» VX304x, VX3035 & VM6050 «

Release Notes Fedora 16 on VX304x, VX3035 & VM6050 Version 3.1 - ID 13049

SD.DT.G11-3e - October 2013

Revision History

Publication Title:	Fedora 16 Release Notes	
Doc. ID:	SD.DT.G11-3e	
Rev.	Brief Description of Changes	Date of Issue
3e	Chapter 8 - RC boards added Section 8.3 RC Specifications added	10-2013
2e	Chapters 8 and 9 added	08-2013
1e	VX3035 and VM6050 boards added	06-2013
0e	Initial Version	04-2013

Proprietary Note

This document contains information proprietary to Kontron. It may not be copied or transmitted by any means, disclosed to others, or stored in any retrieval system or media without the prior written consent of Kontron or one of its authorized agents.

The information contained in this document is, to the best of our knowledge, entirely correct. However, Kontron cannot accept liability for any inaccuracies or the consequences thereof, or for any liability arising from the use or application of any circuit, product, or example shown in this document.

Kontron reserves the right to change, modify, or improve this document or the product described herein, as seen fit by Kontron without further notice.

Trademarks

This document may include names, company logos and trademarks, which are registered trademarks and, therefore, proprietary to their respective owners.

Environmental Protection Statement

This product has been manufactured to satisfy environmental protection requirements where possible. Many of the components used (structural parts, printed circuit boards, connectors, batteries, etc.) are capable of being recycled.

Final disposition of this product after its service life must be accomplished in accordance with applicable country, state, or local laws or regulations.

Environmental protection is a high priority with Kontron.

Kontron follows the DEEE/WEEE directive.

You are encouraged to return our products for proper disposal.

The Waste Electrical and Electronic Equipment (WEEE) Directive aims to:

- reduce waste arising from electrical and electronic equipment (EEE)
- make producers of EEE responsible for the environmental impact of their products, especially when they become waste
- encourage separate collection and subsequent treatment, reuse, recovery, recycling and sound environmental disposal of EEE
- improve the environmental performance of all those involved during the lifecycle of EEE

Conventions

This guide uses several types of notice: Note, Caution, ESD.

Note: this notice calls attention to important features or instructions.


Caution: this notice alert you to system damage, loss of data, or risk of personal injury.


ESD: This banner indicates an Electrostatic Sensitive Device.

All numbers are expressed in decimal, except addresses and memory or register data, which are expressed in hexadecimal. The prefix `0x' shows a hexadecimal number, following the `C' programming language convention.

The multipliers `k', `M' and `G' have their conventional scientific and engineering meanings of $*10^3$, $*10^6$ and $*10^9$ respectively. The only exception to this is in the description of the size of memory areas, when `K', `M' and `G' mean $*2^{10}$, $*2^{20}$ and $*2^{30}$ respectively.

When describing transfer rates, `k' `M' and `G' mean $*10^3$, $*10^6$ and $*10^9$ *not $*2^{10}$ $*2^{20}$ and $*2^{30}$.*

In PowerPC terminology, multiple bit fields are numbered from 0 to n, where 0 is the MSB and n is the LSB. PCI and CompactPCI terminology follows the more familiar convention that bit 0 is the LSB and n is the MSB.

Signal names ending with an asterisk (*) or a hash (#) denote active low signals; all other signals are active high.

Signal names follow the PICMG 2.0 R3.0 CompactPCI Specification and the PCI Local Bus 2.3 Specification.

For Your Safety

Your new Kontron product was developed and tested carefully to provide all features necessary to ensure its compliance with electrical safety requirements. It was also designed for a long fault-free life. However, the life expectancy of your product can be drastically reduced by improper treatment during unpacking and installation. Therefore, in the interest of your own safety and of the correct operation of your new Kontron product, you are requested to conform with the following guidelines.

High Voltage Safety Instructions


Warning!
All operations on this device must be carried out by sufficiently skilled personnel only.


Caution, Electric Shock!
Before installing a not hot-swappable Kontron product into a system always ensure that your mains power is switched off. This applies also to the installation of piggybacks. Serious electrical shock hazards can exist during all installation, repair and maintenance operations with this product. Therefore, always unplug the power cable and any other cables which provide external voltages before performing work.

Special Handling and Unpacking Instructions

ESD Sensitive Device!
Electronic boards and their components are sensitive to static electricity. Therefore, care must be taken during all handling operations and inspections of this product, in order to ensure product integrity at all times

Do not handle this product out of its protective enclosure while it is not used for operational purposes unless it is otherwise protected.

Whenever possible, unpack or pack this product only at EOS/ESD safe work stations. Where a safe work station is not guaranteed, it is important for the user to be electrically discharged before touching the product with his/her hands or tools. This is most easily done by touching a metal part of your system housing.

It is particularly important to observe standard anti-static precautions when changing piggybacks, ROM devices, jumper settings etc. If the product contains batteries for RTC or memory backup, ensure that the board is not placed on conductive surfaces, including anti-static plastics or sponges. They can cause short circuits and damage the batteries or conductive circuits on the board.

General Instructions on Usage

In order to maintain Kontron's product warranty, this product must not be altered or modified in any way. Changes or modifications to the device, which are not explicitly approved by Kontron and described in this manual or received from Kontron's Technical Support as a special handling instruction, will void your warranty.

This device should only be installed in or connected to systems that fulfill all necessary technical and specific environmental requirements. This applies also to the operational temperature range of the specific board version, which must not be exceeded. If batteries are present, their temperature restrictions must be taken into account.

In performing all necessary installation and application operations, please follow only the instructions supplied by the present manual.

Keep all the original packaging material for future storage or warranty shipments. If it is necessary to store or ship the board, please re-pack it as nearly as possible in the manner in which it was delivered.

Special care is necessary when handling or unpacking the product. Please consult the special handling and unpacking instruction on the previous page of this manual.

Table Of Contents

Chapter 1 - Overview

Functional changes that differ from previous version of the document are identified by a vertical bar in the margin.

Linux, the Open Source Operating System is now taking a significant share of the OS market in Defense and Aerospace, after having taken ground initially in the enterprise server sector.

The goal of this document is to help you through the installation process of the Fedora 16 BSP distribution on the Kontron VX304x, VX3035 & VM6050 boards.

In this document, the terms VX304x, VX3035 & VM6050 are used for the VX304x, VX3035 & VM6050 boards in standard or rugged conduction-cooled version:

VX304x 3U VPX Computing Node

- VX304x-SA Standard Commercial version
- VX304x-RC Rugged Conduction-Cooled version

VX3035 Single-slot 3U VPX board

- VX3035-SA Standard Commercial version
- VX3035-RC Rugged Conduction-Cooled version

VM6050 Single-slot 6U VME board

- VM6050-SA Standard Commercial version
- VM6050-RC Rugged Conduction-Cooled version

In this document, the terms VX304x-RTM, VX3035-RTM & VM6050-RTM are associated to the VX304x, VX3035 & VM6050 Rear Transition Module (RTM):

VX304x-RTM Rear Transition Module for the single-slot 3U VPX board

- PB-VX3-4xx

VX3035-RTM Rear Transition Module for the single-slot 6U VPX board

- PB-VX3-011

VM6050-RTM Rear Transition Module for the single-slot 6U VPX board

- PBV36-P0-VM6-00

Chapter 2 - Release Content

In this document the term Ki7 is used to refer to a board among the VX304x, VX3035 & VM6050 boards.

The release is made of:

> **Fedora 16 x86_64 DVD for Ki7 boards**

This distribution includes the standard Fedora 16 Linux release as well as the BSP packages related to the Ki7 boards.

You can choose to install this distribution in a graphical configuration or in a serial console configuration.

> **Fedora 16 x86_64 LiveDVD for Ki7 boards**

The content of this distribution is identical to the previous one but the usage is different.

This distribution has a double role: at first it can be used in order to evaluate this new release without impacting the content of onboard SATA disk.

In a second time, it allows to install a LiveUSB flavor of Linux on onboard SATA flash SSD or USB flash (stick).

The main interest of a LiveUSB installation is to preserve as much as possible the flash devices by limiting the writing cycles to them.

The Board Support Package (BSP) provides support for some specific features of the board:

> Kernel:	Update of the kernel of the Fedora 16 to support Ki7 boards specifics features and to fix issues. Note that this kernel is required by the following features.
> Sensors:	CPU Cores and Board temperatures and voltages.
> Vital Product Data (VPD) Tool:	Get board's serial number, order code, E.C. Level, ...
> LEDs:	Four Front Panel Tri-color LEDs
> GPIO:	Driver to support the GPIOs of the Ki7 boards
> Watchdog:	Drivers to setup the Watchdogs of the board.
> BIOS Update tool:	A command and script to update the BIOS of the board.
> CPLD register Tool (cpldtool):	Tool to deal with hardware registers of the onboard CPLD
> FRAM support:	Driver and special API file to read/write from/to the FRAM
> Diskless:	Tool to configure, boot and manage diskless boards.
> PBIT report:	sysvartool gives the report of the PBIT.
> I2C buses drivers:	i2c bus drivers for the local i2c bus and the two backplane i2c busses.

The VXFabric product is delivered apart from the BSP.

More information on Ki7 boards BSP in Chapter 7 “BSP Specific Features” page 22.

Information on Fedora 16 is available at http://fedoraproject.org/

Chapter 3 - Associated Documentation

Kontron Documentation

> Hardware

- VX304x 3U VPX Computing Node User's Guide CA.DT.A98
- VX304x Hardware Release Notes CA.DT.A99

- VX3035 3U VPX SBC User's Guide CA.DT.A95
- VX3035 Hardware Release Notes CA.DT.A96

- VM6050 6U VME SBC User's Guide CA.DT.A93
- VM6050 Hardware Release Notes CA.DT.A94

> Firmware

- VX304x BIOS User Manual SD.DT.F96

- VX3035 BIOS User Manual SD.DT.F97

- VM6050 BIOS User Manual SD.DT.F89

Fedora 16 Documentation

> Documentation available at http://fedoraproject.org/

Chapter 4 - Required Configuration

4.1 Hardware Requirements

4.1.1 Hardware Requirements for VX304x Boards

> A Kontron VX304x board.

> The Fedora 16 release may be installed on one of the following bootable disks:

- a SATA disk connected to the SATA connectors available on VX304x-RTM board.
- an optional onboard USB Flash Disk.
- an optional onboard SSD Flash Disk

> A USB DVD-ROM device (for installation from DVD-ROM).

> - A graphical display (with mini DisplayPort interface), USB keyboard and USB mouse (for a graphics install)

OR

- A console on serial line (text or VNC install).

4.1.2 Hardware Requirements for VX3035 Board

> A Kontron VX3035 board.

> The Fedora 16 release may be installed on one of the following bootable disks:

- a SATA disk connected to the SATA connectors available on VX3035-RTM board.
- Optional onboard USB Flash Disk.

> A USB DVD-ROM device (for installation from DVD-ROM).

> Graphical display, USB keyboard and USB mouse (for a graphics install) or a console on serial line (text or VNC install).

4.1.3 Hardware Requirements for VM6050 Board

> A Kontron VM6050 board.

> The Fedora 16 release may be installed on one of the following bootable disks:

- a SATA disk connected to the SATA connectors available on VM6050-RTM board.
- Optional onboard USB Flash Disk.

> A USB DVD-ROM device (for installation from DVD-ROM).

> A console on serial line (text or VNC install).

For a graphic configuration on VM6050 a specific order code and a specific graphic module are required: VM6050-2SA34-12110 and MOD-GX-SA-00. The module provides two DP ports and a VGA connector.

4.2 Firmware Requirements

4.2.1 VX304x Boards

The version of the BIOS firmware must be at least:

> 12332

This version is displayed in the BIOS Setup.

4.2.2 VX3035 Board

The version of the BIOS firmware must be at least:

> 12174

This version is displayed in the BIOS Setup.

4.2.3 VM6050 Board

The version of the BIOS firmware must be at least:

> 11332

This version is displayed in the BIOS Setup.

4.3 Software Requirements

> The DVD-ROMs:

- Fedora 16 x86_64 DVD for Ki7 boards
- Fedora 16 x86_64 LiveDVD for Ki7 boards

After the release is installed, the release version is saved in `/etc/ki7-release`:

```
# cat /etc/ki7-release Ki7 Board Support Package 3.1 [13049]
```

4.4 DVD-ROM Installation Example

4.4.1 DVD-ROM Installation Example for VX304x Board

1. Plug a DVD-ROM USB device to the front panel USB connector using a standard USB cable.
2. Plug the HDD SATA device on the SATA0-A connector of the Rear Transition Module (VX304x-RTM).

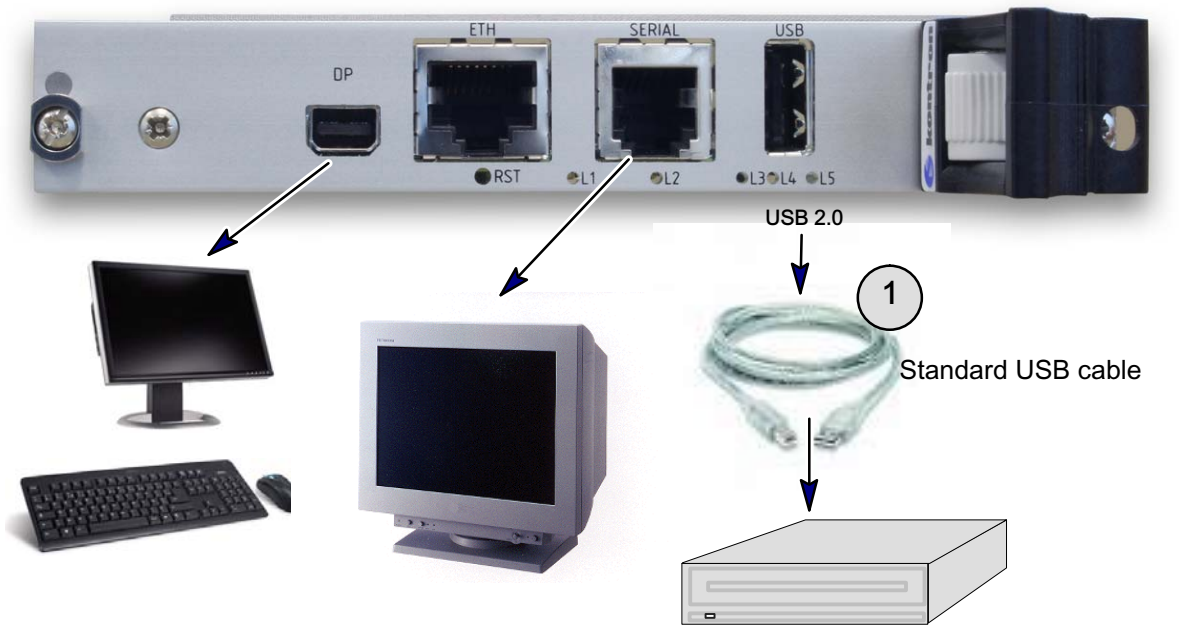

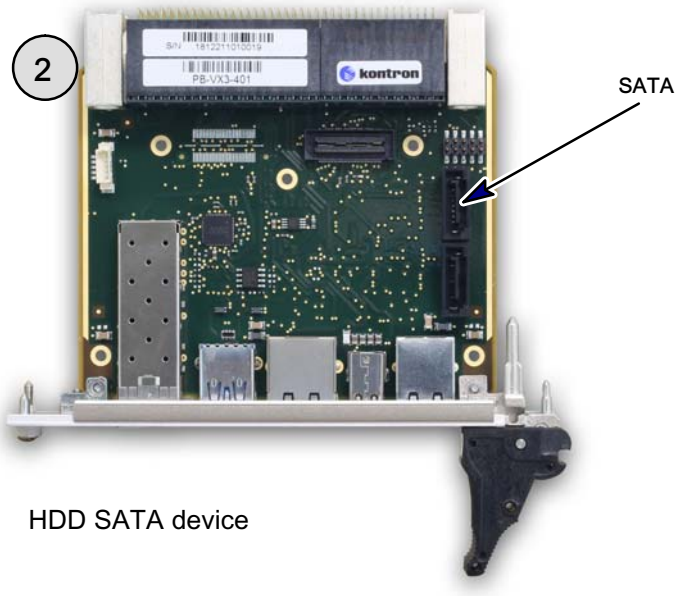

There are 2 USB ports available: one on the front panel and the other one on the RTM board if a RTM is available. So in order to use a USB keyboard and a USB mouse plus a USB DVD-ROM, a USB HUB is required.

4.4.2 DVD-ROM Installation Example for VX3035

1. Plug a DVD-ROM USB device to the front panel USB connector using a standard USB cable.
2. Plug the HDD SATA device on the SATA0-A connector of the Rear Transition Module (VX3035-RTM).

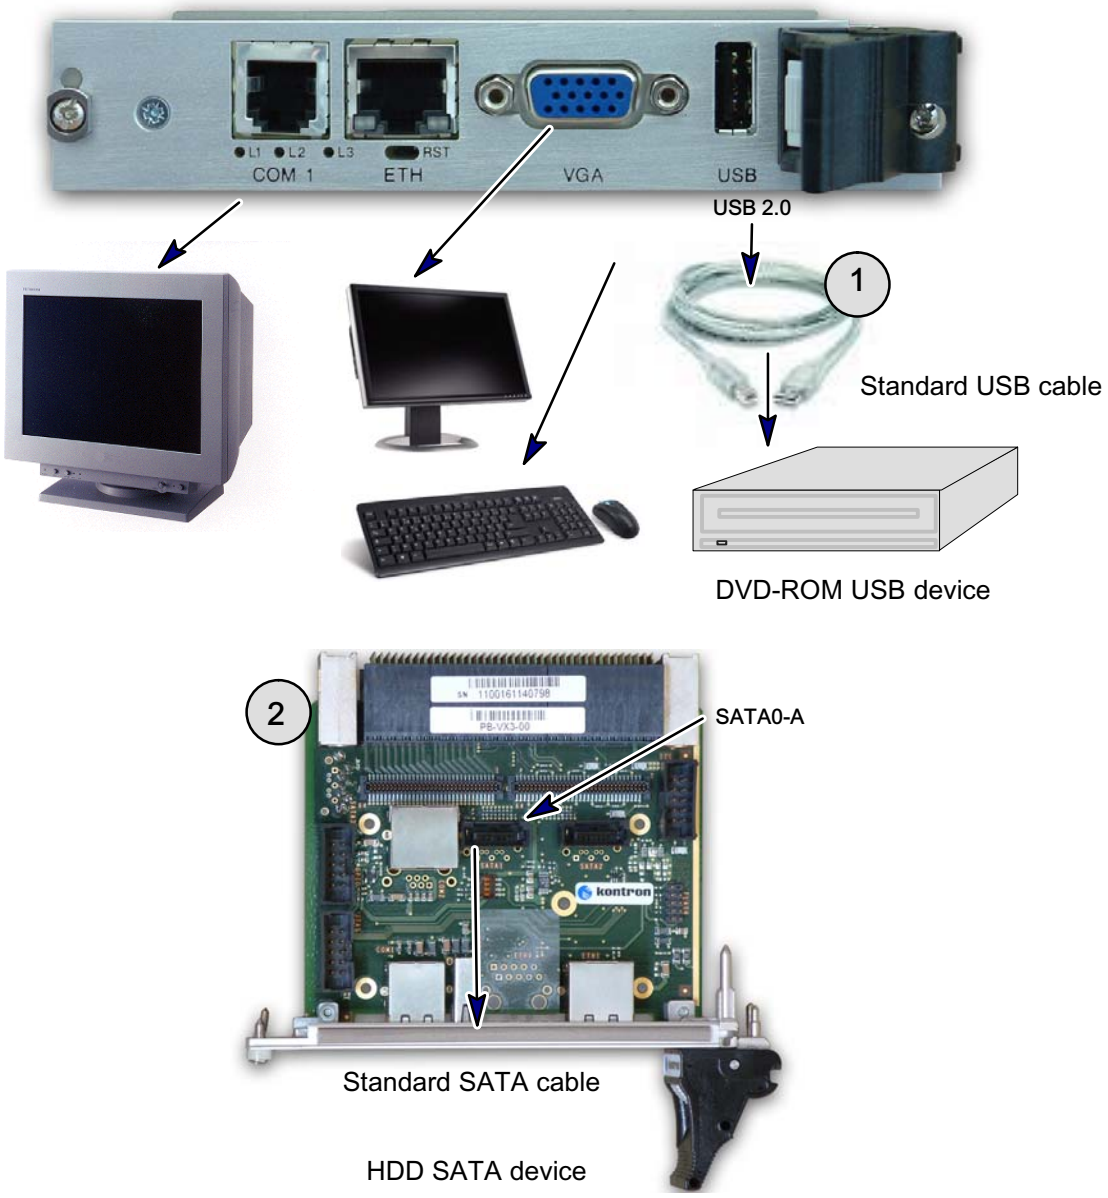

There are 2 USB ports available: one on the front panel and the other one on the RTM board if a RTM is available. So in order to use a USB keyboard and a USB mouse plus a USB DVD-ROM, a USB HUB is required.

4.4.3 DVD-ROM Installation Example for VM6050

1. Plug a DVD-ROM USB device to the front panel USB connector using a standard USB cable.
2. Plug the HDD SATA device on the SATA0-A connector of the Rear Transition Module (VM6050-RTM).

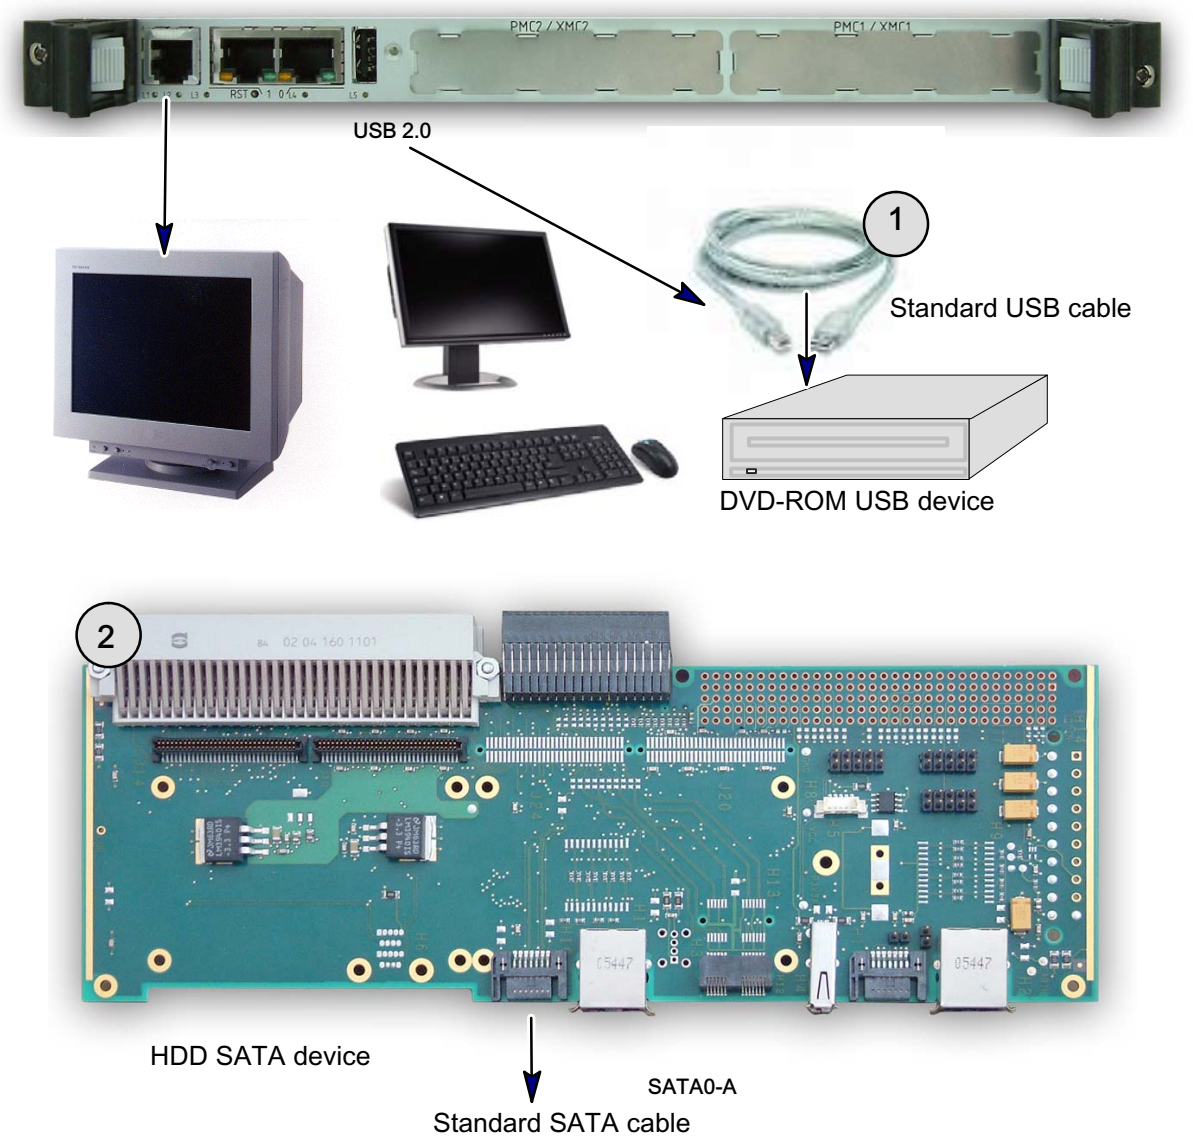

Graphic configuration on VM6050 requires a specific order code and a specific graphic module: VM6050-2SA34-12110 and MOD-GX-SA-00. The module provides two DP ports and a VGA connector.

There are 2 USB ports available: one on the front panel and the other one on the RTM board if a RTM is available. So in order to use a USB keyboard and a USB mouse plus a USB DVD-ROM, a USB HUB is required.

Chapter 5 - Installation

This chapter describes the specific steps of the installation process of the Fedora 16 on Ki7 boards.

5.1 Disk Installation

5.1.1 Fedora 16 DVD Menu

This section describes the options added to the installation menu on the Fedora 16 x86_64 DVD for Ki7 boards media.

When booting from this media, the following menu appears on the serial line (and also on the graphics display if present):

```
*                  Welcome to Fedora 16 ki7                    *
************************************************************
* Install or upgrade using GRAPHICS mode                     *
* Install or upgrade using SERIAL console                    *
* Install or upgrade using VNC                               *
* Install system with basic video driver                     *
* Rescue installed system using GRAPHICS mode                *
* Rescue installed system using SERIAL console               *
* Boot from local drive                                      *
* Memory test                                                *
*                                                            *
*                                                            *
*                                                            *
*                                                            *
************************************************************
```

Select one of these options :

> "`Install or upgrade using SERIAL console`": to do all the installation using the serial console. However in this mode, the disk partitionning cannot be customized and only a fixed very minimal set of 200 packages is installed. Depending on the needed services and libraries, some post installations may be required, so this is not the recommended method. If possible, install using VNC (see below) or using a graphics display.

> "`Install or upgrade using VNC`": to start the installation using the serial console, but then use a VNC client on another system to do the rest of the installation in graphics mode with all installation options available. This requires the network to be connected to the board.

When installing using the "`Install or upgrade using VNC`", the default behaviour is to use a DHCP server to provide the network settings to the board.

If no DHCP server is available, it is possible to provide these settings manually as follows:

> Move the cursor to the "`Install or upgrade using VNC`"

> Press <tab> and append the network settings; at least the IP address :

```
> vmlinuz initrd=initrd.img console=ttyS0,115200 text stage2=hd:LABEL="Fedora" vnc
ip=172.20.144.95
```

or more settings :

```
> vmlinuz initrd=initrd.img console=ttyS0,115200 text stage2=hd:LABEL="Fedora" vnc
ip=172.20.161.201 netmask=255.255.255.0 gateway=172.20.161.46 dns=172.20.144.1
```

> and press <ENTER>

The following menu should be displayed (the Ethernet device list is an example):

```
Welcome to Fedora for x86_64

              ***************** Networking Device *****************
              *                                                   *
              * You have multiple network devices on this system. *
              * Which would you like to install through?          *
              *                                                   *
              *  eth0 - Ethernet device eth0 - 00:00:de:40:39:b5  *
              *  eth1 - Ethernet device eth1 - 00:00:de:40:39:b6  *
              *  eth2 - Ethernet device eth2 - 00:00:de:40:39:b7  *
              *  eth3 - Ethernet device eth3 - 00:00:de:40:39:b8  *
              *  eth4 - Ethernet device eth4 - 00:30:f7:98:2b:00  *
              *  eth5 - Ethernet device eth5 - 00:30:f7:98:2b:01  *
              *                                                   *
              *    ******        ************        ********     *
              *    * OK *        * Identify *        * Back *     *
              *    ******        ************        ********     *
              *                                                   *
              *                                                   *
              *****************************************************

  <Tab>/<Alt-Tab> between elements  | <Space> selects | <F12> next screen
```

Select the interface to use during installation to connect with the VNC client

```
Welcome to Fedora for x86_64

              ************* Disc Found *************
              *                                    *
              * To begin testing the media before  *
              * installation press OK.             *
              *                                    *
              * Choose Skip to skip the media test *
              * and start the installation.        *
              *                                    *
              *     ******           ********      *
              *     * OK *           * Skip *      *
              *     ******           ********      *
              *                                    *
              *                                    *
              ************************************

<Tab>/<Alt-Tab> between elements  | <Space> selects | <F12> next screen
```

Select `Skip` in case you want to verify the validity of the media. At the end, the following messages are displayed:

```
Running anaconda 14.22, the Fedora system installer - please wait.
19:01:12 Starting VNC...
19:01:14 The VNC server is now running.
19:01:14

WARNING!!! VNC server running with NO PASSWORD!
You can use the vncpassword=<password> boot option
if you would like to secure the server.

19:01:14 Please manually connect your vnc client to 172.20.144.95:1 to begin the install.
Press <enter> for a shell
19:01:18 Starting graphical installation.
```

Then on a workstation, launch a VNC client to connect to the server running on the board. In the example above, the VNC server is 172.20.144.95:1

Then continue the installation process.

On a Fedora workstation, you can use the "TigerVNC Viewer" VNC client supplied with the distribution (can be launched from menu Application / Internet)

5.1.2 Fedora 16 Installation on Ki7 Boards

This section describes the installation procedure from an USB DVD-ROM drive.

There is no major difference between the installation on the Ki7 boards and the standard Fedora 16 installation, so refer to the Fedora documentation to get more details on all the Fedora installation menus.

Nevertheless, the Ki7 boards are a graphic board, so the installation may be done in graphic mode or in text mode on the serial port 0.

1. Insert the DVD Fedora 16 in the DVD-ROM driver.
2. After a board reset, or a board power-on, type **<F7>** to get the the Boot Manager Menu or **<F2>** to get the Setup from the BIOS and to select the DVD-ROM device as the boot device.
3. Select the installation method using the menu, as explained above.
4. Then the standard Fedora 16 Installer menus from anaconda should be displayed.
5. Proceed as a standard Fedora 16 installation.
6. Note that the mezzanine USB Flash device, if present, is probed as another SATA drive.
7. At the end of the installation, reboot on the installed disk drive through the Boot Manager Menu **<F7>**. The BIOS Setup menu should be used to set the boot devices priorities. Refer to the BIOS Manual - SD.DT.F96.

It is strongly recommended to disable the swap partition if the installation is done on a USB or SATA flash device.

5.1.3 BSP Installation on Ki7 Boards

The BSP packages are present on the media: **Fedora 16 x86_64 DVD for Ki7 boards.**

Follow the procedure below to install the BSP onto an existing bootable media (SATA disk of USB flash) or to access to the source packages:

1. Insert the Kontron media Fedora 16 x86_64 DVD for Ki7 boards in the USB DVD-ROM drive.
2. Boot the Ki7 board on the SATA disk previously installed with the Fedora 16.
3. Login as root and run the following commands:
   - At first, create a repo file in order to acces to the packages present on the DVD.

```
# mount /dev/sr0 /mnt
# cat >/etc/yum.repos.d/ki7_bsp.repo <<EOF
[ki7_bsp]
name=ki7_bsp
gpgcheck=0
enabled=0
baseurl=file:///mnt/
EOF
```

   - Then install the BSP with the command:

```
#yum -y --disablerepo=* --enablerepo=ki7_bsp install ki7_bsp
```

   The installation process will take few minutes.
4. Reboot the machine to take into account the new version of the BSP.

When the BSP is installed, the kernel is updated.

At this step you should configure the system to use the network service instead of the NetworkManager, and also do some other settings to configure the network interfaces to meet the requirements of an embedded system. See important information about that at section "6.1 - Network page 16"

5.1.4 LiveDVD Installation on SSD SATA Flash or USB Flash

The LiveDVD media allows to evaluate a new version of Linux in an easy way without degrading data present on local SATA disk or flash SSD disk.

An another usage of the LiveDVD is the creation of liveUSB on flash devices.

The main advantage of the liveUSB is to prevent, by default, all writing accesses to the flash disk (or USB stick). It is especially important for some USB flash (stick) or SSD SATA flash devices which do not tolerate a lot of writing cycles. However it is needed to be able to preserve some changes from one boot to the other one. The LiveUSB tool features this capability.

Creating a LiveUSB on the SSD flash disk implies the presence of the LiveDVD into the USB DVD-ROM drive.

At first, boot on the LiveDVD in the following way:

- Insert the Kontron LiveDVD into the USB DVD-ROM drive. You should also have a console connected to the serial port `ttyS0`.
- Under BIOS context, select the USB DVD-ROM drive as the first boot device.
- At the beginning of the boot step, a first menu should occur.
- Select the option:

  `Start FEDORA16 Ki7`
- At the end of the boot step notice the presence of a login on the console.

Using the `livecd-iso-to-disk` is the easiest way to create a LiveUSB image on a USB stick/SSD disk.

In case you use the media for the first time, you have to re-partition and format your media. To reach this purpose, use `fdisk`.

For example:

```
# fdisk /dev/sdX
n
Return
Return
Return
Return
t
6
p
w!
```

should create a unique partition which covers the whole media.

Then, reformat the media:

```
# mkfs.vfat -F 32 /dev/sdb1
```

Write the ISO image to the USB/SSD SATA flash disk using `livecd-iso-to-disk`:

```
# ./livecd-iso-to-disk --overlay-size-mb <OVERLAY_SIZE> --force --format /dev/sr0
/dev/sdX1
```

- Replace the reference sdX1 by the right one in your particular case.
- Replace the field <OVERLAY_SIZE> by the value which fit your own configuration (1024 for example will reserve a 1 Gigabyte writable area).
- Press <Enter> to launch the process.

Then reboot the board and select the SSD device as new boot device this time.

A menu is displayed proposing 2 choices:

```
- Start FEDORA16 Ki7
- Troubleshooting
```

Choose Start FEDORA16 Ki7. As a result, you have access to a complete Fedora 16 x86_64 system including the BSP specific to the Ki7 board.

Of course, you might customize this system with your own changes. All these changes will be preserved into the overlay layer.

Chapter 6 - Fedora System Configuration

In this chapter, informations related to some specific configuration items of the Fedora 16 system are detailed.

6.1 Network

6.1.1 Network Manager

With Fedora 16, the network interfaces are managed by the NetworkManager service by default.

For an embedded system, it is recommended to use the older network service instead which is easier to configure through configuration files.

For this:

> Disable the NetworkManager service:

```
[root@ki7]# chkconfig NetworkManager off
```

> Enable the network service:

```
[root@ki7]# chkconfig network on
```

> Stop the network manager: `service NetworkManager stop`

> Check the configuration files and modify then if needed :

- `/etc/sysconfig/network-scripts/ifcfg-ethx` files
- `/etc/resolv.conf`
- `/etc/sysconfig/network`

> Start the network service: `service network start`

> Reboot if `/etc/sysconfig/network` has been modified

6.1.2 MAC Address

By default, the MAC address is stored with the configuration parameters of each interface. If the MAC address of a device is found different from the one expected (board changed for example), the interface is not brought up. This is not suitable for an embedded system when boards must be changed for maintenance without requiring additional configuration.

To workaround this behavior, do not bind an Ethernet interface to a MAC address:

> Run `system-config-network`

> For each interface:

- click on Edit
- click on Hardware Device tab
- unselect Bind to MAC address

> Exit from `system-config-network` saving changes

This can be done also by editing the `/etc/sysconfig/network-scripts/ifcfg-eth*` files and removing the `HWADDR` lines.

Removing `HWADDR` only works if the service network (and not NetworkManager) is used.

6.1.3 Firewall

If the firewall must be disabled but has been enabled during the installation:

> Run `system-config-firewall`, click on disable and exit

OR run

```
[root@ki7]# service iptables stop
[root@ki7]# service ip6tables stop
```

> Make sure to disable the iptables service by running:

```
[root@ki7]# chkconfig iptables off
[root@ki7]# chkconfig ip6tables off
```

> Reboot

6.2 SELinux

SELinux stands for Security-Enhanced Linux. The Security-Enhanced Linux kernel enforces mandatory access control policies that confine user programs and system servers to the minimum amount of privilege they require to do their jobs.

If you experience some trouble running some services or have some permission issues, try to set the System Default Policy to Permissive instead of Enforcing by running the `system-config-selinux` tool, or from command line doing as follows:

- disable on boot by editing `/etc/selinux/config` to set `SELINUX=permissive` instead of `SELINUX=enforcing`
- disable now: `setenforce 0`

6.3 GRUB Boot Loader

If your console is on the serial line and that access to the grub boot menu is not required on the graphics console, you should comment the following line in `/boot/grub/grub.conf` if present :

```
terminal --timeout=5 serial console
```

This will prevent GRUB from waiting 5 seconds to let you select a console by typing a character on one console to select it.

This setting has nothing to do with the timeout on the boot menu that is set by the line

```
timeout=<value>
```

6.4 SATA Speed

With some specific environmental constraints or with some SATA devices, it may be required to reduce the SATA speed of a specific SATA port.

The SATA speed may be up to 6 Gbps for Ki7 boards.

To properly manage the SATA speed, first of all check at the BIOS setup that AHCI mode is enabled. Furthermore, AHCI mode should allow the access to Hotplug option (refer to next chapter).

The current speed of the SATA ports may be checked at boot time by:

```
[root@ki7]# dmesg | egrep ata[1-9]:
[    3.480624] ata1: SATA max UDMA/133 abar m2048@0xf1c02000 port 0xf1c02100 irq 45
[    3.487984] ata2: SATA max UDMA/133 irq_stat 0x00400040, connection status changed irq 45
[    3.496119] ata3: DUMMY
[    3.498552] ata4: DUMMY
[    3.500985] ata5: DUMMY
[    3.503419] ata6: DUMMY
[    3.809480] ata1: SATA link down (SStatus 0 SControl 300)
[    4.226743] ata2: SATA link up 3.0 Gbps (SStatus 123 SControl 300)
```

Here the port 2 is set at 3.0 Gbps.

If the hard drive is finaly not supporting this rate, it is possible to decrease it by adding `libata.force=2:1.5` (for example) to the kernel command boot line to force at boot time the SATA speed limit of the port 2 at 1.5 Gbps.

```
[root@ki7]# KERNEL=`grubby --default-kernel` ; grubby --update-kernel=$KERNEL --args=libata.force=2:1.5
```

Then after a reboot:

```
[root@ki7]# dmesg | egrep ata[1-9]:
[    3.480245] ata1: SATA max UDMA/133 abar m2048@0xf1c02000 port 0xf1c02100 irq 45
[    3.487603] ata2: FORCE: PHY spd limit set to 1.5Gbps
[    3.492627] ata2: SATA max UDMA/133 irq_stat 0x00400040, connection status changed irq 45
[    3.500762] ata3: DUMMY
[    3.503195] ata4: DUMMY
[    3.505627] ata5: DUMMY
[    3.508059] ata6: DUMMY
[    3.814088] ata1: SATA link down (SStatus 0 SControl 300)
[    4.231351] ata2: SATA link up 1.5 Gbps (SStatus 113 SControl 310)
```

It is required for the Ki7 to setup the SATA Speed of the onboard SDD Flash device to 1.5 Gbps. This device is located at port 4. So, if it is not already present, add the following argument to the kernel command line using:

```
[root@ki7]# KERNEL=`grubby --default-kernel` ; grubby
--update-kernel=$KERNEL --args=libata.force=5:1.5
```

6.5 SATA Hotplug

In AHCI mode, the SATA controller of the Ki7 boards provides with hotplug function. First of all this has to be setup at BIOS menus (Chipset->South Bridge Configuration->SATA Configuration). After the Hot Plug option is enabled for the SATA ports, boot the system.

» To remove a SATA device from the system:

Close all users of the device and backup device data as needed. Use `umount` to unmount any file systems that mounted the device.

Remove the device from any md and LVM volume using it. If the device is a member of an LVM Volume group, then it may be necessary to move data off the device using the `pvmove` command, then use the `vgreduce` command to remove the physical volume, and (optionally) `pvremove` to remove the LVM metadata from the disk.

If the device uses multipathing, run `multipath -l` and note all the paths to the device. Afterwards, remove the multipathed device using `multipath -f` device.

Run `blockdev -flushbufs` device to flush any outstanding I/O to all paths to the device. This is particularly important for raw devices, where there is no umount or vgreduce operation to cause an I/O flush.

Remove any reference to the device's path-based name, like `/dev/sd`, `/dev/disk/by-path` or the major:minor number, in applications, scripts, or utilities on the system. This is important in ensuring that different devices added in the future will not be mistaken for the current device.

Finally, remove each path to the device from the SCSI subsystem. To do so, use the command:

```
[root@ki7]# echo 1 > /sys/block/device-name/device/delete
```

where device-name may be `sde`, for example.

```
Nov 10 15:16:24 Ki7 kernel: [10018.256462] sd 1:0:0:0: [sdb] Synchronizing SCSI cache
Nov 10 15:16:24 Ki7 kernel: [10018.534156] sd 1:0:0:0: [sdb] Stopping disk
Nov 10 15:16:24 Ki7 kernel: [10018.934519] ata2.00: disabled
```

Then you can shut off the device.

» To add a SATA device:

When the system is up and running, power on the hotplugeable SATA device and the system should be warned that a new SATA device is available:

```
[root@ki7]# dmesg
...
Nov 10 15:22:40 Ki7 kernel: [10394.408164] ata2: irq_stat 0x00400040, connection status changed
Nov 10 15:22:40 Ki7 kernel: [10394.414149] ata2: SError: { RecovComm PHYRdyChg CommWake DevExch }
Nov 10 15:22:40 Ki7 kernel: [10394.420310] ata2: hard resetting link
Nov 10 15:22:43 Ki7 kernel: [10396.923493] ata2: SATA link up 3.0 Gbps (SStatus 123 SControl 300)
Nov 10 15:22:43 Ki7 kernel: [10396.925408] ata2.00: ATA-8: ST9160314AS, 0001SDM1, max UDMA/133
Nov 10 15:22:43 Ki7 kernel: [10396.925414] ata2.00: 312581808 sectors, multi 16: LBA48 NCQ (depth 31/32)
Nov 10 15:22:43 Ki7 kernel: [10396.927836] ata2.00: configured for UDMA/133
Nov 10 15:22:43 Ki7 kernel: [10396.927846] ata2: EH complete
Nov 10 15:22:43 Ki7 kernel: [10396.927991] scsi 1:0:0:0: Direct-Access ATA ST9160314AS\0001 PQ: 0 ANSI: 5
Nov 10 15:22:43 Ki7 kernel: [10396.928246] sd 1:0:0:0: Attached scsi generic sg1 type 0
Nov 10 15:22:43 Ki7 kernel: [10396.928299] sd 1:0:0:0: [sdb] 312581808 512-byte logical blocks: \
                                                                                (160 GB/149 GiB)
Nov 10 15:22:43 Ki7 kernel: [10396.928398] sd 1:0:0:0: [sdb] Write Protect is off
Nov 10 15:22:43 Ki7 kernel: [10396.928443] sd 1:0:0:0: [sdb] Write cache: enabled, \
                                                        read cache: enabled, doesn´t support DPO or FUA
Nov 10 15:22:43 Ki7 kernel: [10396.928772]  sdb: sdb1 sdb2
Nov 10 15:22:43 Ki7 kernel: [10396.955136] sd 1:0:0:0: [sdb] Attached SCSI disk
```

If automount services are enabled, the partitions should be mounted too.

6.6 VITA 57

Using VX3830 IO cards for VPX boards or VM6050, it is possible to use the VITA 57 option. Kontron is delivering a toolkit to help the setup and development of application using the onboard FPGA and the FMC modules defined though the VITA 57 standard. Refer to your representative to get more information about this option.

6.7 VXFabric

Kontron VXFabric™ is an open infrastructure which implements efficient inter board communication at hardware speed. The architecture is compliant with the OpenVPX standard (VITA 65) which defines two main hardware topologies of the backplane: distributed and centralized topologies.

To get more information about this software, go to the www.kontron.com web site and enter the key word "vxfabric" into the search engine.

Chapter 7 - BSP Specific Features

7.1 Sensors

The BSP contains an RPM named `sensors_addons` that configures the standard `lm_sensor` software for the Ki7 boards.

To display sensors information:

```
[root@ki7]# sensors
```

» Example on VX304x

```
# sensors
acpitz-virtual-0
Adapter: Virtual device
temp1:        +37.0°C  (crit = +106.0°C)
temp2:        +36.0°C  (crit = +106.0°C)

nct7802y-i2c-22-28
Adapter: I2C CPLD adapter
 3V3 SB (A/D REF):    +3.24 V  (min =  +0.18 V, max =  +0.38 V)
 Processor Vcore  :   +0.86 V  (min =  +0.00 V, max =  +0.00 V)
 1V5 Memory       :   +1.52 V  (min =  +0.40 V, max =  +0.04 V)
 12V VPX VS1      :  +12.10 V  (min =  +0.13 V, max =  +1.34 V)
 5V  VPX VS3      :   +4.99 V  (min =  +0.78 V, max =  +1.15 V)
 VPX 3V3 aux      :   +0.00 V
nct7802y local temp:  +30.0°C  (low  = -40.0°C, high = +85.0°C)

coretemp-isa-0000
Adapter: ISA adapter
Physical id 0: +38.0°C  (high = +87.0°C, crit = +105.0°C)
Core 0:         +35.0°C  (high = +87.0°C, crit = +105.0°C)
Core 1:         +31.0°C  (high = +87.0°C, crit = +105.0°C)
Core 2:         +34.0°C  (high = +87.0°C, crit = +105.0°C)
Core 3:         +32.0°C  (high = +87.0°C, crit = +105.0°C)

pchtemp-pci-00fe
Adapter: PCI adapter
temp1:        +43.5°C
```

The sensor command reveals the presence of low, high and critical thresholds. In case the temperature temp1 go beyond the critical threshold, an automatic reset of the board will occur.

In case the temperature or the voltage go beyond one of the limits low and high, an explicit alarm message will occur in the sensors command output. So, in order to track down this kind of event, run the following command:

```
# sensors | grep ALARM
```

On the other hand, it is possible to synchronize the execution of an application on a high limit temperature event.

For example, at first, create a file name /tmp/test.sh containing:

```
cd /sys/devices/platform/i2c-cpld.6/i2c-22/22-0028
cat < temp1_max_alarm_intr
echo ´Nuvoton local temp went beyond the high limit´
shutdown now
```

Make it executable and run it:

```
# chmod +x /tmp/test.sh
# /tmp/test.sh
```

As a consequence, in case the nuvoton temperature go beyond the high limit temperature, this shell script will be automatically called, causing a shutdown of the board.

7.2 Watchdog

The BSP 3.0 provides a driver for the watchdog of the CPLD.

The `cpld_wdt` is used instead of the standard kernel iTCO_wdt module added to the blacklist by the file `/etc/modprobe.d/cpld_wdt.conf`. Additional cpld_wdt module options can be added to this file if needed.

To check that the driver is loaded:

```
[root@ki7]# lsmod | grep cpld_wdt
cpld_wdt                3138  0
...
```

This driver can be loaded with two options: timeout of the watchdog (timeout) and the mode (trigger_mode).

By default, timeout is 21s and the trigger_mode is reset mode which means that if the watchdog falls, the board is reset.

To start or service (restart) the watchdog :

```
[root@ki7]# echo > /dev/watchdog
```

To stop the watchdog :

```
[root@ki7]# echo V > /dev/watchdog
```

To change the timeout:

```
[root@ki7]# rmmod cpld_wdt
[root@ki7]# modprobe cpld_wdt timeout=11
```

Only odd timeout values are available.

To change the mode to get an interruption instead of a reset:

```
[root@ki7]# rmmod cpld_wdt
[root@ki7]# modprobe cpld_wdt timeout=11 trigger_mode=2
```

Then start the watchdog by writing to `/dev/watchdog` (echo), wait for the interrupt reading the watchdog (cat), and when the interrupt occurs disable the watchdog that is restarted by the termination of "cat" (close of the watchdog device) :

```
[root@ki7]# echo > /dev/watchdog ; time cat /dev/watchdog ; echo V > /dev/watchdog
```

If no one is waiting on the interruption (no "cat"), an emergency reboot system call is run when the watchdog expires.

7.3 VPD Tool

`vpdtool` command is useful to get information related to the VX304x boards.

```
# vpdtool
Board type   : VX3044-SA48-2010M00
EC Level     : 20000BU
Serial Number: 1812251080014
Variant      : 1041B00401008800
Keylist                                                                                   :
/PCB_B/SACLASS/MEZZUSBSATA/BHDUAL/BHQUADMODE/ndef4_CPU_type/ndef3_DDR3_type/DDR_1B/DDR3-1333/
DDR3-Dual-Rank/IRTC/NOXDP/STDCLK/STD_EARTH/BATOFF/DPBDP/PWRMAGOFF/FP422OFF/TMLVDSOFF/I2CSTD/1
INCHES/1SLOT/P80ON/VS2OFF/NODCLK/IBOMOFF/ITIN/LFPROCESS/LFCOMP/
```

To get the features list:

```
# vpdtool --human
Board type   : VX3044-SA48-2010M00
EC Level     : 20000BU
Serial Number: 1812251080014
Variant      : 1041B00401008800
Keylist                                                                                   :
/PCB_B/SACLASS/MEZZUSBSATA/BHDUAL/BHQUADMODE/ndef4_CPU_type/ndef3_DDR3_type/DDR_1B/DDR3-1333/
DDR3-Dual-Rank/IRTC/NOXDP/STDCLK/STD_EARTH/BATOFF/DPBDP/PWRMAGOFF/FP422OFF/TMLVDSOFF/I2CSTD/1
INCHES/1SLOT/P80ON/VS2OFF/NODCLK/IBOMOFF/ITIN/LFPROCESS/LFCOMP/

Features     :
ERROR: PCB B not defined
SA Class
USB/SATA mezzanine slot configuration
Dual link chip NH82580DB
Quad link chip NH52580EB not forced in Dual mode
ERROR: ndef4_CPU_type
ERROR: DDR3 type not defined
SDRAM 1 bank only
DDR3-1333 dram speed
DDR3 Dual Rank (=Twin die)
PCH internal RTC configuration
No XDP port available
Standard clock option
Standard EARTH connection
Battery super cap not present
DPB port set to DP mode
Internal magnetic power option off
Front panel RS-422 option off
MLVDS terminations off
CPLD I2C equipment: one F-RAM, one EEPROM
1 inch VPX pitch
One VPX slot version for SA
P80 debug option available on debug connector
VS2 PWR pins not used
No PCIe switch dual clocking mode
Normal BOM generation
PCB plating: Immersion Tin
Lead Free Process
All components Lead Free
```

Note: Do not pay attention to the message `ERROR: PCB B not defined`. It will be fixed in next release.

7.4 LEDs

The driver `leds_cpld` allows to deal with the front panel LEDs for a user mode.

The driver `leds_cpld` creates a list of special files and classes in `/sys`:

```
[root@ki7]# ls /sys/class/leds/
led2:amber  led2:OFF  led3:amber  led3:OFF
led2:green  led2:red  led3:green  led3:red
```

For each LED, there are three different colors available: green, red, amber that can be set by addressing the related file.

For each color (which are exclusive), there are four different modes:

- ON (echo 0, see following example)
- low blinking (echo 1, see following example)
- fast blinking (echo 2, see following example)
- OFF (echo 1, see following example)

Example, to set those different modes on the LED 2 in amber:

```
[root@ki7]# echo 0 > led2:amber/brightness The LED 2 is ON in AMBER
[root@ki7]# echo 1 > led2:amber/brightness The LED 2 blinks low in AMBER
[root@ki7]# echo 2 > led2:amber/brightness The LED 2 blinks fast in AMBER
[root@ki7]# echo 0 > led2:OFF/brightness The LED 2 is OFF
```

When setting led2 or led3, both led2 and led3 are set to user mode.

To exit from this mode, set led<n>:OFF/brightness to a non null value; for example

```
[root@ki7]# echo 1 > led2:OFF/brightness
```

> VX304x: On the front panel, the name of the led1 is L1, led2 is L2 and led3 is L3.

The led1 (L1 and L4) is not manageable at user level, so no special file for led1 is available.

7.5 Multinodes Diskless

Multinodes Diskless is a powerfull set of python scripts to help the configuration of diskless root filesystems and to boot distant nodes through the Ethernet network (PXE).

The main command is `diskless`:

```
[root@ki7] diskless --help
```

There are two xml files to configure the diskless system:

`/etc/diskless/images.xml`: describes the layers, the nodes and main configuration setup.

`/etc/diskless/diskless_fs.xml`: lists the rpms used to build the root filesystem of the distant nodes;

The user manual is available at: `/usr/local/share/doc/diskless/pdf/manual.pdf`

If diskless is not needed, you can disable this service at boot time :

```
[root@ki7]# chkconfig diskless off
```

7.5.1 Introduction

In case you have several Ki7 boards into your machine, you can use one of the board as a diskless server.

Most of the time, it will be the first board present into a rack but it is not mandatory. This one must be equipped with a SATA hard disk (or SATA SSD).

Of course, you won't need any disk media present on the other targets providing that the targets will boot on the diskless server through the network (that is the point).

The package multinodes-diskless is available to setup and configure a NFS exportable root filesystem and to get a kernel and a initrd bootable through PXE.

The diskless product is uncompatible with SELINUX enforcing policy as well as iptables services.

So run the following commands in order to disable these services:

- Stop the iptables service and disable the SELINUX by running the commands:

```
[root@ki7]# chkconfig iptables off
[root@ki7]# service iptables stop
```

- Set the variable SELINUX=disabled into the file: `/etc/selinux/config`

The following steps describe how to initialize a diskless configuration and how to boot it on a given number of targets.

7.5.2 Get the whole Fedora 16 Distribution on a Local Repository

The very first step is to get a local repository on the diskless server of the Fedora 16 RPM packages.

To reach this purpose, insert the Fedora 16 DVD and run these commands :

```
[root@ki7]# mkdir -p /dist/Fedora_16_ki7_boards/{BSP,RPMS}
[root@ki7]# mount /dev/dvd /mnt
[root@ki7]# cp /mnt/Packages/* /dist/Fedora_16_ki7_boards/RPMS
[root@ki7]# cd /dist/Fedora_16_ki7_boards
[root@ki7]# createrepo .
```

Then code the right value into the following field of the /etc/diskless/diskless_fs.xml:

```
<baseurl>file:///dist/Fedora_16_ki7_boards/</baseurl>
```

In this way the diskless tool will know the location of the Fedora 16 RPMs just installed.

The next step is to build a diskless root file system based on these Fedora 16 RPMs.

7.5.3 Build the diskless RFS (root file system)

Force the layer building even if it already exists by running:

```
# diskless -B
```

The command will take few minutes to complete.

The content of the RFS is defined into the section named <rpms> present into the xml file /etc/diskless_fs.xml

If you want change the content of the RFS, for example in order to add your own package, simply add the package name to the list.

Be careful when you change one of the xml files : /etc/diskless/images.xml or /etc/diskless/diskless_fs.xml . It is suggested to preserve the initial version into an another file.

You should notice the presence of two layers under /diskless/layers

```
# ls /diskless/layers
fedora_base fedora_config
```

with

- fedora_base: Containing the reference root file system.
- fedora_config: A customizable layer you can change as you want.

For example, if you want to use your own /etc/hosts file, simply install it under /diskless/layers/fedora_config by running:

```
# cp /etc/hosts /diskless/layers/fedora_config
```

In this way, the target(s) after booting throught the network will "see" your own /etc/hosts file instead of the reference one (in fact, from the diskless client point of view, the layer fedora_config overloads the underneath fedora_base layer).

The RFS is now ready. The next step is to configure the DHCP server as well as the PXE boot.

7.5.4 Configure the DHCP and PXE Services in MAC Ethernet Address Mode

This step is mandatory to authorize remote targets to boot over the LAN on the diskless server.

Two modes are available for this purpose:

- MAC Ethernet mode: in this mode the IP address delivered by the DHCP server depends on the MAC Ethernet address of the target board.
- GEOID mode: in this mode the IP address delivered by the DHCP server depends on the location of the target board into the rack.

This section presents the MAC Ethernet adress mode. The next one presents the GEOID Ethernet addres mode. Of course both of them are exclusive.

At first, uncomment if needed the XML structure called `<nodemap id="MAC_ADDR-diskless-cluster"..>` into the file `/etc/diskless/image.xml` and comment the one related to the GEOID mode.

Then code the right MAC Ethernet address (related to the network interface used during the boot step) into this field:

```
<match criterium="macaddress">
<value>00:00:xx:xx:xx:xx</value>
</match>
```

From the server side, configure the network interface `em1` (for example) with the right address with the command:

```
 # ifconfig em1 192.168.1.1
```

If you want to use a different network base address, simply code this one into the field `<ip>192.168.1.%(1+order.value)</ip>` of the structure `<nodemap>`.

Start the required service `tftpd` by running the commands:

```
 [root@ki7]# chkconfig tftp on
```

Finally run the command:

```
 # diskless -x
```

Verify the right files have been created under `/diskless/tftp`:

```
# cd /diskless/tftp
# find .
```

The created files should be:

```
./pxelinux.0
./pxelinux.cfg
./pxelinux.cfg/default
```

A another impact of the `diskless -x` command is the starting of the DHCPD daemon.

So verify all worked well in this area by running:

```
 # service dhcp status
```

The output should include: `active` (running)

At this step, unless you want to configure diskless in GEOID mode, go directly to the section 7.5.6 - Export the Root File System.

7.5.5 Configure the DHCP and PXE Services in GEOID Ethernet Address Mode

An alternative to the MAC address mode is the GEOID mode. In this mode the IP address delivered by the DHCP server depends on the location of the target board into the rack.

This mode is interesting in the sense that it allows to replace a board by another one without changing the configuration file: `image.xml` (because the MAC Ethernet address, specific to each board, is not coded into the image.xml file in this case).

Indeed, in this case, it is required to code the GEOID range into the `image.xml` file instead of the explicit MAC addresses related to the boards plugged into the rack.

So edit the file `/etc/diskless/image.xml`, uncomment the right XML structure called `<nodemap id="GEOID-diskless-cluster"..>`... and comment the one related to the MAC-ADDR mode.

Then, code the right GEOID value into this file (field `<match criterium="geoid">`).

In order to take into account the previous change, run the following command:

```
# diskless -x
```

At this step, the DHCPD daemon should be ready as well as the TFTPD one.

7.5.6 Export the Root File System

The diskless RFS based on aufs (which authorized to have different stackable layers) must be exported now.

To reach this purpose, use this command:

```
# diskless -e
```

Finally start the NFSD service with the command:

```
# service nfs-server start
```

and verify the status is correct.

```
# service nfs-server status | grep Active
```

At this step, the remote target should be ready to be booted in diskless MAC-Ethernet mode.

Reboot the server to take into account the previous changes.

7.5.7 Boot the Target through the Network

On the diskless target and after a power-on, take the hand under CPUB by pressing <F2>.

Under BIOS Setup, select successively the menus: **Chipset ->South Bridge -> Ibex Peak Option** and enable the PXE boot (example given on VM6050).

Save the changes and after reset, press again **<F2>**.

Select the menu "**boot**" and put the network device at the top of the list.

Save the change and exit. As a result the PXE boot will happen and the following should be displayed:

```
CLIENT MAC ADDR: 00 00 DE 40 36 F9
GUID: 00020003 0004 0005 0006
000700080009
DHCP.|
```

You should notice the occurence of a login at the end of the boot step.

Log as root and verify the system is operational.

7.5.8 Diskless Service

In case you want to make automatic the launching of the diskless at the boot step, simply enable the related service by running :

```
chkconfig diskless on
```

In this case, the next time you will boot the server, the diskless service `diskless` will be launched.

As result the commands `diskless -x` and `diskless -e` will be run automatically, authorizing the diskless boot of the target.

7.5.9 Initiate a Reset of the Remote Targets

To reach this purpose:

```
# diskless -aR
```

Should cause a reset of all the remote targets.

7.6 Sysvartool

To get the report of the PBIT (Power Built In Test), run the command:

```
[root@ki7]# sysvartool -A pbit -l
VX3O4x detected
 area = 2, arch = 2
POSTs configured to run from command line:
       mem_data: PASSED
       mem_addr: PASSED
   mem_pattern1: PASSED
   mem_pattern2: PASSED
   mem_pattern3: PASSED
   mem_pattern4: PASSED
    ether_loop0: PASSED (FAILED ONCE)
    ether_loop1: PASSED (FAILED ONCE)
    ether_loop2: PASSED
         system: PASSED

PASSED  : 10
FAILED  : 0
NOT RUN : 0
TOTAL   : 10

POSTs configured to run automatically from RAM:

PASSED  : 0
FAILED  : 0
NOT RUN : 0
TOTAL   : 0

POSTs configured to run automatically from ROM:

PASSED  : 0
FAILED  : 0
NOT RUN : 0
TOTAL   : 0
```

7.7 GPIOs

There are 5 GPIOs on Ki7 boards. When the GPIO driver is loaded, a collection of special sys files is created for each GPIO:

`/sys/class/cpld-gpio/gpioX/direction` should be set to `in` or `out`:

```
[root@ki7]# echo out > /sys/class/cpld-gpio/gpio4/direction
[root@ki7]# cat /sys/class/cpld-gpio/gpio4/direction
out
```

`/sys/class/cpld-gpio/gpioX/value` to get or set (depending on the direction) the value of the GPIO: `1` or `0`

```
[root@ki7]# echo 0 > /sys/class/cpld-gpio/gpio4/value
[root@ki7]# cat /sys/class/cpld-gpio/gpio3/value
1
```

`/sys/class/cpld-gpio/gpioX/value_it` get or set the value. A read to this file is waiting for an interruption on the GPIOs (refer below).

`/sys/class/cpld-gpio/gpioX/interrupt` enables or disables interruption on the GPIO: `on` or `off`

```
[root@ki7]# echo on > /sys/class/cpld-gpio/gpio3/interrupt
```

`/sys/class/cpld-gpio/gpioX/polarity` to set the polarity of the interruption: `hi` (high level or rising edge) or `lo` (low level or falling edge).

`/sys/class/cpld-gpio/gpioX/mode`: `edge` or `level` interruption mode.

```
[root@ki7]# cat /sys/class/cpld-gpio/gpio5/mode
edge
```

`/sys/class/cpld-gpio/gpio3/toggle` set toggle mode means the interruption is raised when the GPIOs state changes.

```
[root@ki7]# echo on > /sys/class/cpld-gpio/gpio3/toggle
```

`/sys/class/cpld-gpio/gpio3/intr_stat` is a counter of the interruption on that GPIOs.

```
[root@ki7]# cat /sys/class/cpld-gpio/gpio5/intr_stat
12
```

The GPIO2 is dual function because it is also a Masquable Reset. This is configurable through CPLD registers:

To set the GPIO2 as a standard GPIO (without issuing a local reset):

- on a VPX board :

```
[root@ki7]# cpldtool -f VPX_CONTROL MSKR2LOC 0
```

- on a VME board :

```
[root@ki7]# cpldtool -f VME_CONTROL MSKR2LOC 0
```

- or on all boards :

```
[root@ki7]# cpldtool -f 0x70 MSKR2LOC 0
```

This can also be done with the lower level command (but take care not to change some other bits) :

```
[root@ki7]# port 0x870
@0x870 = 0x1d
[root@ki7]# port 0x870 0xd
@0x870 <- 0x0d
[root@ki7]# port 0x870
@0x870 = 0x0d
```

7.8 cpldtool

`cpldtool` may be used to get some high level information about the setting of the CPLD.

This tool also has a write mode to change the settings, but this mode should not be used, except for a very small set of settings and with a lot of caution.

```
# cpldtool -a
Reg 0x0 - CPLD_ID = 0x05
              CPLD_ID=0x0
              CPLD_Debug=0x0
              CPLD_Version=0x5
Reg 0x1 - PCB_ID = 0x44
Reg 0x3 - PWON_STATUS = 0x00
              DPOST=0x0
              POST_RTC=0x0
Reg 0x4 - PWR_RST_CONFIG = 0x0E
              PWRON_MODE=0x0
              PLTRST_to_PERST3U_INHIB=0x0
              Alarm_inhib=0x0
              PLTRST_to_PERST6U_INHIB=0x0
              Reset_Mode=0x3
              PERSTb_Control=0x1
              Software_Cross_Reset=0x0
Reg 0x5 - VPD_BC1 = 0x00
Reg 0x6 - VPD_BC2 = 0x00
Reg 0x9 - MEM_PROTECT = 0x00
              Boot_flash_CS_swap_DIP=0x0
              Boot_flash_CS_swap_Valid#=0x0
              Boot_both_flash=0x0
              Onboard SSD hardware protection=0x0
              USER_WP=0x0
              SYS_WP=0x0
              VPD_WP=0x0
              NVMRO=0x0
Reg 0xa - COM_TX = 0x00
              Semaphore=0x0
              Data=0x0
Reg 0xb - COM_RX = 0x00
Reg 0xc - BOARD_CONFIGURATION_VX3040 = 0x00
Reg 0xd - PCI_MODE_VX3040 = 0xD2
              DRAM_CHANNEL=0x1
              DRAM_RANK=0x1
              DRAM_SIZE=0x1
              XMC_PRESENT=0x0
Reg 0xe - DIP_SWITCH_STATE_VX3040 = 0x00
              DBG_MD=0x0
              BKPPCI_CFG0=0x0
              XPCIE_MD=0x0
              RPCIE_MD=0x0
              BKPPCI_CFG1=0x0
              TDP_MD=0x0
              FACTORY_MODE=0x0
```

```
Reg 0xf - SERIRQ_CONTROL = 0xE0
              SMB_MSK=0x1
              SMB_TEMP_MSK=0x1
              TIP_MSK=0x1
              GPIOs_INT=0x0
              SMB_INT=0x0
              SMB_TEMP_INT=0x0
              TIP_INT=0x0
              WDG_INT=0x0
Reg 0x5b - ALERT_STATUS = 0x08
              CATERR=0x0
              THERMTRIP=0x0
              THERM_PROT=0x0
              PROCHOT=0x0
              TEMP_ALERT=0x1
              6U_PCIe_ERR=0x0
              3U_PCIe_ERR=0x0
              DDR_Throttling=0x0
Reg 0x5b - ALERT_STATUS = 0x08
              CATERR=0x0
              THERMTRIP=0x0
              THERM_PROT=0x0
              PROCHOT=0x0
              TEMP_ALERT=0x1
              PCIE_ERR=0x0
Reg 0x6a - GEO_ADD = 0x3E
              SYSCON=0x0
              Error=0x0
              GAP=0x1
              GA=0x1E
Reg 0x70 - VPX_CONTROL = 0xDE
              GDiscrete1_Ctl=0x1
              SMB_Alert_Status=0x1
              MSKR2LOC=0x1
              VPX_RST=0x1
              VPX2LOC=0x1
              LOC2VPX=0x0
Reg 0x71 - PCIe_SWITCH_VPX = 0x10
              VPX_REFCLK=0x0
              VPX_AUXCLK=0x0
Reg 0x72 - I2C_BOARD_STATUS = 0xF4
              Power_Status=0x1
              Reset_Source=0x3
              Reset_Status=0x1
              Boot_Status=0x4
Reg 0x73 - I2C_BOARD_CONTROL = 0x81
              Board_Id=0x8
              Check_Errors=0x0
              Reset_3UB=0x0
              Reset_3UA=0x0
              Power_OnOff=0x1
```

7.9 I2C Busses

The `cpld_i2c` driver is supporting the local i2c bus (I2C bus number 22) which is local to the board and used to address sensors and system eeproms.

There are also two other I2C busses routed to the backplane; the bus numbers are 23 and 24:

- `/dev/i2c-23` is the SMB bus connecting it to the I2C devices in the chassis (if any).
- `/dev/i2c-24` is the IPMB bus connecting all boards through the backplane

By default, those two busses are available only on the system controller.

It is possible to enable those busses on other slots by adding the option `force_i2c_extend=1` at the load of the module "`cpld`". Add this argument "`cpld.force_i2c_extend=1`" at the kernel command line to enable this option.

7.10 BIOS Update

The `flashrom` package provides a set of commands and scripts to update the BIOS of the boards.

To update the BIOS of the Ki7 boards, use the command `ki7updbios`:

Help command:

```
[root@ki7]# ki7updbios -h
usage:

/usr/local/sbin/ki7updbios <options> <bios image>
Script to update the BIOS on boards using flashrom tools.\n
Options:\n
-r Read the current BIOS in bios image file.\n
-w Update the BIOS image using the given bios image file.\n
-a Update the all BIOS eeprom.\n
[root@ki7]#
```

Read the current BIOS:

```
[root@ki7]# ki7updbios -r BIOS.bin
flashrom v0.9.4 on Linux 2.6.32.14-11035.vx304x.fc12.i686.PAE (i686), built with libpci 3.1.7,
GCC 4.4.2 20091027 (Red Hat 4.4.2-7), little endian
flashrom is free software, get the source code at http://www.flashrom.org

Calibrating delay loop... OK.
No coreboot table found.
Found chipset "Intel QM57", enabling flash write... OK.
This chipset supports the following protocols: FWH,SPI.
Found chip "SST SST25VF032B" (4096 KB, SPI) at physical address 0xffc00000.
Reading flash... done.
[root@ki7]# ls -ltr BIOS.bin
-rw-r--r-- 1 root root 4194304 2011-03-03 15:54 BIOS.bin
```

Write a new BIOS file:

```
[root@ki7]# ki7updbios -w vx304x_ID12xxx.bin
flashrom v0.9.4 on Linux 2.6.32.14-11035.vx304x.fc12.i686.PAE (i686), built with libpci 3.1.7,
GCC 4.4.2 20091027 (Red Hat 4.4.2-7), little endian
flashrom is free software, get the source code at http://www.flashrom.org

Looking for "bios"... found.
Calibrating delay loop... OK.
No coreboot table found.
Found chipset "Intel QM57", enabling flash write... OK.
This chipset supports the following protocols: FWH,SPI.
Found chip "SST SST25VF032B" (4096 KB, SPI) at physical address 0xffc00000.
Flash image seems to be a legacy BIOS. Disabling checks.
Erasing and writing flash chip... Done.
Verifying flash... VERIFIED.
```

Take care that if the system is booted from the RESCUE BIOS flash (refer to the BIOS user manual), the update of the BIOS using the **ki7updbios** command will update the RESCUE BIOS Flash which is not recommended.

The update of the BIOS through the **ki7updbios** command does not preserve the setup parameters of the BIOS. For deployement of a BIOS version with its own setup parameters, preset one board with the desired parameters, backup this using **ki7updbios -r**, and use this version of BIOS+ setup to be deployed on other similar boards.

7.11 FMRAM Example

The `fmram` packages gives a example of how to access to the FMRAM device which may be used to save some customer data which need to be backuped.

```
[root@ki7]# fmram -h
```

Usage `fmram` [options]
a tool to read or write the FerroMagnetic RAM

Options are :

`-h`	: this help
`-r <value>`	: read at offset <value> (default 0)
`-w <value>`	: write at offset <value> (default 0)
`-s <value>`	: data size to read or write (default 4)
`-f <filename>`	: file name used to store (optional) or read (mandatory) data

Example:
`fmram -w 0x10 -s 0x11 -f data_file` :
store 17 bytes read from data_file to the ferromagnetic RAM at offset 16

With src package, the C code of this command is delivered as an example.

Chapter 8 - Specificities Related to the RC Boards

8.1 How to Manage to the Lack of RTC Battery

Some hardware constraints may prevent the use of the RTC's battery on the boards. This implies to workaround or disable some standard behaviour of Fedora services. The very first one is "`fsck`" which is run at boot time. `fsck` checks that the date of the last mount of the checked partition is prior than the current date. If not, the boot is stopped in a maintenance mode and this could happen after few times with the power off and without the RTC's battery. To workaround this, simply create a file named `/etc/e2fsck.conf` with the following content:

```
 [option]
broken_system_clock=true
```

The same file has to be added also into the initramfs too. In order to reach this goal, simply run the command:

```
# dracut --force
```

Without the right date for the system, some other services and functions may present dysfunctions. For example, the command "`make`" warns if there is some previous build detected in the future and this could lead to bad recompilation. If the network is available, it is recommended to setup the service `ntpd` to adjust the system date with a date server.

8.2 External Devices Connection

One of the impact of the RC board is the missing of front panel.

As a consequence, the devices required to access to the system such as:

- Serial line console.
- USB mouse
- USB keyboard
- Display monitor (through the mini DP or VGA)
- Ethernet

will be plugged from the rear side through a Rear Transition Module (RTM) Paddle board.

The Kontron PB-VX3-4xx is a 3U VPX Rear Transition Module compliant with the definition of the Rear Transition Module on VPX Standard –VITA 46.10.

It provides rear I/O peripherals connectivity for Kontron VX30xx Single Board Computers.

The Kontron VM6050-RTM (Order Code: PBV36-P0-VM6-00) rear transition module is compliant to PMC I/O Module Standard VITA 36 - 199x Draft 0.1 July 19, 1999 (mechanical and PIM format) and is available for VM6050 boards.

For example, in order to perform a graphical linux installation stage on a VX3040-RC, plug the required devices to the paddle board:

- USB mouse and USB keyboard through a USB HUB.
- The display monitor (through the mini DP).

After turning on the system, you should notice all devices have been recognized correctly allowing to take control of the graphical environment.

8.3 RC Specifications

The RC version of the boards is designed to work in different environmental constraints and for different levels of temperature and power dissipation.

This may implies specific setups (cpu frequencies, hardware parameters,…) to guaranty the specified behavior.

Those setup are described in the “Hardware User’s Guide” of each boards. Please check this document for your board and for the required environment.

Chapter 9 - Power Management

9.1 Introduction

At the core of power management is an understanding of how to effectively optimize energy consumption of each system component.

By studying the different tasks that your system performs, and configuring each component to ensure that its performance is just sufficient for the job, you can save energy, generate less heat.

Many of the principles for analysis and tuning of a system in regard to power consumption are similar to those for performance tuning.

To some degree, power management and performance tuning are opposite approaches to system configuration, because systems are usually optimized either towards performance or power.

Two types of tools are available:

> The ones which allow to set-up the power management configuration:
  > BIOS menu and `cpufreq` (under linux).
> The ones which allow to evaluate the impact of this setting on the system:
  > `powertop` under linux.

9.2 Power Management Setting

9.2.1 Under BIOS

A way to reduce the power consuming (the drawback beeing the decreasing of the performance) is to disable the **Hyper-Threading** mode as well as the **Turbo-Mode**.

To disable the **Hyper-Threading**:

> Take the hand under bios and select the submenu **Advanced -> CPU Configuration**
> Then set the option **hyper threading** to **Disabled**.

To disable the Turbo Mode.

Take the hand under bios and select the submenu **Advanced -> CPU PPM Configuration**

Then set the option Turbo mode to Disabled.

An alternative is to use the BIOS TDP menu:

> This option is specific to the VX3044 board.

The thermal design power (TDP), sometimes called thermal design point, refers to the maximum amount of power the cooling system in a computer is required to dissipate.

The move to a 22nm process and Tri-Gate transistors alone should already account for some pretty significant power savings. But there are a few other changes in Ivy Bridge meant to optimize power consumption.

An important addition brought to mobile Ivy Bridge processors is the inclusion of a configurable TDP that allows them to switch between three different ratings:

nominal, a lower configurable TDP and an upper configurable TDP.

The lower configurable TDP implies the lowest power comsumption of course.The BIOS TDP menu is accessible under `enhanced/CPU Configuration/CPU PPM` Configuration.

9.2.2 Under Linux

The main power management tool available under linux is CPUFreq.

CPUfreq allows the clock speed of the processor to be adjusted on the fly.

This enables the system to run at a reduced clock speed to save power.

Different types of CPUfreq governors are available:

```
cpufreq_performance
```

The Powersave governor forces the CPU to use the highest possible clock frequency (no power saving benefit at all).

```
cpufreq_powersave
```

By contrast, the Powersave governor forces the CPU to use the lowest possible clock frequency.

```
cpufreq_ondemand
```

The ondemand governor is a dynamic governor that allows the CPU to achieve maximum clock frequency when system load is high and also minimum clock frequency when system is idle.

This is the default mode (best compromize between heat emission,power consumption,performance and manageability).

```
cpufreq_userspace
```

Allows userspace program to set the frequency. Used normally in conjunction with the cpuspeed daemon.

```
cpufreq_conservative
```

Similar to the cpufreq_ondemand mode but this mode switches between frequencies more gradually. Boot the target under linux and log as root.

At first, you can view which governor the system is currently using with the command:

```
 #cat /sys/devices/system/cpu/cpu0/cpufreq/scaling_governor
```

The result will be probably:

ondemand because this is the default value.

You can also view which governors are available for the CPUs:

```
#cat/sys/devices/system/cpu/cpu0/cpufreq/scaling_available_go
```

At this step, select the governor mode wich fit the best your need with the command:

```
 #echo xxxx     > /sys/devices/system/cpu/cpu0/cpufreq/scaling_governor
```

Control the result by:

```
#cat /sys/devices/system/cpu/cpu0/cpufreq/scaling_governor
```

An alternative to that consists of selecting an explicit frequency by using a command such as:

```
 # cd /sys/devices/system/cpu
# for i in 0 1 2 3
do ; cd cpu$i ; echo 1300000 > cpufreq/scaling_max_freq; cd ..; done
```

In this particular case, setting of the frequency 1.3 GHZ for all the CPUs. If you want to know the set of available frequencies:

```
# cd /sys/devices/system/cpu
# for i in 0 1 2 3
do ; cd cpu$i ; cat cpufreq/scaling_available_frequencies; cd ..; done
```

9.3 Impact of the Power Management Policy

Boot linux and login as root. The powertop tool identifies specific components of kernel and use-space applications that frequently wake up the CPU. At first, install `powertop` package with the command:

```
# yum install powertop
```

Next run `powertop` with the command

```
# powertop
```

while it runs, `powertop` gathers statistics from the system.

9.3.1 Hyper-Threading + Turbo Mode Enabled

The purpose of this configuration is to obtain as much as performance as possible. This is the default configuration (Except on VX3040RC board which implies to disable the turbo mode in order to fit the maximal temperature constraint).

Run PowerTop with the command:

```
# powertop
```

And move to the `Idle Stats` option.

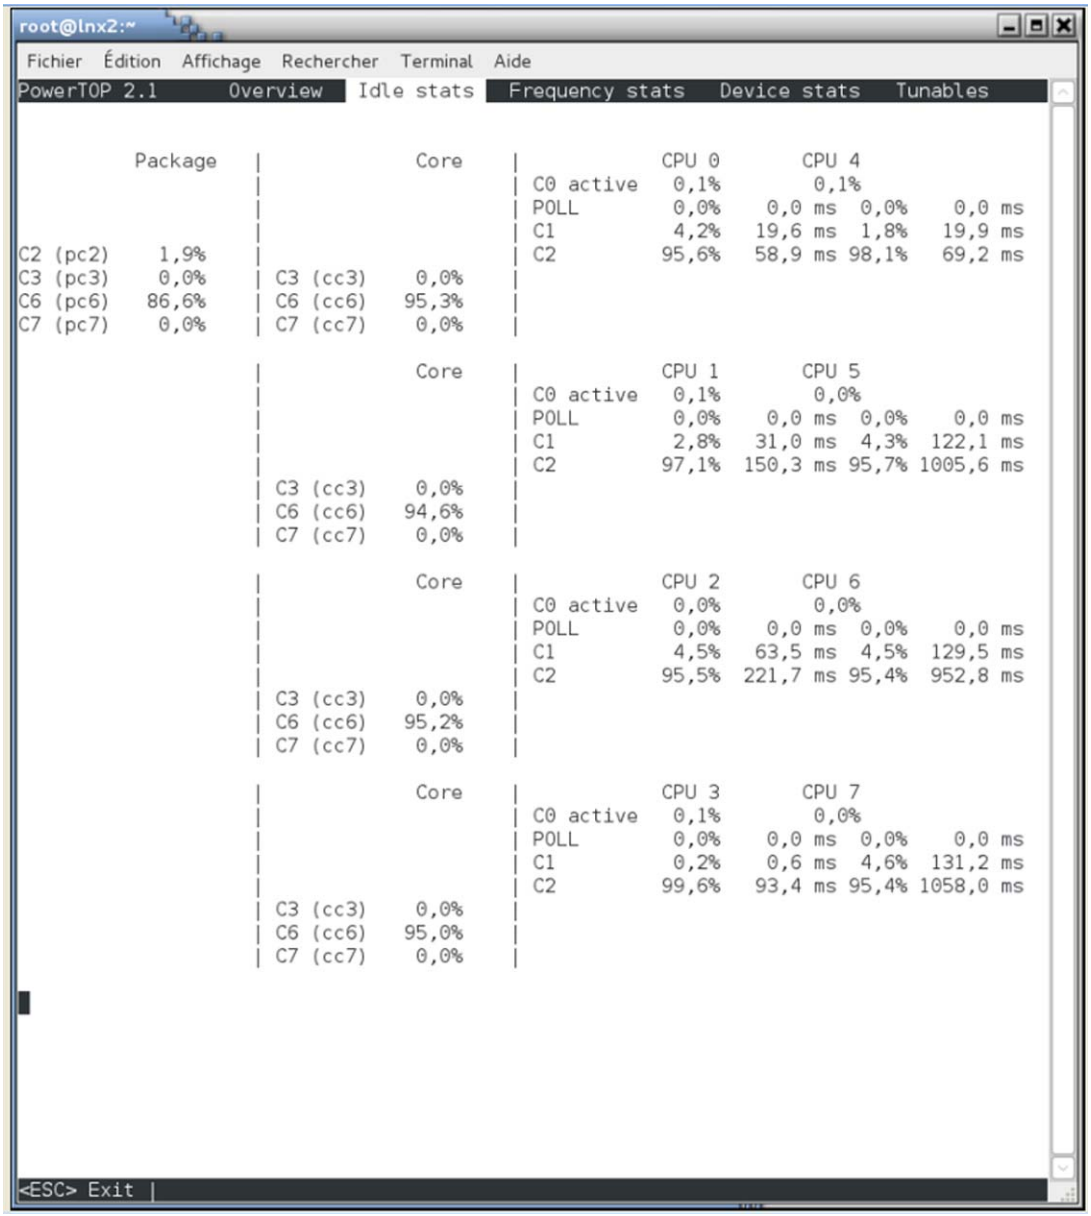
root@lnx2:~
Fichier Édition Affichage Rechercher Terminal Aide
PowerTOP 2.1 Overview Idle stats Frequency stats Device stats Tunables
Package
C2 (pc2) 1,9%
C3 (pc3) 0,0%
C6 (pc6) 86,6%
C7 (pc7) 0,0%
Core
C3 (cc3) 0,0%
C6 (cc6) 95,3%
C7 (cc7) 0,0%
CPU 0
C0 active 0,1%
POLL 0,0% 0,0 ms
C1 4,2% 19,6 ms
C2 95,6% 58,9 ms
CPU 4
0,1%
0,0% 0,0 ms
1,8% 19,9 ms
98,1% 69,2 ms
Core
C3 (cc3) 0,0%
C6 (cc6) 94,6%
C7 (cc7) 0,0%
CPU 1
C0 active 0,1%
POLL 0,0% 0,0 ms
C1 2,8% 31,0 ms
C2 97,1% 150,3 ms
CPU 5
0,0%
0,0% 0,0 ms
4,3% 122,1 ms
95,7% 1005,6 ms
Core
C3 (cc3) 0,0%
C6 (cc6) 95,2%
C7 (cc7) 0,0%
CPU 2
C0 active 0,0%
POLL 0,0% 0,0 ms
C1 4,5% 63,5 ms
C2 95,5% 221,7 ms
CPU 6
0,0%
0,0% 0,0 ms
4,5% 129,5 ms
95,4% 952,8 ms
Core
C3 (cc3) 0,0%
C6 (cc6) 95,0%
C7 (cc7) 0,0%
CPU 3
C0 active 0,1%
POLL 0,0% 0,0 ms
C1 0,2% 0,6 ms
C2 99,6% 93,4 ms
CPU 7
0,0%
0,0% 0,0 ms
4,6% 131,2 ms
95,4% 1058,0 ms
<ESC> Exit

The different C-states of the CPU are defined as follows:

> C0

The operating of running state.

The CPU is working and no `idle` at all.

> C1, Halt

A state where the processor is not executing any instruction but is typically not in a lower power state.

> C2 Stop clock

A state where the clock is frozen for the processor but it keeps the complete state for its registers and caches.

> C3 sleep

A state where the processor really goes to sleep.

> C6 (specific to Nehalem architecture) the power supply of the CPU can be reduced up to 0.

Then, run the following process in order to overload the CPUs:

```
# while true; do rm -rf squashfs-root ; unsquashfs squashfs.img; done &
```

Analyse the impact on the `Idle Stats`:

`Run Powertop` with the command:

```
# powertop
```

And move to the `Idle Stats` option.

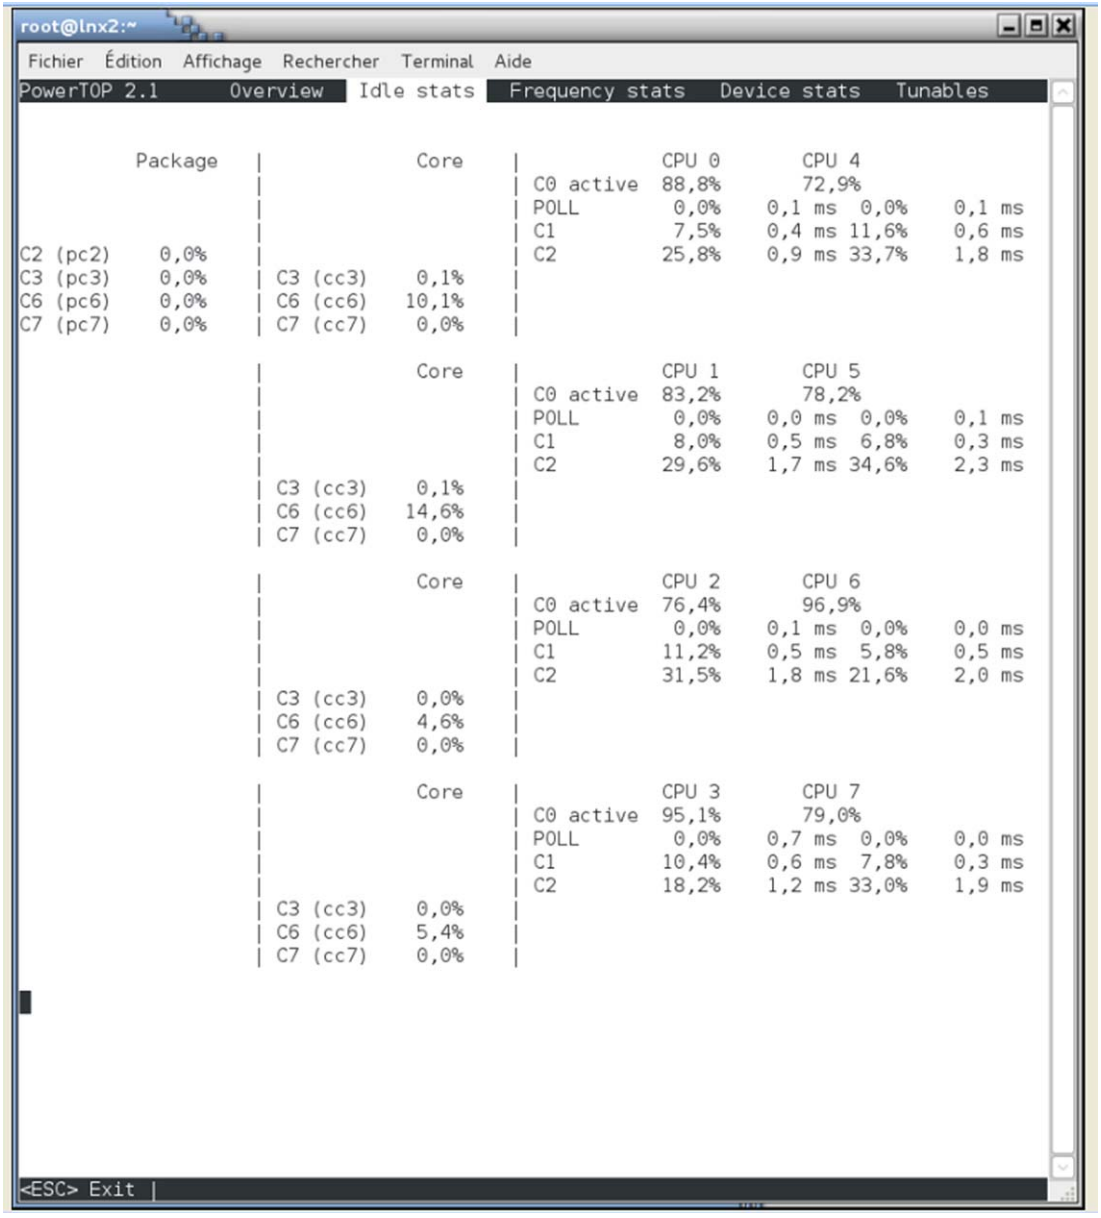

Notice the CPU passed most of the time in the C0 state revealing the CPU is very busy.

Analyse the frequencies of the different CPUs with the command:

```
# powertop
```

And move to the `Frequencies Stats` option.

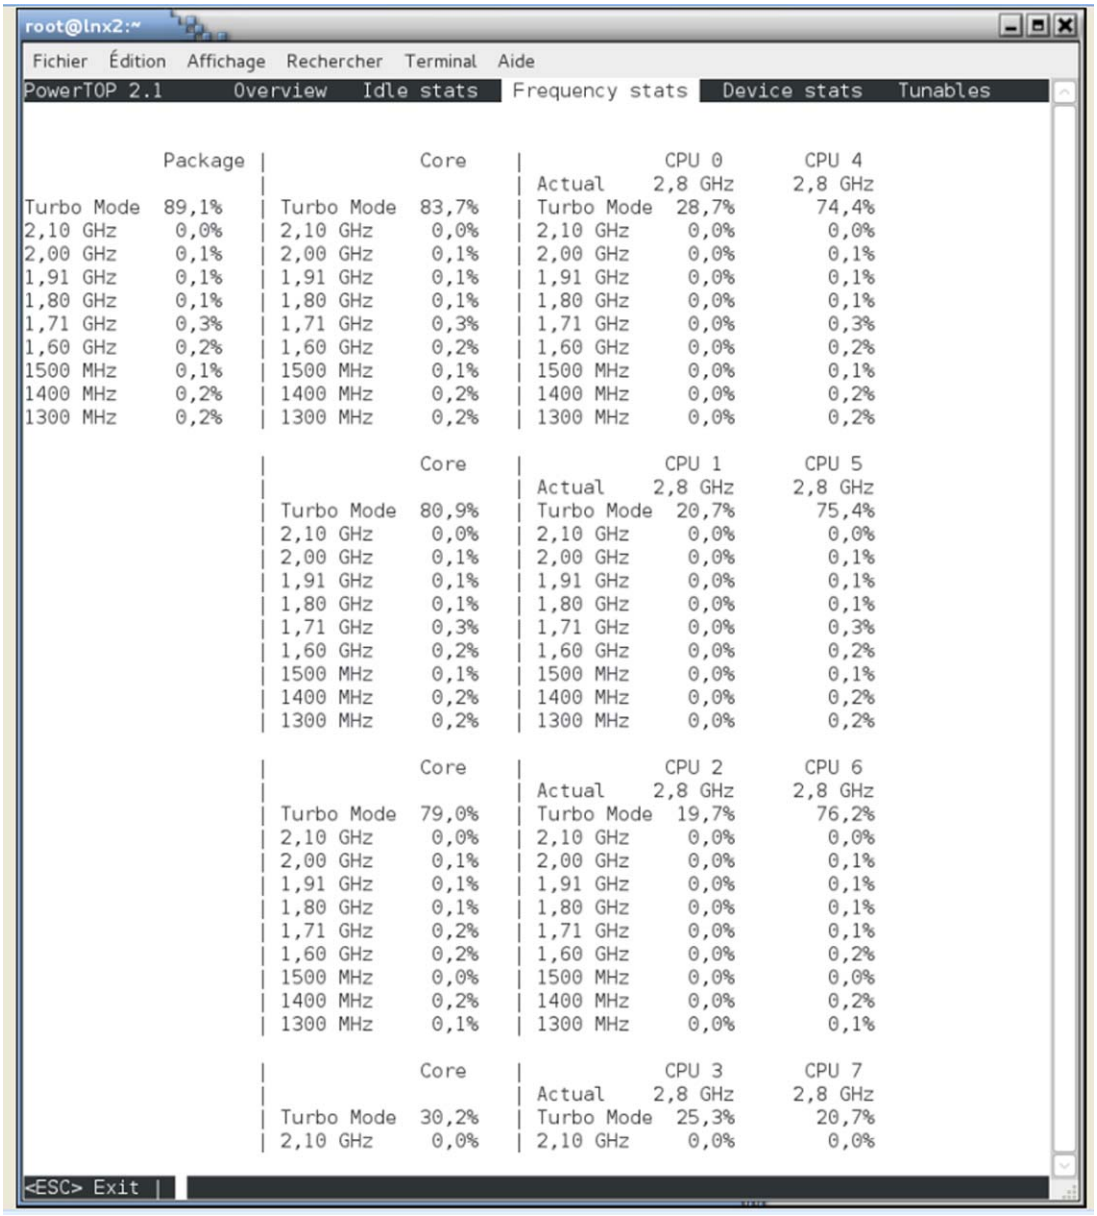

LOG related to the `powertop Freq Stats` in turbo mode with load

Notice the frequencies remains at the top level around 2.8 GHZ.

9.3.2 Turbo mode disabled + CpuFreq policy=Powersave

Select the `CpuFreq powersave` mode by running:

```
# cd /sys/devices/system/cpu
# for i in 0 1 2 3 4 5 6 7
do ; cd cpu$i ; echo powersave > cpufreq/scaling_governor; cd ..; done
```

Then run this process in order to overload the CPU:

```
# while true; do rm -rf squashfs-root ; unsquashfs squashfs.img; done &
```

Notice the impact using the `powertop` tool.

```
# powertop
```

And move to the `Idle Stats` option.

The result should look like:

Then move to the `Frequency Stats` option.

9.3.3 TDP LOW mode (low power comsumption oriented)

Select the appropriate option in the BIOS TDP menu.

Then run this loop in order to load the system as much as possible:

```
# while true; do rm -rf squashfs-root ; unsquashfs squashfs.img; done &
```

And analyse the result with `powertop`

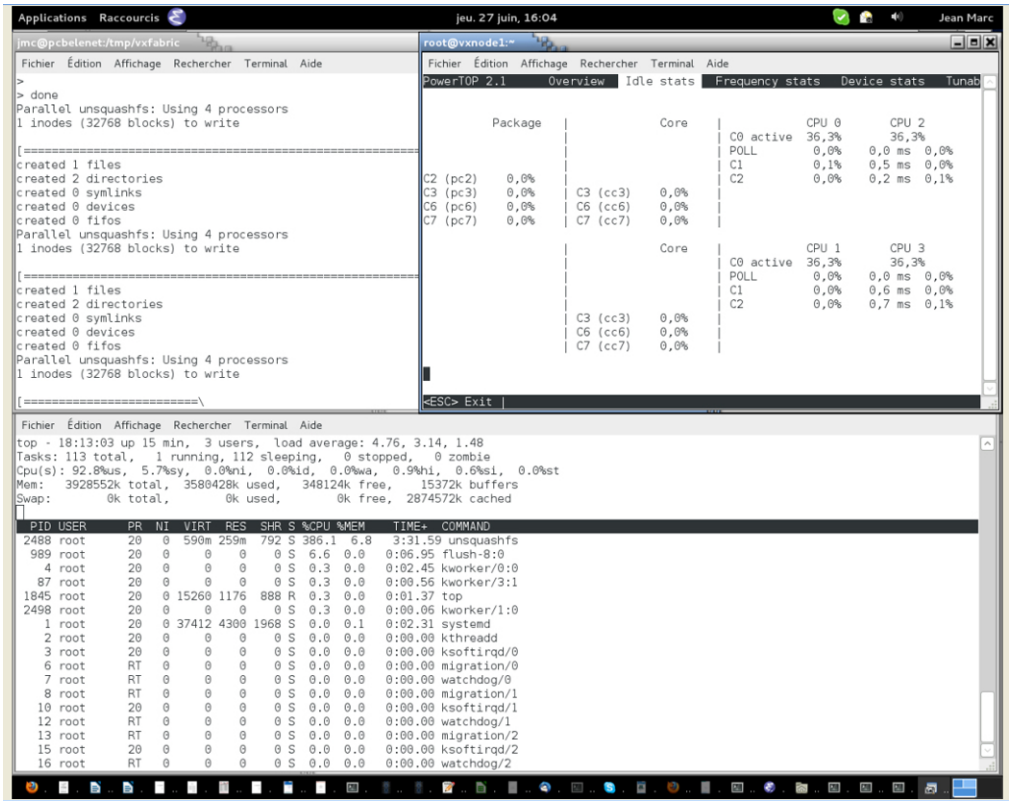

The CPU is used only at 36 % of its capacity.

Chapter 10 - Additional Information

10.1 Known Limitations

SUSPEND Mode not supported

The SUSPEND mode is not supported by the current graphic chipset hardware.

So, follow this procedure in order to disable it:

- Create a new configuration file under `/etc/polkit-1/localauthority`

```
# vi /etc/polkit-1/localauthority/50-local.d/50-disable-suspend.pkla
```

- Copy the following code into that file:

```
[Disable Suspend]
Identity=unix-user:*
Action=org.freedesktop.upower.hibernate;org.freedesktop.upower.suspend
ResultAny=no
ResultInactive=no
ResultActive=no
```

- Finally check that the permission was successfully revoked

```
$ pkcheck --action-id org.freedesktop.upower.suspend --process $$

Not authorized.
```

If you get no output from the `pkcheck` command (make sure to run it as a normal user, not root) the permission is still there.

That is it, after re-login the annoying Suspend will be replaced by a friendly Power Off.

MAILING ADDRESS
Kontron Modular Computers S.A.S.
150 rue Marcelin Berthelot - BP 244
ZI TOULON EST
83078 TOULON CEDEX - France

TELEPHONE AND E-MAIL
+33 (0) 4 98 16 34 00
Sales: Order-MAR-Toulon@kontron.com
Support: GSS-MAR-Toulon@kontron.com

For further information about other Kontron products, please visit our Internet web site: www.kontron.com.