Subject: Re: [PATCH] Make common helpers for seq_files that work with list_head-s
Posted by Andrew Morton on Thu, 17 May 2007 17:36:55 GMT
View Forum Message <> Reply to Message

On Thu, 17 May 2007 19:19:40 +0400 Pavel Emelianov <xemul@sw.ru> wrote:

> Many places in kernel use seq_file API to iterate over a regular
> list_head. The code for such iteration is identical in all the
> places, so it's worth introducing a common helpers.
>
> This makes code more than 300 lines smaller.
>
> Cc-ed are people, who maintain the code altered by the patch.
>
> Signed-off-by: Pavel Emelianov <xemul@openvz.org>
>
> block/genhd.c | 40 +++-----
> crypto/proc.c | 17 ---
> drivers/char/misc.c | 18 ----
> drivers/input/input.c | 29 ------
> drivers/isdn/capi/kcapi_proc.c | 28 ------
> fs/afs/proc.c | 81 ++----------------
> fs/namespace.c | 14 ---
> fs/nfs/client.c | 54 +----------
> fs/proc/proc_tty.c | 15 ---
> include/linux/seq_file.h | 35 +++++++
> kernel/module.c | 17 ---
> mm/slab.c | 28 +-----
> net/atm/br2684.c | 22 ----
> net/core/sock.c | 39 --------
> net/ipv4/netfilter/nf_conntrack_l3proto_ipv4_compat.c | 27 ------
> net/netfilter/nf_conntrack_expect.c | 27 ------
> net/rxrpc/ar-proc.c | 48 +---------
> 17 files changed, 116 insertions(+), 423 deletions(-)

Can't complain about the diffstat. Please experiment with uninlining
seq_list_start(), see if that reduces overall text size.

And as Dmitry indicated, it would be less disruptive if we could have the
one core patch then a stream of per-subsystem patches, please.