author	Dominique Martinet <asmadeus@codewreck.org>	2025-03-19 20:20:15 +0900
committer	Greg Kroah-Hartman <gregkh@linuxfoundation.org>	2025-05-02 07:59:20 +0200
commit	c548f95688e2b5ae0e2ae43d53cf717156c7d034 (patch)	
tree	be59daad2367c9392291d3cc727af9dda95dd530	
parent	a3b8d8cf519693b9db81c6c86a2a201985404c10 (diff)	
download	linux-c548f95688e2b5ae0e2ae43d53cf717156c7d034.tar.gz	

diff options

context: 3

space: include

mode: unified

9p/net: fix improper handling of bogus negative read/write replies

```
[ Upstream commit d0259a856afca31d699b706ed5e2adf11086c73b ]

In p9_client_write() and p9_client_read_once(), if the server
incorrectly replies with success but a negative write/read count then we
would consider written (negative) <= rsize (positive) because both
variables were signed.

Make variables unsigned to avoid this problem.

The reproducer linked below now fails with the following error instead
of a null pointer deref:
9pnet: bogus RWRITE count (4294967295 > 3)

Reported-by: Robert Morris <rtm@mit.edu>
Closes: https://lore.kernel.org/16271.1734448631@26-5-164.dynamic.csail.mit.edu
Message-ID: <20250319-9p_unsigned_rw-v3-1-71327f1503d0@codewreck.org>
Reviewed-by: Christian Schoenebeck <linux_oss@crudebyte.com>
Signed-off-by: Dominique Martinet <asmadeus@codewreck.org>
Signed-off-by: Sasha Levin <sashal@kernel.org>
```

Diffstat

-rw-r--r-- net/9p/client.c 30

1 files changed, 16 insertions, 14 deletions

```
diff --git a/net/9p/client.c b/net/9p/client.c
index 09f8ced9f8bb7f..52a5497cfca794 100644
--- a/net/9p/client.c
+++ b/net/9p/client.c
@@ -1548,7 +1548,8 @@ p9_client_read_once(struct p9_fid *fid, u64 offset, struct iov_iter *to,
        struct p9_client *clnt = fid->clnt;
        struct p9_req_t *req;
        int count = iov_iter_count(to);
-       int rsize, received, non_zc = 0;
+       u32 rsize, received;
+       bool non_zc = false;
        char *dataptr;

        *err = 0;
@@ -1571,7 +1572,7 @@ p9_client_read_once(struct p9_fid *fid, u64 offset, struct iov_iter *to,
                                       0, 11, "dqd", fid->fid,
                                       offset, rsize);
        } else {
-               non_zc = 1;
+               non_zc = true;
                req = p9_client_rpc(clnt, P9_TREAD, "dqd", fid->fid, offset,
```

```
                                     rsize);
        }
@@ -1592,11 +1593,11 @@ p9_client_read_once(struct p9_fid *fid, u64 offset, struct iov_iter *to,
                return 0;
        }
        if (rsize < received) {
-               pr_err("bogus RREAD count (%d > %d)\n", received, rsize);
+               pr_err("bogus RREAD count (%u > %u)\n", received, rsize);
                received = rsize;
        }

-       p9_debug(P9_DEBUG_9P, "<<< RREAD count %d\n", received);
+       p9_debug(P9_DEBUG_9P, "<<< RREAD count %u\n", received);

        if (non_zc) {
                int n = copy_to_iter(dataptr, received, to);
@@ -1623,9 +1624,9 @@ p9_client_write(struct p9_fid *fid, u64 offset, struct iov_iter *from, int *err)
        *err = 0;

        while (iov_iter_count(from)) {
-               int count = iov_iter_count(from);
-               int rsize = fid->iounit;
-               int written;
+               size_t count = iov_iter_count(from);
+               u32 rsize = fid->iounit;
+               u32 written;

                if (!rsize || rsize > clnt->msize - P9_IOHDRSZ)
                        rsize = clnt->msize - P9_IOHDRSZ;
@@ -1633,7 +1634,7 @@ p9_client_write(struct p9_fid *fid, u64 offset, struct iov_iter *from, int *err)
                if (count < rsize)
                        rsize = count;

-               p9_debug(P9_DEBUG_9P, ">>> TWRITE fid %d offset %llu count %d (/%d)\n",
+               p9_debug(P9_DEBUG_9P, ">>> TWRITE fid %d offset %llu count %u (/%zu)\n",
                         fid->fid, offset, rsize, count);

                /* Don't bother zerocopy for small IO (< 1024) */
@@ -1659,11 +1660,11 @@ p9_client_write(struct p9_fid *fid, u64 offset, struct iov_iter *from, int *err)
                        break;
                }
                if (rsize < written) {
-                       pr_err("bogus RWRITE count (%d > %d)\n", written, rsize);
+                       pr_err("bogus RWRITE count (%u > %u)\n", written, rsize);
                        written = rsize;
                }

-               p9_debug(P9_DEBUG_9P, "<<< RWRITE count %d\n", written);
+               p9_debug(P9_DEBUG_9P, "<<< RWRITE count %u\n", written);

                p9_req_put(clnt, req);
                iov_iter_revert(from, count - written - iov_iter_count(from));
@@ -2098,7 +2099,8 @@ EXPORT_SYMBOL_GPL(p9_client_xattrcreate);

 int p9_client_readdir(struct p9_fid *fid, char *data, u32 count, u64 offset)
 {
-       int err, rsize, non_zc = 0;
+       int err, non_zc = 0;
+       u32 rsize;
        struct p9_client *clnt;
        struct p9_req_t *req;
        char *dataptr;
@@ -2107,7 +2109,7 @@ int p9_client_readdir(struct p9_fid *fid, char *data, u32 count, u64 offset)

        iov_iter_kvec(&to, ITER_DEST, &kv, 1, count);
```

```
-       p9_debug(P9_DEBUG_9P, ">>> TREADDIR fid %d offset %llu count %d\n",
+       p9_debug(P9_DEBUG_9P, ">>> TREADDIR fid %d offset %llu count %u\n",
                 fid->fid, offset, count);

        clnt = fid->clnt;
@@ -2142,11 +2144,11 @@ int p9_client_readdir(struct p9_fid *fid, char *data, u32 count, u64 offset)
                goto free_and_error;
        }
        if (rsize < count) {
-               pr_err("bogus RREADDIR count (%d > %d)\n", count, rsize);
+               pr_err("bogus RREADDIR count (%u > %u)\n", count, rsize);
                count = rsize;
        }

-       p9_debug(P9_DEBUG_9P, "<<< RREADDIR count %d\n", count);
+       p9_debug(P9_DEBUG_9P, "<<< RREADDIR count %u\n", count);

        if (non_zc)
                memmove(data, dataptr, count);
```