

dld

COLLABORATORS

	<i>TITLE :</i> dld	
<i>ACTION</i>	<i>NAME</i>	<i>DATE</i>
WRITTEN BY		December 7, 2024
<i>SIGNATURE</i>		

REVISION HISTORY

NUMBER	DATE	DESCRIPTION	NAME

Contents

1	dld	1
1.1	dld.guide	1
1.2	dld.guide/dld_init	1
1.3	dld.guide/dld_link	3
1.4	dld.guide/dld_unlink	3
1.5	dld.guide/dld_get_func	4
1.6	dld.guide/Executability	5
1.7	dld.guide/dld_list_undefined_sym	6
1.8	dld.guide/dld_create_reference	6
1.9	dld.guide/Explicit Definition	7
1.10	dld.guide/C++ support	7
1.11	dld.guide/dld_perror	9
1.12	dld.guide/Error Codes	10

Chapter 1

dld

1.1 dld.guide

What is dld?

"Dld" is a library package of C functions that performs "dynamic link editing". Programs that use dld can add compiled object code to or remove such code from a process anytime during its execution. Loading modules, searching libraries, resolving external references, and allocating storage for global and static data structures are all performed at run time.

Dld is now available for VAX, Sun 3, SPARCstation, Sequent Symmetry, Atari ST, and Linux.

This text describes how the dld functions can be called and some technical details that should be aware of. For the internals of dld and sample applications, please refer to 'An Approach to Genuine Dynamic Linking', Software-Practice and Experirnce, Vol. 21(4), 375-390 (April 1991). An early draft of that paper is included with the original distribution (prep.ai.mit.edu:pub/gnu/dld-3.2.3.tar.gz).

dld_init	Initializing Dld
dld_link	Dynamically Linking in New Modules
dld_unlink	Unlinking a Module
dld_get_func	Invoking Dynamically Linked Functions
Executability	Determining If a Function is Executable
dld_list_undefined_sym	Listing the Undefined Symbols
dld_create_reference	Explicitly Referencing a Symbol
Explicit Definition	Explicitly Defining a Symbol
C++ support	
dld_perror	Printing out the Error Messages
Error Codes	Definition of Error Codes

1.2 dld.guide/dld_init

Initializing Dld

=====

To use any of the dld functions, you must include the header file 'dld.h' for declaration of the functions and definition of the error code constants.

The function 'dld_init' must be called before any other dld functions.

- Function: int dld_init (char *PATH)
where PATH is a string containing the path name of the executable file of the executing process.

This function initializes internal data structures of dld and loads into memory symbol definitions of the executing process. By doing so, other dynamically loaded functions can reference symbols already defined or share functions already exist in the executing process.

'dld_init' returns 0 when successful; otherwise, it returns an error code that is non-zero (see Definition of Error Codes).

Locating the Executable File

The path name of the executing process as required by 'dld_init' might not be easily obtained all the time. Not all systems pass the entire path name of the executable file as the first argument ('argv[0]') to 'main'. In order to obtain the full path of the executable file, the 'dld_find_executable' function can be used.

- Function: char *dld_find_executable (char *COMMAND)
'dld_find_executable' returns the absolute path name of the file that would be executed if COMMAND were given as a command. It looks up the environment variable PATH, searches in each of the directory listed for COMMAND, and returns the absolute path name for the first occurrence. Thus, it is advisable to invoke 'dld_init' as:

```
main (int argc, char **argv)
{
    ...
    if (dld_init (dld_find_executable (argv[0]))) {
        ...
    }
    ...
}
```

Note: If the current process is executed using the 'execve' call without passing the correct path name as argument 0, 'dld_find_executable (argv[0])' will also fail to locate the executable file.

'dld_find_executable' returns zero if 'command' is not found in any of the directories listed in 'PATH'.

1.3 dld.guide/dld_link

Dynamically Linking in New Modules

=====

The function 'dld_link' dynamically links in the named relocatable object or library file into memory.

- Function: int dld_link (char *FILENAME)
 where FILENAME is the path name of the file to be linked.
 Specifically, if the named file is a relocatable object file, it is completely loaded into memory. If it is a library file, only those modules defining an unresolved external reference are loaded. Since a module in the library may itself reference other routines in the library, loading it may generate more unresolved external references. Therefore, a library file is searched repeatedly until a scan through all library members is made without having to load any new modules.

Storage for the text and data of the dynamically linked modules is allocated using 'malloc'. In other words, they are kept in the *heap* of the executing process.

After all modules are loaded, 'dld_link' resolves as many external references as possible. Note that some symbols might still be undefined at this stage, because the modules defining them have not yet been loaded.

If the specified module is linked successfully, 'dld_link' returns 0; otherwise, it returns a non-zero error code (see Definition of Error Codes).

1.4 dld.guide/dld_unlink

Unlinking a Module

=====

The major difference between dld and other dynamic linker is that dld allows object modules to be removed from the process anytime during execution. Unlinking a module is simply the reverse of the link operation (see Important Points in Using Unlink). The specified module is removed and the memory allocated to it is reclaimed. Additionally, resolution of external references must be undone.

There are two unlink functions:

- Function: int dld_unlink_by_file (char *PATH, int HARD)
 - Function: int dld_unlink_by_symbol (char *ID, int HARD)
- The two unlink functions are basically the same except that

'dld_unlink_by_file' takes as argument the path name (PATH) of a file corresponding to a module previously linked in by 'dld_link', but 'dld_unlink_by_symbol' unlinks the module that defines the specified symbol (ID).

Both functions take a second argument HARD. When HARD is non-zero ("hard unlink"), the specified module is removed from memory unconditionally. On the other hand, if HARD is zero ("soft unlink"), this module is removed from memory only if it is not referenced by any other modules. Furthermore, if unlinking a module results in leaving some other modules being unreferenced, these unreferenced modules are also removed.

Hard unlink is usually used when you want to explicitly remove a module and probably replace it by a different module with the same name. For example, you may want to replace the system's 'printf' by your own version. When you link in your version of 'printf', dld will automatically redirect all references to 'printf' to the new version.

Soft unlink should be used when you are not sure if the specified module is still needed. If you just want to clean up unnecessary functions, it is always safe to use soft unlink.

Both unlink functions returns 0 if the specified object file or symbol is previously loaded. Otherwise, they return a non-zero error code (see Definition of Error Codes).

Important Points in Using Unlink

When a module is being unlinked, dld tries to clean up as much as it can to restore the executing process to a state as if this module has never been linked. This clean up includes removing and reclaiming the memory for storing the text and data segment of the module, and "un-defining" any global symbols defined by this module.

However, side effects--such as modification of global variables, input/output operations, and allocations of new memory blocks--caused by the execution of any function in this module are not reversed. Thus, it is the responsibility of the programmer to explicitly carry out all necessary clean up operations before unlinking a module.

1.5 dld.guide/dld_get_func

Invoking Dynamically Linked Functions

=====

Dynamically linked functions may still be invoked from modules (e.g., 'main') that do not contain references to such functions.

- Function: unsigned long dld_get_symbol (char *ID)
Returns the entry point of the function named ID if found, 0 if not found. Non-zero returned values can be used as pointers to the

functions.

- Function: unsigned long dld_get_func (char *FUNC)
Returns the address of the global variable named FUNC if found, 0 if not found.

A typical use of 'dld_get_func' would be:

```

{
    void (*func) ();
    int error_code;

    ...

    /* First, link in the object file "my_object_file.o". Proceed
       only if the link operation is successful, i.e. it returns 0.
       "my_new_func" is a function defined in "my_object_file.o".
       Set func to point at the entry point of this function and
       then invoke it indirectly through func. */

    if ((error_code = dld_link ("my_object_file.o")) == 0) {
        if ((func = (void (*) ()) get_func ("my_new_func")) != 0)
            (*func) ();
        ...
    } else {

        ...
    }
}

```

1.6 dld.guide/Executability

Determining If a Function is Executable

=====

Since dld allows modules to be added to or removed from an executing process dynamically, some global symbols may not be defined. As a result, an invocation of a function might reference an undefined symbol. We say that a function is "executable" if and only if all its external references have been fully resolved and all functions that it might call are executable.

- Function: int dld_function_executable_p (char *FUNC)
The predicate function 'dld_function_executable_p' helps solve this problem by tracing the cross references between modules and returns non-zero only if the named function is executable.

Note that the implementation of 'dld_function_executable_p' is not complete according to the (recursive) definition of executability. External references through pointers are not traced. That is, 'dld_function_executable_p' will still return non-zero if the named function uses a pointer to indirectly call another function which has already been unlinked. Furthermore, if one external reference of a object module is unresolved, all functions defined in this

module are considered unexecutable. Therefore, `'dld_function_executable_p'` is usually too conservative.

However, it is advisable to use `'dld_function_executable_p'` to check if a function is executable before its invocation. In such a dynamic environment where object modules are being added and removed, a function that is executable at one point in time might not be executable at another. Under most circumstances, `'dld_function_executable_p'` is accurate. Also, the implementation of this function has been optimized and it is relatively cheap to use.

1.7 dld.guide/dld_list_undefined_sym

Listing the Undefined Symbols

=====

- Function: `char **dld_list_undefined_sym ()`
 The function `'dld_list_undefined_sym'` returns an array of undefined global symbol names.

The list returned contains all the symbols that have been referenced by some modules but have not been defined. This function is designed for debugging, especially in the case when a function is found to be not executable but you do not know what the missing symbols are.

The length of the array is given by the global variable `'dld_undefined_sym_count'`, which always holds the current total number of undefined global symbols. Note that all C symbols are listed in their internal representation--i.e., they are prefixed by the underscore character `'_'`.

Storage for the array returned is allocated by `'malloc'`. It is the programmer's responsibility to release this storage by `'free'` when it is not needed anymore.

1.8 dld.guide/dld_create_reference

Explicitly Referencing a Symbol

=====

Normally, a library module is loaded only when it defines one of more symbols that has been referenced. To force a library routine to be loaded, one need to explicitly create a reference to a symbol defined by that library routine. The function `'dld_create_reference'` is designed for this purpose:

- Function: `int dld_create_reference (char *NAME)`
 Usually `NAME` is the name of the library routine that should be

loaded, but it can be any symbol defined by that routine. After such a reference has been created, linking the appropriate library by `'dld_link'` would cause the required library routine to be loaded.

If the call is successful, `'dld_create_reference'` returns 0; otherwise, it returns a non-zero error code (see Definition of Error Codes).

The library routine loaded by this method can be unlinked by `'dld_unlink_by_symbol (NAME)'`. Once it has been unlinked, the corresponding reference created by `'dld_create_reference'` is also removed so that this routine will not be loaded in again by subsequent linking of the library.

1.9 dld.guide/Explicit Definition

Explicitly Defining a Symbol

=====

Dld allows a programmer to explicitly define global symbols. That is, a programmer can force a symbol to have storage assigned for it. This is especially useful in incremental program testing where the function being tested needs to access some global variables which are defined by another function not yet linked in (or even not yet written). There are two functions related to explicit definition:

- Function: `int dld_define_sym (char *NAME, unsigned int SIZE)`
`'dld_define_sym'` forces dld to allocate SIZE bytes for symbol NAME. It can be called before or after a reference to NAME is made. If references to NAME already exist when it is defined, all such references are directed to point to the correct address allocated for NAME.

`'dld_define_sym'` returns 0 if successful. Otherwise, it returns a non-zero error code (see Definition of Error Codes). The typical error is a multiple definition of NAME.

- Function: `void dld_remove_defined_symbol (char *NAME)`
When the definition of NAME is no longer needed, it can be removed by `'dld_remove_define_symbol'`.

1.10 dld.guide/C++ support

C++ support

=====

From: Thomas Hiller, Hiller@tu-harburg.d400.de

The original DLD supports only C function linking. When using C++

there is a problem with the global constructors and destructors. The new DLD version contains code which mimics the original "collect2" supplied by gcc (see ??). The way GNU 'ld' (see ??) does this is better (more direct). But I'm not familiar with the internals of GNU 'ld' and DLD.

The file 'gxxload.cc' contains two C functions ('dyn_load' and 'dyn_unload'), which function as 'dld_load' and 'dld_unload', but they take account that the C++ objects or libraries may contain global constructors and destructors.

- Function: void dyn_load (char *NAME)
Behaves like 'dld_link(name)' but constructors are called before control returns to the program.
- Function: void dyn_unload (char *NAME)
Behaves like 'dld_unlink' but the destructors are called before the unlink takes place.
- Global Array: char* dyn_libraries
The caller needs to define the null terminated array of strings 'dyn_libraries'. The strings in the array 'dyn_libraries', give the pathnames of the files which should be used to resolve link references after a call is made to 'dyn_load'. As in the example below, 'dyn_libraries' might contain:


```
foo "/usr/lib/libg++.a"

foo "/usr/lib/libm.a"

foo "/usr/lib/libc.a"
```

A small demonstration program is included in the 'test' sub-directory's files 'gxxtest.cc' and 'sub.cc'. When invoked, 'gxxtest' should print (1) :

```
DLD function:
dyn_load:
Hello, this is A
dyn_unload:
Hello, this was A
---
```

The constructor (which prints the first 'Hello' message) is called by 'dyn_load'. 'dyn_unload' calls the destructor (second 'Hello' message). The final '---' message demonstrates that the destructor is called by 'dyn_unload' and not the 'exit()' function.

sub.cc

```
#include <iostream.h>

class A {
public:
  A() { cout<<"Hello, this is A"<<endl; }
  ~A() { cout<<"Hello, this was A"<<endl; }
};
```

```
A a;
```

```
gxxtest.cc
```

```
-----
```

```
#include <iostream.h>

#include <dld.h>

char* dyn_libraries[] = {
"/usr/lib/libg++.a", /* link in libg++ */
"/usr/lib/libm.a", /* link in libm */
"/usr/lib/libc.a", /* link in libc */
LIBGCC, /* link in libgcc */
0
};

main()
{
  cout<<"DLD function: "<<endl;
  dld_link("sub.o");
  dld_unlink_by_file("sub.o",0);
  cout<<"dyn_load: "<<endl;
  dyn_load("sub.o");
  cout<<"dyn_unload: "<<endl;
  dyn_unload("sub.o");
  cout<<"---"<<endl;
}

```

```
----- Footnotes -----
```

(1) Running this example on my machine does not print the 'Hello' lines. But I don't know enough C++ to fix it.

1.11 dld.guide/dld_perror

Printing out the Error Messages

```
=====
```

- Function: void dld_perror (char *USER_MESG)
where USER_MESG is a user-supplied string prepended to the error message. The function 'dld_perror' prints out a short message explaining the error returns by the last dld functions.
- Function: char * dld_strerror (int CODE)
The function 'dld_strerror' returns the error message string corresponding to the given error code (from 'dld_errno').

1.12 dld.guide/Error Codes

Definition of Error Codes

=====

The dld functions return a non-zero error code when they fail. The definitions of these error codes are:

```
`DLD_ENOFILE'
    cannot open file.

`DLD_EBADMAGIC'
    bad magic number.

`DLD_EBADHEADER'
    failure reading header.

`DLD_ENOTEXT'
    premature eof in text section.

`DLD_ENOSYMBOLS'
    premature eof in symbols.

`DLD_ENOSTRINGS'
    bad string table.

`DLD_ENOTXTRELOC'
    premature eof in text relocation.

`DLD_ENODATA'
    premature eof in data section.

`DLD_ENODATRELOC'
    premature eof in data relocation.

`DLD_EMULTDEFS'
    multiple definitions of symbol.

`DLD_EBADLIBRARY'
    malformed library archive.

`DLD_EBADCOMMON'
    common block not supported.

`DLD_EBADOBJECT'
    malformed input file (not object file or archive).

`DLD_EBADRELOC'
    bad relocation info.

`DLD_ENOMEMORY'
    virtual memory exhausted.

`DLD_EUNDEFSYM'
    undefined symbol.
```
