
Subject: [RFC][PATCH] Static init struct pid for swapper
Posted by [Sukadev Bhattiprolu](#) on Sat, 13 Jan 2007 21:57:00 GMT
[View Forum Message](#) <> [Reply to Message](#)

From: Sukadev Bhattiprolu <sukadev@us.ibm.com>
Subject: Statically initialize struct pid for swapper

Statically initialize a struct pid for the swapper process (pid_t == 0)
and attach it to init_task. This is needed so task_pid(), task_pgrp()
and task_session() interfaces work on the swapper process also.

Signed-off-by: Sukadev Bhattiprolu <sukadev@us.ibm.com>
Cc: Cedric Le Goater <clg@fr.ibm.com>
Cc: Dave Hansen <haveblue@us.ibm.com>
Cc: Serge Hallyn <serue@us.ibm.com>
Cc: containers@lists.osdl.org

arch/alpha/kernel/init_task.c		1 +
arch/arm/kernel/init_task.c		1 +
arch/arm26/kernel/init_task.c		1 +
arch/avr32/kernel/init_task.c		1 +
arch/frv/kernel/init_task.c		1 +
arch/h8300/kernel/init_task.c		1 +
arch/i386/kernel/init_task.c		1 +
arch/ia64/kernel/init_task.c		1 +
arch/m32r/kernel/init_task.c		1 +
arch/m68knommu/kernel/init_task.c		1 +
arch/mips/kernel/init_task.c		1 +
arch/parisc/kernel/init_task.c		1 +
arch/powerpc/kernel/init_task.c		1 +
arch/s390/kernel/init_task.c		1 +
arch/sh/kernel/init_task.c		1 +
arch/sh64/kernel/init_task.c		1 +
arch/sparc/kernel/init_task.c		1 +
arch/sparc64/kernel/init_task.c		1 +
arch/um/kernel/init_task.c		1 +
arch/v850/kernel/init_task.c		1 +
arch/x86_64/kernel/init_task.c		1 +
include/linux/init_task.h		33 ++++++

22 files changed, 54 insertions(+)

Index: lx26-20-rc2-mm1/arch/x86_64/kernel/init_task.c

=====

--- lx26-20-rc2-mm1.orig/arch/x86_64/kernel/init_task.c 2007-01-13 04:38:35.892740080 -0800
+++ lx26-20-rc2-mm1/arch/x86_64/kernel/init_task.c 2007-01-13 04:57:58.930931400 -0800
@@ -14,6 +14,7 @@ static struct fs_struct init_fs = INIT_F
static struct files_struct init_files = INIT_FILES;

```
static struct signal_struct init_signals = INIT_SIGNALS(init_signals);
static struct sighand_struct init_sighand = INIT_SIGHAND(init_sighand);
+static struct pid init_struct_pid = INIT_STRUCT_PID;
struct mm_struct init_mm = INIT_MM(init_mm);
```

```
EXPORT_SYMBOL(init_mm);
```

```
Index: lx26-20-rc2-mm1/include/linux/init_task.h
```

```
=====
--- lx26-20-rc2-mm1.orig/include/linux/init_task.h 2007-01-13 04:38:35.893739928 -0800
+++ lx26-20-rc2-mm1/include/linux/init_task.h 2007-01-13 05:25:52.178559096 -0800
@@ -88,6 +88,19 @@ extern struct nsproxy init_nsproxy;
```

```
extern struct group_info init_groups;
```

```
+#define INIT_STRUCT_PID { \
+ .count = ATOMIC_INIT(1), \
+ .nr = 0, \
+ /* Do we need to put this struct pid in pid_hash ? */ \
+ .pid_chain = { .next = NULL, .pprev = NULL }, \
+ .tasks = { \
+ { .first = &init_task.pids[PIDTYPE_PID].node }, \
+ { .first = &init_task.pids[PIDTYPE_PGID].node }, \
+ { .first = &init_task.pids[PIDTYPE_SID].node }, \
+ }, \
+ .rcu = RCU_HEAD_INIT, \
+}
+
/*
 * INIT_TASK is used to set up the first task table, touch at
 * your own risk!. Base=0, limit=0x1ffff (=2MB)
@@ -139,6 +152,26 @@ extern struct group_info init_groups;
 .cpu_timers = INIT_CPU_TIMERS(tsk.cpu_timers), \
 .fs_excl = ATOMIC_INIT(0), \
 .pi_lock = SPIN_LOCK_UNLOCKED, \
+ .pids = { \
+ { .node = { \
+ .next = NULL, \
+ .pprev = &init_struct_pid.tasks[PIDTYPE_PID].first,\
+ }, \
+ .pid = &init_struct_pid, \
+ }, \
+ { .node = { \
+ .next = NULL, \
+ .pprev = &init_struct_pid.tasks[PIDTYPE_PGID].first,\
+ }, \
+ .pid = &init_struct_pid, \
+ }, \
+ { .node = { \
```

```

+ .next = NULL, \
+ .pprev = &init_struct_pid.tasks[PIDTYPE_SID].first,\
+ }, \
+ .pid = &init_struct_pid, \
+ }, \
+ } \
  INIT_TRACE_IRQFLAGS \
  INIT_LOCKDEP \
}

```

Index: lx26-20-rc2-mm1/arch/alpha/kernel/init_task.c

```

=====
--- lx26-20-rc2-mm1.orig/arch/alpha/kernel/init_task.c 2006-11-29 13:57:37.000000000 -0800
+++ lx26-20-rc2-mm1/arch/alpha/kernel/init_task.c 2007-01-13 12:41:56.745940256 -0800
@@ -12,6 +12,7 @@ static struct fs_struct init_fs = INIT_F
static struct files_struct init_files = INIT_FILES;
static struct signal_struct init_signals = INIT_SIGNALS(init_signals);
static struct sighand_struct init_sighand = INIT_SIGHAND(init_sighand);
+static struct pid init_struct_pid = INIT_STRUCT_PID;
struct mm_struct init_mm = INIT_MM(init_mm);
struct task_struct init_task = INIT_TASK(init_task);

```

Index: lx26-20-rc2-mm1/arch/arm/kernel/init_task.c

```

=====
--- lx26-20-rc2-mm1.orig/arch/arm/kernel/init_task.c 2006-11-29 13:57:37.000000000 -0800
+++ lx26-20-rc2-mm1/arch/arm/kernel/init_task.c 2007-01-13 12:42:18.614615712 -0800
@@ -16,6 +16,7 @@ static struct fs_struct init_fs = INIT_F
static struct files_struct init_files = INIT_FILES;
static struct signal_struct init_signals = INIT_SIGNALS(init_signals);
static struct sighand_struct init_sighand = INIT_SIGHAND(init_sighand);
+static struct pid init_struct_pid = INIT_STRUCT_PID;
struct mm_struct init_mm = INIT_MM(init_mm);

```

```
EXPORT_SYMBOL(init_mm);
```

Index: lx26-20-rc2-mm1/arch/arm26/kernel/init_task.c

```

=====
--- lx26-20-rc2-mm1.orig/arch/arm26/kernel/init_task.c 2006-11-29 13:57:37.000000000 -0800
+++ lx26-20-rc2-mm1/arch/arm26/kernel/init_task.c 2007-01-13 12:42:12.116603560 -0800
@@ -19,6 +19,7 @@ static struct fs_struct init_fs = INIT_F
static struct files_struct init_files = INIT_FILES;
static struct signal_struct init_signals = INIT_SIGNALS(init_signals);
static struct sighand_struct init_sighand = INIT_SIGHAND(init_sighand);
+static struct pid init_struct_pid = INIT_STRUCT_PID;
struct mm_struct init_mm = INIT_MM(init_mm);

```

```
EXPORT_SYMBOL(init_mm);
```

Index: lx26-20-rc2-mm1/arch/avr32/kernel/init_task.c

```

=====
--- lx26-20-rc2-mm1.orig/arch/avr32/kernel/init_task.c 2006-11-29 13:57:37.000000000 -0800

```

```
+++ lx26-20-rc2-mm1/arch/avr32/kernel/init_task.c 2007-01-13 12:42:25.664543960 -0800
@@ -17,6 +17,7 @@ static struct fs_struct init_fs = INIT_F
static struct files_struct init_files = INIT_FILES;
static struct signal_struct init_signals = INIT_SIGNALS(init_signals);
static struct sighand_struct init_sighand = INIT_SIGHAND(init_sighand);
+static struct pid init_struct_pid = INIT_STRUCT_PID;
struct mm_struct init_mm = INIT_MM(init_mm);
```

```
EXPORT_SYMBOL(init_mm);
```

```
Index: lx26-20-rc2-mm1/arch/frv/kernel/init_task.c
```

```
=====
--- lx26-20-rc2-mm1.orig/arch/frv/kernel/init_task.c 2006-11-29 13:57:37.000000000 -0800
+++ lx26-20-rc2-mm1/arch/frv/kernel/init_task.c 2007-01-13 12:42:37.605728624 -0800
@@ -14,6 +14,7 @@ static struct fs_struct init_fs = INIT_F
static struct files_struct init_files = INIT_FILES;
static struct signal_struct init_signals = INIT_SIGNALS(init_signals);
static struct sighand_struct init_sighand = INIT_SIGHAND(init_sighand);
+static struct pid init_struct_pid = INIT_STRUCT_PID;
struct mm_struct init_mm = INIT_MM(init_mm);
```

```
EXPORT_SYMBOL(init_mm);
```

```
Index: lx26-20-rc2-mm1/arch/h8300/kernel/init_task.c
```

```
=====
--- lx26-20-rc2-mm1.orig/arch/h8300/kernel/init_task.c 2006-11-29 13:57:37.000000000 -0800
+++ lx26-20-rc2-mm1/arch/h8300/kernel/init_task.c 2007-01-13 12:42:42.575973032 -0800
@@ -16,6 +16,7 @@ static struct fs_struct init_fs = INIT_F
static struct files_struct init_files = INIT_FILES;
static struct signal_struct init_signals = INIT_SIGNALS(init_signals);
static struct sighand_struct init_sighand = INIT_SIGHAND(init_sighand);
+static struct pid init_struct_pid = INIT_STRUCT_PID;
struct mm_struct init_mm = INIT_MM(init_mm);
```

```
EXPORT_SYMBOL(init_mm);
```

```
Index: lx26-20-rc2-mm1/arch/i386/kernel/init_task.c
```

```
=====
--- lx26-20-rc2-mm1.orig/arch/i386/kernel/init_task.c 2006-11-29 13:57:37.000000000 -0800
+++ lx26-20-rc2-mm1/arch/i386/kernel/init_task.c 2007-01-13 12:42:52.308493464 -0800
@@ -14,6 +14,7 @@ static struct fs_struct init_fs = INIT_F
static struct files_struct init_files = INIT_FILES;
static struct signal_struct init_signals = INIT_SIGNALS(init_signals);
static struct sighand_struct init_sighand = INIT_SIGHAND(init_sighand);
+static struct pid init_struct_pid = INIT_STRUCT_PID;
struct mm_struct init_mm = INIT_MM(init_mm);
```

```
EXPORT_SYMBOL(init_mm);
```

```
Index: lx26-20-rc2-mm1/arch/ia64/kernel/init_task.c
```

```
=====
--- lx26-20-rc2-mm1.orig/arch/ia64/kernel/init_task.c 2006-11-29 13:57:37.000000000 -0800
```

```
+++ lx26-20-rc2-mm1/arch/ia64/kernel/init_task.c 2007-01-13 12:43:03.647769632 -0800
@@ -20,6 +20,7 @@ static struct fs_struct init_fs = INIT_F
static struct files_struct init_files = INIT_FILES;
static struct signal_struct init_signals = INIT_SIGNALS(init_signals);
static struct sighand_struct init_sighand = INIT_SIGHAND(init_sighand);
+static struct pid init_struct_pid = INIT_STRUCT_PID;
struct mm_struct init_mm = INIT_MM(init_mm);
```

EXPORT_SYMBOL(init_mm);

Index: lx26-20-rc2-mm1/arch/m32r/kernel/init_task.c

```
=====
--- lx26-20-rc2-mm1.orig/arch/m32r/kernel/init_task.c 2006-11-29 13:57:37.000000000 -0800
+++ lx26-20-rc2-mm1/arch/m32r/kernel/init_task.c 2007-01-13 12:43:08.006107064 -0800
@@ -15,6 +15,7 @@ static struct fs_struct init_fs = INIT_F
static struct files_struct init_files = INIT_FILES;
static struct signal_struct init_signals = INIT_SIGNALS(init_signals);
static struct sighand_struct init_sighand = INIT_SIGHAND(init_sighand);
+static struct pid init_struct_pid = INIT_STRUCT_PID;
struct mm_struct init_mm = INIT_MM(init_mm);
```

EXPORT_SYMBOL(init_mm);

Index: lx26-20-rc2-mm1/arch/m68knommu/kernel/init_task.c

```
=====
--- lx26-20-rc2-mm1.orig/arch/m68knommu/kernel/init_task.c 2006-11-29 13:57:37.000000000
-0800
+++ lx26-20-rc2-mm1/arch/m68knommu/kernel/init_task.c 2007-01-13 12:43:23.268786784 -0800
@@ -16,6 +16,7 @@ static struct fs_struct init_fs = INIT_F
static struct files_struct init_files = INIT_FILES;
static struct signal_struct init_signals = INIT_SIGNALS(init_signals);
static struct sighand_struct init_sighand = INIT_SIGHAND(init_sighand);
+static struct pid init_struct_pid = INIT_STRUCT_PID;
struct mm_struct init_mm = INIT_MM(init_mm);
```

EXPORT_SYMBOL(init_mm);

Index: lx26-20-rc2-mm1/arch/mips/kernel/init_task.c

```
=====
--- lx26-20-rc2-mm1.orig/arch/mips/kernel/init_task.c 2006-11-29 13:57:37.000000000 -0800
+++ lx26-20-rc2-mm1/arch/mips/kernel/init_task.c 2007-01-13 12:43:28.325018120 -0800
@@ -13,6 +13,7 @@ static struct fs_struct init_fs = INIT_F
static struct files_struct init_files = INIT_FILES;
static struct signal_struct init_signals = INIT_SIGNALS(init_signals);
static struct sighand_struct init_sighand = INIT_SIGHAND(init_sighand);
+static struct pid init_struct_pid = INIT_STRUCT_PID;
struct mm_struct init_mm = INIT_MM(init_mm);
```

EXPORT_SYMBOL(init_mm);

Index: lx26-20-rc2-mm1/arch/parisc/kernel/init_task.c

```
=====
```

```

--- lx26-20-rc2-mm1.orig/arch/parisc/kernel/init_task.c 2006-11-29 13:57:37.000000000 -0800
+++ lx26-20-rc2-mm1/arch/parisc/kernel/init_task.c 2007-01-13 12:43:40.204212208 -0800
@@ -37,6 +37,7 @@ static struct fs_struct init_fs = INIT_F
static struct files_struct init_files = INIT_FILES;
static struct signal_struct init_signals = INIT_SIGNALS(init_signals);
static struct sighand_struct init_sighand = INIT_SIGHAND(init_sighand);
+static struct pid init_struct_pid = INIT_STRUCT_PID;
struct mm_struct init_mm = INIT_MM(init_mm);

```

EXPORT_SYMBOL(init_mm);

Index: lx26-20-rc2-mm1/arch/powerpc/kernel/init_task.c

```

=====
--- lx26-20-rc2-mm1.orig/arch/powerpc/kernel/init_task.c 2006-11-29 13:57:37.000000000 -0800
+++ lx26-20-rc2-mm1/arch/powerpc/kernel/init_task.c 2007-01-13 12:43:44.629539456 -0800
@@ -11,6 +11,7 @@ static struct fs_struct init_fs = INIT_F
static struct files_struct init_files = INIT_FILES;
static struct signal_struct init_signals = INIT_SIGNALS(init_signals);
static struct sighand_struct init_sighand = INIT_SIGHAND(init_sighand);
+static struct pid init_struct_pid = INIT_STRUCT_PID;
struct mm_struct init_mm = INIT_MM(init_mm);

```

EXPORT_SYMBOL(init_mm);

Index: lx26-20-rc2-mm1/arch/s390/kernel/init_task.c

```

=====
--- lx26-20-rc2-mm1.orig/arch/s390/kernel/init_task.c 2006-11-29 13:57:37.000000000 -0800
+++ lx26-20-rc2-mm1/arch/s390/kernel/init_task.c 2007-01-13 12:43:54.629019304 -0800
@@ -19,6 +19,7 @@ static struct fs_struct init_fs = INIT_F
static struct files_struct init_files = INIT_FILES;
static struct signal_struct init_signals = INIT_SIGNALS(init_signals);
static struct sighand_struct init_sighand = INIT_SIGHAND(init_sighand);
+static struct pid init_struct_pid = INIT_STRUCT_PID;
struct mm_struct init_mm = INIT_MM(init_mm);

```

EXPORT_SYMBOL(init_mm);

Index: lx26-20-rc2-mm1/arch/sh/kernel/init_task.c

```

=====
--- lx26-20-rc2-mm1.orig/arch/sh/kernel/init_task.c 2006-11-29 13:57:37.000000000 -0800
+++ lx26-20-rc2-mm1/arch/sh/kernel/init_task.c 2007-01-13 12:45:33.954919472 -0800
@@ -11,6 +11,7 @@ static struct fs_struct init_fs = INIT_F
static struct files_struct init_files = INIT_FILES;
static struct signal_struct init_signals = INIT_SIGNALS(init_signals);
static struct sighand_struct init_sighand = INIT_SIGHAND(init_sighand);
+static struct pid init_struct_pid = INIT_STRUCT_PID;
struct mm_struct init_mm = INIT_MM(init_mm);

```

EXPORT_SYMBOL(init_mm);

Index: lx26-20-rc2-mm1/arch/sh64/kernel/init_task.c

```

--- lx26-20-rc2-mm1.orig/arch/sh64/kernel/init_task.c 2006-11-29 13:57:37.000000000 -0800
+++ lx26-20-rc2-mm1/arch/sh64/kernel/init_task.c 2007-01-13 12:45:26.773011288 -0800
@@ -22,6 +22,7 @@ static struct fs_struct init_fs = INIT_F
static struct files_struct init_files = INIT_FILES;
static struct signal_struct init_signals = INIT_SIGNALS(init_signals);
static struct sighand_struct init_sighand = INIT_SIGHAND(init_sighand);
+static struct pid init_struct_pid = INIT_STRUCT_PID;
struct mm_struct init_mm = INIT_MM(init_mm);

```

```

struct pt_regs fake_swapper_regs;

```

```

Index: lx26-20-rc2-mm1/arch/sparc/kernel/init_task.c

```

```

=====
--- lx26-20-rc2-mm1.orig/arch/sparc/kernel/init_task.c 2006-11-29 13:57:37.000000000 -0800
+++ lx26-20-rc2-mm1/arch/sparc/kernel/init_task.c 2007-01-13 12:45:45.048233032 -0800
@@ -11,6 +11,7 @@ static struct fs_struct init_fs = INIT_F
static struct files_struct init_files = INIT_FILES;
static struct signal_struct init_signals = INIT_SIGNALS(init_signals);
static struct sighand_struct init_sighand = INIT_SIGHAND(init_sighand);
+static struct pid init_struct_pid = INIT_STRUCT_PID;
struct mm_struct init_mm = INIT_MM(init_mm);
struct task_struct init_task = INIT_TASK(init_task);

```

```

Index: lx26-20-rc2-mm1/arch/sparc64/kernel/init_task.c

```

```

=====
--- lx26-20-rc2-mm1.orig/arch/sparc64/kernel/init_task.c 2006-11-29 13:57:37.000000000 -0800
+++ lx26-20-rc2-mm1/arch/sparc64/kernel/init_task.c 2007-01-13 12:45:39.008151264 -0800
@@ -12,6 +12,7 @@ static struct fs_struct init_fs = INIT_F
static struct files_struct init_files = INIT_FILES;
static struct signal_struct init_signals = INIT_SIGNALS(init_signals);
static struct sighand_struct init_sighand = INIT_SIGHAND(init_sighand);
+static struct pid init_struct_pid = INIT_STRUCT_PID;
struct mm_struct init_mm = INIT_MM(init_mm);

```

```

EXPORT_SYMBOL(init_mm);

```

```

Index: lx26-20-rc2-mm1/arch/um/kernel/init_task.c

```

```

=====
--- lx26-20-rc2-mm1.orig/arch/um/kernel/init_task.c 2006-11-29 13:57:37.000000000 -0800
+++ lx26-20-rc2-mm1/arch/um/kernel/init_task.c 2007-01-13 12:45:55.279677616 -0800
@@ -19,6 +19,7 @@ struct mm_struct init_mm = INIT_MM(init_
static struct files_struct init_files = INIT_FILES;
static struct signal_struct init_signals = INIT_SIGNALS(init_signals);
static struct sighand_struct init_sighand = INIT_SIGHAND(init_sighand);
+static struct pid init_struct_pid = INIT_STRUCT_PID;
EXPORT_SYMBOL(init_mm);

```

```

/*

```

```

Index: lx26-20-rc2-mm1/arch/v850/kernel/init_task.c

```

```
--- lx26-20-rc2-mm1.orig/arch/v850/kernel/init_task.c 2006-11-29 13:57:37.000000000 -0800
+++ lx26-20-rc2-mm1/arch/v850/kernel/init_task.c 2007-01-13 12:46:03.355449912 -0800
@@ -24,6 +24,7 @@ static struct fs_struct init_fs = INIT_F
static struct files_struct init_files = INIT_FILES;
static struct signal_struct init_signals = INIT_SIGNALS (init_signals);
static struct sighand_struct init_sighand = INIT_SIGHAND (init_sighand);
+static struct pid init_struct_pid = INIT_STRUCT_PID;
struct mm_struct init_mm = INIT_MM (init_mm);

EXPORT_SYMBOL(init_mm);
```

Containers mailing list
Containers@lists.osdl.org
<https://lists.osdl.org/mailman/listinfo/containers>
