

## x86\_64: 32bit emulation problems

**Source:** <http://linux.derkeiler.com/Mailing-Lists/Kernel/2005-02/6518.html>

---

**From:** Bernd Schubert ([bernd.schubert\\_at\\_pci.uni-heidelberg.de](mailto:bernd.schubert_at_pci.uni-heidelberg.de))

**Date:** 02/28/05

To: Linux Kernel Mailing List <[linux-kernel@vger.kernel.org](mailto:linux-kernel@vger.kernel.org)>, Andi Kleen <[ak@muc.de](mailto:ak@muc.de)>

Date: Mon, 28 Feb 2005 21:54:03 +0100

Hi,

I'm just looking into a very strange problem. Some of our systems have athlon64 CPUs. Due to our diskless nfs environment we currently still prefer a 32bit userspace environment, but would like to be able to use a 64-bit chroot environment.

Well, currently there seems to be a stat64() NFS problem when a x86\_64 kernel is booted and stat64() comes from a 32bit libc.

Here's just an example:

```
hitchcock:/home/bernd/src/tests# ./test_stat64 /mnt/test/yp
stat() works fine.
```

```
hitchcock:/home/bernd/src/tests# ./test_stat32 /mnt/test/yp
stat for /mnt/test/yp failed
```

The test program looks rather simple:

```
#include <sys/types.h>
#include <sys/stat.h>
#include <unistd.h>
#include <stdio.h>
#include <errno.h>
#include <string.h>
#include <stdlib.h>

int main(int argc, char **argv)
{
    char *dir;
    struct stat buf;

    dir = argv[1];

    if (stat (dir, &buf) == -1)
        fprintf(stderr, "stat for %s failed \n", dir);
```

## Linux-Kernel: x86\_64: 32bit emulation problems

```
else
    fprintf(stderr, "stat() works fine.\n");
return (0);
}
```

Here are the strace outputs:

=====

32bit:

-----

```
hitchcock:/home/bernd/src/tests# strace32 ./test_stat32 /mnt/test/yp
execve("./test_stat32", ["/test_stat32", "/mnt/test/yp"], [/* 39 vars */]) =
0
uname({sys="Linux", node="hitchcock", ...}) = 0
brk(0) = 0x80ad000
brk(0x80ce000) = 0x80ce000
stat64("/mnt/test/yp", {st_mode=S_IFDIR|0755, st_size=2704, ...}) = 0
write(2, "stat for /mnt/test/yp failed \n", 30stat for /mnt/test/yp failed
) = 30
exit_group(0) = ?
```

64bit:

-----

```
hitchcock:/home/bernd/src/tests# strace ./test_stat64 /mnt/test/yp
execve("./test_stat64", ["/test_stat64", "/mnt/test/yp"], [/* 39 vars */]) =
0
uname({sys="Linux", node="hitchcock", ...}) = 0
brk(0) = 0x572000
brk(0x593000) = 0x593000
stat("/mnt/test/yp", {st_mode=S_IFDIR|0755, st_size=2704, ...}) = 0
write(2, "stat() works fine.\n", 19stat() works fine.
) = 19
_exit(0) = ?
```

Anyone having an idea whats going on? The ethereal capture also looks pretty normal. The kernel of this system is 2.6.9, but it also happens on another system with 2.6.11-rc5.

As usual we are using unfs3 for /etc and /var, but for me that looks like a client problem. I'm even not sure if this is limited to NFS at all.

Thanks in advance,

Bernd

-

To unsubscribe from this list: send the line "unsubscribe linux-kernel" in the body of a message to majordomo@vger.kernel.org

More majordomo info at <http://vger.kernel.org/majordomo-info.html>

Please read the FAQ at <http://www.tux.org/lkml/>