

Unit OS A: Windows Networking

A.2. Windows Sockets Programming

Windows Operating System Internals - by David A. Solomon and Mark E. Russinovich with Andreas Polze

Roadmap for Section A.2

- General Concepts - Berkeley Sockets
- Creating a socket
- Binding an address
- Accepting connections
- Exchanging data
- Closing a connection
- Managing multiple connections with select()

3

Winsock Features

- Support for scatter-gather and asynchronous application I/O
- Quality of service (QoS) conventions so that applications can negotiate latency and bandwidth requirements when the underlying network supports QoS
- Extensibility so that Winsock can be used with protocols other than those Windows requires it to support
- Support for integrated namespaces other than those defined by a protocol an application is using with Winsock. A server can publish its name in Active Directory, for example, and using namespace extensions, a client can look up the server's address in Active Directory
- Support for multipoint messages where messages transmit to multiple receivers simultaneously

4

Windows Socket Programming

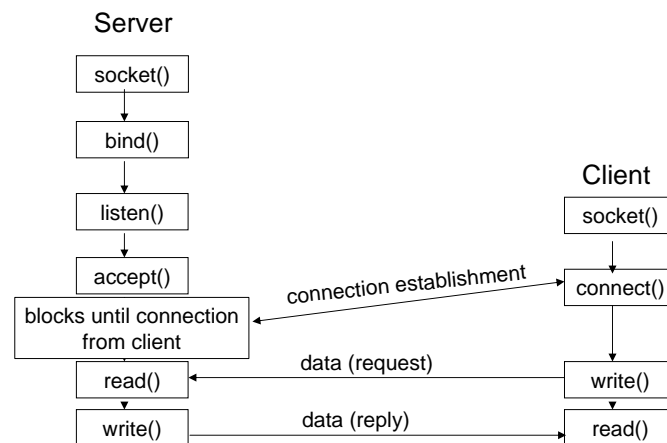
Berkeley Socket programs will port to Window Sockets

Exceptions:

- Call *WSAStartup()* to initialize Windows Socket DLL
- Use *ioctlsocket()* (non-portable) to configure the socket
- *_read()* and *_write()* can be used on sockets, but only after converting the socket descriptor to a file handle via *_open_osfhandle()*
- Use *closesocket()* (non-portable) rather than *close* to close a socket
- Call *WSACleanup()* to shut down the DLL

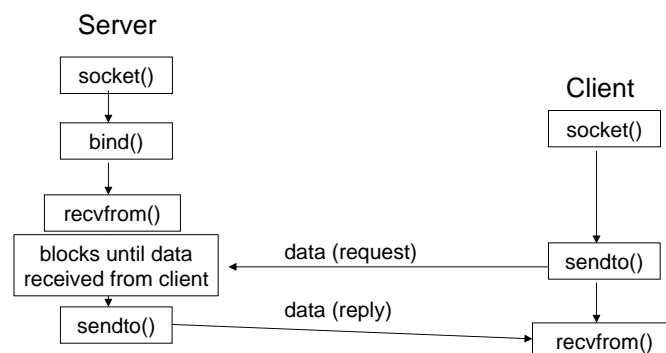
5

Berkeley 4.3 UNIX Sockets – connection-oriented



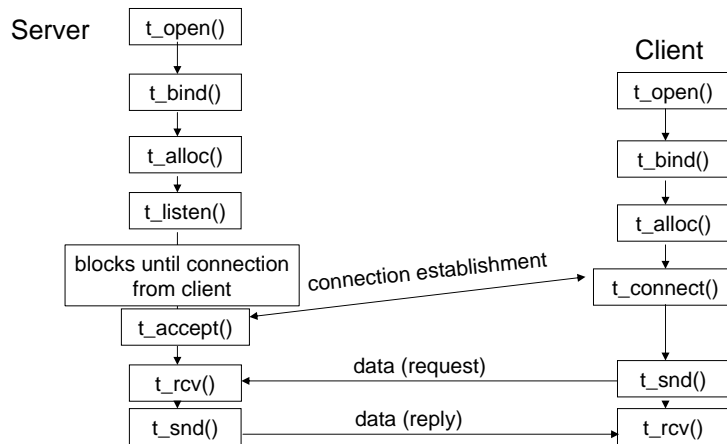
6

Berkeley 4.3 UNIX Sockets - connectionless



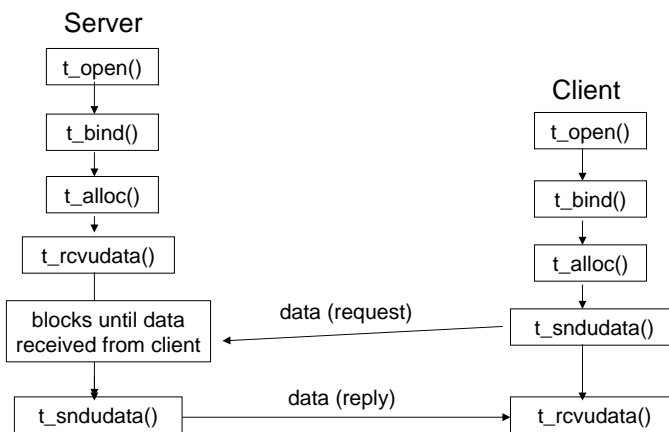
7

Unix SYS V.3 Transport Layer Interface – connection-oriented



8

Unix SYS V.3 Transport Layer Interface - connectionless



9

Create a socket

```
#include <winsock.h>
SOCKET socket (
    int af, int type, int protocol );
```

- **af**: An address format specification. The only format currently supported is AF_INET, which is the ARPA Internet address format.
- **type**: A type specification for the new socket.
- **protocol**: A particular protocol to be used with the socket, or 0 if the caller does not wish to specify a protocol.

10

Accept a connection on a socket

```
#include <winsock.h>
SOCKET accept (
    SOCKET s, struct sockaddr FAR * addr,
    int FAR * addrlen );
```

- **s**: A descriptor identifying a socket which is listening for connections after a listen().
- **addr**: An optional pointer to a buffer which receives the address of the connecting entity, as known to the communications layer. The exact format of the addr argument is determined by the address family established when the socket was created.
- **addrlen**: An optional pointer to an integer which contains the length of the address addr.

11

struct sockaddr

- From winsock.h (Windows) or /usr/include/sys/socket.h (UNIX)

```
/*
 * Structure used by kernel to store most
 * addresses.
 */
struct sockaddr {
    u_char  sa_len;          /* total length */
    u_char  sa_family;      /* address family */
    char    sa_data[14];    /* actually longer; address value*/
};
#define SOCK_MAXADDRLEN 255 /* longest possible addresses */
```

12

Associate a local address with a socket

```
#include <winsock.h>
int bind (
    SOCKET s, const struct sockaddr FAR * name,
    int namelen );
```

- **s**: A descriptor identifying an unbound socket.
- **name**: The address to assign to the socket.
- **namelen**: length of the name

```
struct sockaddr {
    u_short sa_family;
    char sa_data[14];
};
```

13

Internet address family

- **In the Internet address family, a name consists of several components.**
- **For SOCK_DGRAM and SOCK_STREAM, the name consists of three parts:**
 - a host address, the protocol number (set implicitly to UDP or TCP, respectively), and a port number which identifies the application.
 - If an application does not care what address is assigned to it, it may specify an Internet address equal to INADDR_ANY, a port equal to 0, or both.
 - If the Internet address is equal to INADDR_ANY, any appropriate network interface will be used; this simplifies application programming in the presence of multi-homed hosts.
 - If the port is specified as 0, the Windows Sockets implementation will assign a unique port to the application with a value between 1024 and 5000.
- **The application may use getsockname() after bind() to learn the address that has been assigned to it**
 - getsockname() will not necessarily fill in the Internet address until the socket is connected; several Internet addresses may be valid if the host is multi-homed.

14

Example: bind to an reserved port

```
SOCKADDR_IN sin;
SOCKET s;
u_short alport = IPPORT_RESERVED; /* 1024 */
sin.sin_family = AF_INET;
sin.sin_addr.s_addr = 0;
for (;;) {
    sin.sin_port = htons(alport);
    if (bind(s, (LPSOCKADDR)&sin, sizeof (sin)) == 0) {
        /* it worked */
    }
    if ( GetLastError() != WSAEADDRINUSE) {
        /* fail */
    }
    alport--;
    if (alport == IPPORT_RESERVED/2 ) {
        /* fail--all unassigned reserved ports are in use.*/
    }
}
```

15

Close a socket

```
#include <winsock.h>
int closesocket ( SOCKET s );
```

- This function closes a socket.
 - releases the socket descriptor `s`, so that further references to `s` will fail with the error `WSAENOTSOCK`.
 - If this is the last reference to the underlying socket, the associated naming information and queued data are discarded.
- Semantics influenced by socket options:

Option	Interval	Type of close	Wait for close?
<code>SO_DONTLINGER</code>	Don't care	Graceful	No
<code>SO_LINGER</code>	Zero	Hard	No
<code>SO_LINGER</code>	Non-zero	Graceful	Yes

16

Establish a connection to a peer

```
#include <winsock.h>
int connect ( SOCKET s,
             const struct sockaddr * name,
             int namelen );
```

- **s**: A descriptor identifying an unconnected socket.
- **name**: The name of the peer to which the socket is to be connected.
- **namelen**: The length of the name.
- create a connection to the specified foreign association. The parameter `s` specifies an unconnected datagram or stream socket

17

Establish a socket to listen for incoming connection

```
#include <winsock.h>
int listen ( SOCKET s, int backlog );
```

- **s**: A descriptor identifying a bound, unconnected socket.
- **backlog**: The maximum length to which the queue of pending connections may grow.
- typically used by servers that could have more than one connection request at a time:
 - if a connection request arrives with the queue full, the client will receive an error with an indication of WSAECONNREFUSED

18

Receiving data from a socket (connection-oriented)

```
#include <winsock.h>
int recv ( SOCKET s,
          char * buf, int len, int flags );
```

- **s**: A descriptor identifying a connected socket.
- **buf**: A buffer for the incoming data.
- **len**: The length of buf.
- **flags**: Specifies the way in which the call is made.

19

Receive a datagram and store the source address (connectionless)

```
#include <winsock.h>
int recvfrom ( SOCKET s,
              char * buf, int len, int flags,
              struct sockaddr * from, int * fromlen );
```

- **s**: A descriptor identifying a bound socket.
- **buf**: A buffer for the incoming data.
- **len**: The length of buf.
- **flags**: Specifies the way in which the call is made.
- **from**: An optional pointer to a buffer which will hold the source address upon return.
- **fromlen**: An optional pointer to the size of the from buffer.

20

Determine the status of one or more sockets, waiting if necessary.

```
#include <winsock.h>
int select ( int nfds, fd_set * readfds,
            fd_set * writefds, fd_set * exceptfds,
            const struct timeval * timeout );
```

- **nfds**:
This argument is ignored and included only for the sake of compatibility.
- **readfds**:
An optional pointer to a set of sockets to be checked for readability.
- **writefds**:
An optional pointer to a set of sockets to be checked for writability
- **exceptfds**:
An optional pointer to a set of sockets to be checked for errors.
- **timeout**:
The maximum time for select() to wait, or NULL for blocking operation.

21

Send data on a connected socket (connection-oriented)

```
#include <winsock.h>
int send ( SOCKET s,
          const char * buf, int len, int flags );
```

- **s**: A descriptor identifying a connected socket.
- **buf**: A buffer containing the data to be transmitted.
- **len**: The length of the data in buf.
- **flags**: Specifies the way in which the call is made.

22

Send data to a specific destination (connectionless)

```
#include <winsock.h>
int sendto ( SOCKET s,
            const char * buf, int len, int flags,
            const struct sockaddr * to, int tolen );
```

- **s**: A descriptor identifying a socket.
- **buf**: A buffer containing the data to be transmitted.
- **len**: The length of the data in buf.
- **flags**: Specifies the way in which the call is made.
- **to**: An optional pointer to the address of the target socket.
- **tolen**: The size of the address in to.

23

Further Reading

- Mark E. Russinovich and David A. Solomon, Microsoft Windows Internals, 4th Edition, Microsoft Press, 2004;
 - Windows Sockets (from pp. 791)
- Abraham Silberschatz, Peter B. Galvin, Operating System Concepts, John Wiley & Sons, 6th Ed., 2003;
 - Chapter 15 - Distributed System Structures
- W. Richard Stevens, Unix Network Programming, Prentice Hall Software Series, 1990; (The Book)
 - Chapter 6 - Berkeley Sockets