

Under the Hood of AFD.sys Part 4: Receiving TCP packets

Mateusz Lewczak : : 06/08/2025

A hands-on foray into the IOCTL_AFD_RECEIVE Fast-I/O path: stalking AfdFastConnectionReceive in WinDbg, decoding the AFD_SENDRECV_INFO / WSABUF triad, flipping TDI flags for peek-and-poke tricks, and slurping raw TCP responses straight out of AFD.sys—zero Winsock, pure kernel-level packet sorcery.

Posted Aug 6, 2025

By [Mateusz Lewczak](#)

11 min read

Under the Hood of AFD.sys Part 4: Receiving TCP packets

Introduction

Ok, the time has come, we can finally receive some data in our sockets. If you haven't seen the previous batches, I encourage you to check them out, so far we've managed to create a socket and send TCP packets. We still have a long way to go to fully understand how networking works by communicating directly with the AFD.sys driver, but there will be time for that yet. No need to procrastinate, let's go!

Looking for recv

As before, `recv` is handled as *Fast I/O* (unless we set the flags differently). A good reference for this will be `IOCTL_AFD_RECEIVE` going into our `AfdFastIoDeviceControl` function. By default, as before, our reference will be this code using Winsock:

```
1 void createTCPv4() {
2     const size_t PAYLOAD = 1024;
3     SOCKET s = socket(AF_INET, SOCK_STREAM, IPPROTO_TCP);
4     if (s == INVALID_SOCKET) { std::cerr << "socket: " << WSAGetLastError() << '\n';
5     return; }
6
7     sockaddr_in dst{};
8     dst.sin_family = AF_INET;
9     dst.sin_port = htons(80);
10    InetPtonA(AF_INET, "192.168.1.1", &dst.sin_addr);
11
12    if (connect(s, reinterpret_cast<sockaddr*>(&dst), sizeof(dst)) == SOCKET_ERROR) {
13        std::cerr << "connect: " << WSAGetLastError() << '\n';
14        closesocket(s); return;
15    }
16
17    std::string big(PAYLOAD, 'A');
18
19    size_t sent = 0;
20    while (sent < big.size()) {
```

```

21     int n = send(s, big.data() + sent, static_cast<int>(big.size() - sent), 0);
22     if (n == SOCKET_ERROR) {
23         break;
24     }
25     sent += n;
26 }
27
28 char buf[4096];
29 int n = 0;
30 size_t received = 0;
31 std::string response;
32
33 while ((n = recv(s, buf, static_cast<int>(sizeof(buf)), 0)) > 0) {
34     response.append(buf, n);
35     received += n;
36 }
37
38 if (n == SOCKET_ERROR) {
39     std::cerr << "recv: " << WSAGetLastError() << '\n';
40 }
41
42 std::cout << "Received " << received << " bytes\n";
43 std::cout << "----- RESPONSE BEGIN -----`n"
44     << response << '\n';
45     << "----- RESPONSE END -----`n";
46
47 closesocket(s);
}

```

As we connect to the HTTP server and send garbage we get a `Bad request` in response - a clear case. To confirm that we are indeed `recv` hitting the driver as *Fast I/O* we will use a command like this in WinDbg:

```

14: kd> .foreach /pS 1 (ep { !process 0 0 afd_re.exe }) { bm /p ${ep}
1 afd!AfdFastIoDeviceControl ".printf \"IoControlCode=%p\\n\", @rdi;gc;" }
2 0: fffff801`6d9d4c20 @"!afd!AfdFastIoDeviceControl"
3 Couldn't resolve error at 'SessionId: afd!AfdFastIoDeviceControl ".printf
4 \"IoControlCode=%p\\n\", @rdi;gc;" '
5 14: kd> g
6 IoControlCode=000000000001207b // IOCTL_AFD_GET_INFORMATION
7 IoControlCode=000000000001207b // IOCTL_AFD_GET_INFORMATION
8 IoControlCode=0000000000012047 // IOCTL_AFD_SET_CONTEXT
9 IoControlCode=00000000000120bf // IOCTL_AFD_TRANSPORT_IOCTL
10 IoControlCode=0000000000012047 // IOCTL_AFD_SET_CONTEXT
11 IoControlCode=0000000000012003 // IOCTL_AFD_BIND
12 IoControlCode=0000000000012047 // IOCTL_AFD_SET_CONTEXT
13 IoControlCode=0000000000012007 // IOCTL_AFD_CONNECT
14 IoControlCode=0000000000012047 // IOCTL_AFD_SET_CONTEXT
15 IoControlCode=000000000001201f // IOCTL_AFD_SEND
16 IoControlCode=0000000000012017 // IOCTL_AFD_RECEIVE
    IoControlCode=0000000000012017 // IOCTL_AFD_RECEIVE

```

As we can see our request `IOCTL_AFD_RECEIVE` appears twice. We can explain this by the fact that in our code, the `recv` function is executed in a loop. In practice, we retrieve the response in packets of 4096 bytes until we have received the entire TCP response. The first time we received the entire HTTP response and presumably AFD.sys returned information about how many bytes we actually received. And the second call with which we wanted to retrieve the rest returned us zero bytes, so no more requests were sent - a simple matter.

It's time to find a direct function that is responsible for handling this request, as in `AfdFastConnectionSend`. Let's check this statically using Binary Ninja:

```
1 1c0034be0    int64_t AfdFastIoDeviceControl(struct _FILE_OBJECT* FileObject,
2 1c0034be0        BOOLEAN Wait, PVOID InputBuffer, ULONG InputBufferLength,
3 1c0034be0        PVOID OutputBuffer, ULONG OutputBufferLength, ULONG IoControlCode,
4 1c0034be0        PIO_STATUS_BLOCK IoStatus, struct _DEVICE_OBJECT* DeviceObject) {
5 ...
6 1c00354a6            rbx = (uint64_t)AfdFastConnectionReceive(FsContext, &s,
7 1c00354a6                rax_51, IoStatus);
8 ...
9 1c0034be0        }
```

This time in the code we don't find a condition that directly checks if `IoControlCode == 0x12017`, what's more, before calling our target function we also have a number of checks that for now we don't know what they do. Let's take a breakpoint on this function:

```
12: kd> .foreach /pS 1 (ep { !process 0 0 afd_re.exe }) { bm /p ${ep}
1 afd!AfdFastConnectionReceive ".printf \"HIT!\n\";gc;" }
2 4: fffff801`6d9d3280 @!"afd!AfdFastConnectionReceive"
3 Couldn't resolve error at 'SessionId: afd!AfdFastConnectionReceive ".printf
4 \"HIT!\n\";gc;" '
5 12: kd> g
HIT!
```

We only have one hit despite the fact that two `IOCTL_AFD_RECEIVE` requests went, this could mean that these check functions before calling `AfdFastConnectionReceive` check if, for example, the internal response buffer for the socket is empty.

We now turn to examining what our input buffer looks like for this request. Here, as usual, our invaluable sources ([killvxk](#)), ([unknowncheats.me](#) ICoded post), ([ReactOS Project](#)), ([DynamoRIO / Dr. Memory](#)), ([DeDf](#)), ([diversenok](#)) will help us.

```
1 10: kd> .foreach /pS 1 (ep { !process 0 0 afd_re.exe }) { bm /p ${ep}
2 afd!AfdFastConnectionReceive }
3 6: fffff801`6d9d3280 @!"afd!AfdFastConnectionReceive"
4 Couldn't resolve error at 'SessionId: afd!AfdFastConnectionReceive '
5 10: kd> g
6 Breakpoint 6 hit
7 afd!AfdFastConnectionReceive:
8 fffff801`6d9d3280 4c894c2420      mov      qword ptr [rsp+20h],r9
9 4: kd> r
10 rax=0000000000000002  rbx=00000001ac3ae028  rcx=fffff8b05eaffb340
11 rdx=fffff58d13a4ef10  rsi=0000000000000001  rdi=0000000000000000
12 rip=fffff8016d9d3280  rsp=fffff58d13a4ee88  rbp=fffff58d13a4f4e0
13 r8=00000000000001000  r9=fffff58d13a4f1c8  r10=fffff801d8817c70
14 r11=fffffb1fcfd3800000  r12=fffff8b05eaffb340  r13=0000000000000000
15 r14=00000000000000018  r15=0000000000000af01
16 iopl=0            nv up ei pl zr na po nc
17 cs=0010  ss=0018  ds=002b  es=002b  fs=0053  gs=002b          efl=00040246
18 afd!AfdFastConnectionReceive:
19 fffff801`6d9d3280 4c894c2420      mov      qword ptr [rsp+20h],r9
20 ss:0018:fffff58d`13a4eea8=00000000000000003
21 4: kd> dq 00000001ac3ae028 L3
22 00000001`ac3ae028  00000001`ac3ae108  00000000`00000001
23 00000001`ac3ae038  00000000`00000020
24 4: kd> dq 00000001`ac3ae108 L2
25 00000001`ac3ae108  00000000`00001000  00000001`ac3ae260
26 4: kd> dq 00000001`ac3ae260 L10
27 00000001`ac3ae260  cccccccc`ccccccc cccccccc`ccccccc
```

```

28 00000001`ac3ae270  cccccccc`cccccccc cccccccc`cccccccc
29 00000001`ac3ae280  cccccccc`cccccccc cccccccc`cccccccc
30 00000001`ac3ae290  cccccccc`cccccccc cccccccc`cccccccc
31 00000001`ac3ae2a0  cccccccc`cccccccc cccccccc`cccccccc
32 00000001`ac3ae2b0  cccccccc`cccccccc cccccccc`cccccccc
  00000001`ac3ae2c0  cccccccc`cccccccc cccccccc`cccccccc
  00000001`ac3ae2d0  cccccccc`cccccccc cccccccc`cccccccc

```

We can see that essentially the structure of the input buffer is identical to the one we use to send packets (see part 3). With a slight difference, in our structure we have `AfdFlags`, which are flags describing our buffer. When they are set to `0x00` as in the case of sending then `AFD.sys` treats them as *send buffer*.

Analyzing retrieved data AfdFastConnectionReceive

Earlier we were guided by ([diversenok](#)) to guess what values the flags can take and nowhere there was a value `0x20`. We can instead look at ([unknowncheats.me ICoded post](#)), there we find such definitions:

1 #define TDI_RECEIVE_BROADCAST	0x4
2 #define TDI_RECEIVE_MULTICAST	0x8
3 #define TDI_RECEIVE_PARTIAL	0x10
4 #define TDI_RECEIVE_NORMAL	0x20
5 #define TDI_RECEIVE_EXPEDITED	0x40
6 #define TDI_RECEIVE_PEEK	0x80
7 #define TDI_RECEIVE_NO_RESPONSE_EXP	0x100
8 #define TDI_RECEIVE_COPY_LOOKAHEAD	0x200
9 #define TDI_RECEIVE_ENTIRE_MESSAGE	0x400
10 #define TDI_RECEIVE_AT_DISPATCH_LEVEL	0x800
11 #define TDI_RECEIVE_CONTROL_INFO	0x1000
12 #define TDI_RECEIVE_FORCE_INDICATION	0x2000
13 #define TDI_RECEIVE_NO_PUSH	0x4000

That is, `0x20` would imply a normal reception of data from the driver. But there is a detail, according to ([unknowncheats.me ICoded post](#)), these flags apply to the `TdiFlags` field:

```

NTSTATUS AfdRecv(HANDLE SocketHandle, PVOID Buffer, ULONG_PTR BufferLength, PULONG_PTR
1 pBytes)
2 {
3     NTSTATUS Status;
4     IO_STATUS_BLOCK IoStatus;
5     AFD_SENDRECV_INFO RecvInfo;
6     HANDLE Event;
7     AFD_WSABUF AfdBuffer;
8
9     Status = NtCreateEvent(&Event, EVENT_ALL_ACCESS, NULL, NotificationEvent,
10 FALSE);
11     if (NT_SUCCESS(Status))
12     {
13         /**
14             AfdBuffer.len = (ULONG)BufferLength;
15
16             RecvInfo.BufferArray = &AfdBuffer;
17             RecvInfo.BufferCount = 1;
18             RecvInfo.TdiFlags = TDI_RECEIVE_NORMAL;
19             RecvInfo.AfdFlags = 0;
20         /**
21     }
22     return Status;
}

```

Which in our case is not quite true. The buffer sent to `AFD.sys` is `0x18` in size, i.e. it has three fields of `0x8` bytes. And this third field (in our case `AfdFlags`) is just set to `0x20`. I have experimentally checked and our version is the one that works. Of course, I am not saying that the ([unknowncheats.me |Coded post](https://unknowncheats.me/ICoded/post)) version does not work, it just does not apply in our case.

With all this in mind, let us create a working proof-of-concept using everything we already have:

```

1 #include <stdint.h>
2 #include <Windows.h>
3 #include <winternl.h>
4 #include <iostream>
5 #include "afd_defs.h"
6 #include "afd_ioctl.h"
7 #pragma comment(lib, "ntdll.lib")
8
9 NTSTATUS createAfdSocket(PHANDLE socket) {/**/}
10 NTSTATUS bindAfdSocket(HANDLE socket) {/**/}
11 NTSTATUS connectAfdSocket(HANDLE socket) {/**/}
12 NTSTATUS sendAfdPacketTCP(HANDLE socket) {/**/}
13
14 NTSTATUS receiveAfdPacketTCP(HANDLE socket) {
15     const int BUF_NUM = 1;
16     const int BUF_SIZE = 1000;
17     AFD_BUFF* payload = new AFD_BUFF[BUF_NUM];
18     for (int i = 0; i < BUF_NUM; i++) {
19         payload[i].buf = (uint8_t*)malloc(BUF_SIZE);
20         memset(payload[i].buf, 0x00, BUF_SIZE);
21         payload[i].len = BUF_SIZE;
22     }
23
24     AFD_SEND_PACKET* afdSendPacket = new AFD_SEND_PACKET;
25     afdSendPacket->buffersArray = payload;
26     afdSendPacket->buffersCount = BUF_NUM;
27     afdSendPacket->afdFlags = 0x20; // RECEIVE_NORMAL
28
29     IO_STATUS_BLOCK ioStatus;
30     NTSTATUS status = NtDeviceIoControlFile(socket, NULL, NULL, NULL, &ioStatus,
31 IOCTL_AFD_RECEIVE,
32                                         afdSendPacket, sizeof(AFD_SEND_PACKET),
33                                         nullptr, 0);
34     if (status == STATUS_PENDING) {
35         WaitForSingleObject(socket, INFINITE);
36         status = ioStatus.Status;
37     }
38
39     std::cout << "[+] SERVER RESPONSE: " << std::endl;
40     std::cout << payload[0].buf << std::endl;
41
42     return status;
43 }
44
45 int main() {
46     HANDLE socket;
47     // 1. Create socket
48     // 2. Bind socket
49     // 3. Connect to remote host
50     // 4. Send 1000x'A'
51
52     status = receiveAfdPacketTCP(socket);
53     if (!NT_SUCCESS(status)) {
54         std::cout << "[-] Could not receive TCP packet: " << std::hex << status <<
55 std::endl;
56         return 1;
57     }

```

```

58     std::cout << "[+] Received!" << std::endl;
59
60     return 0;
61 }
```

What we do. We allocate our buffers to store the received response somewhere, set `AfdFlags` to `0x20` (normal reception), and then send the request to `AFD.sys`. What more do you need?

Well, it would be useful to somehow find out how much of this data we have received. At first, I thought that maybe `AFD.sys` would modify the buffer structure and change its size. However, this did not happen. The answer was much simpler. The `IO_STATUS_BLOCK` structure has an `Information` field:

```

1 // ref: https://learn.microsoft.com/en-us/windows-hardware/drivers/ddi/wdm/ns-wdm-
2 _io_status_block
3 typedef struct _IO_STATUS_BLOCK {
4     union {
5         NTSTATUS Status;
6         PVOID     Pointer;
7     };
8     ULONG_PTR Information;
9 } IO_STATUS_BLOCK, *PIO_STATUS_BLOCK;
```

And it is in the `Information` field that we get a return on how many bytes have been read, but we do not know how many are actually left to read. To check this we would now have to send another request to `AFD.sys` and check if `Information` is equal to `0x0`. Then only then would we know if this is all there is.

Other receive flags

This question may be best answered by [documentation](#) from Microsoft. All the flags are very nicely described there. Although some of them are explained in terminology familiar to driver developers. I tried to reproduce some of them in Winsock and 'make up' my own explanation:

```

1 // Receive normal packets
2 afdSendPacket->afdFlags = TDI_RECEIVE_NORMAL;
3 // Receive normal packet, but don't clear AFD.sys input queue
4 afdSendPacket->afdFlags = TDI_RECEIVE_NORMAL | TDI_RECEIVE_PEEK;
5 // Receive normal packet, but wait for all data, equivalent of MSG_WAITALL in Winsock
6 afdSendPacket->afdFlags = TDI_RECEIVE_NORMAL | TDI_RECEIVE_NO_PUSH;
7 // Receive packets with tcp.flags.urg == 1
8 afdSendPacket->afdFlags = TDI_RECEIVE_EXPEDITED;
9 // Receive packets with tcp.flags.urg == 1, but don't clear AFD.sys input queue
10 afdSendPacket->afdFlags = TDI_RECEIVE_EXPEDITED | TDI_RECEIVE_PEEK;
11 // Receive packets with tcp.flags.urg == 1, but wait for all data, equivalent of
12 MSG_WAITALL in Winsock
13 afdSendPacket->afdFlags = TDI_RECEIVE_EXPEDITED | TDI_RECEIVE_NO_PUSH;
```

Some of our flags relate to UDP and we will certainly look at this when the opportunity arises.

Next steps

At this point, we already have the necessary functionality to be able to create a simple TCP client. In the next batches we will look more at socket operations. How to change its parameters, how to close a connection, how to close a socket, how to handle other types of TCP messages.

Final code

Essentially, the content of the final code is no different from what you can find in the proof-of-concept above. Of course, the code for IPv6 will look identical.

References

1. Vittitoe, Steven. "Reverse Engineering Windows AFD.sys: Uncovering the Intricacies of the Ancillary Function Driver." *Proceedings of REcon 2015*, 2015, <https://doi.org/10.5446/32819>.
2. killvxk. *CVE-2024-38193 Nephster PoC*. 2024, <https://github.com/killvxk/CVE-2024-38193-Nephster/blob/main/Poc/poc.h>.
3. unknowncheats.me ICoded post. *Native TCP Client Socket*. n.d., <https://www.unknowncheats.me/forum/c-and-c-/500413-native-tcp-client-socket.html>.
4. ReactOS Project. *Afd.h*. n.d., <https://github.com/reactos/reactos/blob/master/drivers/network/afd/include/afd.h>.
5. DynamoRIO / Dr. Memory. *afd_shared.h*. n.d., https://github.com/DynamoRIO/drmemory/blob/master/wininc/afd_shared.h.
6. Dr. Memory - GH issue#376. *Issue #376: AFD Support Improvements*. n.d., <https://github.com/DynamoRIO/drmemory/issues/376>.
7. Microsoft. *NtCreateFile Function (Winternl.h)*. n.d., <https://learn.microsoft.com/en-us/windows/win32/api/winternl/nf-winternl-ntcreatefile>.
8. ---. *x64 Calling Convention*. n.d., <https://learn.microsoft.com/en-us/cpp/build/x64-calling-convention?view=msvc-170>.
9. ---. *x64 Calling Convention*. n.d., <https://learn.microsoft.com/pl-pl/windows/win32/api/winsock2/nf-winsoc2-wsasocketa>.
10. DeDf. *AFD Repository*. n.d., <https://github.com/DeDf/afd/tree/master>.
11. Allievi, Andrea, et al. *Windows® Internals Part 2 - 6th Edition*. 6th ed., Microsoft Press (Pearson Education), 2022, <https://learn.microsoft.com/sysinternals/resources/windows-internals>.
12. diversenok. *\Text\ntntafdf.h – Ancillary Function Driver Definitions*. commit 2dda0dd, Hunt & Hackett, April 2025, <https://github.com/winsiderss/systeminformer/blob/master/phnt/include/ntafdf.h>.