

---

Subject: Re: [PATCH 0/16] Pid namespaces

Posted by [Badari Pulavarty](#) on Thu, 12 Jul 2007 18:55:59 GMT

[View Forum Message](#) <> [Reply to Message](#)

---

On Wed, 2007-07-11 at 20:19 -0700, sukadev@us.ibm.com wrote:

> Pavel Emelianov [xemul@openvz.org] wrote:

> | sukadev@us.ibm.com wrote:

> |..

> My x86\_64 system boots fine but crashes as below, when I run my

> 'pidns\_exec' test with a simple program that prints getpid(), getppid()

> etc of the process in the child pid ns.

>

> Pls see

>

> <http://www.geocities.com/sukadevb/Pidspace/2.6.22-rc6-mm1-pavel-1.tgz>

>

> for the patches I currently have applied and let me know if I need more

> on top.

>

> And see

>

> <http://www.geocities.com/sukadevb/Pidspace/Test1/>

>

> for the test programs. You may need to run the 'mypid-loop.x' script

> to repro the crash. The pidns\_exec.c program calls clone() with CLONE\_NEWPID

> and execs the given program (it was included in Patch 0 of the patchset I

> posted to Containers).

>

> Suka

>

> login: Unable to handle kernel NULL pointer dereference at 00000000000002fc RIP:

Yes. I am noticing different set of problems on my ppc64 with above tests and patchset.

Thanks,

Badari

Not cloning container for unused subsystem ns

Not cloning container for unused subsystem ns

Not cloning container for unused subsystem ns

Not cloning container for unused subsystem ns

Not cloning container for unused subsystem ns

Not cloning container for unused subsystem ns

Not cloning container for unused subsystem ns

Not cloning container for unused subsystem ns

-----[ cut here ]-----

kernel BUG at kernel/workqueue.c:258!

```

Oops: Exception in kernel mode, sig: 5 [#11]
SMP NR_CPUS=32 NUMA pSeries
Modules linked in:
NIP: c00000000062e6c LR: c00000000062e54 CTR: 0000000000000000
REGS: c00000000d157af0 TRAP: 0700 Tainted: G    D (2.6.22-rc6-mm1)
MSR: 800000000029032 <EE,ME,IR,DR> CR: 24000084 XER: 2000000f
TASK = c00000000d0f7460[33] 'events/6' THREAD: c00000000d154000 CPU: 6
GPR00: 0000000000000001 c00000000d157d70 c00000000060fed8 0000000000000001
GPR04: c0000000005c2938 0000000000000000 c00000000064438 c00000000d005600
GPR08: c000000000684f20 7fffffff0000000003506570 c0000000f204fe88
GPR12: 0000000000004000 c000000000529980 0000000000000000 c00000000045e928
GPR16: 4000000001c00000 c00000000045d340 0000000000000000 0000000000000000
GPR20: c00000000050e1d8 000000000210e1d8 000000000210e448 c00000000050e448
GPR24: 0000000001a1fb80 c00000000050c2d8 c00000000d154000 c0000000005c2968
GPR28: c0000000f204fe80 c0000000f2665008 c00000000540eb8 c0000000f2665000
NIP [c00000000062e6c] .run_workqueue+0xe0/0x208
LR [c00000000062e54] .run_workqueue+0xc8/0x208
Call Trace:
[c00000000d157d70] [c00000000062ea4] .run_workqueue+0x118/0x208 (unreliable)
[c00000000d157e10] [c000000000639bc] .worker_thread+0x114/0x138
[c00000000d157f00] [c00000000068ac8] .kthread+0x78/0xc4
[c00000000d157f90] [c000000000231b4] .kernel_thread+0x4c/0x68
Instruction dump:
980d01cc 7c2004ac 38000000 38600001 901c0000 4bfa80f9 60000000 e81dfff8
78000764 7c00e278 3120ffff 7c090110 <0b000000> 38000001 7d20f8a8 7d290078
Not cloning container for unused subsystem ns
Not cloning container for unused subsystem ns
Unable to handle kernel paging request for data at address 0x00000008
Faulting instruction address: 0xc000000000d9990
Oops: Kernel access of bad area, sig: 11 [#12]
SMP NR_CPUS=32 NUMA pSeries
Modules linked in:
NIP: c000000000d9990 LR: c000000000d9a28 CTR: c00000000048060
REGS: c00000000de9b870 TRAP: 0300 Tainted: G    D (2.6.22-rc6-mm1)
MSR: 8000000000009032 <EE,ME,IR,DR> CR: 24000428 XER: 2000000f
DAR: 0000000000000008, DSISR: 0000000042000000
TASK = c00000000d15cfa0[4871] 'pidns_exec' THREAD: c00000000de98000 CPU: 2
GPR00: c000000000d9a28 c00000000de9baf0 c00000000060fed8 0000000000000000
GPR04: c00000000065d708 c0000000100b3f00 c000000000d8eac c000000000631980
GPR08: c0000000100b3f00 0000000000000000 c000000000966158 c00000000de9bc00
GPR12: 0000000000000001 c000000000529180 00000000ffa0d094 0000000000000000
GPR16: 00000000100a0000 00000000100a92c8 00000000100a0000 0000000010080000
GPR20: 0000000000000000 0000000000000000 0000000000000000 00000000100a9568
GPR24: 00000000f7fddb60 c00000000065e2c8 c00000000de9bc00 c000000000634500
GPR28: 0000000000000000 c0000000d27a200 c00000000055ec28 c0000000d27a200
NIP [c000000000d9990] .release_mounts+0x3c/0x108
LR [c000000000d9a28] .release_mounts+0xd4/0x108
Call Trace:

```

```

[c00000000de9baf0] [c000000000d9a28] .release_mounts+0xd4/0x108 (unreliable)
[c00000000de9bb90] [c000000000d9b38] .__put_mnt_ns+0xdc/0x114
[c00000000de9bc50] [c0000000006cd74] .free_nsproxy+0x54/0xe8
[c00000000de9bce0] [c0000000005272c] .do_exit+0x948/0xa08
[c00000000de9bda0] [c000000000528c0] .sys_exit_group+0x0/0x8
[c00000000de9be30] [c0000000000872c] syscall_exit+0x0/0x40
Instruction dump:
fb61ffd8 fb81ffe0 fba1ffe8 fbe1fff8 f8010010 f821ff61 ebc2c998 7c7a1b78
480000a8 e97f0008 e93f0000 f92b0000 <f9690008> fbff0000 fbff0008 e81f0010
Not cloning container for unused subsystem ns
Not cloning container for unused subsystem ns
Not cloning container for unused subsystem ns
Not cloning container for unused subsystem ns
Unable to handle kernel paging request for data at address 0x00000020
Faulting instruction address: 0xc00000000051e88
Oops: Kernel access of bad area, sig: 11 [#13]
SMP NR_CPUS=32 NUMA pSeries
Modules linked in:
NIP: c00000000051e88 LR: c00000000051e1c CTR: 800000000013f270
REGS: c00000000de9b3a0 TRAP: 0300 Tainted: G D (2.6.22-rc6-mm1)
MSR: 8000000000009032 <EE,ME,IR,DR> CR: 28000444 XER: 00000002
DAR: 0000000000000020, DSISR: 0000000040000000
TASK = c00000000d15cfa0[4871] 'pidns_exec' THREAD: c00000000de98000 CPU: 2
GPR00: 0000000000000000 c00000000de9b620 c00000000060fed8 0000000000000000
GPR04: 0000000000000000 c00000000d15cfa0 ffffffff 0000000000000000
GPR08: 0000000000000000 0000000000000000 c00000000063e580 c00000000063e580
GPR12: 0000000000004000 c000000000529180 00000000ffa0d094 0000000000000000
GPR16: 00000000100a0000 00000000100a92c8 00000000100a0000 0000000010080000
GPR20: 0000000000000000 0000000000000000 0000000000000000 00000000100a9568
GPR24: 00000000f7fddb60 c00000000065e2c8 0000000000000000b 0000000000000000b
GPR28: c00000000045fad8 c00000000d15cfa0 c000000000540998 0000000000000001
NIP [c00000000051e88] .do_exit+0xa4/0xa08
LR [c00000000051e1c] .do_exit+0x38/0xa08
Call Trace:
[c00000000de9b620] [c00000000051e1c] .do_exit+0x38/0xa08 (unreliable)
[c00000000de9b6e0] [c000000000021478] .die+0x20c/0x210
[c00000000de9b780] [c00000000002971c] .bad_page_fault+0xb8/0xd4
[c00000000de9b800] [c000000000004e98] handle_page_fault+0x3c/0x58
--- Exception: 300 at .release_mounts+0x3c/0x108
LR = .release_mounts+0xd4/0x108
[c00000000de9bb90] [c000000000d9b38] .__put_mnt_ns+0xdc/0x114
[c00000000de9bc50] [c0000000006cd74] .free_nsproxy+0x54/0xe8
[c00000000de9bce0] [c0000000005272c] .do_exit+0x948/0xa08
[c00000000de9bda0] [c000000000528c0] .sys_exit_group+0x0/0x8
[c00000000de9be30] [c0000000000872c] syscall_exit+0x0/0x40
Instruction dump:
801d0198 2f800000 40be0010 e87e8050 4bffb5c9 60000000 e80d01a0 7fa00278
3120ffff 7c090110 0b000000 e93d04a8 <e9290020> e8090828 7fbd0000 40be0048

```

```

Not cloning container for unused subsystem ns
Not cloning container for unused subsystem ns
Not cloning container for unused subsystem ns
Not cloning container for unused subsystem ns
Not cloning container for unused subsystem ns
Not cloning container for unused subsystem ns
Unable to handle kernel paging request for data at address 0x00000020
Faulting instruction address: 0xc000000000051e88
Oops: Kernel access of bad area, sig: 11 [#14]
SMP NR_CPUS=32 NUMA pSeries
Modules linked in:
NIP: c000000000051e88 LR: c000000000051e1c CTR: 800000000013f270
REGS: c00000000de9aed0 TRAP: 0300 Tainted: G D (2.6.22-rc6-mm1)
MSR: 8000000000009032 <EE,ME,IR,DR> CR: 28000444 XER: 00000002
DAR: 0000000000000020, DSISR: 0000000040000000
TASK = c00000000d15cfa0[4871] 'pidns_exec' THREAD: c00000000de98000 CPU: 2
GPR00: 0000000000000000 c00000000de9b150 c00000000060fed8 0000000000000000
GPR04: 0000000000000000 c00000000d15cfa0 ffffffff ffffffff 0000000000000000
GPR08: 0000000000000000 0000000000000000 c00000000063e580 c00000000063e580
GPR12: 0000000000004000 c000000000529180 00000000ffa0d094 0000000000000000
GPR16: 00000000100a0000 00000000100a92c8 00000000100a0000 0000000010080000
GPR20: 0000000000000000 0000000000000000 0000000000000000 00000000100a9568
GPR24: 00000000f7fdbb60 c00000000065e2c8 0000000000000000b 0000000000000000b
GPR28: c00000000045fad8 c00000000d15cfa0 c000000000540998 00000000000000001
NIP [c000000000051e88] .do_exit+0xa4/0xa08
LR [c000000000051e1c] .do_exit+0x38/0xa08
Call Trace:
[c00000000de9b150] [c000000000051e1c] .do_exit+0x38/0xa08 (unreliable)
[c00000000de9b210] [c000000000021478] .die+0x20c/0x210
[c00000000de9b2b0] [c00000000002971c] .bad_page_fault+0xb8/0xd4
[c00000000de9b330] [c000000000004e98] handle_page_fault+0x3c/0x58
--- Exception: 300 at .do_exit+0xa4/0xa08
    LR = .do_exit+0x38/0xa08
[c00000000de9b6e0] [c000000000021478] .die+0x20c/0x210
[c00000000de9b780] [c00000000002971c] .bad_page_fault+0xb8/0xd4
[c00000000de9b800] [c000000000004e98] handle_page_fault+0x3c/0x58
--- Exception: 300 at .release_mounts+0x3c/0x108
    LR = .release_mounts+0xd4/0x108
[c00000000de9bb90] [c0000000000d9b38] .__put_mnt_ns+0xdc/0x114
[c00000000de9bc50] [c00000000006cd74] .free_nsproxy+0x54/0xe8
[c00000000de9bce0] [c00000000005272c] .do_exit+0x948/0xa08
[c00000000de9bda0] [c0000000000528c0] .sys_exit_group+0x0/0x8
[c00000000de9be30] [c00000000000872c] syscall_exit+0x0/0x40
Instruction dump:
801d0198 2f800000 40be0010 e87e8050 4bffb5c9 60000000 e80d01a0 7fa00278
3120ffff 7c090110 0b000000 e93d04a8 <e9290020> e8090828 7fbd0000 40be0048
Not cloning container for unused subsystem ns
Not cloning container for unused subsystem ns

```

```

Not cloning container for unused subsystem ns
Not cloning container for unused subsystem ns
Not cloning container for unused subsystem ns
Not cloning container for unused subsystem ns
Unable to handle kernel paging request for data at address 0x00000020
Faulting instruction address: 0xc000000000051e88
Oops: Kernel access of bad area, sig: 11 [#15]
SMP NR_CPUS=32 NUMA pSeries
Modules linked in:
NIP: c000000000051e88 LR: c000000000051e1c CTR: 800000000013f270
REGS: c00000000de9aa00 TRAP: 0300 Tainted: G      D (2.6.22-rc6-mm1)
MSR: 8000000000009032 <EE,ME,IR,DR> CR: 28000444 XER: 00000002
DAR: 0000000000000020, DSISR: 0000000040000000
TASK = c00000000d15cfa0[4871] 'pidns_exec' THREAD: c00000000de98000 CPU: 2
GPR00: 0000000000000000 c00000000de9ac80 c00000000060fed8 0000000000000000
GPR04: 0000000000000000 c00000000d15cfa0 ffffffff0000000000000000
GPR08: 0000000000000000 0000000000000000 c00000000063e580 c00000000063e580
GPR12: 0000000000004000 c000000000529180 00000000ffa0d094 0000000000000000
GPR16: 00000000100a0000 00000000100a92c8 00000000100a0000 0000000010080000
GPR20: 0000000000000000 0000000000000000 0000000000000000 00000000100a9568
GPR24: 00000000f7fdb60 c00000000065e2c8 000000000000000b 000000000000000b
GPR28: c0000000045fad8 c00000000d15cfa0 c000000000540998 0000000000000001
NIP [c000000000051e88] .do_exit+0xa4/0xa08
LR [c000000000051e1c] .do_exit+0x38/0xa08
Call Trace:
[c00000000de9ac80] [c000000000051e1c] .do_exit+0x38/0xa08 (unreliable)
[c00000000de9ad40] [c000000000021478] .die+0x20c/0x210
[c00000000de9ade0] [c00000000002971c] .bad_page_fault+0xb8/0xd4
[c00000000de9ae60] [c000000000004e98] handle_page_fault+0x3c/0x58
--- Exception: 300 at .do_exit+0xa4/0xa08
    LR = .do_exit+0x38/0xa08
[c00000000de9b210] [c000000000021478] .die+0x20c/0x210
[c00000000de9b2b0] [c00000000002971c] .bad_page_fault+0xb8/0xd4
[c00000000de9b330] [c000000000004e98] handle_page_fault+0x3c/0x58
--- Exception: 300 at .do_exit+0xa4/0xa08
    LR = .do_exit+0x38/0xa08
[c00000000de9b6e0] [c000000000021478] .die+0x20c/0x210
[c00000000de9b780] [c00000000002971c] .bad_page_fault+0xb8/0xd4
[c00000000de9b800] [c000000000004e98] handle_page_fault+0x3c/0x58
--- Exception: 300 at .release_mounts+0x3c/0x108
    LR = .release_mounts+0xd4/0x108
[c00000000de9bb90] [c000000000d9b38] .__put_mnt_ns+0xdc/0x114
[c00000000de9bc50] [c00000000006cd74] .free_nsproxy+0x54/0xe8
[c00000000de9bce0] [c00000000005272c] .do_exit+0x948/0xa08
[c00000000de9bda0] [c0000000000528c0] .sys_exit_group+0x0/0x8
[c00000000de9be30] [c00000000000872c] syscall_exit+0x0/0x40
Instruction dump:
801d0198 2f800000 40be0010 e87e8050 4bffb5c9 60000000 e80d01a0 7fa00278

```

```

3120ffff 7c090110 0b000000 e93d04a8 <e9290020> e8090828 7fbd0000 40be0048
Not cloning container for unused subsystem ns
Not cloning container for unused subsystem ns
Not cloning container for unused subsystem ns
Not cloning container for unused subsystem ns
Not cloning container for unused subsystem ns
Not cloning container for unused subsystem ns
Unable to handle kernel paging request for data at address 0x00000020
Faulting instruction address: 0xc0000000000051e88
Oops: Kernel access of bad area, sig: 11 [#16]
SMP NR_CPUS=32 NUMA pSeries
Modules linked in:
NIP: c0000000000051e88 LR: c0000000000051e1c CTR: 800000000013f270
REGS: c00000000de9a530 TRAP: 0300 Tainted: G D (2.6.22-rc6-mm1)
MSR: 8000000000009032 <EE,ME,IR,DR> CR: 28000444 XER: 00000002
DAR: 0000000000000020, DSISR: 0000000040000000
TASK = c0000000d15cfa0[4871] 'pidns_exec' THREAD: c0000000de98000 CPU: 2
GPR00: 0000000000000000 c0000000de9a7b0 c0000000060fed8 0000000000000000
GPR04: 0000000000000000 c0000000d15cfa0 ffffffff0000000000000000
GPR08: 0000000000000000 0000000000000000 c0000000063e580 c0000000063e580
GPR12: 0000000000004000 c00000000529180 00000000ffa0d094 0000000000000000
GPR16: 00000000100a0000 00000000100a92c8 00000000100a0000 0000000010080000
GPR20: 0000000000000000 0000000000000000 0000000000000000 00000000100a9568
GPR24: 00000000f7fddb60 c0000000065e2c8 0000000000000000b 0000000000000000b
GPR28: c0000000045fad8 c0000000d15cfa0 c00000000540998 0000000000000001
NIP [c0000000000051e88] .do_exit+0xa4/0xa08
LR [c0000000000051e1c] .do_exit+0x38/0xa08
Call Trace:
[c0000000de9a7b0] [c0000000000051e1c] .do_exit+0x38/0xa08 (unreliable)
[c0000000de9a870] [c0000000000021478] .die+0x20c/0x210
[c0000000de9a910] [c000000000002971c] .bad_page_fault+0xb8/0xd4
[c0000000de9a990] [c000000000004e98] handle_page_fault+0x3c/0x58
--- Exception: 300 at .do_exit+0xa4/0xa08
    LR = .do_exit+0x38/0xa08
[c0000000de9ad40] [c0000000000021478] .die+0x20c/0x210
[c0000000de9ade0] [c000000000002971c] .bad_page_fault+0xb8/0xd4
[c0000000de9ae60] [c000000000004e98] handle_page_fault+0x3c/0x58
--- Exception: 300 at .do_exit+0xa4/0xa08
    LR = .do_exit+0x38/0xa08
[c0000000de9b210] [c0000000000021478] .die+0x20c/0x210
[c0000000de9b2b0] [c000000000002971c] .bad_page_fault+0xb8/0xd4
[c0000000de9b330] [c000000000004e98] handle_page_fault+0x3c/0x58
--- Exception: 300 at .do_exit+0xa4/0xa08
    LR = .do_exit+0x38/0xa08
[c0000000de9b6e0] [c0000000000021478] .die+0x20c/0x210
[c0000000de9b780] [c000000000002971c] .bad_page_fault+0xb8/0xd4
[c0000000de9b800] [c000000000004e98] handle_page_fault+0x3c/0x58
--- Exception: 300 at .release_mounts+0x3c/0x108

```

```
LR = .release_mounts+0xd4/0x108
[c00000000de9bb90] [c0000000000d9b38] .__put_mnt_ns+0xdc/0x114
[c00000000de9bc50] [c00000000006cd74] .free_nsproxy+0x54/0xe8
[c00000000de9bce0] [c00000000005272c] .do_exit+0x948/0xa08
[c00000000de9bda0] [c0000000000528c0] .sys_exit_group+0x0/0x8
[c00000000de9be30] [c00000000000872c] syscall_exit+0x0/0x40
Instruction dump:
801d0198 2f800000 40be0010 e87e8050 4bffb5c9 60000000 e80d01a0 7fa00278
3120ffff 7c090110 0b000000 e93d04a8 <e9290020> e8090828 7fbd0000 40be0048
...
```

---

Containers mailing list  
Containers@lists.linux-foundation.org  
<https://lists.linux-foundation.org/mailman/listinfo/containers>

---