
Subject: [PATCH -mm] remove CONFIG_UTS_NS and CONFIG_IPC_NS
Posted by [Cedric Le Goater](#) on Fri, 08 Jun 2007 11:48:38 GMT

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

CONFIG_UTS_NS and CONFIG_IPC_NS have very little value as they only deactivate the unshare of the uts and ipc namespaces and do not improve performance.

Signed-off-by: Cedric Le Goater <clg@fr.ibm.com>
Acked-by: "Serge E. Hallyn" <serue@us.ibm.com>
Cc: Eric W. Biederman <ebiederm@xmission.com>
Cc: Herbert Poetzl <herbert@13thfloor.at>
Cc: Pavel Emelianov <xemul@openvz.org>

include/linux/ipc.h | 11 +++-----
include/linux/utsname.h | 13 -----
init/Kconfig | 17 -----
ipc/msg.c | 4 +---
ipc/sem.c | 4 +---
ipc/shm.c | 4 +---
ipc/util.c | 11 +-----
ipc/util.h | 8 ++-----
kernel/Makefile | 4 ++--
kernel/nsproxy.c | 10 -----
kernel/utsname_sysctl.c | 5 +----
11 files changed, 12 insertions(+), 79 deletions(-)

Index: 2.6.22-rc4-mm2/include/linux/utsname.h

=====

--- 2.6.22-rc4-mm2.orig/include/linux/utsname.h

+++ 2.6.22-rc4-mm2/include/linux/utsname.h

@@ -48,7 +48,6 @@ static inline void get_uts_ns(struct uts
kref_get(&ns->kref);
}

-#ifdef CONFIG_UTS_NS

extern struct uts_namespace *copy_utsname(int flags, struct uts_namespace *ns);
extern void free_uts_ns(struct kref *kref);

@@ -56,18 +55,6 @@ static inline void put_uts_ns(struct uts

{
kref_put(&ns->kref, free_uts_ns);
}

-#else

-static inline struct uts_namespace *copy_utsname(int flags,
- struct uts_namespace *ns)

-{
- return ns;

```

-}
-
-static inline void put_uts_ns(struct uts_namespace *ns)
-{
-}
-#endif
-
static inline struct new_utsname *utsname(void)
{
    return &current->nsproxy->uts_ns->name;
}
Index: 2.6.22-rc4-mm2/init/Kconfig
=====
--- 2.6.22-rc4-mm2.orig/init/Kconfig
+++ 2.6.22-rc4-mm2/init/Kconfig
@@ -142,15 +142,6 @@ config SYSVIPC
     section 6.4 of the Linux Programmer's Guide, available from
     <http://www.tldp.org/guides.html>.

-config IPC_NS
-bool "IPC Namespaces"
-depends on SYSVIPC
-default n
-help
-    Support ipc namespaces. This allows containers, i.e. virtual
-    environments, to use ipc namespaces to provide different ipc
-    objects for different servers. If unsure, say N.
-
config SYSVIPC_SYSCTL
bool
depends on SYSVIPC
@@ -240,14 +231,6 @@ config TASK_IO_ACCOUNTING

    Say N if unsure.

-config UTS_NS
-bool "UTS Namespaces"
-default n
-help
-    Support uts namespaces. This allows containers, i.e.
-    vservers, to use uts namespaces to provide different
-    uts info for different servers. If unsure, say N.
-
config AUDIT
bool "Auditing support"
depends on NET
Index: 2.6.22-rc4-mm2/kernel/Makefile
=====
--- 2.6.22-rc4-mm2.orig/kernel/Makefile

```

```

+++ 2.6.22-rc4-mm2/kernel/Makefile
@@ -8,7 +8,8 @@ obj-y    = sched.o fork.o exec_domain.o
    signal.o sys.o kmod.o workqueue.o pid.o \
    rcupdate.o extable.o params.o posix-timers.o \
    kthread.o wait.o kfifo.o sys_ni.o posix-cpu-timers.o mutex.o \
-   hrtimer.o rwsem.o latency.o nsproxy.o srcu.o die_notifier.o
+   hrtimer.o rwsem.o latency.o nsproxy.o srcu.o die_notifier.o \
+   utsname.o

obj-$(CONFIG_STACKTRACE) += stacktrace.o
obj-y += time/
@@ -53,7 +54,6 @@ obj-$(CONFIG_DEBUG_SYNCHRO_TEST) += sync
obj-$(CONFIG_RCU_TORTURE_TEST) += rcutorture.o
obj-$(CONFIG_RELAY) += relay.o
obj-$(CONFIG_SYSCTL) += utsname_sysctl.o
-obj-$(CONFIG_UTS_NS) += utsname.o
obj-$(CONFIG_TASK_DELAY_ACCT) += delayacct.o
obj-$(CONFIG_TASKSTATS) += taskstats.o tsacct.o

```

Index: 2.6.22-rc4-mm2/include/linux/ipc.h

```

=====
--- 2.6.22-rc4-mm2.orig/include/linux/ipc.h
+++ 2.6.22-rc4-mm2/include/linux/ipc.h
@@ -93,6 +93,7 @@ extern struct ipc_namespace init_ipc_ns;

#ifdef CONFIG_SYSVIPC
#define INIT_IPC_NS(ns) .ns = &init_ipc_ns,
+extern void free_ipc_ns(struct kref *kref);
extern struct ipc_namespace *copy_ipcs(unsigned long flags,
    struct ipc_namespace *ns);
#else
@@ -104,13 +105,9 @@ static inline struct ipc_namespace *copy
}
#endif

-#ifndef CONFIG_IPC_NS
-extern void free_ipc_ns(struct kref *kref);
-#endif
-
static inline struct ipc_namespace *get_ipc_ns(struct ipc_namespace *ns)
{
-#ifndef CONFIG_IPC_NS
+﻿#ifndef CONFIG_SYSVIPC
    if (ns)
        kref_get(&ns->kref);
    #endif
@@ -119,7 +116,7 @@ static inline struct ipc_namespace *get_

```

```

static inline void put_ipc_ns(struct ipc_namespace *ns)
{
-#ifdef CONFIG_IPC_NS
+#ifdef CONFIG_SYSVIPC
    kref_put(&ns->kref, free_ipc_ns);
#endif
}
@@ -127,5 +124,3 @@ static inline void put_ipc_ns(struct ipc
#endif /* __KERNEL__ */

#endif /* _LINUX_IPC_H */

```

Index: 2.6.22-rc4-mm2/ipc/msg.c

=====

```

--- 2.6.22-rc4-mm2.orig/ipc/msg.c
+++ 2.6.22-rc4-mm2/ipc/msg.c
@@ -87,7 +87,7 @@ static int newque (struct ipc_namespace
static int sysvipc_msg_proc_show(struct seq_file *s, void *it);
#endif

```

```

-static void __ipc_init __msg_init_ns(struct ipc_namespace *ns, struct ipc_ids *ids)
+static void __msg_init_ns(struct ipc_namespace *ns, struct ipc_ids *ids)
{
    ns->ids[IPC_MSG_IDS] = ids;
    ns->msg_ctlmax = MSGMAX;
@@ -96,7 +96,6 @@ static void __ipc_init __msg_init_ns(str
    ipc_init_ids(ids, ns->msg_ctlmni);
}

```

```

-#ifdef CONFIG_IPC_NS
int msg_init_ns(struct ipc_namespace *ns)
{
    struct ipc_ids *ids;
@@ -128,7 +127,6 @@ void msg_exit_ns(struct ipc_namespace *n
    kfree(ns->ids[IPC_MSG_IDS]);
    ns->ids[IPC_MSG_IDS] = NULL;
}
-#endif

```

```

void __init msg_init(void)
{

```

Index: 2.6.22-rc4-mm2/ipc/sem.c

=====

```

--- 2.6.22-rc4-mm2.orig/ipc/sem.c
+++ 2.6.22-rc4-mm2/ipc/sem.c
@@ -121,7 +121,7 @@ static int sysvipc_sem_proc_show(struct
#define sc_semopm sem_ctls[2]

```

```

#define sc_semmni sem_ctls[3]

-static void __ipc_init __sem_init_ns(struct ipc_namespace *ns, struct ipc_ids *ids)
+static void __sem_init_ns(struct ipc_namespace *ns, struct ipc_ids *ids)
{
    ns->ids[IPC_SEM_IDS] = ids;
    ns->sc_semmns = SEMMSL;
@@ -132,7 +132,6 @@ static void __ipc_init __sem_init_ns(str
    ipc_init_ids(ids, ns->sc_semmni);
}

-#ifdef CONFIG_IPC_NS
int sem_init_ns(struct ipc_namespace *ns)
{
    struct ipc_ids *ids;
@@ -164,7 +163,6 @@ void sem_exit_ns(struct ipc_namespace *n
    kfree(ns->ids[IPC_SEM_IDS]);
    ns->ids[IPC_SEM_IDS] = NULL;
}
-#endif

void __init sem_init (void)
{
Index: 2.6.22-rc4-mm2/ipc/shm.c
=====
--- 2.6.22-rc4-mm2.orig/ipc/shm.c
+++ 2.6.22-rc4-mm2/ipc/shm.c
@@ -77,7 +77,7 @@ static void shm_destroy (struct ipc_name
static int sysvipc_shm_proc_show(struct seq_file *s, void *it);
-#endif

-static void __ipc_init __shm_init_ns(struct ipc_namespace *ns, struct ipc_ids *ids)
+static void __shm_init_ns(struct ipc_namespace *ns, struct ipc_ids *ids)
{
    ns->ids[IPC_SHM_IDS] = ids;
    ns->shm_ctlmax = SHMMAX;
@@ -98,7 +98,6 @@ static void do_shm_rmid(struct ipc_names
    shm_destroy(ns, shp);
}

-#ifdef CONFIG_IPC_NS
int shm_init_ns(struct ipc_namespace *ns)
{
    struct ipc_ids *ids;
@@ -130,7 +129,6 @@ void shm_exit_ns(struct ipc_namespace *n
    kfree(ns->ids[IPC_SHM_IDS]);
    ns->ids[IPC_SHM_IDS] = NULL;
}

```

```
-#endif
```

```
void __init shm_init (void)
```

```
{
```

```
Index: 2.6.22-rc4-mm2/ipc/util.c
```

```
=====
```

```
--- 2.6.22-rc4-mm2.orig/ipc/util.c
```

```
+++ 2.6.22-rc4-mm2/ipc/util.c
```

```
@@ -50,7 +50,6 @@ struct ipc_namespace init_ipc_ns = {
```

```
},
```

```
};
```

```
-#ifdef CONFIG_IPC_NS
```

```
static struct ipc_namespace *clone_ipc_ns(struct ipc_namespace *old_ns)
```

```
{
```

```
int err;
```

```
@@ -110,14 +109,6 @@ void free_ipc_ns(struct kref *kref)
```

```
shm_exit_ns(ns);
```

```
kfree(ns);
```

```
}
```

```
-#else
```

```
-struct ipc_namespace *copy_ipcs(unsigned long flags, struct ipc_namespace *ns)
```

```
-{
```

```
- if (flags & CLONE_NEWIPC)
```

```
- return ERR_PTR(-EINVAL);
```

```
- return ns;
```

```
-}
```

```
-#endif
```

```
/**
```

```
 * ipc_init - initialise IPC subsystem
```

```
@@ -145,7 +136,7 @@ __initcall(ipc_init);
```

```
 * array itself.
```

```
*/
```

```
-void __ipc_init ipc_init_ids(struct ipc_ids* ids, int size)
```

```
+void ipc_init_ids(struct ipc_ids* ids, int size)
```

```
{
```

```
int i;
```

```
Index: 2.6.22-rc4-mm2/ipc/util.h
```

```
=====
```

```
--- 2.6.22-rc4-mm2.orig/ipc/util.h
```

```
+++ 2.6.22-rc4-mm2/ipc/util.h
```

```
@@ -41,12 +41,8 @@ struct ipc_ids {
```

```
};
```

```
struct seq_file;
```

```

-#ifdef CONFIG_IPC_NS
-#define __ipc_init
-#else
-#define __ipc_init __init
-#endif
-void __ipc_init ipc_init_ids(struct ipc_ids *ids, int size);
+
+void ipc_init_ids(struct ipc_ids *ids, int size);
#ifdef CONFIG_PROC_FS
void __init ipc_init_proc_interface(const char *path, const char *header,
    int ids, int (*show)(struct seq_file *, void *));
Index: 2.6.22-rc4-mm2/kernel/utsname_sysctl.c
=====
--- 2.6.22-rc4-mm2.orig/kernel/utsname_sysctl.c
+++ 2.6.22-rc4-mm2/kernel/utsname_sysctl.c
@@ -18,10 +18,7 @@
static void *get_uts(ctl_table *table, int write)
{
    char *which = table->data;
-#ifdef CONFIG_UTS_NS
- struct uts_namespace *uts_ns = current->nsproxy->uts_ns;
- which = (which - (char *)&init_uts_ns) + (char *)uts_ns;
-#endif
+
+ if (!write)
+     down_read(&uts_sem);
+ else
Index: 2.6.22-rc4-mm2/kernel/nsproxy.c
=====
--- 2.6.22-rc4-mm2.orig/kernel/nsproxy.c
+++ 2.6.22-rc4-mm2/kernel/nsproxy.c
@@ -157,16 +157,6 @@ int unshare_nsproxy_namespaces(unsigned
    if (!(unshare_flags & (CLONE_NEWNS | CLONE_NEWUTS | CLONE_NEWIPC)))
        return 0;

-#ifndef CONFIG_IPC_NS
- if (unshare_flags & CLONE_NEWIPC)
-     return -EINVAL;
-#endif
-
-#ifndef CONFIG_UTS_NS
- if (unshare_flags & CLONE_NEWUTS)
-     return -EINVAL;
-#endif
-
    if (!capable(CAP_SYS_ADMIN))
        return -EPERM;

```

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