
Subject: [PATCH COMMIT] diff-merge-2.6.15.5-20060413

Posted by [xemul](#) on Thu, 13 Apr 2006 11:42:38 GMT

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

Added to 026test008

Patch from OpenVZ team <devel@openvz.org>
Merge /linux/kernel/git/stable/linux-2.6.16.y

```
diff --git a/arch/i386/kernel/cpu/cpufreq/Kconfig b/arch/i386/kernel/cpu/cpufreq/Kconfig
index 26892d2..16f2e35 100644
```

```
--- a/arch/i386/kernel/cpu/cpufreq/Kconfig
+++ b/arch/i386/kernel/cpu/cpufreq/Kconfig
@@ -203,6 +203,7 @@ config X86_LONGRUN
config X86_LONGHAUL
tristate "VIA Cyrix III Longhaul"
select CPU_FREQ_TABLE
+ depends on BROKEN
help
```

This adds the CPUFreq driver for VIA Samuel/CyrixIII,
VIA Cyrix Samuel/C3, VIA Cyrix Ezra and VIA Cyrix Ezra-T

```
diff --git a/arch/i386/kernel/cpu/cpufreq/p4-clockmod.c
b/arch/i386/kernel/cpu/cpufreq/p4-clockmod.c
```

```
index cc73a7a..ebe1848 100644
```

```
--- a/arch/i386/kernel/cpu/cpufreq/p4-clockmod.c
+++ b/arch/i386/kernel/cpu/cpufreq/p4-clockmod.c
@@ -244,7 +244,7 @@ static int cpufreq_p4_cpu_init(struct cp
for (i=1; (p4clockmod_table[i].frequency != CPUFREQ_TABLE_END); i++) {
    if ((i<2) && (has_N44_O17_errata[policy->cpu]))
        p4clockmod_table[i].frequency = CPUFREQ_ENTRY_INVALID;
-   else if (has_N60_errata[policy->cpu] && p4clockmod_table[i].frequency < 2000000)
+   else if (has_N60_errata[policy->cpu] && ((stock_freq * i)/8) < 2000000)
        p4clockmod_table[i].frequency = CPUFREQ_ENTRY_INVALID;
    else
        p4clockmod_table[i].frequency = (stock_freq * i)/8;
```

```
diff --git a/arch/i386/kernel/cpu/cpufreq/speedstep-smi.c
b/arch/i386/kernel/cpu/cpufreq/speedstep-smi.c
```

```
index 28cc5d5..fcf4276 100644
```

```
--- a/arch/i386/kernel/cpu/cpufreq/speedstep-smi.c
+++ b/arch/i386/kernel/cpu/cpufreq/speedstep-smi.c
@@ -75,7 +75,9 @@ static int speedstep_smi_ownership (void
 __asm__ __volatile__(
 "out %%al, (%%dx)\n"
 : "=D" (result)
- : "a" (command), "b" (function), "c" (0), "d" (smi_port), "D" (0), "S" (magic)
+ : "a" (command), "b" (function), "c" (0), "d" (smi_port),
+ "D" (0), "S" (magic)
```

```

+ : "memory"
);

dprintk("result is %x\n", result);
diff --git a/arch/i386/kernel/dmi_scan.c b/arch/i386/kernel/dmi_scan.c
index 6a93d75..ca2a0cb 100644
--- a/arch/i386/kernel/dmi_scan.c
+++ b/arch/i386/kernel/dmi_scan.c
@@ -106,7 +106,7 @@ static void __init dmi_save_devices(stru
 struct dmi_device *dev;

 for (i = 0; i < count; i++) {
- char *d = ((char *) dm) + (i * 2);
+ char *d = (char *)(dm + 1) + (i * 2);

 /* Skip disabled device */
 if ((*d & 0x80) == 0)
diff --git a/arch/powerpc/kernel/pci_64.c b/arch/powerpc/kernel/pci_64.c
index ba92bab..4c4449b 100644
--- a/arch/powerpc/kernel/pci_64.c
+++ b/arch/powerpc/kernel/pci_64.c
@@ -78,6 +78,7 @@ int global_phb_number; /* Global phb co

 /* Cached ISA bridge dev. */
 struct pci_dev *ppc64_isabridge_dev = NULL;
+EXPORT_SYMBOL_GPL(ppc64_isabridge_dev);

 static void fixup_broken_pcnet32(struct pci_dev* dev)
{
diff --git a/arch/x86_64/kernel/entry.S b/arch/x86_64/kernel/entry.S
index e9c9b8a..a15161e 100644
--- a/arch/x86_64/kernel/entry.S
+++ b/arch/x86_64/kernel/entry.S
@@ -180,6 +180,10 @@ rff_trace:
*
* XXX if we had a free scratch register we could save the RSP into the stack frame
*      and report it properly in ps. Unfortunately we haven't.
+*
+* When user can change the frames always force IRET. That is because
+* it deals with uncanonical addresses better. SYSRET has trouble
+* with them due to bugs in both AMD and Intel CPUs.
*/
ENTRY(system_call)
@@ -254,7 +258,10 @@ sysret_signal:
 xorl %esi,%esi # oldset -> arg2
 call ptregscall_common
 1: movl $_TIF_NEED_RESCHED,%edi

```

```
- jmp sysret_check
+ /* Use IRET because user could have changed frame. This
+   works because ptregscall_common has called FIXUP_TOP_OF_STACK. */
+ cli
+ jmp int_with_check
```

badsys:

```
    movq $-ENOSYS,RAX-ARGOFFSET(%rsp)
@@ -280,7 +287,8 @@ tracesys:
    call syscall_trace_leave
    RESTORE_TOP_OF_STACK %rbx
    RESTORE_REST
- jmp ret_from_sys_call
+ /* Use IRET because user could have changed frame */
+ jmp int_ret_from_sys_call
    CFI_ENDPROC
```

```
/*
@@ -408,25 +416,9 @@ ENTRY(stub_execve)
    CFI_ADJUST_CFA_OFFSET -8
```

```
    CFI_REGISTER rip, r11
    SAVE_REST
- movq %r11, %r15
- CFI_REGISTER rip, r15
    FIXUP_TOP_OF_STACK %r11
    call sys_execve
- GET_THREAD_INFO(%rcx)
- bt $TIF_IA32,threadinfo_flags(%rcx)
- CFI_REMEMBER_STATE
- jc exec_32bit
    RESTORE_TOP_OF_STACK %r11
- movq %r15, %r11
- CFI_REGISTER rip, r11
- RESTORE_REST
- pushq %r11
- CFI_ADJUST_CFA_OFFSET 8
- CFI_REL_OFFSET rip, 0
- ret
-
```

```
-exec_32bit:
- CFI_RESTORE_STATE
    movq %rax,RAX(%rsp)
    RESTORE_REST
```

```
    jmp int_ret_from_sys_call
diff --git a/drivers/base/cpu.c b/drivers/base/cpu.c
index 07a7f97..29f3d75 100644
--- a/drivers/base/cpu.c
+++ b/drivers/base/cpu.c
```

```

@@ -141,7 +141,7 @@ int __devinit register_cpu(struct cpu *c
    return error;
}

-struct sys_device *get_cpu_sysdev(int cpu)
+struct sys_device *get_cpu_sysdev(unsigned cpu)
{
    if (cpu < NR_CPUS)
        return cpu_sys_devices[cpu];
diff --git a/drivers/base/firmware_class.c b/drivers/base/firmware_class.c
index e97e911..4723182 100644
--- a/drivers/base/firmware_class.c
+++ b/drivers/base/firmware_class.c
@@ -211,18 +211,20 @@ static int
fw_realloc_buffer(struct firmware_priv *fw_priv, int min_size)
{
    u8 *new_data;
+ int new_size = fw_priv->alloc_size;

    if (min_size <= fw_priv->alloc_size)
        return 0;

- new_data = vmalloc(fw_priv->alloc_size + PAGE_SIZE);
+ new_size = ALIGN(min_size, PAGE_SIZE);
+ new_data = vmalloc(new_size);
    if (!new_data) {
        printk(KERN_ERR "%s: unable to alloc buffer\n", __FUNCTION__);
        /* Make sure that we don't keep incomplete data */
        fw_load_abort(fw_priv);
        return -ENOMEM;
    }
- fw_priv->alloc_size += PAGE_SIZE;
+ fw_priv->alloc_size = new_size;
    if (fw_priv->fw->data) {
        memcpy(new_data, fw_priv->fw->data, fw_priv->fw->size);
        vfree(fw_priv->fw->data);
diff --git a/drivers/base/node.c b/drivers/base/node.c
index 16c513a..c80c3ae 100644
--- a/drivers/base/node.c
+++ b/drivers/base/node.c
@@ -106,7 +106,7 @@ static ssize_t node_read_numastat(struct
    other_node = 0;
    for (i = 0; i < MAX_NR_ZONES; i++) {
        struct zone *z = &pg->node_zones[i];
- for (cpu = 0; cpu < NR_CPUS; cpu++) {
+ for_each_online_cpu(cpu) {
        struct per_cpu_pageset *ps = zone_pcp(z,cpu);
        numa_hit += ps->numa_hit;

```

```

    numa_miss += ps->numa_miss;
diff --git a/drivers/block/cciss.c b/drivers/block/cciss.c
index 0d65394..c149d57 100644
--- a/drivers/block/cciss.c
+++ b/drivers/block/cciss.c
@@ -3269,8 +3269,8 @@ clean2:
 unregister_blkdev(hba[i]->major, hba[i]->devname);
clean1:
 release_io_mem(hba[i]);
- free_hba(i);
 hba[i]->busy_initializing = 0;
+ free_hba(i);
 return(-1);
}

diff --git a/drivers/char/Kconfig b/drivers/char/Kconfig
index 05ba410..8b72a61 100644
--- a/drivers/char/Kconfig
+++ b/drivers/char/Kconfig
@@ -187,6 +187,7 @@ config MOXA_SMARTIO
config ISI
 tristate "Multi-Tech multiport card support (EXPERIMENTAL)"
 depends on SERIAL_NONSTANDARD
+ select FW_LOADER
help
This is a driver for the Multi-Tech cards which provide several
serial ports. The driver is experimental and can currently only be
diff --git a/drivers/char/tlclk.c b/drivers/char/tlclk.c
index 4c27218..2546637 100644
--- a/drivers/char/tlclk.c
+++ b/drivers/char/tlclk.c
@@ -767,6 +767,7 @@ static int __init tlclk_init(void)
 printk(KERN_ERR "tlclk: can't get major %d.\n", tlclk_major);
 return ret;
}
+ tlclk_major = ret;
alarm_events = kzalloc( sizeof(struct tlclk_alarms), GFP_KERNEL);
if (!alarm_events)
 goto out1;
diff --git a/drivers/ieee1394/sbp2.c b/drivers/ieee1394/sbp2.c
index eca92eb..d83248e 100644
--- a/drivers/ieee1394/sbp2.c
+++ b/drivers/ieee1394/sbp2.c
@@ -495,22 +495,17 @@ static struct sbp2_command_info *sbp2uti
/*
 * This function finds the sbp2_command for a given outstanding SCpnt.
 * Only looks at the inuse list.
+ * Must be called with scsi_id->sbp2_command_orb_lock held.

```

```

*/
-static struct sbp2_command_info *sbp2util_find_command_for_SCpnt(struct
scsi_id_instance_data *scsi_id, void *SCpnt)
+static struct sbp2_command_info *sbp2util_find_command_for_SCpnt(
+ struct scsi_id_instance_data *scsi_id, void *SCpnt)
{
    struct sbp2_command_info *command;
- unsigned long flags;

- spin_lock_irqsave(&scsi_id->sbp2_command_orb_lock, flags);
- if (!list_empty(&scsi_id->sbp2_command_orb_inuse)) {
- list_for_each_entry(command, &scsi_id->sbp2_command_orb_inuse, list) {
- if (command->Current_SCpnt == SCpnt) {
-     spin_unlock_irqrestore(&scsi_id->sbp2_command_orb_lock, flags);
+ if (!list_empty(&scsi_id->sbp2_command_orb_inuse))
+ list_for_each_entry(command, &scsi_id->sbp2_command_orb_inuse, list)
+ if (command->Current_SCpnt == SCpnt)
    return command;
- }
- }
- }
- spin_unlock_irqrestore(&scsi_id->sbp2_command_orb_lock, flags);
    return NULL;
}

```

@@ -579,17 +574,15 @@ static void sbp2util_free_command_dma(st

```

/*
 * This function moves a command to the completed orb list.
+ * Must be called with scsi_id->sbp2_command_orb_lock held.
 */
-static void sbp2util_mark_command_completed(struct scsi_id_instance_data *scsi_id,
- struct sbp2_command_info *command)
+static void sbp2util_mark_command_completed(
+ struct scsi_id_instance_data *scsi_id,
+ struct sbp2_command_info *command)
{
- unsigned long flags;
-
- spin_lock_irqsave(&scsi_id->sbp2_command_orb_lock, flags);
    list_del(&command->list);
    sbp2util_free_command_dma(command);
    list_add_tail(&command->list, &scsi_id->sbp2_command_orb_completed);
- spin_unlock_irqrestore(&scsi_id->sbp2_command_orb_lock, flags);
}

/*
@@ -2177,7 +2170,9 @@ static int sbp2_handle_status_write(stru
```

```

 * Matched status with command, now grab scsi command pointers and check status
 */
SCpnt = command->Current_SCpnt;
+ spin_lock_irqsave(&scsi_id->sbp2_command_orb_lock, flags);
    sbp2util_mark_command_completed(scsi_id, command);
+ spin_unlock_irqrestore(&scsi_id->sbp2_command_orb_lock, flags);

if (SCpnt) {

@@ -2513,6 +2508,7 @@ static int sbp2scsi_abort(struct scsi_cm
 (struct scsi_id_instance_data *)SCpnt->device->host->hostdata[0];
 struct sbp2scsi_host_info *hi = scsi_id->hi;
 struct sbp2_command_info *command;
+ unsigned long flags;

SBP2_ERR("aborting sbp2 command");
scsi_print_command(SCpnt);
@@ -2523,6 +2519,7 @@ static int sbp2scsi_abort(struct scsi_cm
 * Right now, just return any matching command structures
 * to the free pool.
 */
+ spin_lock_irqsave(&scsi_id->sbp2_command_orb_lock, flags);
    command = sbp2util_find_command_for_SCpnt(scsi_id, SCpnt);
    if (command) {
        SBP2_DEBUG("Found command to abort");
@@ -2540,6 +2537,7 @@ static int sbp2scsi_abort(struct scsi_cm
     command->Current_done(command->Current_SCpnt);
    }
}
+ spin_unlock_irqrestore(&scsi_id->sbp2_command_orb_lock, flags);

/*
 * Initiate a fetch agent reset.
diff --git a/drivers/md/dm.c b/drivers/md/dm.c
index 745ca1f..d559569 100644
--- a/drivers/md/dm.c
+++ b/drivers/md/dm.c
@@ -533,30 +533,35 @@ static void __clone_and_map(struct clone

} else {
/*
- * Create two copy bios to deal with io that has
- * been split across a target.
+ * Handle a bvec that must be split between two or more targets.
*/
    struct bio_vec *bv = bio->bi_io_vec + ci->idx;
+ sector_t remaining = to_sector(bv->bv_len);
+ unsigned int offset = 0;

```

```

- clone = split_bvec(bio, ci->sector, ci->idx,
-         bv->bv_offset, max);
- __map_bio(ti, clone, tio);
-
- ci->sector += max;
- ci->sector_count -= max;
- ti = dm_table_find_target(ci->map, ci->sector);
-
- len = to_sector(bv->bv_len) - max;
- clone = split_bvec(bio, ci->sector, ci->idx,
-         bv->bv_offset + to_bytes(max), len);
- tio = alloc_tio(ci->md);
- tio->io = ci->io;
- tio->ti = ti;
- memset(&tio->info, 0, sizeof(tio->info));
- __map_bio(ti, clone, tio);
+ do {
+     if (offset) {
+         ti = dm_table_find_target(ci->map, ci->sector);
+         max = max_io_len(ci->md, ci->sector, ti);
+
+         tio = alloc_tio(ci->md);
+         tio->io = ci->io;
+         tio->ti = ti;
+         memset(&tio->info, 0, sizeof(tio->info));
+     }
+
+     len = min(remaining, max);
+
+     clone = split_bvec(bio, ci->sector, ci->idx,
+             bv->bv_offset + offset, len);
+
+     __map_bio(ti, clone, tio);
+
+     ci->sector += len;
+     ci->sector_count -= len;
+     offset += to_bytes(len);
+ } while (remaining -= len);

- ci->sector += len;
- ci->sector_count -= len;
- ci->idx++;
}
}

diff --git a/drivers/media/video/Kconfig b/drivers/media/video/Kconfig
index d82c8a3..ef42a26 100644
--- a/drivers/media/video/Kconfig

```

```

+++ b/drivers/media/video/Kconfig
@@ -349,6 +349,7 @@ config VIDEO_AUDIO_DECODER
config VIDEO_DECODER
tristate "Add support for additional video chipsets"
depends on VIDEO_DEV && I2C && EXPERIMENTAL
+ select FW_LOADER
---help---
Say Y here to compile drivers for SAA7115, SAA7127 and CX25840
video decoders.
diff --git a/drivers/media/video/tuner-types.c b/drivers/media/video/tuner-types.c
index 6fe7817..5f3d46d 100644
--- a/drivers/media/video/tuner-types.c
+++ b/drivers/media/video/tuner-types.c
@@ -1087,8 +1087,8 @@ static struct tuner_params tuner_tnf_533
/* ----- TUNER_SAMSUNG_TCPN_2121P30A - Samsung NTSC ----- */
static struct tuner_range tuner_samsung_tcpn_2121p30a_ntsc_ranges[] = {
- { 16 * 175.75 /*MHz*/, 0x01, },
- { 16 * 410.25 /*MHz*/, 0x02, },
+ { 16 * 130.00 /*MHz*/, 0x01, },
+ { 16 * 364.50 /*MHz*/, 0x02, },
{ 16 * 999.99 , 0x08, },
};

diff --git a/drivers/net/irda/irda-usb.c b/drivers/net/irda/irda-usb.c
index 8936058..6e2ec56 100644
--- a/drivers/net/irda/irda-usb.c
+++ b/drivers/net/irda/irda-usb.c
@@ -740,7 +740,7 @@ static void irda_usb_receive(struct urb
    struct sk_buff *newskb;
    struct sk_buff *dataskb;
    struct urb *next_urb;
- int docopy;
+ unsigned int len, docopy;

    IRDA_DEBUG(2, "%s(), len=%d\n", __FUNCTION__, urb->actual_length);

@@ -851,10 +851,11 @@ static void irda_usb_receive(struct urb
    dataskb->dev = self->netdev;
    dataskb->mac.raw = dataskb->data;
    dataskb->protocol = htons(ETH_P_IRDA);
+ len = dataskb->len;
    netif_rx(dataskb);

    /* Keep stats up to date */
- self->stats.rx_bytes += dataskb->len;
+ self->stats.rx_bytes += len;
    self->stats.rx_packets++;

```

```

self->netdev->last_rx = jiffies;

diff --git a/drivers/net/wireless/Kconfig b/drivers/net/wireless/Kconfig
index ef85d76..8101657 100644
--- a/drivers/net/wireless/Kconfig
+++ b/drivers/net/wireless/Kconfig
@@ -239,7 +239,8 @@ config IPW2200_DEBUG

config AIRO
    tristate "Cisco/Aironet 34X/35X/4500/4800 ISA and PCI cards"
- depends on NET_RADIO && ISA_DMA_API && CRYPTO && (PCI || BROKEN)
+ depends on NET_RADIO && ISA_DMA_API && (PCI || BROKEN)
+ select CRYPTO
---help---
    This is the standard Linux driver to support Cisco/Aironet ISA and
    PCI 802.11 wireless cards.
@@ -374,6 +375,7 @@ config PCMCIA_HERMES
config PCMCIA_SPECTRUM
    tristate "Symbol Spectrum24 Trilogy PCMCIA card support"
    depends on NET_RADIO && PCMCIA && HERMES
+ select FW_LOADER
---help---
    This is a driver for 802.11b cards using RAM-loadable Symbol
@@ -387,6 +389,7 @@ config PCMCIA_SPECTRUM
config AIRO_CS
    tristate "Cisco/Aironet 34X/35X/4500/4800 PCMCIA cards"
    depends on NET_RADIO && PCMCIA && (BROKEN || !M32R)
+ select CRYPTO
---help---
    This is the standard Linux driver to support Cisco/Aironet PCMCIA
    802.11 wireless cards. This driver is the same as the Aironet
diff --git a/drivers/net/wireless/hostap/hostap_80211_tx.c
b/drivers/net/wireless/hostap/hostap_80211_tx.c
index 4a85e63..5f398bd 100644
--- a/drivers/net/wireless/hostap/hostap_80211_tx.c
+++ b/drivers/net/wireless/hostap/hostap_80211_tx.c
@@ -469,7 +469,7 @@ int hostap_master_start_xmit(struct sk_b
}

if (local->ieee_802_1x && meta->ethertype == ETH_P_PAE && tx.crypt &&
- !(fc & IEEE80211_FCTL_VERS)) {
+ !(fc & IEEE80211_FCTL_PROTECTED)) {
    no_encrypt = 1;
    PDEBUG(DEBUG_EXTRA2, "%s: TX: IEEE 802.1X - passing "
           "unencrypted EAPOL frame\n", dev->name);
diff --git a/drivers/net/wireless/ipw2200.c b/drivers/net/wireless/ipw2200.c
index 287676a..aa6f3a4 100644

```

```

--- a/drivers/net/wireless/ipw2200.c
+++ b/drivers/net/wireless/ipw2200.c
@@ -9956,9 +9956,8 @@ static int ipw_ethtool_set_eeprom(struct
    return -EINVAL;
    down(&p->sem);
    memcpy(&p->eeprom[eeprom->offset], bytes, eeprom->len);
-   for (i = IPW EEPROM DATA;
-        i < IPW EEPROM DATA + IPW EEPROM IMAGE_SIZE; i++)
-       ipw_write8(p, i, p->eeprom[i]);
+   for (i = 0; i < IPW EEPROM IMAGE_SIZE; i++)
+       ipw_write8(p, i + IPW EEPROM DATA, p->eeprom[i]);
    up(&p->sem);
    return 0;
}
diff --git a/drivers/pcmcia/ds.c b/drivers/pcmcia/ds.c
index bb96ce1..a4333a8 100644
--- a/drivers/pcmcia/ds.c
+++ b/drivers/pcmcia/ds.c
@@ -546,7 +546,7 @@ static int pcmcia_device_query(struct pc
    tmp = vers1->str + vers1->ofs[i];

    length = strlen(tmp) + 1;
-   if ((length < 3) || (length > 255))
+   if ((length < 2) || (length > 255))
        continue;

    p_dev->prod_id[i] = kmalloc(sizeof(char) * length,
diff --git a/drivers/scsi/sata_mv.c b/drivers/scsi/sata_mv.c
index 2770005..b00af08 100644
--- a/drivers/scsi/sata_mv.c
+++ b/drivers/scsi/sata_mv.c
@@ -1102,6 +1102,7 @@ static u8 mv_get_crbp_status(struct ata_
    void __iomem *port_mmio = mv_ap_base(ap);
    struct mv_port_priv *pp = ap->private_data;
    u32 out_ptr;
+   u8 ata_status;

    out_ptr = readl(port_mmio + EDMA_RSP_Q_OUT_PTR_OF);

```

@ @ -1109,6 +1110,8 @ @ static u8 mv_get_crbp_status(struct ata_

```

    assert(((out_ptr >> EDMA_RSP_Q_PTR_SHIFT) & MV_MAX_Q_DEPTH_MASK) ==
           pp->rsp_consumer);

+   ata_status = pp->crpb[pp->rsp_consumer].flags >> CRPB_FLAG_STATUS_SHIFT;
+
/* increment our consumer index... */
    pp->rsp_consumer = mv_inc_q_index(&pp->rsp_consumer);

```

```

@@ -1123,7 +1126,7 @@ static u8 mv_get_crpb_status(struct ata_
 writelfl(out_ptr, port_mmio + EDMA_RSP_Q_OUT_PTR_OFS);

 /* Return ATA status register for completed CRPB */
- return (pp->crpb[pp->rsp_consumer].flags >> CRPB_FLAG_STATUS_SHIFT);
+ return ata_status;
}

/***
@@ -1192,7 +1195,6 @@ static void mv_host_intr(struct ata_host
 u32 hc_irq_cause;
 int shift, port, port0, hard_port, handled;
 unsigned int err_mask;
- u8 ata_status = 0;

 if (hc == 0) {
    port0 = 0;
@@ -1210,6 +1212,7 @@ static void mv_host_intr(struct ata_host
 hc,relevant,hc_irq_cause);

 for (port = port0; port < port0 + MV_PORTS_PER_HC; port++) {
+ u8 ata_status = 0;
    ap = host_set->ports[port];
    hard_port = port & MV_PORT_MASK; /* range 0-3 */
    handled = 0; /* ensure ata_status is set if handled++ */
diff --git a/drivers/usb/core/message.c b/drivers/usb/core/message.c
index 7135e54..96cabeb 100644
--- a/drivers/usb/core/message.c
+++ b/drivers/usb/core/message.c
@@ -1388,11 +1388,13 @@ free_interfaces:
 if (dev->state != USB_STATE_ADDRESS)
    usb_disable_device (dev, 1); // Skip ep0

- i = dev->bus_mA - cp->desc.bMaxPower * 2;
- if (i < 0)
- dev_warn(&dev->dev, "new config #%"d exceeds power "
- "limit by %dmA\n",
- configuration, -i);
+ if (cp) {
+ i = dev->bus_mA - cp->desc.bMaxPower * 2;
+ if (i < 0)
+ dev_warn(&dev->dev, "new config #%"d exceeds power "
+ "limit by %dmA\n",
+ configuration, -i);
+ }

if ((ret = usb_control_msg(dev, usb_sndctrlpipe(dev, 0),
    USB_REQ_SET_CONFIGURATION, 0, configuration, 0,

```

```

diff --git a/drivers/usb/host/ehci-sched.c b/drivers/usb/host/ehci-sched.c
index ebcca97..88419c6 100644
--- a/drivers/usb/host/ehci-sched.c
+++ b/drivers/usb/host/ehci-sched.c
@@ -707,6 +707,7 @@ iso_stream_init (
} else {
    u32 addr;
    int think_time;
+   int hs_transfers;

    addr = dev->ttport << 24;
    if (!ehci_is_TDI(ehci)
@@ -719,6 +720,7 @@ iso_stream_init (
        think_time = dev->tt ? dev->tt->think_time : 0;
        stream->tt_usecs = NS_TO_US (think_time + usb_calc_bus_time (
            dev->speed, is_input, 1, maxp));
+   hs_transfers = max (1u, (maxp + 187) / 188);
    if (is_input) {
        u32 tmp;

@@ -727,12 +729,11 @@ iso_stream_init (
        stream->usecs = HS_USECS_ISO (1);
        stream->raw_mask = 1;

-   /* pessimistic c-mask */
-   tmp = usb_calc_bus_time (USB_SPEED_FULL, 1, 0, maxp)
-   / (125 * 1000);
-   stream->raw_mask |= 3 << (tmp + 9);
+   /* c-mask as specified in USB 2.0 11.18.4 3.c */
+   tmp = (1 << (hs_transfers + 2)) - 1;
+   stream->raw_mask |= tmp << (8 + 2);
} else
-   stream->raw_mask = smask_out [maxp / 188];
+   stream->raw_mask = smask_out [hs_transfers - 1];
    bandwidth = stream->usecs + stream->c_usecs;
    bandwidth /= 1 << (interval + 2);

```

```

diff --git a/drivers/video/cfbimgblt.c b/drivers/video/cfbimgblt.c
index 910e233..8ba6152 100644
--- a/drivers/video/cfbimgblt.c
+++ b/drivers/video/cfbimgblt.c
@@ -169,7 +169,7 @@ static inline void slow_imageblit(const

    while (j--) {
        l--;
-       color = (*s & 1 << (FB_BIT_NR(l))) ? fgcolor : bgcolor;
+       color = (*s & (1 << l)) ? fgcolor : bgcolor;
        val |= FB_SHIFT_HIGH(color, shift);

```

```

/* Did the bitshift spill bits to the next long? */
diff --git a/drivers/video/i810/i810_main.c b/drivers/video/i810/i810_main.c
index d8467c0..788297e 100644
--- a/drivers/video/i810/i810_main.c
+++ b/drivers/video/i810/i810_main.c
@@ -1508,7 +1508,7 @@ static int i810fb_cursor(struct fb_info
    int size = ((cursor->image.width + 7) >> 3) *
        cursor->image.height;
    int i;
- u8 *data = kmalloc(64 * 8, GFP_KERNEL);
+ u8 *data = kmalloc(64 * 8, GFP_ATOMIC);

    if (data == NULL)
        return -ENOMEM;
diff --git a/fs/9p/vfs_inode.c b/fs/9p/vfs_inode.c
index 3ad8455..651a9e1 100644
--- a/fs/9p/vfs_inode.c
+++ b/fs/9p/vfs_inode.c
@@ -614,6 +614,7 @@ static struct dentry *v9fs_vfs_lookup(st

    sb = dir->i_sb;
    v9ses = v9fs_inode2v9ses(dir);
+   dentry->d_op = &v9fs_dentry_operations;
    dirfid = v9fs_fid_lookup(dentry->d_parent);

    if (!dirfid) {
@@ -681,8 +682,6 @@ static struct dentry *v9fs_vfs_lookup(st
        goto FreeFcall;

    fid->qid = fcall->params.rstat.stat.qid;
-
-   dentry->d_op = &v9fs_dentry_operations;
    v9fs_stat2inode(&fcall->params.rstat.stat, inode, inode->i_sb);

    d_add(dentry, inode);
diff --git a/fs/nfsd/nfs3proc.c b/fs/nfsd/nfs3proc.c
index 6d2dfed..f61142a 100644
--- a/fs/nfsd/nfs3proc.c
+++ b/fs/nfsd/nfs3proc.c
@@ -682,7 +682,7 @@ static struct svc_procedure nfsd_proced
    PROC(lookup, dirop, dirop, fhandle2, RC_NOCACHE, ST+FH+pAT+pAT),
    PROC(access, access, access, fhandle, RC_NOCACHE, ST+pAT+1),
    PROC(readlink, readlink, readlink, fhandle, RC_NOCACHE,
ST+pAT+1+NFS3_MAXPATHLEN/4),
-   PROC(read, read, read, fhandle, RC_NOCACHE, ST+pAT+4+NFSSVC_MAXBLKSIZE),
+   PROC(read, read, read, fhandle, RC_NOCACHE, ST+pAT+4+NFSSVC_MAXBLKSIZE/4),
    PROC(write, write, write, fhandle, RC_REPLBUFF, ST+WC+4),

```

```

PROC(create, create, create, fhandle2, RC_REPLBUFF, ST+(1+FH+pAT)+WC),
PROC(mkdir, mkdir, create, fhandle2, RC_REPLBUFF, ST+(1+FH+pAT)+WC),
diff --git a/fs/nfsd/nfs4proc.c b/fs/nfsd/nfs4proc.c
index 6d63f1d..ca8a4c4 100644
--- a/fs/nfsd/nfs4proc.c
+++ b/fs/nfsd/nfs4proc.c
@@ -975,7 +975,7 @@ struct nfsd4_voidargs { int dummy; };
 */
static struct svc_procedure nfsd_procedures4[2] = {
    PROC(null, void, void, void, RC_NOCACHE, 1),
-   PROC(compound, compound, compound, compound, RC_NOCACHE, NFSD_BUFSIZE)
+   PROC(compound, compound, compound, compound, RC_NOCACHE, NFSD_BUFSIZE/4)
};

struct svc_version nfsd_version4 = {
diff --git a/fs/nfsd/nfsproc.c b/fs/nfsd/nfsproc.c
index 3e6b75c..06cd0db 100644
--- a/fs/nfsd/nfsproc.c
+++ b/fs/nfsd/nfsproc.c
@@ -553,7 +553,7 @@ static struct svc_procedure nfsd_proced
    PROC(None, void, void, none, RC_NOCACHE, ST),
    PROC(lookup, diropargs, diropres, fhandle, RC_NOCACHE, ST+FH+AT),
    PROC(readlink, readlinkargs, readlinkres, none, RC_NOCACHE,
ST+1+NFS_MAXPATHLEN/4),
-   PROC(read, readargs, readres, fhandle, RC_NOCACHE, ST+AT+1+NFSSVC_MAXBLKSIZE),
+   PROC(read, readargs, readres, fhandle, RC_NOCACHE,
ST+AT+1+NFSSVC_MAXBLKSIZE/4),
    PROC(None, void, void, none, RC_NOCACHE, ST),
    PROC(write, writeargs, attrstat, fhandle, RC_REPLBUFF, ST+AT),
    PROC(create, createargs, diropres, fhandle, RC_REPLBUFF, ST+FH+AT),
diff --git a/fs/proc/proc_misc.c b/fs/proc/proc_misc.c
index bedd830..2040b52 100644
--- a/fs/proc/proc_misc.c
+++ b/fs/proc/proc_misc.c
@@ -351,7 +351,7 @@ static void *devinfo_next(struct seq_fil
    case BLK_HDR:
        info->state = BLK_LIST;
        (*pos)++;
-    break;
+ /*fallthrough*/
    case BLK_LIST:
        if (get_blkdev_info(info->blkdev,&idummy,&ndummy)) {
            /*
diff --git a/fs/proc/vmcore.c b/fs/proc/vmcore.c
index 4063fb3..164a7d0 100644
--- a/fs/proc/vmcore.c
+++ b/fs/proc/vmcore.c
@@ -103,8 +103,8 @@ static ssize_t read_vmcore(struct file *

```

```

    size_t buflen, loff_t *fpos)
{
    ssize_t acc = 0, tmp;
- size_t tsz, nr_bytes;
- u64 start;
+ size_t tsz;
+ u64 start, nr_bytes;
    struct vmcore *curr_m = NULL;

    if (buflen == 0 || *fpos >= vmcore_size)
diff --git a/fs/sysfs/dir.c b/fs/sysfs/dir.c
index ddc6717..cf94bda 100644
--- a/fs/sysfs/dir.c
+++ b/fs/sysfs/dir.c
@@ -309,6 +309,7 @@ void sysfs_remove_dir(struct kobject * k
    * Drop reference from dget() on entrance.
    */
    dput(dentry);
+ kobj->dentry = NULL;
}

int sysfs_rename_dir(struct kobject * kobj, const char *new_name)
diff --git a/fs/sysfs/file.c b/fs/sysfs/file.c
index 0af020f..a8df0fb 100644
--- a/fs/sysfs/file.c
+++ b/fs/sysfs/file.c
@@ -183,7 +183,7 @@ fill_write_buffer(struct sysfs_buffer *
    return -ENOMEM;

    if (count >= PAGE_SIZE)
- count = PAGE_SIZE;
+ count = PAGE_SIZE - 1;
    error = copy_from_user(buffer->page,buf,count);
    buffer->needs_read_fill = 1;
    return error ? -EFAULT : count;
diff --git a/fs/sysfs/inode.c b/fs/sysfs/inode.c
index d384e91..8018161 100644
--- a/fs/sysfs/inode.c
+++ b/fs/sysfs/inode.c
@@ -226,12 +226,16 @@ void sysfs_drop_dentry(struct sysfs_dire
void sysfs_hash_and_remove(struct dentry * dir, const char * name)
{
    struct sysfs_dirent * sd;
- struct sysfs_dirent * parent_sd = dir->d_fsd;
+ struct sysfs_dirent * parent_sd;
+
+ if (!dir)
+     return;

```

```

if (dir->d_inode == NULL)
/* no inode means this hasn't been made visible yet */
return;

+ parent_sd = dir->d_fsdata;
mutex_lock(&dir->d_inode->i_mutex);
list_for_each_entry(sd, &parent_sd->s_children, s_sibling) {
    if (!sd->s_element)
diff --git a/fs/sysfs/symlink.c b/fs/sysfs/symlink.c
index 85ec936..bd75039 100644
--- a/fs/sysfs/symlink.c
+++ b/fs/sysfs/symlink.c
@@ -66,6 +66,7 @@ static int sysfs_add_link(struct dentry
    if (!error)
        return 0;

+ kobject_put(target);
    kfree(sl->link_name);
exit2:
    kfree(sl);
diff --git a/fs/xfs/linux-2.6/xfs_aops.c b/fs/xfs/linux-2.6/xfs_aops.c
index 74d8be8..a980736 100644
--- a/fs/xfs/linux-2.6/xfs_aops.c
+++ b/fs/xfs/linux-2.6/xfs_aops.c
@@ -616,7 +616,7 @@ xfs_is_delayed_page(
    acceptable = (type == IOMAP_UNWRITTEN);
    else if (buffer_delay(bh))
        acceptable = (type == IOMAP_DELAY);
-   else if (buffer_mapped(bh))
+   else if (buffer_dirty(bh) && buffer_mapped(bh))
        acceptable = (type == 0);
    else
        break;
diff --git a/include/asm-powerpc/floppy.h b/include/asm-powerpc/floppy.h
index e258778..608164c 100644
--- a/include/asm-powerpc/floppy.h
+++ b/include/asm-powerpc/floppy.h
@@ -35,6 +35,7 @@ 
#endif CONFIG_PCI

#include <linux/pci.h>
+#include <asm/ppc-pci.h> /* for ppc64_isabridge_dev */

#define fd_dma_setup(addr,size,mode,io) powerpc_fd_dma_setup(addr,size,mode,io)

@@ -52,12 +53,12 @@ static __inline__ int powerpc_fd_dma_set
    if (bus_addr

```

```

    && (addr != prev_addr || size != prev_size || dir != prev_dir) {
    /* different from last time -- unmap prev */
- pci_unmap_single(NULL, bus_addr, prev_size, prev_dir);
+ pci_unmap_single(ppc64_isabridge_dev, bus_addr, prev_size, prev_dir);
    bus_addr = 0;
}

if (!bus_addr) /* need to map it */
- bus_addr = pci_map_single(NULL, addr, size, dir);
+ bus_addr = pci_map_single(ppc64_isabridge_dev, addr, size, dir);

/* remember this one as prev */
prev_addr = addr;
diff --git a/include/linux/cpu.h b/include/linux/cpu.h
index 0ed1d48..d612b89 100644
--- a/include/linux/cpu.h
+++ b/include/linux/cpu.h
@@ -32,7 +32,7 @@ struct cpu {
};

extern int register_cpu(struct cpu *, int, struct node *);
-extern struct sys_device *get_cpu_sysdev(int cpu);
+extern struct sys_device *get_cpu_sysdev(unsigned cpu);
#ifndef CONFIG_HOTPLUG_CPU
extern void unregister_cpu(struct cpu *, struct node *);
#endif
diff --git a/include/linux/fb.h b/include/linux/fb.h
index 2cb19e6..2fdd8ae 100644
--- a/include/linux/fb.h
+++ b/include/linux/fb.h
@@ -839,12 +839,10 @@ struct fb_info {
#define FB_LEFT_POS(bpp)      (32 - bpp)
#define FB_SHIFT_HIGH(val, bits) ((val) >> (bits))
#define FB_SHIFT_LOW(val, bits) ((val) << (bits))
-#define FB_BIT_NR(b)         (7 - (b))
#else
#define FB_LEFT_POS(bpp)      (0)
#define FB_SHIFT_HIGH(val, bits) ((val) << (bits))
#define FB_SHIFT_LOW(val, bits) ((val) >> (bits))
-#define FB_BIT_NR(b)         (b)
#endif

/*
diff --git a/include/linux/proc_fs.h b/include/linux/proc_fs.h
index 4e48793..092f894 100644
--- a/include/linux/proc_fs.h
+++ b/include/linux/proc_fs.h
@@ -78,7 +78,7 @@ struct kcore_list {

```

```

struct vmcore {
    struct list_head list;
    unsigned long long paddr;
-   unsigned long size;
+   unsigned long long size;
    loff_t offset;
};

diff --git a/include/linux/raid/raid1.h b/include/linux/raid/raid1.h
index 9d5494a..3009c81 100644
--- a/include/linux/raid/raid1.h
+++ b/include/linux/raid/raid1.h
@@ -130,6 +130,6 @@ struct r1bio_s {
    * with failure when last write completes (and all failed).
    * Record that bi_end_io was called with this flag...
 */
#define R1BIO_Returned 4
+#define R1BIO_Returned 6

#endif
diff --git a/include/linux/rtc.h b/include/linux/rtc.h
index 0b2ba67..b739ac1 100644
--- a/include/linux/rtc.h
+++ b/include/linux/rtc.h
@@ -11,8 +11,6 @@ 
#ifndef __LINUX_RTC_H__
#define __LINUX_RTC_H__

#include <linux/interrupt.h>
-
/*
 * The struct used to pass data via the following ioctl. Similar to the
 * struct tm in <time.h>, but it needs to be here so that the kernel
@@ -95,6 +93,8 @@ struct rtc_pll_info {

#endif /* __KERNEL__ */

diff --git a/kernel/exec_domain.c b/kernel/exec_domain.c
index 867d6db..c01cead 100644
--- a/kernel/exec_domain.c
+++ b/kernel/exec_domain.c
@@ -140,6 +140,7 @@ __set_personality(u_long personality)
    ep = lookup_exec_domain(personality);

```

```

if (ep == current_thread_info()->exec_domain) {
    current->personality = personality;
+ module_put(ep->module);
    return 0;
}

diff --git a/kernel/sched.c b/kernel/sched.c
index 9c561e2..fc46033 100644
--- a/kernel/sched.c
+++ b/kernel/sched.c
@@ -240,6 +240,7 @@ struct runqueue {

    task_t *migration_thread;
    struct list_head migration_queue;
+ int cpu;
#endif

#ifndef CONFIG_SCHEDSTATS
@@ -1965,6 +1966,9 @@ EXPORT_SYMBOL(nr_iowait_ve);
/*
 * double_rq_lock - safely lock two runqueues
 *
+ * We must take them in cpu order to match code in
+ * dependent_sleeper and wake_dependent_sleeper.
+ *
 * Note this does not disable interrupts like task_rq_lock,
 * you need to do so manually before calling.
 */
@@ -1976,7 +1980,7 @@ static void double_rq_lock(runqueue_t *rq1, runqueue_t *rq2)
    spin_lock(&rq1->lock);
    __acquire(rq2->lock); /* Fake it out ; */
} else {
- if (rq1 < rq2) {
+ if (rq1->cpu < rq2->cpu) {
    spin_lock(&rq1->lock);
    spin_lock(&rq2->lock);
} else {
@@ -2012,7 +2016,7 @@ static void double_lock_balance(runqueue_t *busiest, runqueue_t *this_rq)
    __acquires(this_rq->lock)
{
    if (unlikely(!spin_trylock(&busiest->lock))) {
- if (busiest < this_rq) {
+ if (busiest->cpu < this_rq->cpu) {
    spin_unlock(&this_rq->lock);
    spin_lock(&busiest->lock);
    spin_lock(&this_rq->lock);
@@ -6391,6 +6395,7 @@ void __init sched_init(void)
    rq->push_cpu = 0;
}

```

```

rq->migration_thread = NULL;
INIT_LIST_HEAD(&rq->migration_queue);
+ rq->cpu = i;
#endif
atomic_set(&rq->nr_iowait, 0);

diff --git a/kernel/signal.c b/kernel/signal.c
index 86d9a29..d58a823 100644
--- a/kernel/signal.c
+++ b/kernel/signal.c
@@ -992,7 +992,6 @@ __group_complete_signal(int sig, struct
    if (t == NULL)
        /* restart balancing at this thread */
        t = p->signal->curr_target = p;
- BUG_ON(t->tgid != p->tgid);

    while (!wants_signal(sig, t)) {
        t = next_thread(t);
diff --git a/net/core/sock.c b/net/core/sock.c
index 489d1f0..18839a2 100644
--- a/net/core/sock.c
+++ b/net/core/sock.c
@@ -408,8 +408,9 @@ set_rcvbuf:
    if (!valbool) {
        sk->sk_bound_dev_if = 0;
    } else {
-    if (optlen > IFNAMSIZ)
-        optlen = IFNAMSIZ;
+    if (optlen > IFNAMSIZ - 1)
+        optlen = IFNAMSIZ - 1;
+    memset(devname, 0, sizeof(devname));
        if (copy_from_user(devname, optval, optlen)) {
            ret = -EFAULT;
            break;
    }
diff --git a/net/ipv4/fib_trie.c b/net/ipv4/fib_trie.c
index e320b32..24009be 100644
--- a/net/ipv4/fib_trie.c
+++ b/net/ipv4/fib_trie.c
@@ -314,11 +314,6 @@ static void __leaf_free_rcu(struct rcu_head
    kfree(container_of(head, struct leaf, rcu));
}

-static inline void free_leaf(struct leaf *leaf)
-{
-    call_rcu(&leaf->rcu, __leaf_free_rcu);
-}
-
static void __leaf_info_free_rcu(struct rcu_head *head)

```

```

{
    kfree(container_of(head, struct leaf_info, rcu));
@@ -357,7 +352,12 @@ static void __tnode_free_rcu(struct rcu_
static inline void tnode_free(struct tnode *tn)
{
- call_rcu(&tn->rcu, __tnode_free_rcu);
+ if(IS_LEAF(tn)) {
+     struct leaf *l = (struct leaf *) tn;
+     call_rcu_bh(&l->rcu, __leaf_free_rcu);
+ }
+     else
+     call_rcu(&tn->rcu, __tnode_free_rcu);
}

static struct leaf *leaf_new(void)
diff --git a/net/ipv4/netfilter/ip_conntrack_netlink.c b/net/ipv4/netfilter/ip_conntrack_netlink.c
index a7167e0..56b3f1e 100644
--- a/net/ipv4/netfilter/ip_conntrack_netlink.c
+++ b/net/ipv4/netfilter/ip_conntrack_netlink.c
@@ -1631,7 +1631,7 @@ static void __exit ctnetlink_exit(void)
    printk("ctnetlink: unregistering from nfnetlink.\n");

#endif CONFIG_IP_NF_CONNTRACK_EVENTS
- ip_conntrack_unregister_notifier(&ctl_nNotifier_exp);
+ ip_conntrack_expect_unregister_notifier(&ctl_nNotifier_e xp);
    ip_conntrack_unregister_notifier(&ctl_nNotifier);
#endif

diff --git a/net/netfilter/nf_conntrack_netlink.c b/net/netfilter/nf_conntrack_netlink.c
index 9ff3463..40eedef 100644
--- a/net/netfilter/nf_conntrack_netlink.c
+++ b/net/netfilter/nf_conntrack_netlink.c
@@ -1641,7 +1641,7 @@ static void __exit ctnetlink_exit(void)
    printk("ctnetlink: unregistering from nfnetlink.\n");

#endif CONFIG_NF_CONNTRACK_EVENTS
- nf_conntrack_unregister_notifier(&ctl_nNotifier_exp);
+ nf_conntrack_expect_unregister_notifier(&ctl_nNotifier_e xp);
    nf_conntrack_unregister_notifier(&ctl_nNotifier);
#endif

diff --git a/security/keys/key.c b/security/keys/key.c
index 99781b7..0e2584e 100644
--- a/security/keys/key.c
+++ b/security/keys/key.c
@@ -785,6 +785,10 @@ key_ref_t key_create_or_update(key_ref_t

```

```

key_check(keyring);

+ key_ref = ERR_PTR(-ENOTDIR);
+ if (keyring->type != &key_type_keyring)
+     goto error_2;
+
down_write(&keyring->sem);

/* if we're going to allocate a new key, we're going to have
diff --git a/security/keys/keyring.c b/security/keys/keyring.c
index d65a180..bffa924 100644
--- a/security/keys/keyring.c
+++ b/security/keys/keyring.c
@@ -437,6 +437,7 @@ EXPORT_SYMBOL(keyring_search);
*/
/* search the given keyring only (no recursion)
 * - keyring must be locked by caller
+ * - caller must guarantee that the keyring is a keyring
 */
key_ref_t __keyring_search_one(key_ref_t keyring_ref,
                               const struct key_type *ktype,
diff --git a/sound/isa/opti9xx/opti92x-ad1848.c b/sound/isa/opti9xx/opti92x-ad1848.c
index 63d96be..65b28cb 100644
--- a/sound/isa/opti9xx/opti92x-ad1848.c
+++ b/sound/isa/opti9xx/opti92x-ad1848.c
@@ -2088,9 +2088,11 @@ static int __init alsa_card_opti9xx_init
int error;
struct platform_device *device;

+ifdef CONFIG_PNP
pnp_register_card_driver(&opti9xx_pnpc_driver);
if (snd_opti9xx_pnp_is_probed)
    return 0;
+endif
if (!isapnp_selected()) {
    error = platform_driver_register(&snd_opti9xx_driver);
    if (error < 0)
@@ -2102,7 +2104,9 @@ static int __init alsa_card_opti9xx_init
    }
    platform_driver_unregister(&snd_opti9xx_driver);
}
+ifdef CONFIG_PNP
pnp_unregister_card_driver(&opti9xx_pnpc_driver);
+endif
#endif
printk(KERN_ERR "no OPTi " CHIP_NAME " soundcard found\n");
#endif
@@ -2115,7 +2119,9 @@ static void __exit alsa_card_opti9xx_exi

```

```

platform_device_unregister(snd_opti9xx_platform_device);
platform_driver_unregister(&snd_opti9xx_driver);
}
#endif CONFIG_PNP
pnp_unregister_card_driver(&opti9xx_pnpc_driver);
#endif
}

module_init(alsa_card_opti9xx_init)
diff --git a/sound/pci/hda/patch_realtek.c b/sound/pci/hda/patch_realtek.c
index b767552..d5cd3a1 100644
--- a/sound/pci/hda/patch_realtek.c
+++ b/sound/pci/hda/patch_realtek.c
@@ -2948,6 +2948,8 @@ static struct hda_board_config alc260_cf
{ .modelname = "basic", .config = ALC260_BASIC },
{ .pci_subvendor = 0x104d, .pci_subdevice = 0x81bb,
  .config = ALC260_BASIC }, /* Sony VAIO */
+ { .pci_subvendor = 0x152d, .pci_subdevice = 0x0729,
+   .config = ALC260_BASIC }, /* CTL Travel Master U553W */
{ .modelname = "hp", .config = ALC260_HP },
{ .pci_subvendor = 0x103c, .pci_subdevice = 0x3010, .config = ALC260_HP },
{ .pci_subvendor = 0x103c, .pci_subdevice = 0x3011, .config = ALC260_HP },

```

File Attachments

-
- 1) [diff-merge-2.6.15.5-20060413](#), downloaded 395 times
-