

# Sockets

---

Sockets are a common interface that make TCP streams look like file streams. The socket interface shown here has been implemented for many operating systems and languages. This handout shows the Perl versions, but the socket functions look very similar in other languages. A Perl program needs the statement "use Socket;" up near its top to import these functions. In C, the headers are in "sys/socket.h".

## Error Handling

When writing Internet socket code, it's important to do all the error handling, because in a network situation the errors will actually happen with some regularity (host name not found, connection refused...). You need to put in error handling right from the start so you can distinguish the network errors from your own programming errors. Most of the Perl networking functions return false on error. In addition, in the case of an error, they will put a somewhat descriptive string in the global variable \$!.

## Buffering -- Autoflush

Many I/O systems do "buffering" automatically for efficiency. When you do a print(), the data is not written immediately. The system buffers up the characters until there is enough data to send all at once. That can be a real problem if the print is part of a protocol where you send, say, one line of text and then wait for a response. In Perl, the FileHandle module provides the "autoflush" function -- sets a file handle so it does not buffer...

```
use FileHandle;

....

## Suppose the SOCK file handle has been opened,
## set it to be unbuffered
autoflush SOCK, 1;
```

## 1. Hostname Lookup

### **\$ip = inet\_aton(\$hostname)**

Tries to look up the IP address of a hostname such as "www.yahoo.com". This will probably do a quick transaction with the local DNS server to get the IP address. Fails if the hostname cannot be looked up either because the hostname is wrong, or the local DNS system is down. Returns false on error and sets \$!.

## 2. Create Socket Address

### **\$sockaddr = sockaddr\_in(\$port , \$ip);**

Given a port number and an ipaddr, combine the two into a socket address that can be passed to connect(). We'll assume this step doesn't fail.

### 3. Allocate Socket

**socket(SOCK, PF\_INET, SOCK\_STREAM, 0);**

Given a file handle (e.g. SOCK, or could be a string like "HANDLE" stored in a variable \$sock), allocate a TCP socket and connect it to the file handle. The other arguments are constants. This just allocates the socket resources; it does not try to connect the socket to anything. The constants shown create a TCP stream connection. Other constants can be used to create other sorts of sockets. Fails if there are not enough resources to create another socket on the local machine. Most OS's have some kernel limit on how many simultaneous sockets can be maintained -- probably a few thousand. Returns false on error and sets \$!.

### 4. Connect

**connect(SOCK, \$sockaddr)**

Given a socket file handle (from step 3) and a socket address (from step 2), try to connect to the given port on that machine. On success, the given socket can now be used for reading and writing to the remote machine. Reading will block if there is no data available. Fails if the connection cannot be made to the remote machine because it is unreachable or it is not listening on that port number. Returns false on error and sets \$!.

### 5. Read

The simplest way to read text data is just with the standard <SOCK> file reading operator. Returns undef on error or the EOF.

```
...
connect(SOCK, $sockaddr);
$line = <SOCK>;    ## read a line of text
```

Most protocols involve taking turns exchanging a few lines of text back and forth. The code for such a protocol will need to know when it is supposed to read and when it is supposed to write.

Reading all the data until the EOF is like reading from a file.

```
while ($line = <SOCK>) {
    ...
}
```

This while-loop form is only appropriate if you want to continually read data until the other end does a close().

The function `sysread(SOCK, str, length)` is an alternative that tries to read the given number of bytes, and puts them in the given string. Do not use both <> and `sysread()` -- use one or the other. `sysread()` can be more efficient, since it does not have to read the data one line at a time...

```
$result = "";
while(sysread(SOCK, $buff, 1000)) { ## attempt to read 1000 bytes at a time
    $result = $result . $buff;
}
```

## Irregular Timing

The data will not all show up at once. The reading code will block (stop running) when there is no data. Usually, the data will show up in irregularly timed and sized bursts. In reality, the CPU is much, much faster than the network, so the code will spend most of its time blocked, waking up for the occasional burst of network activity.

## Irregular Sizing

Suppose that one end of the socket makes three separate write calls of 100 bytes each. This does not mean that the data will arrive in three 100 byte chunks. They may show up as a 150 byte chunk, a 37 byte chunk, and a 113 byte chunk. The reading code is responsible for accumulating the read bytes to build up meaningful chunk..

## 6. Write

The simplest way to write text is data is with print.

```
print SOCK "Hello there, how are you?\012";
```

The function `syswrite(SOCK, str, length)` is an alternative. Do not use both `print` and `syswrite()` -- use one or the other. Most Internet protocols end their text lines with carriage-return + linefeed. In Perl that's `"\015\012"` (decimal 13 and 10). Specifying the characters in this way avoids the problem where `"\n"` is mapped to the local end-of-line character, which changes with the platform you are running on.

## 7 Close

`close(SOCK)` -- closes the reading and writing capabilities of this end of the socket. The close on the writing end here will show up as an EOF on the other end.

## finger.pl example

```
#!/usr/bin/perl -w
## A simple example client-side sockets
##
## The "finger" service listens on port 79
## You send it a user name on a line by itself,
## and it sends back some lines of text
## It used to be that finger worked on most hosts.
## Now however, it's rare, since it allows a bad-guy
## to figure out user names on the host.
## finger.pl nick whois.stanford.edu works
## finger.pl nick elaine.stanford.edu rejected
## finger.pl president whitehouse.gov -- hangs ?!

use strict 'vars';    ## enforce global/local declarations
use Socket;
use FileHandle;

## The end-of-line we'll use
## (return+linefeed)
my($EOL);
$EOL = "\015\012";

## Because of 'use strict' -- we declare all vars,
## even these globals
my($hostname, $user);

(scalar(@ARGV) == 2) || die "usage: <user> <host>";
$user = shift(@ARGV) || die "need username first";
$hostname = shift(@ARGV) || die "need hostname second";

my ($sport, $err, $ipaddr, $sockaddr);

$sport = 79;## port for "finger" service
## could do the service->port num lookup dynamically

## Look up the hostname to get its IP addr
$ipaddr = inet_aton($hostname) || die "cannot find '$hostname'";

## Make a socket addr out of the IP+port
$sockaddr = sockaddr_in($sport , $ipaddr);

## Create the socket
socket(SOCK, PF_INET, SOCK_STREAM, 0) || die "cannot create socket";

## Connect the socket
connect(SOCK, $sockaddr) || die "cannot connect";

## Set the socket to be unbuffered
autoflush SOCK, 1;

## Send the username
print SOCK "$user$EOL";
```

```

## Read back and print what they send back until EOF

my($line);

while ($line = <SOCK>) {
    chomp($line);
    print $line, "\n";
}

## Alternately, we could use the perl function
## sysread(SOCK, str, desired_len) for reading

close(SOCK);

```

## Running Finger

```
elaine0:~/my193i> finger.pl nick whois.stanford.edu
```

Stanford University Whois Service

Try "whois help" for more information.

For directory service on the web: <http://www.stanford.edu/services/stanford.who>

```

-----
STANFORD-ONLY view - not public
The information below may not be available to the public. People in
the Stanford directory may choose whether information about them is
"public", "Stanford-only" or "private". Since this is the Stanford view,
please do not share the information below with the public. To see the
public view, go to http://www.stanford.edu/services/stanford.who
-----

```

```

Nick Parlante                <DS186W641>    <Phone Unavail>
  nick@cs.stanford.edu
  Computer Science, Academic staff-Lecturer

```

```

Nick Petalas                 <DS450S829>    (650) 555-5555
  nick_spam@pangea.Stanford.EDU
  Petroleum Engineering (2220), Academic staff-Engr Res Assoc

```

```

elaine0:~/my193i> finger.pl nick nosuchhost.stanford.edu
cannot find 'nosuchhost.stanford.edu' at finger.pl line 37.
elaine0:~/my193i> finger.pl nick elaine27.stanford.edu
cannot connect at finger.pl line 46.

```