



Four Ways of Inter Process Messaging

Well done! You have your first Linux program running. Now you are ready to code the second one but stop! You noticed that they have to communicate and send/receive messages. At that moment, I will tell you about four possible ways of **Inter Process Communication** on Linux 😎

1 Shared Memory

1.1 Creating Shm

1.2 Messaging over Mapped Shared Memory in C++

2 First In First Out Pipes

2.1 Creating FIFO Messaging Object

2.2 Messaging over Linux FIFO in C++

3 Unix Domain Socket

3.1 Creating a Unix Domain Socket

3.2 Connected Server-Client Messaging

3.2.1 Server Side

3.2.2 Client Side

3.2.3 IPC Example via Unix Domain Socket in C++


3.3 Connectionless Datagram Messaging

3.3.1 Simple IPC Unix Domain Socket Messaging in C++

4 Final Words

Shared Memory

On a Linux, the easiest way of communicating to another process is via shared memory. Operating system creates a special file for you and all the processes that wants to communicate to each other only have to access that file. Yet the kernel does much more for you by mapping it into the memory of your process and you only have to use it like any other memory region.



Creating Shm

It is very simple. The first process has to create shared memory object at a predefined location with the correct permissions. The following line creates a file named `/dev/shm/ipc` with full permission.

```
1 int fd = shm_open("ipc", O_RDWR | O_CREAT, 0777);
```

The file descriptor `fd` will be non-negative on success. The second step is mapping the shared memory area to RAM so that it has an address in memory for reading and writing. Use `mmap()` for this.

```
1 mmap(0, SIZE_OF_MEMORY, PROT_READ | PROT_WRITE, MAP_SHARED, fd, 0);
```

This line maps the file into memory with read and write permissions and returns the address on success.

Messaging over Mapped Shared Memory in C++

I prepared a short demo that sends and receives a simple message made of `messageId` and `messageBody`.

```
shared_memory_ipc.cpp
1  /**
2   * @file      shared_memory_ipc.cpp
3   * @author    Atakan S.
4   * @date      01/01/2019
5   * @version   1.0
6   * @brief     Example project for Linux Shared Memory IPC Messaging tutorial.
7   *
8   * @copyright Copyright (c) 2019 Atakan SARIOGLU ~ www.atakansarioglu.com
9   *
10  * Permission is hereby granted, free of charge, to any person obtaining a
11  * copy of this software and associated documentation files (the "Software"),
12  * to deal in the Software without restriction, including without limitation
13  * the rights to use, copy, modify, merge, publish, distribute, sublicense,
14  * and/or sell copies of the Software, and to permit persons to whom the
15  * Software is furnished to do so, subject to the following conditions:
16  *
17  * The above copyright notice and this permission notice shall be included in
18  * all copies or substantial portions of the Software.
19  *
20  * THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR
21  * IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY,
22  * FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE
23  * AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER
24  * LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING
25  * FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER
26  * DEALINGS IN THE SOFTWARE.
27  */
28
29 #include <stdlib.h>
30 #include <stdio.h>
31 #include <fcntl.h>
32 #include <sys/mman.h>
33 #include <unistd.h>
34 #include <iostream>
35 #include <mutex>
36
37 struct Message {
38     std::mutex mutex;
39     int messageId;
40     int messageBody;
41 };
42
43 int main() {
44     // Remove the old file if exists.
45     shm_unlink("ipc");
46
47     // Start 2nd process and initialize random number generator.
48     if(fork()) sleep(1);
49     srand(getpid());
50
51     // Create shared memory object.
52     int shm;
53     if((shm = shm_open("ipc", O_RDWR | O_CREAT, 0777)) < 0) {
54         std::cout << "Cannot open shm" << std::endl;
55         return -1;
56     }
```

```


57
58 // Create an empty message.
59 if(lseek(shm, 0, SEEK_END) < sizeof(Message)) { // Fill with zero
60     lseek(shm, 0, SEEK_SET);
61     Message dummy{};
62     if(write(shm, &dummy, sizeof(Message)) < sizeof(Message)) {
63         std::cout << "Cannot initialize shm" << std::endl;
64         return -1;
65     }
66 }
67
68 // Map the file to memory and obtain a pointer to that region.
69 Message * message;
70 if((message = (Message *)mmap(0, sizeof(Message), PROT_READ | PROT_WRITE, MAP_
71     std::cout << "Cannot create mapping" << std::endl;
72     return -1;
73 }
74
75 // Receive and send messages.
76 while(true) {
77     // Use mutex to send and receive the message inact.
78     std::cout << "Waiting for lock... " << std::endl;
79     while(!message->mutex.try_lock());
80     std::cout << "Has lock..." << std::endl;
81
82     // Read from and write to the memory as usual.
83     std::cout << "Process " << getpid() << " Received message" << std::endl;
84     std::cout << "    Got messageId=" << message->messageId << std::endl;
85     std::cout << "    Got messageBody=" << message->messageBody << std::endl;
86     std::cout << "Process " << getpid() << " Sending message" << std::endl;
87     std::cout << "    Set messageId=" << (message->messageId = getpid()) << std
88     std::cout << "    Set messageBody=" << (message->messageBody = rand()) << s
89
90     // Release the lock and wait.
91     message->mutex.unlock();
92     sleep(5);
93 }
94
95 // Unmap and unlink the shared memory.
96 munmap(message, sizeof(Message));
97 close(shm);
98 shm_unlink("ipc");
99
100 // Exit.
101 return 0;
102 }

```

You can compile with `g++ shared_memory_ipc.cpp -o shared_memory_ipc -lrt -std=c++11` as noted [here](#). In the example code you will find two processes, one is `fork()`'ed from the parent process, and both share the same memory area to read/write the message. Usage of mutex is important since there is no other mechanism to prevent concurrent access to the data.

First In First Out Pipes

A **FIFO** object in Linux is a special type of named pipe which provides first-in-first-out buffering which makes it ideal for inter process messaging.



Creating FIFO Messaging Object

First created by calling `mkfifo()` and later `open()` it and use it for `read()` and `write()` operations.

Messaging over Linux FIFO in C++

The example project creates two processes and the parent process sends messages while the child process receives and prints them. Note that `write()` is non-blocking as long as the FIFO has space but `read()` blocks the current thread. That's why I create an `alarm(1)` that creates `SIGALRM` every second to interrupt ongoing or blocking system calls.

```
linux_fifo_ipc.cpp
1  /**
2   * @file      linux_fifo_ipc.cpp
3   * @author    Atakan S.
4   * @date      01/01/2019
5   * @version   1.0
6   * @brief     Example for Linux IPC using Named Pipe FIFO Messaging.
7   *
8   * @copyright Copyright (c) 2019 Atakan SARIOGLU ~ www.atakansarioglu.com
9   *
10  * Permission is hereby granted, free of charge, to any person obtaining a
11  * copy of this software and associated documentation files (the "Software"),
12  * to deal in the Software without restriction, including without limitation
13  * the rights to use, copy, modify, merge, publish, distribute, sublicense,
14  * and/or sell copies of the Software, and to permit persons to whom the
15  * Software is furnished to do so, subject to the following conditions:
16  *
17  * The above copyright notice and this permission notice shall be included in
18  * all copies or substantial portions of the Software.
19  *
20  * THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR
21  * IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY,
22  * FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE
23  * AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER
24  * LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING
25  * FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER
26  * DEALINGS IN THE SOFTWARE.
27  */
28
29 #include <stdlib.h>
30 #include <stdio.h>
```

```

31 #include <fcntl.h>
32 #include <unistd.h>
33 #include <iostream>
34 #include <chrono>
35 #include <sys/stat.h>
36 #include <signal.h>
37
38 // Alarm signal handler.
39 void alarmHandler(int) {}
40
41 // Simple message type.
42 typedef std::chrono::time_point<std::chrono::high_resolution_clock> Data;
43
44 int main() {
45     // Create FIFO PIPE.
46     int fd;
47     struct stat buffer;
48     if(stat("fifo", &buffer) != 0 && mkfifo("fifo", 0777) < 0) {
49         std::cout << "Cannot create fifo" << std::endl;
50         return -1;
51     }
52
53     // Parent process.
54     if(fork()) {
55         // Open FIFO in write only mode.
56         if((fd = open("fifo", O_WRONLY)) < 0) {
57             std::cout << "Parent: Cannot open fifo" << std::endl;
58             return -1;
59         }
60
61         // Always send message.
62         while(true) {
63             // Send current time, write to the pipe.
64             Data data = std::chrono::system_clock::now();
65             std::cout << "Parent: Sending: " << data.time_since_epoch().count() <<
66
67             // This will send immediately.
68             if(write(fd, &data, sizeof(Data)) < sizeof(Data)) {
69                 std::cout << "Parent: Cannot write to fifo" << std::endl;
70                 return -1;
71             }
72             std::cout << "Parent: Sent." << std::endl;
73             sleep(5);
74         }
75
76         // Child process.
77     } else {
78         // Register signal handler for SIGALRM.
79         struct sigaction signalConfiguration{};
80         signalConfiguration.sa_handler = alarmHandler;
81         sigaction(SIGALRM, &signalConfiguration, 0);
82
83         // Open FIFO in read only mode.
84         if((fd = open("fifo", O_RDONLY)) < 0) {
85             std::cout << "Child: Cannot open fifo" << std::endl;
86             return -1;
87         }
88
89         // Always receive message.
90         while(true) {
91             // Receive message.
92             Data data;
93             std::cout << "Child: Waiting..." << std::endl;
94
95             // Set alarm that will interrupt the read if no message arrives for lo
96             alarm(1);
97

```

```

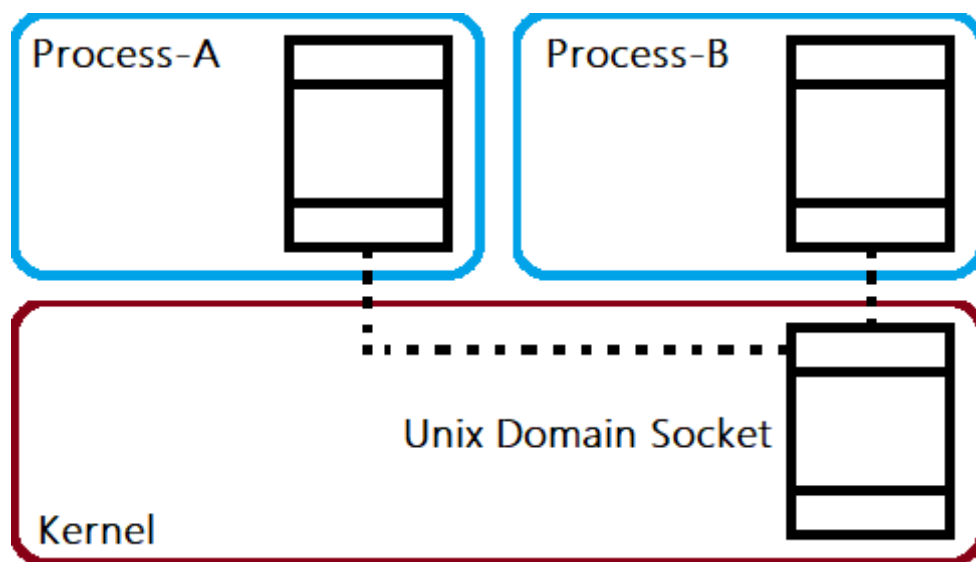
98         // This will block until the parent writes to the pipe.
99         if(read(fd, &data, sizeof(Data)) == sizeof(Data)) {
100             std::cout << "Child: Received: " << data.time_since_epoch().count()
101         }
102     }
103 }
104
105 // Close the pipe.
106 close(fd);
107
108 // Exit.
109 return 0;
110 }

```

Please compile with `g++ linux_fifo_ipc.cpp -o linux_fifo_ipc -std=c++11` command.

Unix Domain Socket

The most complex but the most professional way of Linux IPC is socket communication. **Unix Domain Socket** is not so different than server/client network socket communication but it is intended for local file system use. I will only cover the basics here to showcase its usage.



Creating a Unix Domain Socket

The socket is created just like a file in Linux returning a file handle.

```

1 int fd = socket(AF_UNIX, SOCK_TYPE_SEE_BELOW, 0);

```

Socket file path is passed with `sockaddr_un` struct. Here I use file named `socket`.

```

1 struct sockaddr_un address;

```

```
2 address.sun_family = AF_UNIX;
3 strcpy(address.sun_path, "socket");
```

Connected Server-Client Messaging

Server Side

The server is responsible for socket binding. This will create the socket file on the disk.

```
1 bind(fd, (struct sockaddr *)&address, sizeof(sockaddr_un));
```

Later the server starts listening, allowing multiple connections and accepts connections to serve them.

```
1 listen(fd, NUM_MAX_CONNECTIONS);
2 int connection = accept(fd, (struct sockaddr*)NULL, NULL);
```

Now the returned object from `accept()` is a new file descriptor for `read()` and `write()` operations between the server and the newly connected client. If the server is intended to serve multiple connections, it is a good practice to `fork()` and serve in the child process while the parent process can accept other clients. For IPC purpose, this is rarely necessary.

Client Side

Clients has to create socket and connect to a server. Rest of the operation is `read()` and `write()` just like any other file operation. To close the connection use `shutdown()` function.

IPC Example via Unix Domain Socket in C++

Here is a short example which summarizes the above concept.

```
domain_socket_ipc.cpp
1 /**
2  * @file      domain_socket_ipc.cpp
3  * @author    Atakan S.
4  * @date      01/01/2019
5  * @version   1.0
6  * @brief     Example IPC Messaging using Unix Domain Socket Connection.
7  *
8  * @copyright Copyright (c) 2019 Atakan SARIOGLU ~ www.atakansarioglu.com
9  *
10 * Permission is hereby granted, free of charge, to any person obtaining a
11 * copy of this software and associated documentation files (the "Software"),
12 * to deal in the Software without restriction, including without limitation
13 * the rights to use, copy, modify, merge, publish, distribute, sublicense,
```



```

14  * and/or sell copies of the Software, and to permit persons to whom the
15  * Software is furnished to do so, subject to the following conditions:
16  *
17  * The above copyright notice and this permission notice shall be included in
18  * all copies or substantial portions of the Software.
19  *
20  * THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR
21  * IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY,
22  * FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE
23  * AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER
24  * LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING
25  * FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER
26  * DEALINGS IN THE SOFTWARE.
27  */
28
29 #include <stdlib.h>
30 #include <stdio.h>
31 #include <fcntl.h>
32 #include <unistd.h>
33 #include <iostream>
34 #include <chrono>
35 #include <sys/stat.h>
36 #include <sys/socket.h>
37 #include <sys/un.h>
38 #include <string>
39 #include <memory.h>
40
41 int main() {
42     // Unix domain socket file address.
43     struct sockaddr_un address;
44     address.sun_family = AF_UNIX;
45     strcpy(address.sun_path, "socket");
46
47     // Server process.
48     if(fork()) {
49         // Delete the old socket file.
50         unlink("socket");
51
52         // Create a unix domain socket.
53         int fd;
54         if((fd = socket(AF_UNIX, SOCK_STREAM, 0)) < 0) {
55             std::cout << "Server: Cannot create socket" << std::endl;
56             return -1;
57         }
58
59         // Bind the socket to the address.
60         if(bind(fd, (struct sockaddr *)&address, sizeof(sockaddr_un)) != 0) {
61             std::cout << "Server: Cannot bind socket" << std::endl;
62             return -1;
63         }
64
65         // Listen up to 100 connections.
66         if(listen(fd, 100) != 0) {
67             std::cout << "Server: Cannot listen" << std::endl;
68             return -1;
69         }
70
71         // Serve.
72         while(true) {
73             // Accept incoming connections.
74             std::cout << "Server: Waiting connection..." << std::endl;
75             int connection;
76             if((connection = accept(fd, (struct sockaddr*)&NULL, NULL)) < 0) {
77                 std::cout << "Server: Connection cannot be accepted" << std::endl;
78                 return -1;
79             }
80             std::cout << "Server: Connected." << std::endl;

```

```

81
82 // After accpting a connection, fork to child process to serve it.
83 if(!fork()) {
84     // Read as long as the connection is alive.
85     while(true) {
86         std::cout << "Server: Waiting..." << std::endl;
87
88         // Receive data.
89         char buffer[100];
90         if(read(connection, buffer, 100) > 0) {
91             std::cout << "Server: Received: " << buffer << std::endl;
92         } else break;
93     }
94
95     // Close the connection.
96     close(connection);
97 }
98
99 // Close the socket.
100 close(fd);
101
102 // Client process.
103 } else {
104     // Initial delay until the server is ready.
105     sleep(1);
106
107     // Create a unix domain socket.
108     int fd;
109     if((fd = socket(AF_UNIX, SOCK_STREAM, 0)) < 0) {
110         std::cout << "Client: Cannot create socket" << std::endl;
111         return -1;
112     }
113
114     // Connect to the server.
115     while(connect(fd, (struct sockaddr *)&address, sizeof(sockaddr_un)) != 0)
116         std::cout << "Client: Connecting... " << std::endl;
117         sleep(1);
118     }
119
120     // Send timestamp.
121     int counter = 5;
122     while(counter-- > 0) {
123         auto now = std::to_string(std::chrono::system_clock::now().time_since_epoch());
124         std::cout << "Client: Requesting with: " << now << std::endl;
125
126         // Write data to the server.
127         if(write(fd, now.c_str(), now.size()) < now.size()) {
128             std::cout << "Client: Cannot send request" << std::endl;
129             return -1;
130         }
131
132         // Delay.
133         sleep(1);
134     }
135
136     // Shutdown and close the connection.
137     shutdown(fd, SHUT_RDWR);
138     close(fd);
139 }
140
141 // Exit.
142 return 0;
143 }
144

```

The client connects to the server and sends timestamp but also the server can send data back. The above code compiles with `g++ domain_socket_ipc.cpp -o domain_socket_ipc -std=c++11` command.

Connectionless Datagram Messaging

Different than the above approach, here there is no connection and let's say no server/client. One of the pairs has to `bind()` the socket and that is the only difference between them. If you are messaging full-duplex, be careful not to read the messages that you have sent yourself (take a look at `MSG_PEEK` flag).

Simple IPC Unix Domain Socket Messaging in C++

This is much simpler than the connected method I mentioned above. Using `sendto()` and `recvfrom()` methods, you can pass messages between processes. I won't call the processes server and client instead they have nearly identical roles.

```
socket_datagram_ipc.cpp
1  /**
2   * @file      socket_datagram_ipc.cpp
3   * @author    Atakan S.
4   * @date      01/01/2019
5   * @version   1.0
6   * @brief     Example IPC Messaging using Unix Domain Socket Datagram.
7   *
8   * @copyright Copyright (c) 2019 Atakan SARIOGLU ~ www.atakansarioglu.com
9   *
10  * Permission is hereby granted, free of charge, to any person obtaining a
11  * copy of this software and associated documentation files (the "Software"),
12  * to deal in the Software without restriction, including without limitation
13  * the rights to use, copy, modify, merge, publish, distribute, sublicense,
14  * and/or sell copies of the Software, and to permit persons to whom the
15  * Software is furnished to do so, subject to the following conditions:
16  *
17  * The above copyright notice and this permission notice shall be included in
18  * all copies or substantial portions of the Software.
19  *
20  * THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR
21  * IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY,
22  * FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE
23  * AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER
24  * LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING
25  * FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER
26  * DEALINGS IN THE SOFTWARE.
27  */
28
29 #include <stdlib.h>
30 #include <stdio.h>
31 #include <fcntl.h>
32 #include <unistd.h>
33 #include <iostream>
34 #include <chrono>
35 #include <sys/stat.h>
36 #include <sys/socket.h>
37 #include <sys/un.h>
38 #include <string>
```

```

39 #include <memory.h>
40
41 int main() {
42     // Unix domain socket file address.
43     struct sockaddr_un address;
44     address.sun_family = AF_UNIX;
45     strcpy(address.sun_path, "socket");
46
47     // Receiver process.
48     if(fork()) {
49         // Delete the old socket file.
50         unlink("socket");
51
52         // Create a unix domain socket.
53         int fd;
54         if((fd = socket(AF_UNIX, SOCK_DGRAM, 0)) < 0) {
55             std::cout << "Receiver: Cannot create socket" << std::endl;
56             return -1;
57         }
58
59         // Bind the socket to the address.
60         if(bind(fd, (struct sockaddr *)&address, sizeof(sockaddr_un)) != 0) {
61             std::cout << "Receiver: Cannot bind socket" << std::endl;
62             return -1;
63         }
64
65         // Receive data.
66         while(true) {
67             // Check for incoming data, non-blocking.
68             char buffer[100];
69             if(recvfrom(fd, buffer, 100, MSG_DONTWAIT, NULL, NULL) > 0) {
70                 std::cout << "Receiver: Read: " << buffer << std::endl;
71             }
72
73             // Delay.
74             sleep(1);
75         }
76
77         // Close the socket.
78         close(fd);
79
80         // Sender process.
81     } else {
82         // Initial delay until the socket is ready.
83         sleep(1);
84
85         // Create a unix domain socket.
86         int fd;
87         if((fd = socket(AF_UNIX, SOCK_DGRAM, 0)) < 0) {
88             std::cout << "Sender: Cannot create socket" << std::endl;
89             return -1;
90         }
91
92         // Send timestamp.
93         while(true) {
94             auto now = std::to_string(std::chrono::system_clock::now().time_since_
95             std::cout << "Sender: Writing: " << now << std::endl;
96
97             // Write data to the socket.
98             if(sendto(fd, now.c_str(), now.size(), 0, (struct sockaddr *)&address,
99             std::cout << "Client: Cannot send" << std::endl;
100             return -1;
101         }
102
103         // Delay.
104         sleep(5);
105     }

```

```
106
107     // Close the socket.
108     close(fd);
109 }
110
111 // Exit.
112 return 0;
113 }
```

Compile with `g++ socket_datagram_ipc.cpp -o socket_datagram_ipc -std=c++11` command.

Final Words

I tried to show very basic examples but the IPC topic is huge. Keep this as a starting point and refer to [here](#) for shared memory, [here](#) for named pipes and [here](#) for unix domain socket messaging. Try hard 👍

TAGS: [C++](#), [FEATURED3](#), [FIFO](#), [INTER PROCESS COMMUNICATION](#), [INTER PROCESS MESSAGING](#), [IPC](#),
[LINUX](#), [NAMED PIPE](#), [SHARED MEMORY](#), [SOCKET COMMUNICATION](#), [SYSTEM CALL](#),
[UNIX DOMAIN SOCKET](#)