

Version: 00

Date: 24-Feb-05

Authors: Antoni Munar, Javier Sanchez, Esteban Fullana, Belén Salvachua

Booting installation of the diskless VP-100 Single Board Computer rcc1 and rcc2 at the TileCal Lab in Valencia

Abstract

This note explains the procedures to setup the booting of the diskless Single Board Computers VP-100 rcc1 and rcc2 at the TileCal Lab at Valencia.

1. Introduction

The VP 110/01 VME Pentium Single Board Computers (SBC)¹ for the ATLAS DAQ system are used without hard disk (diskless PC). Therefore, to boot the operating system, file systems and other software resources must be uploaded from somewhere else. Usually such diskless computers are connected to an Ethernet network. Every time it is powered, special instructions stored in the diskless computer ROM are sent in broadcast mode to any PC in the network. If in the network exists a PC with a server running waiting for such instructions exists, an interchange between the two starts, and the diskless computer can upload the operating system from the server². This note describes the steps needed to setup such system between the diskless SBC rcc1, rcc2 and the server tical11.ific.uv.es (a rack mounted Pentium PC) at the Valencia TileCal Lab. Section 2 describes how the diskless computer must be set. Section 3 describes which services must be set in tical11, to download the boot image file and the root file system to the diskless computers rcc1 and rcc2. Section 4 describes how the root file system must be set up in rcc1 and rcc2. After that, rcc2 and rcc1 must be ready to use.

2. Setup of the Diskless Computer

Note: in these note all the instructions refer to rcc2, however the same applies to rcc1.

The SBC rcc2 has two Ethernet interfaces with two different MAC addresses (**Table 1**). The IP addresses have been assigned by the computer services of the University of Valencia. When rcc2 is powered, it automatically sends a broadcast message to any server on the network requesting the download of the boot image file.

SBC Ethernet Port	MAC Address	Assigned IP
rcc2 up	00.40.9E.00.5D.C5	147.156.163.213
rcc2 down	00.40.9E.00.5D.C6	Not assigned
rcc1 up	00.40.9E.00.44.CF	147.156.163.156
rcc1 down	00.40.9E.00.44.D0	Not assigned

Table 2. MAC and IP addresses for the rcc1 and rcc2 diskless computers

To monitor the booting process on rcc2, one can use the serial port on the front panel of rcc2 (COM Ethernet connector). One can connect this port to the serial input port of any other PC (e.g. tical11) and using the utility *minicom* (simply type minicom as superuser), monitor the booting process of rcc2.

¹ Technical Reference Manual for VP 110/01x VME Pentium III-M Single Board Computer, Concurrent Technologies Inc.

² For a very nice and more rigorous description of such mechanism, see the article: <http://www.linuxfocus.org/English/September1998/article63.shtml>

3. Setup of the Server³

The server, tical11.ific.uv.es, must have two services (daemons) running: *dhcpcd* and *tftpd*. The *dhcpcd* server answers the initial request for boot from the rcc2 computer. Once the request is accepted, the control is transferred to the *tftpd* daemon. *tftpd* is an extremely simple ftp protocol, very suitable for diskless computers, because the instructions for file download must be coded in the reduced memory space of a ROM. The *tftpd* daemon initiates the download of the boot image file to the rcc2. The boot image is a binary file that, among other instructions, contains hardwired the location in tical11 of the root file system to be downloaded by rcc2.

3.1 Configuration of the *dhcpcd* server

In this section we explain how to have the *dhcpcd* server running, how to configure it, and how to stop it and restart it if the configuration file is changed.

3.1.1 Setup of the *dhcpcd* server to run automatically when the PC boots:

Firstly, as root user, create the empty file:

```
> touch /var/state/dhcp/dhcpcd.leases
```

Secondly, in the file

```
/etc/rc.d/rc.local
```

include the word “*dhcpcd*” in the last line of the file. After that, the next time the computer is restarted, the *dhcpcd* server will start automatically.

3.1.2 Configuration of the *dhcpcd* server

The *dhcpcd* can be configured modifying the file `/etc/dhcpcd.conf` in tical11.ific.uv.es Presently (25-Feb-05) the existing `dhcpcd.conf` file is

```
ddns-update-style none;
non-authoritative;

##option vendor-class-identifier "PXEClient";
##option vendor-encapsulated-options 06:01:0b:ff;
subnet 147.156.160.0 netmask 255.255.252.0
{
    option routers          147.156.160.1;
    option subnet-mask      255.255.252.0;
    option broadcast-address 147.156.163.255;
    option domain-name      "ific.uv.es";
    option domain-name-servers 147.156.1.1, 147.156.1.3;
```

³ The information in this section is based in the document: Remote Booting on Intel Processors (CERN

```

option time-servers      hal.ific.uv.es;

# option domain-name     "cern.ch";
# option domain-name-servers 137.138.16.5;
# option time-servers     137.138.16.69;
# option dhcp-message     "137.138.131.18:/work_local/linuxOS/vp110_rcc1";

default-lease-time      -1;
# filename "LinuxBootImage.nbi";

host rcc1
{
    hardware ethernet 00:40:9E:00:44:cf;
    fixed-address 147.156.163.205;
    server-name "tical11";
    option host-name "rcc1";
    option dhcp-message "147.156.163.156:/work_local/linuxOS/vp110_rcc1";
    filename "nueva.nbi";
}

host rcc2
{
    hardware ethernet 00:40:9E:00:5D:C5;
    fixed-address 147.156.163.213;
    server-name "tical11";
    option host-name "rcc2";
    option dhcp-message "147.156.163.213:/work_local/linuxOS/vp110_rcc1";
    filename "nueva_rcc2.nbi";
}
}

```

The first parameters are general and refer to the configuration of the network where both PCs are installed. When *rcc2* sends a request for the download of the boot image file, it sends as identifier its MAC address. The *dhcpd* server receives the request, and then looks in the file *dhcpd.conf* for the different diskless computers that are registered (the blocks *host rcc1*, *host rcc2*). If the MAC address and the IP address of the computer that is requesting the download is equal to any of those registered (inside the *host* block, the *hardware ethernet* parameter), then the control is transferred to the *tftpd* server, and the download of the boot image file is started. The option

```
option dhcp-message "147.156.163.213:/work_local/linuxOS/vp110_rcc1";
```

should in principle indicate the place in *tical11* where *rcc2* can find the files of his operating system and download them. However, we found that this was not working. How to indicate such location it is described in Section 3.2 The string

Variant), Markus Joos (Markus.Joos@cern.ch). The quoted document gives a more extensive and detailed (and rather technical) description of how to boot the diskless SBC.

```
filename "nueva_rcc2.nbi";
```

indicates the name of the boot image file. rcc2 needs this file to start the boot. This file is located in the default working directory of the *tftpd* server (see Section 3.2).

3.1.3 Stopping and re-starting the *dhcpcd* server

If one changes the configuration of the *dhcpcd* server while the *dhcpcd* server is running, the *dhcpcd* server must be re-started, so the new configuration can take place. For that:

- 1) Stop the dhcpcd server. Look at the file

```
/var/run/dhcpcd.pid
```

it contains the pid of the dhcpcd running server. After that, as root, issue the command

```
>kill -15 [pid]
```

and the server dhcpcd server will stop.

- 2) Re-start the *dhcpcd* server. Simply type:

```
> dhcpcd
```

3.2 Configuration of the *tftpd* server

The *tftpd* server is already running on tical11. If not, simply as root type

```
> tftpd
```

The configuration file for the *tftpd* server is in

```
/etc/xinetd.d/tftp
```

The file currently in tical11 is:

```

# default: off
# description: The tftp server serves files using the trivial file transfer \
#   protocol. The tftp protocol is often used to boot diskless \
#   workstations, download configuration files to network-aware printers, \
#   and to start the installation process for some operating systems.
service tftp
{
    socket_type      = dgram
    protocol         = udp
    wait             = yes
    user             = root
    server           = /usr/sbin/in.tftpd
    server_args      = -s /tftpboot
    disable          = no
    per_source       = 11
    cps              = 100 2
}

```

The most important parameter here is the *server_args*. `-s /tftpboot`. It indicates the default working directory from where the *tftpd* server can take files and transfer them to rcc2.

In Section 3.1.2 it was explained that the name of the boot image file must be specified in the configuration of the *dhcpd* server. For rcc2 this file is `nueva_rcc2.nbi`. In tical11 this file must be placed in the default working directory of the *tftpd* server:

```
/tftpboot/nueva_rcc2.nbi
```

`nueva_rcc2.nbi` is a binary file compiled to start the booting of the rcc2 computer and which contains hardcoded the location of the root file system that rcc2 will try to download. We did not try to regenerate it, but instead we used a previous version (`/tftpboot/nueva.nbi`) found in tical11⁴. To change the location of the root file system and adapt it to the one that rcc2 needs, we carefully opened the file `nueva_rcc2.nbi` with emacs. Using *insert*, we modified the line

```
/tftpboot/147.156.163.205
```

to

```
/data2/tb/147.156.163.213
```

being very careful not to change the total number of characters, and also making sure that the total size of the file remained unchanged.

As we have seen, when *dhcpd* receives the request from rcc2 to transfer the boot image and the root file system, it redirects the request to *tftpd* to start the download from the `/tftpboot` and `/data2` directories. Since *tftpd* is a very simple ftp server, it has not user authentication, and the any directory besides the one indicate in the *server_args* option must be done accessible to the external world with the `/etc/exports` file. In our case, in `/etc/exports` file you must add the line

⁴ Remote Booting on Intel Processors (CERN Variant), Markus Joos (Markus.Joos@cern.ch).

```
/data2                                rcc2(rw,no_root_squash)
```

and to have the changes to take place, as root type

```
> exportfs -r
```

4. Setup of the rcc2 root file system.

As explained in Section 3, the root file system of rcc2 is downloaded at boot time from tical11 from the directory:

```
/data2/tb/147.156.163.213
```

Once the root file system is downloaded to rcc2, to have it working properly, one has to configure the IP address of rcc2 and its network file system.

1) Configure the Ethernet interface:

Edit the file /data2/tb/147.156.163.213/etc/sysconfig/network

```
NETWORKING=yes  
HOSTNAME="rcc2"
```

Edit the file /data2/tb/147.156.163.213/etc/sysconfig/network-scripts/ifcfg-eth0, and add the line

```
IPADDR=147.156.163.213
```

2) Configure the file system:

Edit the file /data2/tb/147.156.163.213/etc/fstab. It should be like the one below:

```
147.156.163.156:/tftpboot/147.156.163.213 / nfs defaults 1 1  
none /proc proc defaults 0 0  
none /dev/pts devpts gid=5,mode=620 0 0  
147.156.163.156:/home /home nfs defaults 0 0  
#147.156.163.156:/usr /usr nfs defaults 0 0  
147.156.163.156:/data /data nfs defaults 0 0  
147.156.163.156:/work_local /work nfs defaults 0 0  
147.156.163.156:/tdaq /tdaq nfs defaults 0 0  
147.156.163.156:/data2 /data2 nfs defaults 0 0
```

147.156.163.156 is the IP address of tical11. Note that the file /usr is not mounted from tical11. We found that the kernel will issue an error if we mount /usr from tical11. Also,

all the corresponding directories in tical11 must be listed in the file `/etc/exports` of tical11
(Section 3.2)