.
epicsThreadLowestPriorityLevelAbove	Get a priority that is just above the specified priority.
epicsThreadIsEqual	Compares two threadIds and returns (0,1) if they (are not, are) the same.
epicsThreadIsSuspended	BAD NAME. taskwd needs this call. It really means: Is there something wrong with this thread? This could mean suspended or no longer exists or etc. It is a problem because it is implementation dependent.
epicsThreadSleep	Sleep for the specified period of time, i.e. sleep without using the cpu. If delay is >0 then the thread will sleep at least until the next clock tick. The exact time is determined by the underlying architecture. If delay is <= 0 then a delay of 0 is requested of the underlying architecture. What happens is architecture dependent but often it allows other threads of the same priority to run.
epicsThreadSleepQuantum	This function returns the minimum slumber interval obtainable with epicsThreadSleep() in seconds. On most OS there is a system scheduler interrupt interval which determines the value of this parameter. Knowledge of this parameter is used by the various components of EPICS to improve scheduling of software tasks intime when the reduction of average time scheduling errors is important.If this parameter is unknown or is unpredictable for a particular OS then it is safe to return zero.
epicsThreadGetIdSelf	Get the threadId of the calling thread.
epicsThreadGetId	Get the threadId of the specified thread. A return of 0 means that no thread was found with the specified name.
epicsThreadGetNameSelf	Get the name of the calling thread.

<code>epicsThreadGetName</code>	Get the name of the specified thread. The value is copied to a caller specified buffer so that if the thread terminates the caller is not left with a pointer to something that may no longer exist.
<code>epicsThreadIsOkToBlock</code>	Is it OK for a thread to block? This can be called by support code that does not know if it is called in a thread that should not block. For example the <code>errlog</code> system calls this to decide when messages should be displayed on the console.
<code>epicsThreadSetOkToBlock</code>	When a thread is started the default is that it is not allowed to block. This method can be called to change the state. For example <code>iocsh</code> calls this to specify that it is OK to block.
<code>epicsThreadShowAll</code>	Display info about all threads.
<code>epicsThreadShow</code>	Display info about the specified thread.
<code>epicsThreadPrivateCreate</code>	Thread private variables are intended for use by legacy libraries written for a single threaded environment and which uses a global variable to store private data. The only code in base that currently needs this facility is channel access. A library that needs a private variable should make exactly one call to <code>epicsThreadPrivateCreate</code> . Each thread should call <code>epicsThreadPrivateSet</code> when the thread is created. Each library routine can call <code>epicsThreadPrivateGet</code> each time it is called.
<code>epicsThreadPrivateDelete</code>	Delete a thread private variable.
<code>epicsThreadPrivateSet</code>	Set the value for a thread private variable.
<code>epicsThreadPrivateGet</code>	Get the value of a thread private variable, the value is the value set by the call to <code>epicsThreadPrivateSet</code> that was made by the same thread. If called before <code>epicsThreadPrivateSet</code> it returns 0.

`epicsThread` is meant as a somewhat minimal interface for multithreaded applications. It can be implemented on a wide variety of systems with the restriction that the system **MUST** support a multithreaded environment. A POSIX `pthreads` version is provided.

The interface provides the following thread facilities, with restrictions as noted:

- Life cycle - A thread starts life as a result of a call to `epicsThreadCreate`. It terminates when the thread function returns. It should not return until it has released all resources it uses. If a thread is expected to terminate as a natural part of its life cycle then the thread function must return.
- `epicsThreadOnce` - This provides the ability to have an initialization function that is guaranteed to be called exactly once.
- `main` - If a main routine finishes its work but wants to leave other threads running it can call `epicsThreadExitMain`, which should be the last statement in `main`.
- Priorities - Ranges between 0 and 99 with a higher number meaning higher priority. A number of constants are defined for `iocCore` specific threads. The underlying implementation may collapse the range 0 to 99 into a smaller range; even a single priority. User code should never use priorities to guarantee correct behavior.
- Stack Size - `epicsThreadCreate` accepts a stack size parameter. Three generic sizes are available: small, medium, and large. Portable code should always use one of the generic sizes. Some implementation may ignore the stack size request and use a system default instead. Virtual memory systems providing generous stack sizes can be expected to use the system default.
- `epicsThreadId` - This is given a value as a result of a call to `epicsThreadCreate`. A value of 0 always means no thread. If a `threadId` is used for a thread that has terminated the result is not defined (but will normally lead to bad things happening). Thus code that looks after other threads **MUST** be aware of threads terminating.

20.14.2 C++ Interface

```

class epicsThreadRunnable {
public:
    virtual void run() = 0;
    virtual void stop();
    virtual void show(unsigned int level) const;
};

class epicsShareClass epicsThread {
public:
    epicsThread (epicsThreadRunnable &,const char *name, unsigned int stackSize,
                unsigned int priority=epicsThreadPriorityLow);
    virtual ~epicsThread ();
    void start();
    void exitWait ();
    void exitWait (const double delay );
    void exitWaitRelease (); // noop if not called by managed thread
    static void exit ();
    void resume ();
    void getName (char *name, size_t size) const;
    epicsThreadId getId () const;
    unsigned int getPriority () const;
    void setPriority (unsigned int);
    bool priorityIsEqual (const epicsThread &otherThread) const;
    bool isSuspended () const;
    bool isCurrentThread () const;
    bool operator == (const epicsThread &rhs) const;
    /* these operate on the current thread */
    static void suspendSelf ();
    static void sleep (double seconds);
    static epicsThread & getSelf ();
    static const char * getNameSelf ();
    static bool isOkToBlock ();
    static void setOkToBlock(bool isOkToBlock) ;
private:
    ...
};
template <class T>
class epicsThreadPrivate {
public:
    epicsThreadPrivate ();
    ~epicsThreadPrivate ();
    T *get () const;
    void set (T *);
    class unableToCreateThreadPrivate {}; // exception
private:
    ...
};

```

The C++ interface is just a wrapper around the C interface. Two differences are the method `start` and the class `epicsThreadRunnable`.

The method `start` must be called only after the `epicsThread` object is constructed. It in turn calls the `run` method of the `epicsThreadRunnable` object.

Code using the C++ interface code must provide a class that derives from `epicsThreadRunnable`. One way to accomplish this is as follows:

```
class myThread: public epicsThreadRunnable {
public:
    myThread(int arg, const char *name);
    virtual ~myThread();
    virtual void run();
    epicsThread thread;

    ...
}

myThread::myThread(int arg, const char *name) :
    thread(*this, name, epicsThreadGetStackSize(epicsThreadStackSmall), 50)
{
    thread.start();
}

myThread::~~myThread() {}

void myThread::run()
{
    ...
}
```

20.15 epicsTime

`epicsTime.h` contains C++ and C descriptions for time.

20.15.1 Time Related Structures

```
/* epics time stamp for C interface*/
typedef struct epicsTimeStamp {
    epicsUInt32    secPastEpoch; /* seconds since 0000 Jan 1, 1990 */
    epicsUInt32    nsec;         /* nanoseconds within second */
} epicsTimeStamp;

/*TS_STAMP is deprecated */
#define TS_STAMP epicsTimeStamp

struct timespec; /* POSIX real time */
struct timeval; /* BSD */
struct l_fp; /* NTP timestamp */

// extend ANSI C RTL "struct tm" to include nano seconds within a second
// and a struct tm that is adjusted for the local timezone
struct local_tm_nano_sec {
    struct tm ansi_tm; /* ANSI C time details */
    unsigned long nSec; /* nano seconds extension */
}
```



```

};

// extend ANSI C RTL "struct tm" to includes nano seconds within a second
// and a struct tm that is adjusted for GMT (UTC)
struct gm_tm_nano_sec {
    struct tm ansi_tm; /* ANSI C time details */
    unsigned long nSec; /* nano seconds extension */
};

// wrapping this in a struct allows conversion to and
// from ANSI time_t but does not allow unexpected
// conversions to occur
struct time_t_wrapper {
    time_t ts;
};

```

The above structures are for the various time formats.

- `epicsTimeStamp` - This is the structure used by the C interface for epics time stamps. The C++ interface stores this information in private members. The two elements of the class are:
 - `secPastEpoch` - The number of seconds since January 1, 1990 (the epics epoch).
 - `nsec` - nanoseconds within a second

NOTE: `TS_STAMP` is defined for compatibility with existing code.

- `timespec` - This is defined by POSIX Real Time. It requires two mandatory fields:
 - `time_t tv_sec` - Number of seconds since 1970 (The POSIX epoch)
 - `long tv_nsec` - nanoseconds within a second
- `timeval` - BSD and SRV5 Unix timestamp. It has two fields:
 - `time_t tv_sec` - Number of seconds since 1970 (The POSIX epoch)
 - `time_t tv_nsec` - nanoseconds within a second
- `struct l_fp` - Network Time Protocol timestamp. The fields are:
 - `lui` - Number of seconds since 1900 (The NTP epoch)
 - `luf` - Fraction of a second. For example `0x80000000` represents 1/2 second.
- `local_tm_nano_sec` and `gm_tm_nano_sec` - Defined by epics. These just add a nanosecond field to `struct tm`.
- `time_t_wrapper` - This is for converting to/from the ANSI C `time_t`. Since `time_t` is usually an elementary type providing a conversion operator from `time_t` to/from `epicsTime` could cause undesirable implicit conversions. Providing a conversion operator to/from a `time_t_wrapper` prevents implicit conversions.

NOTE on conversion. The epics implementation will properly convert between the various formats from the beginning of the EPICS epoch until at least 2038. Unless the underlying architecture support has defective POSIX, BSD/SRV5, or standard C time support the epics implementation should be valid until 2106.

20.15.2 C++ Interface

```

class epicsTime;

class epicsTimeEvent {

```

```

public:
    epicsTimeEvent (const int &number);
    operator int () const;
private:
    int eventNumber;
};

class epicsTime {
public:
    // exceptions
    class unableToFetchCurrentTime {};
    class formatProblemWithStructTM {};

    epicsTime ();
    epicsTime (const epicsTime &t);

    static epicsTime getEvent (const epicsTimeEvent &event);
    static epicsTime getCurrent ();

    // convert to and from EPICS epicsTimeStamp format
    operator epicsTimeStamp () const;
    epicsTime (const epicsTimeStamp &ts);
    epicsTime & operator = (const epicsTimeStamp &rhs);

    // convert to and from ANSI time_t
    operator time_t_wrapper () const;
    epicsTime (const time_t_wrapper &tv);
    epicsTime & operator = (const time_t_wrapper &rhs);

    // convert to and from ANSI Cs "struct tm" (with nano seconds)
    // adjusted for the local time zone
    operator local_tm_nano_sec () const;
    epicsTime (const local_tm_nano_sec &ts);
    epicsTime & operator = (const local_tm_nano_sec &rhs);

    // convert to and from ANSI Cs "struct tm" (with nano seconds)
    // adjusted for GM time (UTC)
    operator gm_tm_nano_sec () const;
    epicsTime ( const gm_tm_nano_sec & );
    epicsTime & operator = ( const gm_tm_nano_sec & );

    // convert to and from POSIX RT's "struct timespec"
    operator struct timespec () const;
    epicsTime (const struct timespec &ts);
    epicsTime & operator = (const struct timespec &rhs);

    // convert to and from BSD's "struct timeval"
    operator struct timeval () const;
    epicsTime (const struct timeval &ts);
    epicsTime & operator = (const struct timeval &rhs);

    // convert to and from NTP timestamp format
    operator l_fp () const;
    epicsTime (const l_fp &);

```

```

    epicsTime & operator = (const l_fp &rhs);

    // convert to and from WIN32s FILETIME (implemented only on WIN32)
    operator struct _FILETIME () const;
    epicsTime ( const struct _FILETIME & );
    epicsTime & operator = ( const struct _FILETIME & );

    // arithmetic operators
    double operator- (const epicsTime &rhs) const; // returns seconds
    epicsTime operator+ (const double &rhs) const; // add rhs seconds
    epicsTime operator- (const double &rhs) const; // subtract rhs seconds
    epicsTime operator+= (const double &rhs); // add rhs seconds
    epicsTime operator-= (const double &rhs); // subtract rhs seconds

    // comparison operators
    bool operator == (const epicsTime &rhs) const;
    bool operator != (const epicsTime &rhs) const;
    bool operator <= (const epicsTime &rhs) const;
    bool operator < (const epicsTime &rhs) const;
    bool operator >= (const epicsTime &rhs) const;
    bool operator > (const epicsTime &rhs) const;

    // convert current state to user-specified string
    size_t strftime (char *pBuff, size_t bufLength, const char *pFormat) const;

    // dump current state to standard out
    void show (unsigned interestLevel) const;

private:
    ...
};

```

20.15.3 class epicsTimeEvent

```

class epicsShareClass epicsTimeEvent
{
public:
    epicsTimeEvent (const int &number);
    operator int () const;
private:
    int eventNumber;
};

```

Method	Meaning
Convert to/from integer	Does not currently check that the range of the integer is valid, although it might one day.

20.15.4 class epicsTime

epicsTime Method	Meaning
epicsTime()	The default constructor sets the time to the beginning of the epics epoch.

epicsTime Method	Meaning
epicsTime(const epicsTime &t)	Copy constructor, copies the time from its argument.
static getEvent	Returns an epicsTime indicating when the associated event last occurred. See the description of the C routine epicsTimeGetEvent below for details.
static getCurrent	Returns an epicsTime containing the current time. For example: <pre>epicsTime now = epicsTime::getCurrent();</pre>
convert to/from epicsTimeStamp	Three epicsTime methods interface with epicsTimeStamp values: A constructor, an assignment operator, and a conversion operator. For example: <pre>epicsTimeStamp ts1 = {12345, 67890}, ts2; epicsTime t1(ts1), t2; // constructor t2 = ts1; // assignment ts2 = t1; // conversion operator</pre>
convert to/from time_t	The structure time_t_wrapper is used instead of time_t directly to avoid undesired conversions to integer types. Three methods are provided for ANSI time_t: A constructor, an assignment operator, and a conversion operator. Assume the following definitions: <pre>time_t tt; time_t_wrapper ttw; epicsTime time;</pre> <p>An example of the copy constructor is: <pre>ttw.ts = tt; epicsTime time1(ttw);</pre> <p>An example of the assignment operator is: <pre>time = ttw;</pre> <p>An example of the time_t_wrapper operator is: <pre>ttw = time; tt = ttw.ts;</pre> </p></p></p>
convert to/from struct tm in local timezone	The structure local_tm_nano_sec is used instead of struct tm directly to add a nanoseconds field and indicate the time is expressed in the local timezone. Three methods are provided for local_tm_nano_sec: A constructor, an assignment operator, and a conversion operator. Assume the following definitions: <pre>local_tm_nano_sec ltn; epicsTime time;</pre> <p>An example of the copy constructor is: <pre>epicsTime time1(ltn);</pre> <p>An example of the assignment operator is: <pre>time = ltn;</pre> <p>An example of the local_tm_nano_sec operator is: <pre>ltn = time;</pre> </p></p></p>
convert to/from struct tm in UTC	The structure gm_tm_nano_sec is used instead of struct tm directly to add a nanoseconds field and indicate the time is expressed in UTC. Three methods are provided for gm_tm_nano_sec: A constructor, an assignment operator, and a conversion operator. Assume the following definitions: <pre>gm_tm_nano_sec gtn; epicsTime time;</pre> <p>An example of the copy constructor is: <pre>epicsTime time1(gtn);</pre> <p>An example of the assignment operator is: <pre>time = gtn;</pre> <p>An example of the gm_tm_nano_sec operator is: <pre>gtn = time;</pre> </p></p></p>

epicsTime Method	Meaning
convert to/from struct timespec	<p>Three methods are provided for the POSIX struct timespec: A constructor, an assignment operator, and a conversion operator. Assume the following definitions:</p> <pre>struct timespec tts; epicsTime time;</pre> <p>An example of the copy constructor is:</p> <pre>epicsTime time1(tts);</pre> <p>An example of the assignment operator is:</p> <pre>time = tts;</pre> <p>An example of the conversion operator is:</p> <pre>tts = time;</pre>
convert to/from struct timeval	<p>Three methods are provided for BSD's struct timeval: A constructor, an assignment operator, and a conversion operator. Assume the following definitions:</p> <pre>struct timeval ttv; epicsTime time;</pre> <p>An example of the copy constructor is:</p> <pre>epicsTime time1(ttv);</pre> <p>An example of the assignment operator is:</p> <pre>time = ttv;</pre> <p>An example of the conversion operator is:</p> <pre>ttv = time;</pre>
convert to/from NTP timestamps	<p>Three methods are provided for the NTP timestamp structure: A constructor, an assignment operator, and a conversion operator. Assume the following definitions:</p> <pre>l_fp ntp; epicsTime time;</pre> <p>An example of the copy constructor is:</p> <pre>epicsTime time1(ntp);</pre> <p>An example of the assignment operator is:</p> <pre>time = ntp;</pre> <p>An example of the conversion operator is:</p> <pre>ntp = time;</pre>
arithmetic operators -, +, +=, -=	<p>The arithmetic operators allow the difference of two epicsTimes, with the result in seconds. It also allows -, +, +=, and -= where the left hand argument is an epicsTime and the right hand argument is a double. Examples are:</p> <pre>epicsTime time, time1, time2; double t1,t2,t3; ... t1 = time2 - time1; time = time1 + 4.5; time = time2 - t3; time2 += 6.0;</pre>
Comparison operators ==, !=, <=, <, >=, >	<p>Two epics times can be compared:</p> <pre>epicsTime time1, time2; ... if (time1 <= time2) ... </pre>

epicsTime Method	Meaning
strftime	<p>This method is an extension of the standard C library routine <code>strftime</code>. See your OS documentation for details about the standard routine. The <code>epicsTime</code> method adds support for the printing the fractional portion of the time. It searches the format string for the sequence <code>%0nf</code> where <code>n</code> is the desired precision, and uses this format to convert the fractional seconds to the requested precision. For example:</p> <pre>epicsTime time = epicsTime::getCurrent(); char buf[30]; time.strftime(buf, 30, "%Y-%m-%d %H:%M:%S.%06f"); printf("%s\n", buf);</pre> <p>This will print the current time in the format: 2001-01-26 20:50:29.813505</p>
show	Shows the date/time.

20.15.5 C Interface

```
/* All epicsTime routines return (-1,0) for (failure,success) */
#define epicsTimeOK 0
#define epicsTimeERROR (-1)
/*Some special values for eventNumber*/
#define epicsTimeEventCurrentTime 0
#define epicsTimeEventBestTime -1
#define epicsTimeEventDeviceTime -2

/* These are implemented in the "generalTime" framework */
int epicsTimeGetCurrent (epicsTimeStamp *pDest);
int epicsTimeGetEvent (epicsTimeStamp *pDest, int eventNumber);

/* These are callable from an Interrupt Service Routine */
epicsShareFunc int epicsTimeGetCurrentInt(epicsTimeStamp *pDest);
epicsShareFunc int epicsTimeGetEventInt(epicsTimeStamp *pDest, int eventNumber);

/* convert to and from ANSI C's "time_t" */
int epicsTimeToTime_t (time_t *pDest, const epicsTimeStamp *pSrc);
int epicsTimeFromTime_t (epicsTimeStamp *pDest, time_t src);

/*convert to and from ANSI C's "struct tm" with nano seconds */
int epicsTimeToTM (struct tm *pDest, unsigned long *pNSecDest,
    const epicsTimeStamp *pSrc);
int epicsTimeToGMTM (struct tm *pDest, unsigned long *pNSecDest,
    const epicsTimeStamp *pSrc);
int epicsTimeFromTM (epicsTimeStamp *pDest, const struct tm *pSrc,
    unsigned long nSecSrc);
int epicsTimeFromGMTM (epicsTimeStamp *pDest, const struct tm *pSrc,
    unsigned long nSecSrc);

/* convert to and from POSIX RT's "struct timespec" */
int epicsTimeToTimespec (struct timespec *pDest, const epicsTimeStamp *pSrc);
int epicsTimeFromTimespec (epicsTimeStamp *pDest, const struct timespec *pSrc);
```

```

/* convert to and from BSD's "struct timeval" */
int epicsTimeToTimeval (struct timeval *pDest, const epicsTimeStamp *pSrc);
int epicsTimeFromTimeval (epicsTimeStamp *pDest, const struct timeval *pSrc);
/*arithmetic operations */
double epicsTimeDiffInSeconds (
    const epicsTimeStamp *pLeft, const epicsTimeStamp *pRight);
void epicsTimeAddSeconds (
    epicsTimeStamp *pDest, double secondsToAdd); /* adds seconds to *pDest */

/*comparison operations: returns (0,1) if (false,true) */
int epicsTimeEqual(const epicsTimeStamp *pLeft, const epicsTimeStamp *pRight);
int epicsTimeNotEqual(const epicsTimeStamp *pLeft, const epicsTimeStamp *pRight);
int epicsTimeLessThan(const epicsTimeStamp *pLeft, const epicsTimeStamp *pRight);
int epicsTimeLessThanEqual(
    const epicsTimeStamp *pLeft, const epicsTimeStamp *pRight);
int epicsTimeGreaterThan (
    const epicsTimeStamp *pLeft, const epicsTimeStamp *pRight);
int epicsTimeGreaterThanEqual (
    const epicsTimeStamp *pLeft, const epicsTimeStamp *pRight);
/*convert to ASCII string */
size_t epicsTimeToStrftime (
    char *pBuff, size_t bufLength, const char *pFormat, const epicsTimeStamp
    *pTS);

/* dump current state to standard out */
void epicsTimeShow (const epicsTimeStamp *, unsigned interestLevel);
/* OS dependent reentrant versions of the ANSI C interface because */
/* vxWorks gmtime_r interface does not match POSIX standards */
int epicsTime_localtime ( const time_t *clock, struct tm *result );
int epicsTime_gmtime ( const time_t *clock, struct tm *result );

```

The C interface provides most of the features as the C++ interface. The features of the C++ operators are provided as functions.

Note that the `epicsTimeGetCurrent` and `epicsTimeGetEvent` routines are now implemented in `epicsGeneralTime.c`

20.16 osiPoolStatus

`osiPoolStatus.h` contains the following description:

```
int osiSufficientSpaceInPool(void);
```

Method	Meaning
<code>osiSufficientSpaceInPool</code>	Return (true,false) if there is sufficient free memory.

This determines if enough free memory exists to continue.

A vxWorks version returns (true,false) if `memFindMax` returns (>100000, <=100000) bytes.

The default version always returns true.

20.17 osiProcess

osiProcess.h contains the following:

```
typedef enum osiGetUserNameReturn {
    osiGetUserNameFail,
    osiGetUserNameSuccess
}osiGetUserNameReturn;

osiGetUserNameReturn osiGetUserName (char *pBuf, unsigned bufSize);

/*
 * Spawn detached process with named executable, but return
 * osiSpawnDetachedProcessNoSupport if the local OS does not
 * support heavy weight processes.
 */
typedef enum osiSpawnDetachedProcessReturn {
    osiSpawnDetachedProcessFail,
    osiSpawnDetachedProcessSuccess,
    osiSpawnDetachedProcessNoSupport
}osiSpawnDetachedProcessReturn;

osiSpawnDetachedProcessReturn osiSpawnDetachedProcess(
    const char *pProcessName, const char *pBaseExecutableName);
```

Not otherwise documented.

20.18 Ignoring Posix Signals

epicsSignal.h contains the following commented declarations:

```
/*
 * Required to avoid problems with soft IOCs getting SIGHUPs
 * when a Channel Access client disconnects:
 */
void epicsSignalInstallSigHupIgnore ( void );

/*
 * Required to avoid terminating a process which is blocking
 * in a socket send() call when a SIGPIPE signal is generated
 * by the OS:
 */
void epicsSignalInstallSigPipeIgnore ( void );

/*
 * Required only if shutdown() and close() do not interrupt
 * a thread blocking in a socket system call:
 */
void epicsSignalInstallSigAlarmIgnore ( void );
void epicsSignalRaiseSigAlarm ( struct epicsThreadOSD * );
```

Not otherwise documented.

20.19 OS-Independent Socket API

The header file `osiSock.h` provides wrappers around the different socket APIs provided by the supported operating systems. This API was designed to make it possible to write network applications that will compile and run on any OS. See the comments and declarations in the header file for details.

20.20 Device Support Library

NOTE: EPICS Base only provides vxWorks and RTEMS back-end implementations of these routines. Versions of the back-end routines for other operating systems can be added in a support or IOC application.

20.20.1 Overview

`devLib.h` provides definitions for a library of routines useful for device and driver modules, which are primarily intended for accessing VME devices. If all VME drivers register with these routines then addressing conflicts caused by multiple device/drivers trying to use the same VME addresses will be detected.

20.20.2 Location Probing

20.20.2.1 Read Probe

```
long devReadProbe(
    unsigned wordSize,
    volatile const void *ptr,
    void *pValueRead);
```

Performs a bus-error-safe atomic read operation width `wordSize` bytes from the `ptr` location, placing the value read (if successful) at `pValueRead`. The routine returns a failure status (non-zero) if a bus error occurred during the read cycle.

20.20.2.2 Write Probe

```
long devWriteProbe(
    unsigned wordSize,
    volatile void *ptr,
    const void *pValueWritten);
```

Performs a bus-error-safe atomic write operation width `wordSize` which copies the value from `pValueWritten` to the `ptr` location. The routine returns a failure status (non-zero) if a bus error occurred during the write cycle.

20.20.2.3 No Response Probe

```
long devNoResponseProbe(
    epicsAddressType addrType,
    size_t base,
    size_t size);
```

This routine performs a series of read probes for all word sizes from char to long at every naturally aligned location in the range `[base, base+size)` for the given bus address type. It returns an error if any location responds or if any such location cannot be mapped.

20.20.3 Registering VME Addresses

20.20.3.1 Definitions of Address Types

```
typedef enum {
    atVMEA16, atVMEA24, atVMEA32,
    atISA,
    atLast /* atLast is the last enum in this list */
} epicsAddressType;

char *epicsAddressTypeName[] = {
    "VME A16", "VME A24", "VME A32",
    "ISA"
};
```

20.20.3.2 Register Address

```
long devRegisterAddress(
    const char *pOwnerName,
    epicsAddressType addrType,
    size_t logicalBaseAddress,
    size_t size, /* bytes */
    volatile void **pLocalAddress);
```

This routine is called to register a VME address. The routine keeps a list of all VME address ranges requested and returns an error message if an attempt is made to register any addresses that overlap a range that is already being used. `*pLocalAddress` is set equal to the address as seen by the caller.

20.20.3.3 Print Address Map

```
long devAddressMap(void)
```

This routine displays the table of registered VME address ranges, including the owner of each registered address.

20.20.3.4 Unregister Address

```
long devUnregisterAddress(
    epicsAddressType addrType,
    size_t logicalBaseAddress,
    const char *pOwnerName);
```

This routine releases address ranges previously registered by a call to `devRegisterAddress` or `devAllocAddress`.

20.20.3.5 Allocate Address

```
long devAllocAddress(
    const char *pOwnerName,
    epicsAddressType addrType,
    size_t size,
    unsigned alignment, /*number of low zero bits needed in addr*/
    volatile void **pLocalAddress);
```

This routine is called to request the library to allocate an address block of a particular address type. This is useful for devices that appear in more than one address space and can program the base address of one window using registers found in another window.

20.20.4 Interrupt Connection Routines

20.20.4.1 Connect

```
long devConnectInterruptVME(  
    unsigned vectorNumber,  
    void (*pFunction)(void *),  
    void *parameter);
```

Connect ISR `pFunction` up to the VME interrupt `vectorNumber`.

20.20.4.2 Disconnect

```
long devDisconnectInterruptVME(  
    unsigned vectorNumber,  
    void (*pFunction)(void *));
```

Disconnects an ISR from its VME interrupt vector. The parameter `pFunction` should be set to the C function pointer that was connected. It is used as a key to prevent a driver from inadvertently removing an interrupt handler that it didn't install.

20.20.4.3 Check If Used

```
int devInterruptInUseVME(  
    unsigned vectorNumber);
```

Determines if a VME interrupt vector is in use, returning a boolean value.

20.20.4.4 Enable

```
long devEnableInterruptLevelVME(  
    unsigned level);
```

Enable the given VME interrupt level onto the CPU.

20.20.4.5 Disable

```
long devDisableInterruptLevelVME(  
    unsigned level);
```

Disable VME interrupt level. This routine should generally never be used, since it is impossible for a driver to know whether any other active drivers are still making use of a particular interrupt level.

20.20.5 Macros for Normalized Analog Values

20.20.5.1 Convert Digital Value to a Normalized Double Value

```
#define devCreateMask(NBITS) ((1<<(NBITS))-1)
#define devDigToNml(DIGITAL,NBITS) \
    (((double)(DIGITAL))/devCreateMask(NBITS))
```

20.20.5.2 Convert Normalized Double Value to a Digital Value

```
#define devNmlToDig(NORMAL,NBITS) \
    (((long)(NORMAL)) * devCreateMask(NBITS))
```

20.20.6 Deprecated Interrupt Routines

20.20.6.1 Definitions of Interrupt Types (deprecated)

```
typedef enum {intCPU, intVME, intVXI} epicsInterruptType;
```

The routines that use this typedef have all been deprecated, and currently only exist for backwards compatibility purposes. The typedef will be removed in a future release, along with those routines.

20.20.6.2 Connect (deprecated)

```
long devConnectInterrupt(
    epicsInterruptType intType,
    unsigned vectorNumber,
    void (*pFunction)(),
    void *parameter);
```

This routine has been deprecated, and currently only exists for backwards compatibility purposes. Uses of this routine should be converted to call `devConnectInterruptVME` or related routines instead. This routine will be removed in a future release.

20.20.6.3 Disconnect (deprecated)

```
long devDisconnectInterrupt(
    epicsInterruptType intType,
    unsigned vectorNumber);
```

This routine has been deprecated, and currently only exists for backwards compatibility purposes. Uses of this routine should be converted to call `devDisconnectInterruptVME` or related routines instead. This routine will be removed in a future release.

20.20.6.4 Enable Level (deprecated)

```
long devEnableInterruptLevel(
    epicsInterruptType intType,
    unsigned level);
```

This routine has been deprecated, and currently only exists for backwards compatibility purposes. Uses of this routine should be converted to call `devEnableInterruptLevelVME` or related routines instead. This routine will be removed in a future release.

20.20.6.5 Disable Level (deprecated)

```
long devDisableInterruptLevel(
    epicsInterruptType intType,
    unsigned level);
```

This routine has been deprecated, and currently only exists for backwards compatibility purposes. Uses of this routine should be converted to call `devDisableInterruptLevelVME` or related routines instead. This routine will be removed in a future release.

20.21 vxWorks Specific routines and Headers

The routines described in this section are included in a application by the Makfile command:

```
<appl>_OBJS_vxWorks += $(EPICS_BASE_BIN)/vxComLibrary
```

20.21.1 veclist

This routine shows the vxWorks interrupt vector table, but only works properly on 68K family CPUs.

20.21.2 logMsgToErrlog

This traps all calls to `logMsg` and sends them to `errlogPrintf`.

20.21.3 camacLib.h

This was included with 3.13.

20.21.4 epicsDynLink

This provides the routines `symFindByNameEPICS` and `symFindByNameAndTypeEPICS`. It is only provided for modules that have not been converted to use `epicsFindSymbol`. These routines are deprecated.

20.21.5 module_types.h

This is only provided for device/driver support that have not been converted to use OSI features of base. This header is deprecated. Instead of using this, drivers should register a configuration command to obtain the information originally provided by `module_types.h`

20.21.6 task_params.h

This is only provided for device/driver support that have not been converted to use OSI features of base. This header is deprecated.

20.21.7 vxComLibrary

This routine causes epicsDynLink, logMsgToErrlog and veclist to be loaded.

Chapter 21

Registry

Under vxWorks `osiFindGlobalSymbol()` can be used to dynamically bind to record, device, and driver support and functions for use with subroutine records. However on most other systems this routine is not functional, so a registry facility is provided to implement the binding. Any item that is looked up by name at runtime must be registered for it to be visible to other code.

At compile time a perl script reads the IOCs database definition file and produces a C source file containing a routine which registers all record/device/driver/function support defined in that file.

21.1 Registry.h

```
int registryAdd(void *registryID, const char *name, void *data);
void *registryFind(void *registryID, const char *name);
int registrySetTableSize(int size);
void registryFree();
int registryDump(void);
```

This is the code which implements the symbol table. Each different type of symbol has its own unique ID. Everything to be registered is stored in the same gpHash table.

21.2 registryRecordType.h

```
typedef int (*computeSizeOffset)(dbRecordType *pdbRecordType);

typedef struct recordTypeLocation {
    struct rset *prset;
    computeSizeOffset sizeOffset;
} recordTypeLocation;

int registryRecordTypeAdd(const char *name, recordTypeLocation *prtl);
recordTypeLocation *registryRecordTypeFind(const char *name);
```

Provides addresses for both the record support entry table and the routine which computes the size and offset of each field.

21.3 registryDeviceSupport.h

```
int registryDeviceSupportAdd(const char *name, struct dset *pdset)
struct dset *registryDeviceSupportFind(const char *name);
```

This provides addresses for device support entry tables.

21.4 registryDriverSupport.h

```
int registryDriverSupportAdd(const char *name, struct drvet *pdrvet);
struct drvet *registryDriverSupportFind(const char *name);
```

```
int registerRecordDeviceDriver(DBBASE *pdbname);
```

This provides addresses for driver support tables.

21.5 registryFunction.h

```
typedef void (*REGISTRYFUNCTION) (void);
```

```
typedef struct registryFunctionRef {
    const char *name;
    REGISTRYFUNCTION addr;
} registryFunctionRef;
```

```
int registryFunctionAdd(const char *name, REGISTRYFUNCTION func);
REGISTRYFUNCTION registryFunctionFind(const char *name);
int registryFunctionRefAdd(registryFunctionRef ref[], int nfunctions);
```

registryFunctionAdd registers a single function. registryFunctionRefAdd registers several functions.

If you use these routines to register functions directly instead of using a function() statement in a database definition file, the registered functions will not appear in the output from the dbDumpFunction command.

21.6 registerRecordDeviceDriver.c

A version of this is provided for vxWorks. This version makes it unnecessary to use registerRecordDeviceDriver.pl or register other external names. Thus for vxWorks everything can work almost exactly like it did in release 3.13.x

21.7 registerRecordDeviceDriver.pl

This is the perl script which creates a C source file that registers record/device/driver/function support. The following steps are take as part of the standard Make rules:

- Execute this script using a dbd file created by dbExpand
- Compile the resulting C++ file
- Include the object file in the IOC executable

Chapter 22

Database Structures

22.1 Overview

This chapter describes the internal structures describing an IOC database. It is of interest to EPICS system developers but serious application developers may also find it useful. This chapter is intended to make it easier to understand the IOC source listings. It also gives a list of the header files used by IOC Code.

22.2 Include Files

This section lists the files in base/include that are of most interest to IOC Application Developers:

alarm.h alarmString.h - These files contain definitions for all alarm status and severity values.

caDef.h caErr.h caEventMask.h - These files are of interest to anyone writing channel access clients.

callback.h - The definitions for the General Purpose callback system.

dbAccess.h - Definitions for the runtime database access routines.

dbBase.h - Definitions for the structures used to store an EPICS database.

dbDefs.h - A catchall file for definitions that have no other reasonable place to appear.

dbFldTypes.h - Definitions for DBF_XXX and DBR_XXX types.

dbScan.h - Definitions for the scanning system.

dbStaticLib.h - The static databases access system.

db.access.h db_addr.h - Old database access.

devLib.h - The device support library

devSup.h - Device Support Modules

drvSup.h - Driver Support Modules

e11Lib.h - A library that provides the same functions as the vxWorks `1stLib`. All routines start with `e11` instead of `1st`. The `e11Lib` routines work on all architectures.

epicsPrint.h errMdef.h - EPICS error handling system

fast_lock.h - The FASTLOCK routines.

freeList.h - A general purpose free list facility

gpHash.h - A general purpose hash library.

guigroup.h - The guigroup definitions.

initHooks.h - Definitions used by `initHooks.c` routines.

link.h - Link definitions

module_types.h - VME hardware configuration. SHOULD NOT BE USED BY NEW SUPPORT.

recSup.h - The record global routines.


special.h - Definitions for special fields, i.e. `SPC_XXX`.

task_params.h - Definitions for task priorities, stack space, etc.

taskwd.h - Task Watchdog System

tsDefs.h - Time stamp routines. Will also have to look at `base/src/libCom/tsSubr.c`

22.3 Structures



Index

- AB_IO – link field value, 109
- Access Security, 127
 - Configuration File, 129
 - Initialization, 131
- Access Security Group, 131
- Access Security Level, 141
- add_record, 122, 193
- addpath, 96
- addpath – Database Definitions, 98
- adjustment.h, 280
- adjustToWorstCaseAlignment, 281
- algorithm, 266
- alias, 96
- alias_name – record instance definition, 107
- Alloc/Free DBENTRY, 201
- Application Specific Configuration, 37
- archclean, 42
- as_level – field definition, 101
- asAddClient, 138
- asAddMember, 137
- ascar, 159
- asChangeClient, 138
- asChangeGroup, 138
- ascheck, 131
- asCheckGet, 139
- asCheckPut, 139
- asCompute, 140
- asComputeAllAsg, 140
- asComputeAsg, 140
- asdbdump, 143, 157
- asdbdumpFP, 143
- asDbGetAsl, 142
- asDbGetMemberPvt, 142
- asDump, 140
- asDumpFP, 140
- asDumpHag, 140
- asDumpHagFP, 140
- asDumpHash, 141
- asDumpHashFP, 141
- asDumpMem, 141
- asDumpMemFP, 141
- asDumpRules, 141
- asDumpRulesFP, 141
- asDumpUag, 140
- asDumpUagFP, 140
- ASG, 128
 - Access Security configuration, 130
- asGetClientPvt, 139
- asGetMemberPvt, 137
- asInit, 122, 131, 142, 157
- asInitAsyn, 142
- asInitFile, 137
- asInitFP, 137
- asInitialize, 137
- ASL, 128
- asl – field definition rules, 100
- ASL0, 141
- ASL1, 141
- asphag, 143, 157
- asphagFP, 143
- aspmem, 143, 158
- aspmemFP, 143
- asprules, 143, 157
- asprulesFP, 143
- aspuag, 143, 157
- aspuagFP, 143
- asPutClientPvt, 139
- asPutMemberPvt, 138
- asPvt in DBADDR, 217
- asRegisterClientCallback, 139
- asRemoveClient, 138
- asRemoveMember, 137
- asSetFilename, 127, 131, 142, 157
- asSetSubstitutions, 128, 131, 142
- asSubInit, 127, 132, 143
- asSubProcess, 127, 132, 143
- astac, 143
- asTrapWriteAfter, 139, 145
- asTrapWriteBefore, 139, 145
- asTrapWriteId, 145
- asTrapWriteListener, 145
- asTrapWriteMessage, 145
- asTrapWriteRegisterListener, 145
- asTrapWriteUnregisterListener, 145
- asynchronous device support example, 189
- Asynchronous Records, 89
- aToIPAddr, 280
- base – field definition, 102
- base – field definition rules, 100

- base.dbd, 45
- BBGPIB_IO – link field value, 109
- bin directory, 34
- BIN_INSTALLS, 67
- BITBUS_IO – link field value, 109
- breakpoint table – Database Definitions, 106
- Breakpoint Tables, 44, 111
- Breakpoints, 152
- breaktable, 96, 106
- bucketLib.h, 261
- Build Facility, 33
- Build Requirements, 35
- Build software prerequisites, 35
- BuildingR3.13AppsWithR3.14.html, 82
- BuildingR3.13ExtensionsWithR3.14.html, 82

- C++ library, 265
- ca_put_callback, 91
- Cached Puts, 91
- CALC
 - Access Security configuration, 130
- CALC_ERR., 261
- calcArgUsage, 261
- calcErrorStr, 261
- calcPerform, 261
- CALLBACK, 238
- callback.h, 237
- callbackCancelDelayed, 238
- callbackGetUser, 238
- callbackInit, 121, 238
- callbackRequest, 238
- callbackRequestDelayed, 238
- callbackRequestProcessCallback, 238
- callbackRequestProcessCallbackDelayed, 238
- callbackSetCallback, 238
- callbackSetPriority, 238
- callbackSetProcess, 238
- callbackSetQueueSize, 124, 238, 240
- callbackSetUser, 238
- calloc, 281
- callocMustSucceed, 281
- CAMAC_IO – link field value, 109
- camacLib.h, 317
- cantProceed, 281
- cantProceed.h, 281
- casr, 158
- cdCommands, 120
- Cexp, 24
- CFG, 70
- cfg directory, 34
- Channel Access, 25
- channel access link, 83
- Channel Access Monitors, 230
- CHECK_RELEASE, 38
- checkAlarms, 178
- choice, 96
- choice_string – device definition, 104
- class templates, 265
- clean, 42
- ClockTime_Report, 156
- ClockTime_Shutdown, 156
- comment – Database Definitions, 98
- COMPAT_313, 37
- COMPAT_TOOLS_313, 37
- compatibility configuration, 37
- Compiler options, 47
- computeSizeOffset, 319
- Conditionals, 256
- CONFIG, 79
- CONFIG files, user created, 70
- CONFIG.Common.target, 80
- CONFIG.CrossCommon, 78
- CONFIG.gnuCommon, 78
- CONFIG.host.Common, 80
- CONFIG.host.target, 80
- CONFIG.host.vxWorksCommon, 80
- CONFIG.UnixCommon.Common, 80
- CONFIG_ADDONS, 78
- CONFIG_APP_INCLUDE, 78
- CONFIG_BASE, 79
- CONFIG_BASE_VERSION, 79
- CONFIG_COMMON, 79
- CONFIG_COMPAT, 80
- CONFIG_ENV, 79
- CONFIG_FILE_TYPE, 79
- CONFIG_SITE, 36, 79
- CONFIG_SITE.Common.target, 37, 81
- CONFIG_SITE.Common.vxWorksCommon, 37
- CONFIG_SITE.host.Common, 36, 81
- CONFIG_SITE.host.host, 36
- CONFIG_SITE.host.target, 37, 80
- CONFIG_SITE_ENV, 36, 79
- CONFIG_USER, 40
- configure directory, 34
- Configure files, 78
- configure/os File, 80
- CONSTANT – link field value, 108
- constant link, 83
- CONTAINER, 281
- ConvertingR3.13AppsToR3.14.html, 82
- convertRelease, 38
- coreRelease, 121, 160
- cross target, 35
- CROSS_COMPILER_HOST_ARCHS, 35
- CROSS_COMPILER_TARGET_ARCHS, 35, 38, 39
- cvsclean, 43
- cvsclean.pl, 81
- cvt_dbaddr - Record Support Routine, 181

- cvtBitsToUlong, 266
- cvtCharToString, 266
- cvtDoubleToCompactString, 266
- cvtDoubleToExpString, 266
- cvtDoubleToString, 266
- cvtFast.h, 266
- cvtFloatToCompactString, 266
- cvtFloatToExpString, 266
- cvtFloatToString, 266
- cvtLongToHexString, 266
- cvtLongToOctalString, 266
- cvtLongToString, 266
- cvtShortToString, 266
- cvtUcharToString, 266
- cvtUlongToBits, 266
- cvtUlongToString, 266
- cvtUshortToString, 266
- cygwin.bat, 82

- database access routines - List of, 218
- Database Configuration Field Types, 200
- Database Definition, 45, 46, 95
- Database Definition File, 95
- database definitions, 95
- Database Files, 47
- Database Format – Summary, 95
- database link, 83
- Database Link Guidelines, 86
- Database Links, 83
- Database Locking, 85
- Database Scanning, 85
- Database Structures, 321
- DB, 47
- db directory, 34
- DB install files, 47
- DB_INSTALLS, 47
- DB_MAX_CHOICES, 215
- db_post_events, 230
- dba, 150
- dbAccess.h, 215
- dbAdd, 217
- dbAddPath, 202
- DBADDR, 217
- dbAllocBase, 201
- dbAllocEntry, 201
- dbAllocForm, 210
- dbap, 153
- dbb, 152
- dbBkptInit, 122
- dbBufferSize, 226
- dbc, 152
- dbCaAddLink, 233
- dbCaAddLinkCallback, 233
- dbCaGetAlarmLimits, 235
- dbCaGetAttributes, 235
- dbCaGetControlLimits, 235
- dbCaGetGraphicLimits, 235
- dbCaGetLink, 233
- dbCaGetLinkDBFtype, 235
- dbCaGetNelements, 234
- dbCaGetPrecision, 236
- dbCaGetSevr, 235
- dbCaGetTimeStamp, 235
- dbCaGetUnits, 236
- dbCaIsLinkConnected, 234
- dbCaLinkInit, 121, 233
- dbCaPause, 123
- dbCaPutLink, 233
- dbCaPutLinkCallback, 233
- dbcar, 159, 161
- dbCaRemoveLink, 233
- dbCaRun, 123
- dbCopyEntry, 202
- dbCopyEntryContents, 202
- dbCopyRecord, 207
- dbCreateAlias, 207
- dbCreateRecord, 207
- dbCvtLinkToConstant, 210
- dbCvtLinkToPvlink, 210
- DBD, 44–46
- dbd, 152
- dbd directory, 34
- DBD install files, 46
- DBD_INSTALLS, 47
- dbDefs.h, 215, 281
- dbDeleteAliases, 207
- dbDeleteInfo, 212
- dbDeleteRecord, 207
- DBDINC, 44, 45
- dbDumpBreaktable, 213
- dbDumpDevice, 163, 213
- dbDumpDriver, 163, 213
- dbDumpField, 163
- dbDumpFldDes, 213
- dbDumpFunction, 213
- dbDumpMenu, 162, 213
- dbDumpPath, 213
- dbDumpRecord, 213
- dbDumpRecords, 163, 214
- dbDumpRecordType, 163, 213
- dbDumpRegistrar, 213
- dbDumpVariable, 213
- DBE_ALARM, 179
- DBE_LOG, 179
- DBE_VAL, 179
- dbel, 158
- dbExpand, 45, 115, 213
- DBF_CHAR, 216

- DBF_DEVICE, 216
- DBF_DOUBLE, 216
- DBF_ENUM, 216
- DBF_FLOAT, 216
- DBF_FWDLINK, 216
- DBF_INLINK, 216
- DBF_LONG, 216
- DBF_MENU, 216
- DBF_NOACCESS, 216
- DBF_OUTLINK, 216
- DBF_SHORT, 216
- DBF_STRING, 216
- DBF_UCHAR, 216
- DBF_ULONG, 216
- DBF_USHORT, 216
- dbFindBrkTable, 212
- dbFindField, 207
- dbFindInfo, 211
- dbFindMenu, 209
- dbFindRecord, 206
- dbFindRecordType, 204
- dbFinishEntry, 201
- dbFirstField, 204
- dbFirstInfo, 211
- dbFirstRecord, 206
- dbFirstRecordType, 204
- dbFldTypes.h, 215, 216
- dbFoundField, 204, 207
- dbFreeBase, 201
- dbFreeEntry, 201
- dbFreeForm, 210
- dbFreeRecords, 207
- dbGet, 222
- dbGetAlarmLimits, 228
- dbGetControlLimits, 228
- dbGetDefaultName, 205
- dbGetField, 85, 222
- dbGetFieldIndex, 227
- dbGetFieldName, 205
- dbGetFieldType, 204
- dbGetFormPrompt, 210
- dbGetFormValue, 210
- dbGetGraphicLimits, 228
- dbGetInfo, 212
- dbGetInfoName, 211
- dbGetInfoPointer, 212
- dbGetInfoString, 212
- dbGetLink, 84, 86, 222
- dbGetLinkDBFtype, 228
- dbGetLinkField, 209
- dbGetLinkType, 209
- dbGetMenuChoices, 208
- dbGetMenuIndex, 208
- dbGetMenuIndexFromString, 208
- dbGetMenuStringFromIndex, 208
- dbGetNAliases, 206
- dbGetNelements, 227
- dbGetNFields, 204
- dbGetNLinks, 209
- dbGetNMenuChoices, 208
- dbGetNRecords, 206
- dbGetNRecordTypes, 203
- dbGetPdbAddrFromLink, 227
- dbGetPrecision, 228
- dbGetPrompt, 205
- dbGetPromptGroup, 205
- dbGetRange, 208
- dbGetRecordAttribute, 205
- dbGetRecordName, 206
- dbGetRecordTypeName, 204
- dbGetRest, 226
- dbGetRset, 227
- dbGetSevr, 229
- dbGetString, 208
- dbGetTimeStamp, 229
- dbGetUnits, 228
- dbgf, 151
- dbgrep, 150
- dbhcr, 154, 161
- dbInitEntry, 201
- dbInvisibleRecord, 207
- dbior, 154
- dbIsAlias, 206
- dbIsDefaultValue, 208
- dbIsLinkConnected, 227
- dbIsValueField, 227
- dbIsVisibleRecord, 207
- dbl, 149
- dbla, 150
- dbLoadDatabase, 116
- dbLoadRecords, 116
- dbLoadTemplate, 117
- dbLockGetLockId, 231
- dbLockInitRecords, 122, 231
- dbLockSetGblLock, 232
- dbLockSetGblUnlock, 232
- dbLockSetMerge, 231
- dbLockSetRecordLock, 232
- dbLockSetSplit, 232
- dbLockShowLocked, 161
- dblsr, 161
- dbmf.h, 267
- dbmfFree, 268
- dbmfFreeChunks, 268
- dbmfInit, 268
- dbmfMalloc, 268
- dbmfShow, 268
- dbNameToAddr, 221

- dbNextField, 204
- dbNextInfo, 211
- dbNextRecord, 206
- dbNextRecordType, 204
- dbNotifyAdd, 226
- dbNotifyCancel, 226
- dbNotifyCompletion, 226
- dbnr, 152
- dbp, 153
- dbPath, 202
- dbpf, 151
- dbpr, 151
- dbProcess, 230
- dbPut, 224
- dbPutAttribute, 111, 229
- dbPutField, 85, 86, 223
- dbPutForm, 210
- dbPutInfo, 212
- dbPutInfoPointer, 212
- dbPutInfoString, 212
- dbPutLink, 84, 86, 223
- dbPutMenuIndex, 208
- dbPutNotify, 91, 224, 225
- dbPutNotifyInit, 122
- dbPutRecordAttribute, 205
- dbPutString, 208
- dbPvdDump, 164, 213
- dbPvdTableSize, 124
- DBR_AL_DOUBLE, 220
- DBR_AL_LONG, 220
- DBR_CHAR, 219
- DBR_CTRL_DOUBLE, 220
- DBR_CTRL_LONG, 220
- DBR_DOUBLE, 219
- DBR_ENUM, 219
- DBR_ENUM_STRS, 220
- dbr_field_type in DBADDR, 217
- DBR_FLOAT, 219
- DBR_GR_DOUBLE, 220
- DBR_GR_LONG, 220
- DBR_LONG, 219
- DBR_PRECISION, 220
- DBR_PUT_ACKS, 219, 220
- DBR_PUT_ACKT, 219, 220
- DBR_SHORT, 219
- DBR_STATUS, 220
- DBR_STRING, 219
- DBR_TIME, 220
- DBR_UCHAR, 219
- DBR_ULONG, 219
- DBR_UNITS, 220
- DBR_USHORT, 219
- dbReadDatabase, 202
- dbReadDatabaseFP, 202
- dbReadTest, 118
- dbRecordsOnceOnly, 107
- dbRenameRecord, 207
- dbReportDeviceConfig, 213
- dbs, 152
- dbScan.h, 245
- dbScanFwdLink, 229
- dbScanLink, 229
- dbScanLock, 85, 231
- dbScanPassive, 84, 85, 229
- dbScanUnlock, 85, 231
- dbstat, 153
- dbtgif, 160
- dbToMenuH, 112, 113
- dbToRecordtypeH, 112, 113
- dbtspf, 160
- dbtpn, 160
- dbtr, 151
- dbValueSize, 226
- dbVerify, 208
- dbVerifyForm, 210
- dbVisibleRecord, 207
- dbWriteBreaktable, 203
- dbWriteBreaktableFP, 203
- dbWriteDevice, 203
- dbWriteDeviceFP, 203
- dbWriteDriver, 203
- dbWriteDriverFP, 203
- dbWriteFunctionFP, 203
- dbWriteMenu, 203
- dbWriteMenuFP, 203
- dbWriteRecord, 203
- dbWriteRecordFP, 203
- dbWriteRecordType, 203
- dbWriteRecordTypeFP, 203
- dbWriteRegistrarFP, 203
- dbWriteVariableFP, 203
- DCT_FWDLINK, 200
- DCT_INLINK, 200
- DCT_INTEGER, 200
- DCT_LINK_CONSTANT, 209
- DCT_LINK_DEVICE, 209
- DCT_LINK_FORM, 209
- DCT_LINK_PV, 209
- DCT_MENU, 200
- DCT_MENUFORM, 200
- DCT_NOACCESS, 200
- DCT_OUTLINK, 200
- DCT_REAL, 200
- DCT_STRING, 200
- del_record, 194
- devAddressMap, 314
- devAllocAddress, 315
- devAoSoftCallback.c, 233

- devConnectInterrupt, 316
- devConnectInterruptVME, 315
- devCreateMask, 316
- devDisableInterruptLevel, 317
- devDisableInterruptLevelVME, 315
- devDisconnectInterrupt, 316
- devDisconnectInterruptVME, 315
- devEnableInterruptLevel, 317
- devEnableInterruptLevelVME, 315
- device, 96
- device – Database Definitions, 104
- Device Support, 187
- Device Support Entry Table, 173
- devInterruptInUseVME, 315
- devNmlToDig, 316
- devNoResponseProbe, 313
- devReadProbe, 313
- devRegisterAddress, 314
- devUnregisterAddress, 314
- devWriteProbe, 313
- Directory names restriction, 36
- Directory structure, 33
- distclean, 43
- DLL_LIBS, 55
- Doc file, 59
- DOCS, 59
- Documentation Directory, 81
- Documentation Files, 81
- dos2unix.pl, 81
- driver, 96
- driver – database definition, 104
- Driver Support, 195
- Driver Support Entry Table Example, 196
- drvnet_name – driver definition, 104
- DSET, 173
- dset - dbCommon, 187
- dset_name – device definition, 104
- dtyp - dbCommon, 187

- Elements of build system, 34
- ellAdd, 268
- ellConcat, 268
- ellCount, 268
- ellDelete, 268
- ellExtract, 268
- ellFind, 268
- ellFirst, 268
- ellFree, 268
- ellFree2, 268
- ellGet, 268
- ellInit, 268
- ellInsert, 268
- ellLast, 268
- ellLib.h, 268
- ELLIST, 268
- ellNext, 268
- ELLNODE, 268
- ellNStep, 268
- ellNth, 268
- ellPrevious, 268
- ellVerify, 268
- eltc, 153, 168
- eng_value – breakpoint table, 107
- Environment Prerequisites, 35
- Environment Variables, 126
- envPaths, 119
- EPICS, 9, 25
 - Basic Attributes, 26
- EPICS_CA_ADDR_LIST, 126
- EPICS_CA_AUTO_ADDR_LIST, 126
- EPICS_CA_CONN_TMO, 126
- EPICS_CA_MAX_ARRAY_BYTES, 126
- EPICS_CA_REPEATER_PORT, 126
- EPICS_CA_SERVER_PORT, 126
- EPICS_CAS_BEACON_PERIOD, 126
- EPICS_HOST_ARCH, 35, 36
- EPICS_IOC_LOG_FILE_COMMAND, 169
- EPICS_IOC_LOG_FILE_LIMIT, 169
- EPICS_IOC_LOG_FILE_NAME, 169
- EPICS_IOC_LOG_INET, 126
- EPICS_IOC_LOG_PORT, 126, 170
- EPICS_THREAD_ONCE_INIT, 300
- EPICS_TS_NTP_INET, 126
- epicsAddressType, 314
- epicsAddressTypeName, 314
- epicsAlgorithm.h, 266
- epicsAssert, 288
- epicsAtExit, 266
- epicsAtThreadExit, 266
- epicsConvert.h, 281
- epicsConvertDoubleToFloat, 281
- epicsDynLink, 317
- epicsEndian.h, 288
- epicsEnvSet, 126, 253
- epicsEnvShow, 126, 159
- epicsEvent, 289
- epicsEvent.h, 289
- epicsEventCreate, 290
- epicsEventDestroy, 290
- epicsEventId, 290
- epicsEventMustCreate, 290
- epicsEventMustWait, 290
- epicsEventShow, 290
- epicsEventSignal, 290
- epicsEventTryWait, 290
- epicsEventWait, 290
- epicsEventWaitError, 289
- epicsEventWaitOK, 289

- epicsEventWaitStatus, 289
- epicsEventWaitTimeout, 289
- epicsEventWaitWithTimeout, 290
- epicsExit, 266
- epicsExitCallAtExits, 266
- epicsExitCallAtThreadExits, 266
- epicsExport, 17
- epicsExportAddress, 17, 105
- epicsExportRegistrar, 18, 105
- epicsFindSymbol, 290
- epicsFindSymbol.h, 290
- epicsGeneralTime, 291
- epicsGeneralTime.h, 291
- epicsGetStderr, 298
- epicsGetStdin, 298
- epicsGetStdout, 298
- epicsGetThreadStderr, 298
- epicsGetThreadStdin, 298
- epicsGetThreadStdout, 298
- EpicsHostArch, 82
- EpicsHostArch.pl, 36, 82
- epicsInterrupt, 293
- epicsInterrupt.h, 293
- epicsInterruptContextMessage, 294
- epicsInterruptIsInterruptContext, 294
- epicsInterruptLock, 294
- epicsInterruptType, 316
- epicsInterruptUnlock, 294
- epicsMath, 294
- epicsMath.h, 294
- epicsMax, 266
- epicsMemHash, 282
- epicsMessageQueue, 294–296
- epicsMessageQueue.h, 294
- epicsMessageQueueCreate, 296
- epicsMessageQueueDestroy, 296
- epicsMessageQueuePending, 296
- epicsMessageQueueReceive, 296
- epicsMessageQueueReceiveWithTimeout, 296
- epicsMessageQueueSend, 296
- epicsMessageQueueSendWithTimeout, 296
- epicsMessageQueueShow, 296
- epicsMessageQueueTryReceive, 296
- epicsMessageQueueTrySend, 296
- epicsMin, 266
- epicsMutex, 296, 297
- epicsMutex.h, 296
- epicsMutexCreate, 297
- epicsMutexDestroy, 297
- epicsMutexId, 297
- epicsMutexLock, 297
- epicsMutexLockError, 297
- epicsMutexLockOK, 297
- epicsMutexLockStatus, 297
- epicsMutexLockTimeout, 297
- epicsMutexMustCreate, 297
- epicsMutexMustLock, 297
- epicsMutexShow, 297
- epicsMutexShowAll, 297
- epicsMutexTryLock, 297
- epicsMutexUnlock, 297
- epicsParamShow, 159
- epicsPrintf, 167
- epicsPrtEnvParams, 126
- epicsRegisterFunction, 17, 106
- epicsRingBytes.h, 269
- epicsRingBytesCreate, 269
- epicsRingBytesDelete, 269
- epicsRingBytesFlush, 269
- epicsRingBytesFreeBytes, 269
- epicsRingBytesGet, 269
- epicsRingBytesId, 269
- epicsRingBytesIsEmpty, 269
- epicsRingBytesIsFull, 269
- epicsRingBytesPut, 269
- epicsRingBytesSize, 269
- epicsRingBytesUsedBytes, 269
- epicsRingPointer, 269
- epicsRingPointer.h, 269
- epicsRtmsInitHooks.h, 121
- epicsRtmsInitPostSetBootConfigFromNVRAM, 120
- epicsRtmsInitPreSetBootConfigFromNVRAM, 120
- epicsScanDouble, 298
- epicsScanFloat, 298
- epicsSetThreadStderr, 298
- epicsSetThreadStdin, 298
- epicsSetThreadStdout, 298
- epicsShareAPI, 283
- epicsShareFunc, 283
- epicsSignal.h, 312
- epicsSignalInstallSigHupIgnore, 121
- epicsSnprintf, 298
- epicsStdio, 298
- epicsStdio.h, 298
- epicsStdioRedirect, 299
- epicsStdlib, 298
- epicsStdlib.h, 298
- epicsStdoutPrintf, 298
- epicsStrCaseCmp, 282
- epicsStrDup, 282
- epicsStrGlobMatch, 282
- epicsStrHash, 282
- epicsString.h, 281
- epicsStrnCaseCmp, 282
- epicsStrnEscapedFromRaw, 282
- epicsStrnEscapedFromRawSize, 282
- epicsStrnRawFromEscaped, 282
- epicsStrPrintEscaped, 282

- epicsStrtod, 298
- epicsStrtok_r, 282
- epicsSwap, 266
- epicsTempFile, 298
- epicsTempName, 298
- epicsThread, 299, 303
- epicsThread.h, 299
- epicsThreadBooleanStatus, 300
- epicsThreadBooleanStatusFail, 300
- epicsThreadBooleanStatusSuccess, 300
- epicsThreadCreate, 300
- epicsThreadExitMain, 300
- EPICSTHREADFUNC, 300
- epicsThreadGetId, 300
- epicsThreadGetIdSelf, 300
- epicsThreadGetName, 300
- epicsThreadGetNameSelf, 300
- epicsThreadGetPriority, 300
- epicsThreadGetPrioritySelf, 300
- epicsThreadGetStackSize, 300
- epicsThreadHighestPriorityLevelBelow, 300
- epicsThreadId, 300
- epicsThreadInit, 300
- epicsThreadIsEqual, 300
- epicsThreadIsOkToBlock, 121, 300
- epicsThreadIsSuspended, 300
- epicsThreadLowestPriorityLevelAbove, 300
- epicsThreadOnce, 300
- epicsThreadOnceId, 300
- epicsThreadPriorityChannelAccessServer, 300
- epicsThreadPriorityHigh, 300
- epicsThreadPriorityLow, 300
- epicsThreadPriorityMax, 300
- epicsThreadPriorityMedium, 300
- epicsThreadPriorityMin, 300
- epicsThreadPriorityScanHigh, 300
- epicsThreadPriorityScanLow, 300
- epicsThreadPrivateCreate, 300
- epicsThreadPrivateDelete, 300
- epicsThreadPrivateGet, 300
- epicsThreadPrivateId, 300
- epicsThreadPrivateSet, 300
- epicsThreadResume, 300
- epicsThreadRunnable, 303
- epicsThreadSetOkToBlock, 300
- epicsThreadSetPriority, 300
- epicsThreadShow, 300
- epicsThreadShowAll, 300
- epicsThreadSleep, 300
- epicsThreadSleepQuantum, 300
- epicsThreadStackBig, 300
- epicsThreadStackMedium, 300
- epicsThreadStackSizeClass, 300
- epicsThreadStackSmall, 300
- epicsThreadSuspendSelf, 300
- epicsTime, 304, 307, 311
- epicsTime.h, 304
- epicsTimeEvent, 307
- epicsTimeGetCurrent, 293
- epicsTimeGetCurrentInt, 291
- epicsTimeGetEvent, 293
- epicsTimeGetEventInt, 291
- epicsTimer, 271
- epicsTimer.h, 271
- epicsTimerId, 274
- epicsTimerQueueActive, 273
- epicsTimerQueueId, 274
- epicsTimerQueueNotify, 273
- epicsTimerQueuePassive, 273
- epicsTimeStamp, 305
- epicsTypes.h, 282
- epicsUnitTest.h, 283
- epicsVprintf, 167
- epicsVsnprintf, 298
- errlog Listeners, 167
- errlogAddListener, 168
- errlogFatal, 166
- errlogFlush, 166
- errlogGetSevEnumString, 166
- errlogGetSevToLog, 166
- errlogInfo, 166
- errlogInit, 124, 168
- errlogInit2, 124, 168
- errlogListener, 168
- errlogMajor, 166
- errlogMessage, 166
- errlogMinor, 166
- errlogPrintf, 166, 184, 317
- errlogRemoveListener, 168
- errlogSetSevToLog, 166
- errlogSevEnum, 166
- errlogSevPrintf, 166
- errlogSevVprintf, 166
- errlogThread, 168
- errlogVprintf, 166
- errMdef.h, 169
- errMessage, 166
- errPrintf, 166, 167
- errVerbose, 168
- Escape Sequence, 98
- Event, 243
- Event - Scan Type, 243
- Event Scanning, 85, 248
- EVNT - Scan Related Field, 244
- Example IOC Application, 13
- Expanded Database Definition Files, 46
- Expanded DBD files, 45
- expandVars.pl, 80, 81

- Expression Syntax, 262
- Extended device support, 187, 192
- Extension configuration, 37
- extra – field definition rules, 100
- extra_info – field definition, 103

- Features of build system, 34
- field, 96
- field definition rules, 100
- field value – record instance definition, 107
- field_name – field definition, 101
- field_name – record instance definition, 107
- field_size in DBADDR, 217
- field_type – field definition, 101
- field_type in DBADDR, 217
- File Types, 70
- File types, user created, 70
- FILE_TYPE, 70
- filename extension conventions, 98
- filterWarnings.pl, 81
- freeList.h, 277
- freeListCalloc, 277
- freeListCleanup, 277
- freeListFree, 277
- freeListInitPvt, 277
- freeListItemsAvail, 277
- freeListMalloc, 277
- fullpathname.pl, 81
- function, 17, 96
- function – Database Definitions, 106
- function templates, 265
- function_name – function definition, 106
- function_name – registrar definition, 105
- FWDLINK, 84

- generalTime, 291
- generalTime_Init, 291
- generalTimeAddIntCurrentProvider, 292
- generalTimeCurrentProviderName, 291
- generalTimeCurrentTpName, 291
- generalTimeCurrentTpRegister, 292
- generalTimeEventProviderName, 291
- generalTimeEventTpName, 291
- generalTimeEventTpRegister, 292
- generalTimeGetErrorCounts, 291
- generalTimeGetExceptPriority, 292
- generalTimeHighestCurrentName, 291
- generalTimeRegisterCurrentProvider, 292
- generalTimeRegisterEventProvider, 292
- generalTimeRegisterIntEventProvider, 292
- generalTimeReport, 155, 291
- generalTimeResetErrorCounts, 291
- generalTimeSup.h, 292
- get_alarm_double Record Support Routine, 183
- get_array_info - Record Support Routine, 181
- get_control_double - Record Support Routine, 183
- get_enum_str - record Support Routine, 182
- get_enum_strs - record Support Routine, 182
- get_graphic_double - example, 177
- get_graphic_double - Record Support Routine, 183
- get_ioint_info, 247
- get_ioint_info - device support routine, 191
- get_precision - Record Support Routine, 182
- get_units - .example, 177
- get_units - Record Support Routine, 181
- gft, 162
- Global Record Support Routines, 183
- GNU make, 35
- gnumake, 42
- gphAdd, 277
- gphHash.h, 277
- gphDelete, 277
- gphDump, 277
- gphDumpFP, 277
- GPENTRY, 277
- gphFind, 277
- gphFreeMem, 277
- gphInitPvt, 277
- GPIB_IO – link field value, 109
- grecord, 96
- gui_group – field definition, 101

- HAG, 128, 130
- Header dependencies, 43
- Host makefile targets, 40
- hpux, 82
- Html, 59
- html directory, 34
- HTMLS, 59

- I/O Event - Scan Type, 243
- I/O Event scanned, 243
- I/O Event Scanning, 85, 246, 249
- INC, 59
- include, 96
- include – Database Definitions, 99
- include directory, 34
- Include File Generation, 112
- Include files, 59
- Infinite Loop, 90
- INFIX_TO_POSTFIX_SIZE, 261
- info, 96
- info_name – record instance definition, 110
- info_value – record instance definition, 110
- Information item pointer, 111
- Information item, 110
- init - device support routine, 191
- init - Record Support Routine, 180

- init_record, 122
- init_record - device support routine, 191
- init_record - example, 174
- init_record - record support routine, 180
- init_value – field definition, 101
- initDatabase, 122
- initDevSup, 122
- initDrvSup, 122
- initHookAfterCaLinkInit, 121
- initHookAfterCallbackInit, 121
- initHookAfterCaServerInit, 123
- initHookAfterCaServerPaused, 123
- initHookAfterCaServerRunning, 123
- initHookAfterDatabasePaused, 123
- initHookAfterDatabaseRunning, 123
- initHookAfterFinishDevSup, 122
- initHookAfterInitDatabase, 122
- initHookAfterInitDevSup, 122
- initHookAfterInitDrvSup, 122
- initHookAfterInitialProcess, 123
- initHookAfterInitRecSup, 122
- initHookAfterInterruptAccept, 123
- initHookAfterIocBuilt, 123
- initHookAfterIocPaused, 123
- initHookAfterIocRunning, 123
- initHookAfterScanInit, 122
- initHookAtBeginning, 121
- initHookAtEnd, 123
- initHookAtIocPause, 123
- initHookAtIocRun, 123
- initHookFunction, 124
- initHookName, 124
- initHookRegister, 124
- initHooks, 124
- initHookState, 124
- initial – field definition rules, 100
- Initialize Device Processing, 191
- Initialize Logging, 126
- Initialize Specific Record, 191
- initialProcess, 123
- initPeriodic, 251
- initRecSup, 122
- INLINK, 83
- INP
 - Access Security configuration, 130
- Input/Output Controller, 9
 - Software Components, 26
- INST_IO – link field value, 109
- Install Directories, 33
- Install Directory definitions, 37
- INSTALL_LOCATION, 33, 37
- installEpics.pl, 81
- Installing Other Binaries, 67
- Installing Other Libraries, 67
- installLastResortEventProvider, 155, 291
- interest – field definition rules, 100
- interest_level – field definition, 102
- interruptAccept, 123
- IOC, 25
- IOC Error Logging, 165
- IOC Initialization, 119
- Ioc makefile targets, 40
- IOC Shell, 253
- iocBuild, 121, 123
- iocInit, 121
- iocLog, 169
- iocLogClient, 170
- iocLogInit, 126
- iocLogServer, 169
- iocPause, 121, 123
- iocRun, 121, 123
- iocsh, 119, 149, 256
- iocsh.h, 256
- iocshArg, 257
- iocshCmd, 256
- iocshFuncDef, 257
- iocshRegister, 18, 257
- isinf, 294
- isnan, 294
- ISO C++, 265

- JAR, 69
- JAR_INPUT, 69
- JAR_MANIFEST, 69
- JAVA, 68
- Java class files, 68
- Java Example, 68
- Java jar file, 69
- Java native methods, 69
- JAVAINC, 70

- Keywords, 96
- KnownProblems.html, 82

- LAN, 25
- Lex and yac, 59
- lib directory, 34
- LIB_INSTALLS, 67
- LIB_LIBS, 54, 55
- LIB_OBJLIBS, 53
- LIB_OBJJS, 52
- LIB_RCS, 68
- LIB_SRCS, 51
- LIB_SYS_LIBS, 55
- LIBOBSJS, 53
- Libraries, 50
- libraries, 54
- LIBRARY, 50
- Library example, 56

- Library link order, 55
- library name, 50
- Library object file, 52
- Library Source file, 51
- Library version number, 55
- LIBRARY, SCRIPTS, 40
- LIBRARY_HOST, 50
- LIBRARY_IOC, 50
- LIBSRCS, 51
- Linear Conversion, 111
- link options, 47
- link.h, 215
- LINK_ALARM, 84
- link_type – device definition, 104
- LINR, 111
- Loadable libraries, 56, 57
- LOADABLE_LIBRARY, 56
- LOADABLE_LIBRARY_HOST, 56
- LOCAL, 281
- logClientCreate, 278
- logClientFlush, 278
- logClientId, 278
- logClientInit, 278
- logClientSend, 278
- logClientSendMessage, 278
- logClientShow, 278
- logMsg, 317

- macCreateHandle, 280
- macDeleteHandle, 280
- macEnvExpand, 280
- macExpandString, 280
- macGetValue, 280
- macInstallMacros, 280
- macLib.h, 278
- macParseDefns, 280
- macPopScope, 280
- macPushScope, 280
- macPutValue, 280
- macReportMacros, 280
- Macro Substitution, 97
- Macro Substitutions and Include tool, 47
- macSuppressWarning, 280
- Make, 42
- Make commands, 42
- Make targets, 43
- makeBaseApp.pl, 19
- makeBpt, 112
- makeDbDepends.pl, 47, 81
- Makefile, 80
- Makefile contents, 41
- Makefile examples, 41
- Makefile name, 41
- Makefiles, 41
- makeIncludeDbd.pl, 81
- makeMakefile.pl, 81
- makeTestfile.pl, 81
- malloc, 281
- mallocMustSucceed, 281
- Maximize Severity, 84, 88
- menu, 96
- menu – Database Definitions, 99
- menu – field definition rules, 100
- menuConvert, 111
- Menus, 45
- menuScan.dbd, 244
- mkmf.pl, 81
- module_types.h, 317
- monitor - example, 179
- MOTLOAD, 24
- MS, 84
- MSI, 47, 84
- msi, 47
- MSS, 84
- Multiple Definitions, 98
- Multiple host, 35
- munch.pl, 81

- name – breakpoint table, 107
- name_CFLAGS, 49
- name_CPPFLAGS, 49
- name_CXXFLAGS, 49
- name_DLL_LIBS, 55
- name_INCLUDES, 49
- name_LDFLAGS, 49
- name_LDOBJ, 53, 62
- name_LIBS, 54, 55, 63, 64
- name_OBJLIBS, 53, 62
- name_OBJS, 52, 61
- name_RCS, 68
- name_SRCS, 52, 58, 63
- name_SYS_LIBS, 54, 55, 64
- NELEMENTS, 281
- NMS, 84
- no_elements in DBADDR, 217
- NOTRAPWRITE, 129
- NPP, 84
- NTP, 291
- NTPTime_Report, 156

- Object Files, 57
- OBJLIB, 57
- OBJLIB_OBJS, 57
- OBJLIB_SRCS, 57
- OBJS, 40, 57
- OBJS_HOST, 57
- OBJS_IOC, 57
- OFFSET, 281

- OPI, 25
- OS_CLASS, 39
- osclass, 39, 71
- OSI_PATH_LIST_SEPARATOR, 283
- OSI_PATH_SEPARATOR, 283
- osiFindGlobalSymbol, 290
- osiMutex.h, 311
- osiPoolStatus.h, 311
- osiProcess.h, 312
- osiSock.h, 280, 313
- osiSufficientSpaceInPool, 311
- OUTLINK, 83
- override, 35–38, 65, 79
- Overview of Record Processing, 171

- Passive, 243
- Passive - Scan Type, 243
- Passive Scanning, 85
- path, 96
- path – Database Definitions, 98
- Path requirements, 35
- pdbbase, 199
- Periodic - Scan Type, 243
- Periodic Scanning, 85, 250
- periodicTask, 251
- Perl, 35
- pfield in DBADDR, 217
- pfldDes in DBADDR, 217
- pft, 162
- PHAS - Scan Related Field, 244
- POSIX, 44
- Posix C source code, 44
- post_event, 27, 246, 249
- postfix, 261
- postfix.h, 261
- PP, 84
- pp – field definition rules, 100
- pp_value – field definition, 102
- PPCBUG, 23
- precord - DBADDR, 217
- PRIO - Scan Related Field, 244
- process - example, 175
- process - Record Support Routine, 180
- process - record support routine, 86
- Process Passive, 84, 88
- PROD, 40, 60
- PROD_HOST, 60
- PROD_IOC, 60
- PROD_LIBS, 63, 64
- PROD_OBJLIBS, 61
- PROD_OBJS, 61
- PROD_RCS, 68
- PROD_SRCS, 63
- PROD_SYS_LIBS, 64
- PROD_VERSION, 65
- product libraries, 54, 63
- Product link order, 64
- product name, 60
- product object file, 61
- product source file, 62
- product version number, 65
- Products, 60
- prompt – field definition rules, 100
- prompt_value – field definition, 101
- promptgroup – field definition rules, 100
- Psuedo field, 111
- put_array_info - Record Support Routine, 181
- put_enum_str - Record Support Routine, 182
- putNotify, 91
- putNotifyCanceled, 225
- putNotifyInit, 225
- putNotifyOK, 225
- putNotifyStatus, 225
- PV_LINK – link field value, 108
- PVNAME_STRINGSZ, 215
- PVNAME_SZ, 215

- Quoted String, 97

- raw_value – breakpoint table, 107
- RCS, 68
- README.1st, 81
- README.cris, 82
- README.darwin.html, 82
- README.hpux, 82
- README.html, 81
- README.MS_WINDOWS, 82
- README.niCpu030, 82
- README.tru64unix, 82
- realclean, 42
- realuninstall, 43
- rebuild, 42
- recGblDbaddrError, 184
- recGblFwdLink, 186
- recGblGetAlarmDouble, 185
- recGblGetControlDouble, 185
- recGblGetGraphicDouble, 185
- recGblGetPrec, 185
- recGblGetTimeStamp, 185
- recGblInitConstantLink, 186
- recGblRecordError, 184
- recGblRecsupError, 184
- recGblResetAlarms, 184
- recGblSetSevr, 184
- record, 96, 107
- record attribute, 111
- record instance – Database Definitions, 107
- Record Instance File, 95

- Record Processing, 86
- Record Support, 171
- Record Support Entry Table, 172
- record type – Database Definitions, 100
- Record Type Definitions, 44
- record_name – record instance definition, 107
- record_type – device definition, 104
- record_type – record instance definition, 107
- record_type – record type definition, 101
- recordtype, 96
- Registering routines for DBD files, 46
- Registering support routines, 46
- registerRecordDeviceDriver, 320
- registerRecordDeviceDriver.c, 320
- registerRecordDeviceDriver.pl, 320
- registrar, 17, 96, 258
 - iocsh commands, 258
- registrar – Database Defintions, 105
- Registry, 319
- Registry.h, 319
- registryAdd, 319
- registryDeviceSupport.h, 320
- registryDeviceSupportAdd, 320
- registryDeviceSupportFind, 320
- registryDriverSupport.h, 320
- registryDriverSupportAdd, 320
- registryDriverSupportFind, 320
- registryDump, 319
- registryFind, 319
- registryFree, 319
- registryFunction.h, 320
- registryFunctionAdd, 320
- registryFunctionFind, 320
- registryFunctionRef, 320
- registryFunctionRefAdd, 320
- registryRecordTypeAdd, 319
- registrySetTableSize, 319
- RELEASE, 38, 79
- RELEASE_DBDFLAGS, 38
- RELEASE_INCLUDES, 38
- RELEASE_NOTES.html, 82
- ReleaseChecklist.html, 82
- replaceVAR.pl, 81
- report - device support routine, 191
- report - Record Support Routine, 180
- resource files, 67
- resourceLib.h, 267
- RF_IO – link field value, 109
- ringPointer, 270
- ringPointerCreate, 271
- ringPointerDelete, 271
- ringPointerFlush, 271
- ringPointerGetFree, 271
- ringPointerGetSize, 271
- ringPointerGetUsed, 271
- ringPointerId, 271
- ringPointerIsEmpty, 271
- ringPointerIsFull, 271
- ringPointerPop, 271
- ringPointerPush, 271
- RSET, 172
- RSET - example, 174
- rsrv_init, 123
- rsrv_pause, 123
- rsrv_run, 123
- RTEMS, 35
- RULE, 130
- RULES, 79
- RULES files, user created, 70
- RULES.Db, 79
- RULES.ioc, 79
- RULES_ARCHS, 79
- RULES_BUILD, 79
- RULES_DIRS, 79
- RULES_EXPAND, 80
- RULES_FILE_TYPE, 80
- RULES_JAVA, 80
- RULES_TARGET, 80
- RULES_TOP, 80
- runTest, 284
- runTestFunc, 284
- Runtime Database Access, 215
- SCAN - Scan Related Field, 243
- Scan Once - Scan Type, 243
- Scan Related Database Fields, 243
- SCAN_1ST_PERIODIC, 245
- SCAN_ALARM, 90
- scanAdd, 246
- scanDelete, 246
- scanInit, 122, 245
- scanIoInit, 249
- scanIoRequest, 250
- scanOnce, 124, 251
- scanOnceSetQueueSize, 124, 251
- scanPause, 123, 245
- scanpel, 155
- scanPeriod, 246
- scanpiol, 155
- scanppl, 155
- scanRun, 123, 245
- SCRIPTS, 58
- Scripts, 58
- SCRIPTS_HOST, 58
- SCRIPTS_IOC, 58
- secPastEpoch, 305
- SHARED_LIBRARIES, 36, 50, 65
- shareLib.h, 283

- SHRLIB_VERSION, 55
- Site specific Configuration, 36
- Site.cshrc, 82
- Site.profile, 82
- size – field definition rules, 100
- size_value – field definition, 103
- Slope Conversion, 111
- Source file dirs, 44
- SPC_ALARMACK, 102
- SPC_AS, 102
- SPC_CALC, 102
- SPC_DBADDR, 102
- SPC_LINCONV, 102
- SPC_MOD, 102
- SPC_NOMOD, 102
- SPC_RESET, 102
- SPC_SCAN, 102
- special - Record Support Routine, 180
- special – field definition rules, 100
- special in DBADDR, 217
- special_value – field definition, 102
- Specifying libraries, 63
- Specifying libraries dependancies, 54
- spy, 255
- src/tools File, 81
- SRC_DIRS, 44
- SRCS, 51, 62
- standard C++ library, 265
- Startup File Descriptions, 82
- Startup Files, 82
- State Notation Programs, 58
- static builds, 65
- STATIC_BUILD, 65
- status codes, 168
- std::max, 266
- std::min, 266
- std::swap, 266
- strtod, 298
- struct dbAddr, 217
- struct dsxt, 193
- struct putNotify, 225
- substitution file, 47
- substitutions file, 47
- symFindByNameAndTypeEPICS, 317
- symFindByNameEPICS, 317
- synchronous device support example, 189
- Synchronous Records, 89
- SYS_PROD_LIBS, 64

- T_A specific definitions, 40
- Table of Makefile definitions, 71
- TARGETS, 67
- Task Watchdog, 240
- task_params.h, 317
- taskwd.h, 240
- taskwdAnyInsert, 241
- taskwdAnyRemove, 241
- taskwdInit, 121
- taskwdInsert, 240
- taskwdMonitor, 241
- taskwdMonitorAdd, 241
- taskwdMonitorDel, 241
- taskwdRemove, 241
- taskwdShow, 241
- TCL libraries, 68
- TCLINDEX, 68
- TCLLIBNAME, 68
- template file, 47
- TEMPLATES, 59
- Templates, 59
- templates, 265
- templates directory, 34
- Test Products, 65
- Test Scripts, 66
- Test::Harness, 283
- testAbort, 284
- testDiag, 284
- testDone, 284
- testFail, 284
- testHarness, 284
- TESTJAVA, 68
- testMain.h, 285
- testOk, 284
- testOk1, 284
- testOkV, 284
- testPass, 284
- testPlan, 284
- TESTPROD, 40, 65
- TESTPROD_HOST, 65
- TESTPROD_IOC, 65
- TESTSCRIPTS, 66
- TESTSCRIPTS_HOST, 66
- TESTSCRIPTS_IOC, 66
- testSkip, 284
- testTodoBegin, 284
- testTodoEnd, 284
- Time Event, 291
- time provider, 291
- top, 33
- Tornado, 35
- tpn, 162
- TPRO, 153
- Trace Processing, 153
- TRAPWRITE, 129, 145
- truncateFile, 283, 298
- truncateFile.h, 283
- tsBTree.h, 267
- tsDLList.h, 267

tsFreeList.h, [267](#)
tsMinMax.h, [267](#)
tsSLList.h, [267](#)
type – variable definition, [105](#)

UAG, [128](#), [130](#)
UDF, [178](#)
udf, [178](#)
uninstall, [42](#)
Unquoted String, [97](#)
useManifestTool.pl, [81](#)
User created config files, [70](#)
User specific override, [40](#)
USER_DBDFLAGS, [45](#)
USR_CFLAGS, [48](#)
USR_CPPFLAGS, [48](#)
USR_CXXFLAG, [48](#)
USR_INCLUDES, [48](#)
USR_JAVACFLAGS, [68](#)
USR_JAVAHFLAGS, [70](#)
USR_LDFLAGS, [48](#)
USR_LIBS, [54](#), [55](#), [63](#), [64](#)
USR_OBJLIBS, [53](#), [61](#)
USR_OBJS, [52](#), [61](#)
USR_SRCS, [51](#), [62](#)
USR_SYS_LIBS, [55](#), [64](#)

VALID_BUILDS, [40](#)
variable, [17](#), [96](#)
variable – Database Definitions, [105](#)
variable_name – variable definition, [105](#)
veclist, [159](#), [317](#)
VME_IO – link field value, [109](#)
vxComLibrary, [317](#), [318](#)
VXLIO – link field value, [109](#)
vxWorks, [35](#)
vxWorks startup script, [120](#)

win32.bat, [82](#)