

Subsections

- [Socket Creation and Naming](#)
 - [Connecting Stream Sockets](#)
 - [Stream Data Transfer and Closing](#)
 - [Datagram sockets](#)
 - [Socket Options](#)
 - [Example Socket Programs: `socket_server.c`, `socket_client`](#)
 - [socket_server.c](#)
 - [socket_client.c](#)
 - [Exercises](#)
-

IPC:Sockets

Sockets provide point-to-point, two-way communication between two processes. Sockets are very versatile and are a basic component of interprocess and intersystem communication. A socket is an endpoint of communication to which a name can be bound. It has a type and one or more associated processes.

Sockets exist in communication domains. A socket domain is an abstraction that provides an addressing structure and a set of protocols. Sockets connect only with sockets in the same domain. Twenty three socket domains are identified (see `<sys/socket.h>`), of which only the UNIX and Internet domains are normally used. Solaris 2.x Sockets can be used to communicate between processes on a single system, like other forms of IPC.

The UNIX domain provides a socket address space on a single system. UNIX domain sockets are named with UNIX paths. Sockets can also be used to communicate between processes on different systems. The socket address space between connected systems is called the Internet domain.

Internet domain communication uses the TCP/IP internet protocol suite.

Socket types define the communication properties visible to the application. Processes communicate only between sockets of the same type. There are five types of socket.

A stream socket

-- provides two-way, sequenced, reliable, and unduplicated flow of data with no record boundaries. A stream operates much like a telephone conversation. The socket type is `SOCK_STREAM`, which, in the Internet domain, uses Transmission Control Protocol (TCP).

A datagram socket

-- supports a two-way flow of messages. A on a datagram socket may receive messages in a different order from the sequence in which the messages were

sent. Record boundaries in the data are preserved. Datagram sockets operate much like passing letters back and forth in the mail. The socket type is `SOCK_DGRAM`, which, in the Internet domain, uses User Datagram Protocol (UDP).

A sequential packet socket

-- provides a two-way, sequenced, reliable, connection, for datagrams of a fixed maximum length. The socket type is `SOCK_SEQPACKET`. No protocol for this type has been implemented for any protocol family.

A raw socket

provides access to the underlying communication protocols.

These sockets are usually datagram oriented, but their exact characteristics depend on the interface provided by the protocol.

Socket Creation and Naming

`int socket(int domain, int type, int protocol)` is called to create a socket in the specified domain and of the specified type. If a `protocol` is not specified, the system defaults to a protocol that supports the specified socket type. The socket handle (a descriptor) is returned. A remote process has no way to identify a socket until an address is bound to it. Communicating processes connect through addresses. In the UNIX domain, a connection is usually composed of one or two path names. In the Internet domain, a connection is composed of local and remote addresses and local and remote ports. In most domains, connections must be unique.

`int bind(int s, const struct sockaddr *name, int namelen)` is called to bind a path or internet address to a socket. There are three different ways to call `bind()`, depending on the domain of the socket.

- For UNIX domain sockets with paths containing 14, or fewer characters, you can:

```
#include <sys/socket.h>
...
bind (sd, (struct sockaddr *) &addr, length);
```

- If the path of a UNIX domain socket requires more characters, use:

```
#include <sys/un.h>
...
bind (sd, (struct sockaddr_un *) &addr, length);
```

- For Internet domain sockets, use

```
#include <netinet/in.h>
...
bind (sd, (struct sockaddr_in *) &addr, length);
```

In the UNIX domain, binding a name creates a named socket in the file system. Use `unlink()` or `rm ()` to remove the socket.

Connecting Stream Sockets

Connecting sockets is usually not symmetric. One process usually acts as a server and the other process is the client. The server binds its socket to a previously agreed path or address. It then blocks on the socket. For a `SOCK_STREAM` socket, the server calls `int listen(int s, int backlog)`, which specifies how many connection requests can be queued. A client initiates a connection to the server's socket by a call to `int connect(int s, struct sockaddr *name, int namelen)`. A UNIX domain call is like this:

```
struct sockaddr_un server;
...
connect (sd, (struct sockaddr_un *)&server, length);
```

while an Internet domain call would be:

```
struct sockaddr_in;
...
connect (sd, (struct sockaddr_in *)&server, length);
```

If the client's socket is unbound at the time of the connect call, the system automatically selects and binds a name to the socket. For a `SOCK_STREAM` socket, the server calls `accept(3N)` to complete the connection.

`int accept(int s, struct sockaddr *addr, int *addrlen)` returns a new socket descriptor which is valid only for the particular connection. A server can have multiple `SOCK_STREAM` connections active at one time.

Stream Data Transfer and Closing

Several functions to send and receive data from a `SOCK_STREAM` socket. These are `write()`, `read()`, `int send(int s, const char *msg, int len, int flags)`, and `int recv(int s, char *buf, int len, int flags)`. `send()` and `recv()` are very similar to `read()` and `write()`, but have some additional operational flags.

The flags parameter is formed from the bitwise OR of zero or more of the following:

MSG_OOB

-- Send "out-of-band" data on sockets that support this notion. The underlying protocol must also support "out-of-band" data. Only `SOCK_STREAM` sockets created in the `AF_INET` address family support out-of-band data.

MSG_DONTROUTE

-- The `SO_DONTROUTE` option is turned on for the duration of the operation. It is

used only by diagnostic or routing programs.

MSG_PEEK

-- "Peek" at the data present on the socket; the data is returned, but not consumed, so that a subsequent receive operation will see the same data.

A `SOCK_STREAM` socket is discarded by calling `close()`.

Datagram sockets

A datagram socket does not require that a connection be established. Each message carries the destination address. If a particular local address is needed, a call to `bind()` must precede any data transfer. Data is sent through calls to `sendto()` or `sendmsg()`. The `sendto()` call is like a `send()` call with the destination address also specified. To receive datagram socket messages, call `recvfrom()` or `recvmsg()`. While `recv()` requires one buffer for the arriving data, `recvfrom()` requires two buffers, one for the incoming message and another to receive the source address.

Datagram sockets can also use `connect()` to connect the socket to a specified destination socket. When this is done, `send()` and `recv()` are used to send and receive data.

`accept()` and `listen()` are not used with datagram sockets.

Socket Options

Sockets have a number of options that can be fetched with `getsockopt()` and set with `setsockopt()`. These functions can be used at the native socket level (`level = SOL_SOCKET`), in which case the socket option name must be specified. To manipulate options at any other level the protocol number of the desired protocol controlling the option of interest must be specified (see `getprotoent()` in `getprotobyname()`).

Example Socket Programs:

socket_server.c, socket_client

These two programs show how you can establish a socket connection using the above functions.

socket_server.c

```
#include <sys/types.h>
#include <sys/socket.h>
```

```
#include <sys/un.h>
#include <stdio.h>

#define NSTRS      3          /* no. of strings */
#define ADDRESS   "mysocket" /* addr to connect */

/*
 * Strings we send to the client.
 */
char *strs[NSTRS] = {
    "This is the first string from the server.\n",
    "This is the second string from the server.\n",
    "This is the third string from the server.\n"
};

main()
{
    char c;
    FILE *fp;
    int fromlen;
    register int i, s, ns, len;
    struct sockaddr_un saun, fsaun;

    /*
     * Get a socket to work with. This socket will
     * be in the UNIX domain, and will be a
     * stream socket.
     */
    if ((s = socket(AF_UNIX, SOCK_STREAM, 0)) < 0) {
        perror("server: socket");
        exit(1);
    }

    /*
     * Create the address we will be binding to.
     */
    saun.sun_family = AF_UNIX;
    strcpy(saun.sun_path, ADDRESS);

    /*
     * Try to bind the address to the socket. We
     * unlink the name first so that the bind won't
     * fail.
     *
     * The third argument indicates the "length" of
     * the structure, not just the length of the
     * socket name.
     */
    unlink(ADDRESS);
    len = sizeof(saun.sun_family) + strlen(saun.sun_path);
```

```
if (bind(s, &saun, len) < 0) {
    perror("server: bind");
    exit(1);
}

/*
 * Listen on the socket.
 */
if (listen(s, 5) < 0) {
    perror("server: listen");
    exit(1);
}

/*
 * Accept connections.  When we accept one, ns
 * will be connected to the client.  fsaun will
 * contain the address of the client.
 */
if ((ns = accept(s, &fsaun, &fromlen)) < 0) {
    perror("server: accept");
    exit(1);
}

/*
 * We'll use stdio for reading the socket.
 */
fp = fdopen(ns, "r");

/*
 * First we send some strings to the client.
 */
for (i = 0; i < NSTRS; i++)
    send(ns, strs[i], strlen(strs[i]), 0);

/*
 * Then we read some strings from the client and
 * print them out.
 */
for (i = 0; i < NSTRS; i++) {
    while ((c = fgetc(fp)) != EOF) {
        putchar(c);

        if (c == '\n')
            break;
    }
}

/*
 * We can simply use close() to terminate the
 * connection, since we're done with both sides.
 */
```

```
    close(s);

    exit(0);
}
```

socket_client.c

```
#include <sys/types.h>
#include <sys/socket.h>
#include <sys/un.h>
#include <stdio.h>

#define NSTRS      3          /* no. of strings */
#define ADDRESS   "mysocket" /* addr to connect */

/*
 * Strings we send to the server.
 */
char *strs[NSTRS] = {
    "This is the first string from the client.\n",
    "This is the second string from the client.\n",
    "This is the third string from the client.\n"
};

main()
{
    char c;
    FILE *fp;
    register int i, s, len;
    struct sockaddr_un saun;

    /*
     * Get a socket to work with. This socket will
     * be in the UNIX domain, and will be a
     * stream socket.
     */
    if ((s = socket(AF_UNIX, SOCK_STREAM, 0)) < 0) {
        perror("client: socket");
        exit(1);
    }

    /*
     * Create the address we will be connecting to.
     */
    saun.sun_family = AF_UNIX;
    strcpy(saun.sun_path, ADDRESS);

    /*
     * Try to connect to the address. For this to
     * succeed, the server must already have bound
     * this address, and must have issued a listen()
     */
}
```

```
* request.
*
* The third argument indicates the "length" of
* the structure, not just the length of the
* socket name.
*/
len = sizeof(saun.sun_family) + strlen(saun.sun_path);

if (connect(s, &saun, len) < 0) {
    perror("client: connect");
    exit(1);
}

/*
 * We'll use stdio for reading
 * the socket.
 */
fp = fdopen(s, "r");

/*
 * First we read some strings from the server
 * and print them out.
 */
for (i = 0; i < NSTRS; i++) {
    while ((c = fgetc(fp)) != EOF) {
        putchar(c);

        if (c == '\n')
            break;
    }
}

/*
 * Now we send some strings to the server.
 */
for (i = 0; i < NSTRS; i++)
    send(s, strs[i], strlen(strs[i]), 0);

/*
 * We can simply use close() to terminate the
 * connection, since we're done with both sides.
 */
close(s);

exit(0);
}
```

Exercises

Exercise 12776

Configure the above `socket_server.c` and `socket_client.c` programs for your system and compile and run them. You will need to set up `socket ADDRESS` definition.

Dave Marshall

1/5/1999