

# SRM Firmware Howto

---

*David Mosberger* <mailto:davidm@azstarnet.com>, *Rich Payne* <mailto:rdp@alphalinux.org>, and  
*David Huggins-Daines* <mailto:dhuggins@linuxcare.com> v0.6, 3 March 2000

This document describes how to boot Linux/Alpha using the SRM console, which is the console firmware also used to boot Compaq Tru64 Unix (also known as Digital Unix and OSF/1) and OpenVMS.

## Contents

<b>1</b>	<b>About this manual</b>	<b>1</b>
1.1	Who should read this manual . . . . .	1
1.2	Conventions . . . . .	1
<b>2</b>	<b>What is SRM?</b>	<b>2</b>
2.1	Getting to SRM . . . . .	2
2.2	Using the SRM console . . . . .	2
2.3	How Does SRM Boot an OS? . . . . .	3
2.4	Loading The Secondary Bootstrap Loader . . . . .	3
<b>3</b>	<b>The Raw Loader</b>	<b>3</b>
<b>4</b>	<b>The aboot Loader</b>	<b>4</b>
4.1	Getting and Building aboot . . . . .	4
4.2	Floppy Installation . . . . .	5
4.3	Harddisk Installation . . . . .	5
4.4	CD-ROM Installation . . . . .	6
4.5	Building the Linux Kernel . . . . .	6
4.6	Booting Linux . . . . .	6
4.6.1	Device Naming . . . . .	6
4.6.2	Boot Filename . . . . .	7
4.6.3	Boot Flags . . . . .	8
4.6.4	Using aboot interactively . . . . .	8
4.6.5	The aboot.conf configuration file . . . . .	9
4.7	Installing Linux Distributions . . . . .	9
4.7.1	Installation from the Red Hat 6.0 CD . . . . .	9
4.7.2	Installation from the SuSE 6.1 CD . . . . .	10
4.8	Booting Over the Network . . . . .	10
4.9	Partitioning Disks . . . . .	11

4.9.1	What is a disklabel? . . . . .	11
4.9.2	Partitioning the Easy Way: a DOS Disklabel . . . . .	11
4.9.3	Partitioning with a BSD Disklabel . . . . .	11
<b>5</b>	<b>Sharing a Disk With DEC Unix</b>	<b>12</b>
5.1	Partitioning the disk . . . . .	12
5.2	Installing about . . . . .	13
<b>6</b>	<b>Document History</b>	<b>14</b>

## 1 About this manual

### 1.1 Who should read this manual

You should read this manual if you are installing Linux on a new Alpha system that can only boot from the SRM console, or if you are installing Linux on an older Alpha system that can use the SRM console and wish to use SRM to boot your Linux installation.

Because SRM is the only way to boot Linux on modern Alpha systems, and because it provides the proper operating environment for Unix and Unix-like operating systems (such as Linux), it is the recommended way of booting Linux on Alpha when available.

Sometimes, it is preferable to use the ARC, ARCSBIOS, or AlphaBIOS console, such as if you have a machine for which SRM is not available, if you wish to dual-boot with Windows NT without switching consoles, or if you have hardware that is not supported by SRM. On these machines, you will typically use MILO to boot Linux. For more information, refer to the MILO Howto, available from <http://www.alphalinux.org/faq/milo.html> <<http://www.alphalinux.org/faq/milo.html>>.

### 1.2 Conventions

Throughout this manual, we will use the following conventions for commands to be entered by the user:

SRM console commands will be shown with the characteristic SRM '`>>>`' prompt, like this: <sup>1</sup>

```
>>> boot dva0 -fi linux.gz -fl "root=/dev/fd0 load_ramdisk=1"
```

Unix commands will be shown with the '`#`' command prompt if they are to be run as `root`, or '`$`' if they are to be run by a normal user, like this:

```
# swriteboot -f3 /dev/sda /boot/bootlx
```

About commands will be shown with the '`about>`' command prompt, like this:

```
about> b 6/boot/vmlinuz root=/dev/hda6
```

<sup>1</sup>On multiprocessor machines, you will see '`P00 >`' instead, or possibly some other number depending on which processor SRM is running.

## 2 What is SRM?

SRM console is used by Alpha systems as Unix-style boot firmware. Tru64 Unix and OpenVMS depend on it and Linux can boot from it. You can recognize SRM console as a blue screen with a prompt that is presented to you on power-up.

### 2.1 Getting to SRM

Most Alpha systems have both the SRM and ARC/AlphaBIOS console in their firmware. On one of these machines, if your machine starts up with ARC/AlphaBIOS by default, you can switch to SRM through the "Console Selection" option in the Advanced CMOS Setup menu. To make the change permanent, you should set the `os_type` environment variable in SRM to "OpenVMS" or "Unix", like this:

```
>>> set os_type Unix
```

Either one will work to boot Linux. However, if you intend to dual-boot OpenVMS on this machine, you must set `os_type` to "OpenVMS". Conversely, to return to ARC/AlphaBIOS, you can set `os_type` to "NT".

Some older systems may not have both SRM and ARC in firmware as shipped. On these systems, you will have to upgrade your firmware. See <http://ftp.digital.com/pub/DEC/Alpha/firmware/> for the latest firmware updates and instructions.

A few older systems (primarily evaluation boards such as the 164SX and 164LX) are "half-flash" systems, whose firmware can hold SRM or AlphaBIOS, but not both. If you have one of these machines, you will have to reflash your firmware with the SRM console using the AlphaBIOS firmware update utility. Again, see <http://ftp.digital.com/pub/DEC/Alpha/firmware/> for firmware images and instructions. If you wish to return to AlphaBIOS on these machines, you may rerun the firmware update utility from a floppy in SRM using the `fwupdate` command. You can also start AlphaBIOS from a floppy using the `arc` command.

### 2.2 Using the SRM console

The SRM console works very much like a Unix or OpenVMS shell. It views your NVRAM and devices as a pseudo-filesystem. You can see this if you use the `ls` command. Also, it contains a fairly large set of diagnostic, setup, and debugging utilities, the details of which are beyond the scope of this document. As in the Unix shell, you can pipe the output of one command to the input of another, and there is a `more` command that works not unlike the Unix one. To get a full listing of available commands, run:

```
>>> help | more
```

As well, SRM has environment variables, a number of which are pre-defined and correspond to locations in NVRAM. You can view the entire list of environment variables and their values with the `show` command (there are quite a few of them, so you will probably want to pipe its output to `more`). You can also show variables matching a "glob" pattern - for example, `show boot*` will show all the variables starting in "boot".

Environment variables are categorized as either *read-only*, *warm non-volatile*, or *cold non-volatile*. The full listing of pre-defined variables is detailed in the Alpha Architecture Reference Manual. The most useful pre-defined environment variables for the purposes of booting Linux are `BOOTDEF_DEV`, `BOOT_FILE`, `BOOT_FLAGS`, and `AUTO_ACTION`, all of which are cold non-volatile.

To set environment variables, use the `set` command, like this (unlike Unix environment variables, their names are case-insensitive):

```
>>> set bootdef_dev dka0
```

If you set an undefined variable, it will be created for you, however it will not persist across reboots.

The `BOOTDEF_DEV` variable specifies the device (using VMS naming conventions - see 4.6.1 () for an explanation of these) which will be booted from if no device is specified on the `boot` command line, or in an automatic boot. The `BOOT_FILE` variable contains the filename to be loaded by the secondary bootloader, while `BOOT_FLAGS` contains any extra flags. `AUTO_ACTION` specifies the action which the console should take on power-up. By default, it is set to `HALT`, meaning that the machine will start up in the SRM console. Once you have configured your bootloader and the boot-related variables, you can set it to `BOOT` in order to boot automatically on power-up.

Finally, two helpful console keystrokes you should know are Control-C, which, as in the shell, halts a command in progress (such as an automatic boot), and Control-P, which if issued from the `aboot` prompt (or other secondary bootloader) will halt the bootloader and return you to the SRM console.

### 2.3 How Does SRM Boot an OS?

All versions of SRM can boot from SCSI disks and the versions for recent platforms, such as the Noname or AlphaStations can boot from floppy disks as well. Network booting via `bootp` is supported. Note that older SRM versions (notably the one for the Jensen) cannot boot from floppy disks. Booting from IDE devices is supported on newer platforms (DS20, DS10, DP264, UP2000 etc..).

Booting Linux with SRM is a two step process: first, SRM loads and transfers control to the secondary bootstrap loader. Then the secondary bootstrap loader sets up the environment for Linux, reads the kernel image from a disk filesystem and finally transfers control to Linux.

Currently, there are two secondary bootstrap loaders for Linux: the *raw* loader that comes with the Linux kernel and `aboot` which is distributed separately. These two loaders are described in more detail below.

### 2.4 Loading The Secondary Bootstrap Loader

SRM knows nothing about filesystems or disk-partitions. It simply expects that the secondary bootstrap loader occupies a consecutive range of physical disk sector, starting from a given offset. The information on the size of the secondary bootstrap loader and the offset of its first disk sector is stored in the first 512 byte sector. Specifically, the long integer at offset 480 stores the *size* of the secondary bootstrap loader (in 512-byte blocks) and the long at offset 488 gives the *sector number* at which the secondary bootstrap loader starts. The first sector also stores a flag-word at offset 496 which is always 0 and a checksum at offset 504. The checksum is simply the sum of the first 63 long integers in the first sector.

If the checksum in the first sector is correct, SRM goes ahead and reads the *size* sectors starting from the sector given in the *sector number* field and places them in *virtual* memory at address `0x20000000`. If the reading completes successfully, SRM performs a jump to address `0x20000000`.

## 3 The Raw Loader

The sources for this loader can be found in directory `arch/alpha/boot` of the Linux kernel source distribution. It loads the Linux kernel by reading `START_SIZE` bytes starting at disk offset `BOOT_SIZE+512` (also in bytes). The constants `START_SIZE` and `BOOT_SIZE` are defined in `linux/include/asm-alpha/system.h`. `START_SIZE` must be at least as big as the kernel image (i.e., the size of the `.text`, `.data`, and `.bss` segments). Similarly, `BOOT_SIZE` must be at least as big as the image of the raw bootstrap loader. Both constants should be an integer multiple of the sector size, which is 512 bytes. The default values are currently 2MB for `START_SIZE` and 16KB for `BOOT_SIZE`. Note that if you want to boot from a 1.44MB floppy

disk, you have to reduce `START_SIZE` to 1400KB and make sure that the kernel you want to boot is no bigger than that.

To build a raw loader, simply type `make rawboot` in the top directory of your linux source tree (typically `/usr/src/linux`). This should produce the following files in `arch/alpha/boot`:

`tools/lxboot:`

The first sector on the disk. It contains the offset and size of the next file in the format described above.

`tools/bootlx:`

The raw boot loader that will load the file below.

`vmlinux.nh:`

The raw kernel image consisting of the `.text`, `.data`, and `.bss` segments of the object file in `/usr/src/linux/vmlinux`. The extension `.nh` indicates that this file has no object-file header.

The concatenation of these three files should be written to the disk from which you want to boot. For example, to boot from a floppy, insert an empty floppy disk in, say, `/dev/fd0` and then type:

```
# cat tools/lxboot tools/bootlx vmlinux >/dev/fd0
```

You can then shutdown the system and boot from the floppy by issuing the command `boot dva0`.

## 4 The about Loader

When using the SRM firmware, `about` is the preferred way of booting Linux. It supports:

- direct booting from various filesystems (`ext2`, `IS09660`, and `UFS`, the DEC Unix filesystem)
- listing directories and following symbolic links on `ext2`
- booting of executable object files (both `ELF` and `ECOFF`)
- booting compressed kernels
- network booting (using `bootp`)
- partition tables in DEC Unix format (which is compatible with BSD Unix partition tables)
- interactive booting and default configurations for SRM consoles that cannot pass long option strings

### 4.1 Getting and Building about

The latest sources for `about` are available in *this ftp directory* `<ftp://ftp.alphalinux.org.com/about>`. The description in this manual applies to `about` version 0.6 or newer. Please note that many distributions ship `about` with them so downloading `about` from this directory is probably unnecessary.

Once you downloaded and extracted the latest tar file, take a look at the `README` and `INSTALL` files for installation hints. In particular, be sure to adjust the variables in `Makefile` and in `include/config.h` to match your environment. Normally, you won't need to change anything when building under Linux, but it is always a good idea to double check. If you're satisfied with the configuration, simply type `make` to build it (if you're not building under Linux, be advised that `about` requires GNU `make`).

After running `make`, the `about` directory should contain the following files:

**about**

This is the actual `about` executable (either an ECOFF or ELF object file).

**bootlx**

Same as above, but it contains only the text, data and bss segments—that is, this file is not an object file.

**sdisklabel/writeboot**

Utility to install `about` on a hard disk.

**tools/e2writeboot**

Utility to install `about` on an ext2 filesystem (usually used for floppies only).

**tools/isomarkboot**

Utility to install `about` on a iso9660 filesystem (used by CD-ROM distributors).

**tools/aboutconf**

Utility to configure an installed `about`.

## 4.2 Floppy Installation

The bootloader can be installed on a floppy using the `e2writeboot` command (note: this can't be done on a Jensen since its firmware does *not* support booting from floppy). This command requires that the disk is not overly fragmented as it needs to find enough contiguous file blocks to store the entire `about` image (currently about 90KB). If `e2writeboot` fails because of this, reformat the floppy and try again (e.g., with `fdformat(1)`). For example, the following steps install `about` on floppy disk assuming the floppy is in drive `/dev/fd0`:

```
# fdformat /dev/fd0
# mke2fs /dev/fd0
# e2writeboot /dev/fd0 bootlx
```

## 4.3 Harddisk Installation

Since the `e2writeboot` command may fail on highly fragmented disks and since reformatting a harddisk is not without pain, it is generally safer to install `about` on a harddisk using the `swriteboot` command. `swriteboot` requires that the first few sectors are reserved for booting purposes. We suggest that the disk be partitioned such that the first partition starts at an offset of 2048 sectors. This leaves 1MB of space for storing `about`. On a properly partitioned disk, it is then possible to install `about` as follows (assuming the disk is `/dev/sda`):

```
# swriteboot /dev/sda bootlx
```

On systems where partition `c` in the entire disk it will be necessary to 'force' the write of `about`. In this case use the `-f` flag followed by the partition number (in the case of partition `c` this is 3):

```
# swriteboot /dev/sda bootlx -f3
```

On a Jensen, you will want to leave some more space, since you need to write a kernel to this place, too—2MB should be sufficient when using compressed kernels. Use `swriteboot` as described in Section 4.6 () to write `bootlx` together with the Linux kernel.

## 4.4 CD-ROM Installation

To make a CD-ROM bootable by SRM, simply build `about` as described above. Then, make sure that the `bootlx` file is present on the iso9660 filesystem (e.g., copy `bootlx` to the directory that is the filesystem master, then run `mkisofs` on that directory). After that, all that remains to be done is to mark the filesystem as SRM bootable. This is achieved with a command of the form:

```
# isomarkboot filesystem bootlx
```

The command above assumes that `filesystem` is a file containing the iso9660 filesystem and that `bootlx` has been copied into the root directory of that filesystem. That's it!

## 4.5 Building the Linux Kernel

A bootable Linux kernel can be built with the following steps. During the `make config`, be sure to answer "yes" to the question whether you want to boot the kernel via SRM (for certain platforms this is automatically selected). Note that if you build a generic kernel (by selecting "Generic" as the alpha system type), the kernel is able to guess whether it is running under SRM or not.

```
# cd /usr/src/linux
# make config
# make dep
# make boot
# make modules (if applicable)
# make modules_install (if applicable)
```

The last command will build the file `arch/alpha/boot/vmlinux.gz` which can then be copied to the disk from which you want to boot from. In our floppy disk example above, this would entail:

```
# mount /dev/fd0 /mnt
# cp arch/alpha/boot/vmlinux.gz /mnt
# umount /mnt
```

## 4.6 Booting Linux

With the SRM firmware and `about` installed, Linux is generally booted with a command of the form:

```
>>> boot devicename -file filename -flags flags
```

Note that `-file` and `-flags` can be abbreviated to `-fi` and `-fl`, and often are.

The `filename` and `flags` arguments are optional; If they are not specified, SRM uses the default values stored in environment variables `BOOTDEF_DEV`, `BOOT_OSFILE` and `BOOT_OSFLAGS`. The syntax and meaning of these two arguments is described in more detail below. To list the current values of these variables type `show boot*` at the SRM command prompt. This will also show a `BOOT_DEV` variable (among others), this variable is read only and needs to be changed via the `BOOTDEF_DEV` variable.

### 4.6.1 Device Naming

This corresponds to the device from which SRM will attempt to boot. Examples include:

**dva0**

- First floppy drive, `/dev/fd0` under Linux

**dqa0**

- Primary IDE cdrom or hard disk as Master, `/dev/hda` under Linux

**dqa1**

- Primary IDE cdrom or hard disk as Slave, `/dev/hdb` under Linux

**dka0**

- SCSI disk on first bus, Device 0, `/dev/sda` under Linux

**ewa0**

- First Ethernet Device, `/dev/eth0` under Linux

For example to boot from the disk at SCSI id 6, you would enter:

```
>>> boot dka600
```

To list the devices currently installed in the system type `show dev` at the SRM command line. In contrast to Linux device naming, the partition number on a disk device is *not* given as part of the device name (you may see extra numbers after the device names when running `show dev` - these correspond to things like PCI bus and device numbers and are not useful to the user). Remember, as mentioned in 2.3 (), that SRM knows *nothing* about partitions or disklabels - it merely reads a boot block and secondary bootstrap from sectors on a disk. Therefore, the partition number is given as part of the boot filename.

#### 4.6.2 Boot Filename

The filename argument takes the form:

`[n/]filename`

*n* is a single digit in the range 1..8 that gives the partition number from which to boot from. *filename* is the path of the file you want boot. For example to boot a kernel named `vmlinux.gz` from the second partition of SCSI device 6, you would enter:

```
>>> boot dka600 -file 2/vmlinux.gz
```

Or to boot from floppy drive 0, you'd enter:

```
>>> boot dva0 -file vmlinux.gz
```

If a disk has no partition table, `about` pretends the disk contains one `ext2` partition starting at the first diskblock. This allows booting from floppy disks.

As a special case, partition number 0 is used to request booting from a disk that does not (yet) contain a file system. When specifying "partition" number 0, `about` assumes that the Linux kernel is stored right behind the `about` image. Such a layout can be achieved with the `swriteboot` command. For example, to setup a filesystem-less boot from `/dev/sda`, one could use the command:

```
# swriteboot /dev/sda bootlx vmlinux.gz
```

Booting a system in this way is not normally necessary. The reason this feature exists is to make it possible to get Linux installed on a systems that can't boot from a floppy disk (e.g., the Jensen).

### 4.6.3 Boot Flags

A number of bootflags can be specified. The syntax is:

```
-flags "options..."
```

Where "options..." is any combination the following options (separated by blanks). There are many more bootoptions, depending on what drivers your kernel has installed. The options listed below are therefore just examples to illustrate the general idea:

#### **load\_ramdisk=1**

Copy root file system from a (floppy) disk to the RAM disk before starting the system. The RAM disk will be used in lieu of the root device. This is useful to bootstrap Linux on a system with only one floppy drive.

#### **floppy=*str***

Sets floppy configuration to *str*.

#### **root=*dev***

Select device *dev* as the root-file system. The device can be specified as a major/minor hex number (e.g., 0x802 for /dev/sda2) or one of a few canonical names (e.g., /dev/fd0, /dev/sda2).

#### **single**

Boot system in single user mode.

#### **kgdb**

Enable kernel-gdb (works only if CONFIG\_KGDB is enabled; a second Alpha system needs to be connected over the serial port in order to make this work)

Some SRM implementations (e.g., the one for the Jensen) are handicapped and allow only short option strings (e.g., at most 8 characters). In such a case, **about** can be booted with the single-character boot flag "i". With this flag, **about** will enter interactive mode

### 4.6.4 Using about interactively

As of version 0.6, **about** supports a simple command-oriented interactive mode. Note that this is *different* from the prompt which previous versions issued when booted with the "i" flag, or after failing to load a kernel. You can get a summary of the available commands by typing "h" or "?" at the prompt:

```
>>> boot dka0 -fl i
about> ?
h, ?           Display this message
q             Halt the system and return to SRM
p 1-8         Look in partition <num> for configuration/kernel
l            List preconfigured kernels
d <dir>       List directory <dir> in current filesystem
b <file> <args> Boot kernel in <file> (- for raw boot)
              with arguments <args>
0-9          Boot preconfiguration 0-9 (list with 'l')
about> b 3/vmlinux.gz root=/dev/sda3 single
```

### 4.6.5 The `aboot.conf` configuration file

Since booting in that manner quickly becomes tedious, `aboot` allows to define short-hands for frequently used commandlines. In particular, a single digit option (0-9) requests that `aboot` uses the corresponding option string stored in file `/etc/aboot.conf`. A sample `aboot.conf` is shown below:

```
#
# aboot default configurations
#
0:3/vmlinuz.gz root=/dev/sda3
1:3/vmlinuz.gz root=/dev/sda3 single
2:3/vmlinuz.new.gz root=/dev/sda3
3:3/vmlinuz root=/dev/sda3
8:- root=/dev/sda3          # fs-less boot of raw kernel
9:0/vmlinuz.gz root=/dev/sda3 # fs-less boot of (compressed) ECOFF kernel
-
```

With this configuration file, the command

```
>>> boot dka0 -f1 1
```

corresponds exactly to the boot command shown above.

Finally, at the `aboot` prompt, it is possible to enter one of the single character flags ("0"- "9") to get the same effect as if that flag had been specified in the boot command line. As noted in the help text cited above, you can also list the available default configurations with the "l" command.

**Selecting the Partition of `/etc/aboot.conf`** When installed on a harddisk, `aboot` needs to know what partition to search for the `/etc/aboot.conf` file. A newly compiled `aboot` will search the *second* partition (e.g., `/dev/sda2`). Since it would be inconvenient to have to recompile `aboot` just to change the partition number, `abootconf` allows to directly modify an installed `aboot`. Specifically, if you want to change `aboot` to use the *third* partition on disk `/dev/sda`, you'd use the command:

```
# abootconf /dev/sda 3
```

You can verify the current setting by simply omitting the partition number. That is: `abootconf /dev/sda` will print the currently selected partition number. Note that `aboot` does have to be installed already for this command to succeed. As of version 0.6, `swriteboot` will preserve the existing configuration when installing a new `aboot` on a hard disk.

Since `aboot` version 0.5, it is also possible to select the `aboot.conf` partition via the boot command line. This can be done with a command line of the form `a:b` where `a` is the partition that holds `/etc/aboot.conf` and `b` is a single-letter option as described above (0-9, i, or h). For example, if you type `boot -f1 "3:h" dka100` the system boots from SCSI ID 1, loads `/etc/aboot.conf` from the third partition, prints its contents on the screen and waits for you to enter the boot options.

## 4.7 Installing Linux Distributions

### 4.7.1 Installation from the Red Hat 6.0 CD

Red Hat have made their distribution CD bootable from SRM console<sup>2</sup> To start an installation, put the CD in and type the following:

<sup>2</sup>Please note that through the official RedHat CD-ROM is SRM bootable, copies made by various other companies may not be bootable.

```
>>> boot srm-device -file kernels/generic.gz -flags root=linux-device
```

In the above, the SRM device name and Linux device name for your CD-ROM drive are needed. For Example if the machine had an IDE cdrom installed as primary master the command would look like this:

```
>>> boot dqa0 -file kernels/generic.gz -flags "root=/dev/hda"
```

See the section on 4.6.1 () conventions if you don't know what these are.

#### 4.7.2 Installation from the SuSE 6.1 CD

The SuSE 6.1 CD is not bootable from SRM console. SuSE have an alternative approach which involves creating two boot floppies, the images of which are included on the CD. The boot disks can be created in various ways, depending on the systems you have available

Writing the boot disks from a linux system The command to use is dd. From the mount-point of SuSE CD 1, the commands are:

```
# dd if=disks/aboot of=/dev/fd0
# dd if=disks/install of=/dev/fd0
```

For writing the boot disks from a windows system, the command to use is rawrite. It is available on the CD.

```
D:\tools\> rawrite
```

The program then prompts for input disk image and output disk drive. Run this command once for each of the disk images as shown above.

Starting the SuSE installer from the boot disks With the floppy disk made from the aboot image in place, type:

```
>>> boot dva0 -file vmlinux.gz -flags "root=/dev/fd0 load_ramdisk=1"
```

This will start the kernel, prompt you for the second boot disk, and start the installer

## 4.8 Booting Over the Network

Two preliminary steps are necessary before Linux can be booted via a network. First, you need to set the SRM environment variables to enable booting via the bootp protocol and second you need to setup another machine as the your boot server. Please refer to the SRM documentation that came with your machine for information on how to enable bootp. Setting up the boot server is obviously dependent on what operating system that machine is running, but typically it involves starting the program `bootpd` in the background after configuring the `/etc/bootptab` file. The `bootptab` file has one entry describing each client that is allowed to boot from the server. For example, if you want to boot the machine `myhost.cs.arizona.edu`, then an entry of the following form would be needed:

```
myhost.cs.arizona.edu:\
:hd=/remote/:bf=vmlinux.bootp:\
:ht=ethernet:ha=08012B1C51F8:hn:vm=rfc1048:\
:ip=192.12.69.254:bs=auto:
```

This entry assumes that the machine's Ethernet address is 08012B1C51F8 and that its IP address is 192.12.69.254. The Ethernet address can be found with the `show device` command of the SRM console or, if Linux is running, with the `ifconfig` command. The entry also defines that if the client does not specify otherwise, the file that will be booted is `vmlinux.bootp` in directory `/remote`. For more information on configuring `bootpd`, please refer to its man page.

Next, build `about` with with the command `make netboot`. Make sure the kernel that you want to boot has been built already. By default, the `about` Makefile uses the kernel in `/usr/src/linux/arch/alpha/boot/vmlinux.gz` (edit the Makefile if you want to use a different path). The result of `make netboot` is a file called `vmlinux.bootp` which contains `about` and the Linux kernel, ready for network booting.

Finally, copy `vmlinux.bootp` to the bootserver's directory. In the example above, you'd copy it into `/remote/vmlinux.bootp`. Next, power up the client machine and boot it, specifying the Ethernet adapter as the boot device. Typically, SRM calls the first Ethernet adapter `ewa0`, so to boot from that device, you'd use the command:

```
>>> boot ewa0
```

The `-fi` and `-fl` options can be used as usual. In particular, you can ask `about` to prompt for Linux kernel arguments by specifying the option `-fl i`.

## 4.9 Partitioning Disks

### 4.9.1 What is a disklabel?

A disk label is a partition table. Unfortunately, there are several formats the partition table can take, depending on the operating system.

DOS partition tables are the standard used by Linux and Windows. AlphaBIOS systems and every Linux kernel can read DOS partition tables. Unfortunately, the SRM console's boot sector format overlaps with parts of the DOS partition table on disk, and therefore DOS partition tables cannot be used with SRM.

BSD disklabels are used by several variants of Unix, including Tru64. SRM's boot block does not conflict with the BSD disklabel (in fact, the BSD disklabel resides entirely within "reserved" areas of the first sector), and Linux can use a BSD disklabel, provided that support for BSD disklabels has been compiled into the kernel.

To boot from a disk using SRM, a BSD disklabel is required. If the disk is not a boot disk, the BSD disklabel is not required. A BSD disklabel can be created using `fdisk`, the standard Linux disk partitioning tool.

### 4.9.2 Partitioning the Easy Way: a DOS Disklabel

The simplest way to partition your disk is to let your Linux installer do it for you, for example by using Red Hat's disk druid or `fdisk`. On Red Hat 6.1, this will produce a valid BSD disklabel, but *only* if the disk in question previously contained one. In most cases, this will produce a DOS disklabel. It will be readable by Linux, but you will not be able to boot from it via SRM. For this reason, you will probably want to create a BSD disklabel manually in order to boot Linux

### 4.9.3 Partitioning with a BSD Disklabel

1. Start `fdisk` on the disk you're configuring

2. Choose to make a BSD disklabel - option 'b' (newer versions of fdisk will detect existing BSD disklabels and automatically enter disklabel mode)
3. You'll notice some things: Partitions are letters instead of numbers, from a-h Partition 'c' covers the whole of the disk. This is the convention, don't touch it. While you can see it, note down the disk parameters as you'll use them more often than with the DOS-disklabel approach
4. Creating a new partition uses the same procedure as the DOS-disklabel approach, except that the partitions are referred to by letter instead of number. That is, 'n' to make a new partition followed by the partition letter followed by the starting block followed by the end block
5. Setting partition type is slightly different, because the numbering scheme is different (1 is swap, 8 is ext2).
6. When you are finished, write ('w') and quit ('q') as normal.

There are some important catches that you must be aware of when partitioning using a BSD disklabel:

- Partition 'a' should start about 1M into the disk: don't start it at sector 1, try starting at sector 10 (for example). This leaves plenty of space for writing the boot block (see below)
- There is a bug in some versions of fdisk which makes the disk look one sector bigger than it actually is. The listing when you create the BSD disklabel is correct. The last sector of partition 'c' is correct. The default last sector when creating a new partition is 1 sector too big
- Always adjust for this extra sector. This bug exists in the version of fdisk shipped with Red Hat 6.0. Not making an adjustment for this problem almost always leads to "Access beyond end of device" errors from the Linux kernel.

Once you have made a BSD disklabel, continue the installation. After installation, you can write a boot block to your disk to make it bootable from SRM.

## 5 Sharing a Disk With DEC Unix

Unfortunately, DEC Unix doesn't know anything about Linux, so sharing a single disk between the two OSes is not entirely trivial. However, it is not a difficult task if you heed the tips in this section. The section assumes you are using `about` version 0.5 or newer.

### 5.1 Partitioning the disk

First and foremost: *never* use any of the Linux partitioning programs (`minlabel` or `fdisk`) on a disk that is also used by DEC Unix. The Linux `minlabel` program uses the same partition table format as DEC Unix `disklabel`, but there are some incompatibilities in the data that `minlabel` fills in, so DEC Unix will simply refuse to accept a partition table generated by `minlabel`. To setup a Linux `ext2` partition under DEC Unix, you'll have to change the `disktab` entry for your disk. For the purpose of this discussion, let's assume that you have an rz26 disk (a common 1GB drive) on which you want to install Linux. The `disktab` entry under DEC Unix v3.2 looks like this (see file `/etc/disktab`):

```
rz26|RZ26|DEC RZ26 Winchester:\
    :ty=winchester:dt=SCSI:ns#57:nt#14:nc#2570:\
    :oa#0:pa#131072:ba#8192:fa#1024:\
    :ob#131072:pb#262144:bb#8192:fb#1024:\
```



Linux and DEC Unix. Of course, you could also boot Linux from a floppy if speed doesn't matter or over the network, if you have a `bootp`-capable server. But in this section we assume you want to boot Linux from a disk that contains one or more DEC Unix partitions.

*Second big caveat:* installing `about` on a disk shared with DEC Unix renders the first and third partition unusable (since those *must* have a starting offset of 0). For this reason, we recommend that you change the size of partition `a` to something that is just big enough to hold `about` (1MB should be plenty).

Once these two caveats are taken care of, installing `about` is almost as easy as usual: since partition `a` and `c` will overlap with `about`, we need to tell `swriteboot` that this is indeed OK. We can do this under Linux with a command line of the following form (again, assuming we're trying to install `about` on the second SCSI disk):

```
# swriteboot -f1 -f3 /dev/sdb bootlx
```

The `-f1` means that we want to force writing `bootlx` even though it overlaps with partition 1. The corresponding applies for partition 3.

This is it. You should now be able to shutdown the system and boot Linux from the harddisk. In our example, the SRM command line to do this would be:

```
>>> boot dka5 -fi 8/vmlinux.gz -f1 root=/dev/sdb8
```

## 6 Document History

v0.6.1 6 March 2000 A few small fixes:

- Fixed some typos
- Use capitals consistently for SRM environment variables

v0.6 3 March 2000 Changes and information from David Huggins-Daines <dhd@linuxcare.com>

- Moved the notes on MILO vs. SRM to an "About this document" section
- Added sections on switching to SRM, and basic SRM usage
- Added section on the new interactive use of `about`
- Updated the note on DOS partition tables to mention the Red Hat 6.1 installer's behaviour.
- Normalized the markup, and codified the conventions used for user-entered commands.
- Corrected the notes on BSD disklabels (SRM does *not* read BSD disklabels, it's just that they don't conflict with the boot block).

v0.5.2 5 December 1999 Added comments and information from Stig Telfer (stig@alpha-processor.com).

- Added chart on SRM to Linux name mappings
- Added RedHat 6.0 and SuSE 6.1 installation information
- Added Disk Partitioning Information

v0.5.1 (Not Released) 13 November 1999 Took the original 0.5 document and updated several parts:

- Update information on SRM booting from IDE devices
- Fixed URL to about source
- Update toc page to reflect MILO's future
- Included information on BOOTDEF\_DEV and BOOT\_DEV to chapter 3
- Added this section

v0.5 17 August 1996 - Original Document by David Mosberger-Tang