

Re: [PATCH] initramfs: Allow rootfs to use tmpfs instead of ramfs

Source: <http://linux.derkeiler.com/Mailing-Lists/Kernel/2006-07/msg08827.html>

- *From:* Al Boldi <a1426z@xxxxxxxxxx>
 - *Date:* Mon, 31 Jul 2006 00:03:36 +0300
-

Hugh Dickins wrote:

On Sun, 30 Jul 2006, Al Boldi wrote:

Replugs rootfs to use tmpfs instead of ramfs as a Kconfig option.

Why? Looking further down we see what you should have explained here:

- + This option switches rootfs so that it uses tmpfs rather than ramfs
- + for its file storage. This makes rootfs swappable so having a large
- + initrd or initramfs image won't eat up valuable RAM.

Now, I'm far from an expert on initramfs and early userspace, but my understanding is that the "init" of a (properly designed) initramfs would pretty much "rm -rf" everything in the initramfs before passing control to the final "init". So (almost?) no valuable RAM is eaten up, nor the less valuable swap if you do extend this to tmpfs (unless something gets left open, which I think should not be the case).

So I'm inclined to say that this patch is simply unnecessary. But if people who know better think it's a good thing, I've no objection (though I've not tried it): the Kconfigery looks more likely to provoke argument than the tmpfs/ramfs mods.

This patch is based on John Zielinski's
<http://marc.theaimsgroup.com/?l=linux-kernel&m=107013630212011&w=4>
patch.

Modified for 2.6.17.

RunTime tested on i386.

This trivial patch should go into 2.6.18.

Re: [PATCH] initramfs: Allow rootfs to use tmpfs instead of ramfs

Well...

Cc: <stable@xxxxxxxxxx>
Cc: Chris Wright <chrisw@xxxxxxxxxx>

Umm, the -stable tree doesn't usually take enhancements, even if you think it's a very stable enhancement.

Cc: Andrew Morton <akpm@xxxxxxxx>
Cc: John Zielinski <grim@xxxxxxxx>
Signed-off-by: Al Boldi <a1426z@xxxxxxxx>

--- a/drivers/block/Kconfig 2006-07-30 16:44:59.000000000 +0300
+++ b/drivers/block/Kconfig 2006-07-30 16:45:53.000000000 +0300
@@ -412,6 +412,22 @@ config BLK_DEV_INITRD
If RAM disk support (BLK_DEV_RAM) is also included, this also enables initial RAM disk (initrd) support.

+config SHM_ROOTFS

Please say TMPFS_ROOTFS rather than SHM_ROOTFS throughout.

Ok.

+ bool "Use tmpfs (shm) instead of ramfs for rootfs"
+ depends on BLK_DEV_INITRD
+ default n
+ select TMPFS
+ select SHMEM

Selects can be troublesome, but you may well have decided right.

+ help
+ This option switches rootfs so that it uses tmpfs rather than ramfs
+ for it's file storage. This makes rootfs swappable so having a
large

its

Re: [PATCH] initramfs: Allow rootfs to use tmpfs instead of ramfs

Yes.

```
+ initrd or initramfs image won't eat up valuable RAM.
+
+config RAMFS_ROOTFS
+ bool
+ depends on !SHM_ROOTFS
+ default y
+ select RAMFS

config CDROM_PKTCDVD
tristate "Packet writing on CD/DVD media"
--- a/fs/Kconfig 2006-07-30 16:45:25.000000000 +0300
+++ b/fs/Kconfig 2006-07-30 16:46:36.000000000 +0300
@@ -853,7 +853,7 @@ config HUGETLB_PAGE
def_bool HUGETLBFS

config RAMFS
- bool
+ tristate "Ramfs file system support"
default y
---help---
Ramfs is a file system which keeps all files in RAM. It allows
--- a/fs/ramfs/inode.c 2006-07-30 16:45:37.000000000 +0300
+++ b/fs/ramfs/inode.c 2006-07-30 16:46:36.000000000 +0300
@@ -191,22 +191,27 @@ struct super_block *ramfs_get_sb(struct
return get_sb_nODEV(fs_type, flags, data, ramfs_fill_super);
}

+#ifdef CONFIG_RAMFS_ROOTFS
static struct super_block *rootfs_get_sb(struct file_system_type
*fs_type, int flags, const char *dev_name, void *data)
{
return get_sb_nODEV(fs_type, flags|MS_NOUSER, data, ramfs_fill_super);
}
+#endif

static struct file_system_type ramfs_fs_type = {
.name = "ramfs",
.get_sb = ramfs_get_sb,
.kill_sb = kill_litter_super,
};
+
+#ifdef CONFIG_RAMFS_ROOTFS
static struct file_system_type rootfs_fs_type = {
.name = "rootfs",
.get_sb = rootfs_get_sb,
.kill_sb = kill_litter_super,
};
```

Re: [PATCH] initramfs: Allow rootfs to use tmpfs instead of ramfs

```
+#endif

static int __init init_ramfs_fs(void)
{
@@ -221,9 +226,11 @@ static void __exit exit_ramfs_fs(void)
module_init(init_ramfs_fs)
module_exit(exit_ramfs_fs)

+#ifdef CONFIG_RAMFS_ROOTFS
int __init init_rootfs(void)
{
return register_filesystem(&rootfs_fs_type);
}
+#endif

MODULE_LICENSE("GPL");
--- a/mm/shmem.c 2006-07-30 16:45:54.000000000 +0300
+++ b/mm/shmem.c 2006-07-30 16:47:11.000000000 +0300
@@ -2123,7 +2123,7 @@ failed:
return err;
}

-static struct kmem_cache *shmem_inode_cachep;
+static struct kmem_cache *shmem_inode_cachep = NULL;
```

No need to initialize static variables to 0 (or is someone going to raise the spectre of NULL being non-0 on some future architecture?)

Don't know.

```
static struct inode *shmem_alloc_inode(struct super_block *sb)
{
@@ -2156,11 +2156,13 @@ static void init_once(void *foo, struct

static int init_inodecache(void)
{
- shmem_inode_cachep = kmem_cache_create("shmem_inode_cache",
- sizeof(struct shmem_inode_info),
- 0, 0, init_once, NULL);
- if (shmem_inode_cachep == NULL)
- return -ENOMEM;
+ if (!shmem_inode_cachep) {

if (shmem_inode_cachep == NULL) {
is more consistent with the nearby style
```

Re: [PATCH] initramfs: Allow rootfs to use tmpfs instead of ramfs

Re: [PATCH] initramfs: Allow rootfs to use tmpfs instead of ramfs

Ok.

```
+ shmем_inode_cache = kmem_cache_create("shmем_inode_cache",
+ sizeof(struct shmем_inode_info),
+ 0, 0, init_once, NULL);
+ if (shmем_inode_cache == NULL)
+ return -ENOMEM;
+ }
return 0;
}
```

```
@@ -2345,6 +2347,27 @@ put_memory:
return ERR_PTR(error);
}
```

```
+#ifdef CONFIG_SHM_ROOTFS
+static struct super_block *rootfs_get_sb(struct file_system_type
+ *fs_type, + int flags, const char *dev_name, void *data)
+{
+ return get_sb_single(fs_type, flags, data, shmем_fill_super);
+}
+
+static struct file_system_type rootfs_fs_type = {
+ .name = "rootfs",
+ .get_sb = rootfs_get_sb,
+ .kill_sb = kill_litter_super,
+};
+
+int __init init_rootfs(void)
+{
+ if (init_inodecache())
+ panic("Can't initialize shm inode cache");
```

"tmpfs inode cache"
or "shmем inode cache"

Yes, shmем.

```
+ return register_filesystem(&rootfs_fs_type);
+}
+#endif
+
+/*
+ * shmем_zero_setup – setup a shared anonymous mapping
+ */
```

Re: [PATCH] initramfs: Allow rootfs to use tmpfs instead of ramfs

Thanks for your review!

--

Al

-

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

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

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