

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54

PowerPC Reference Platform binding to:

IEEE Std 1275-1994

Standard for Boot

(Initialization, Configuration)

Firmware

Revision: 0.02 DRAFT

Date: July 3, 1995

1. Overview

This document specifies the application of *IEEE Std 1275-1994 Standard for Boot (Initialization, Configuration) Firmware, Core Practices and Requirements* to PowerPC Reference Platform Compliant computer systems, including practices for client program interface and data formats. An implementation of Open Firmware for a PowerPC Reference Platform Compliant system *shall* implement the core requirements as defined in [1], the PowerPC Processor-specific extensions described in [2] and the PowerPC Reference Platform specific extensions described in this binding.

2. References and Terms

2.1. References

This standard shall be used in conjunction with the following publications. When the following standards are superseded by an approved revision, the revision *shall* apply.

[1] *IEEE Std 1275-1994 Standard for Boot (Initialization, Configuration) Firmware, Core Practices and Requirements*.

[2] *PowerPC processor binding to: IEEE Std 1275-1994, Standard for Boot (Initialization, Configuration) Firmware*.

[3] *System V Application Binary Interface*, published by UNIX System Laboratories. This document describes the generic architecture of the ELF (Executable and Linking Format) object file format.

[4] *MS-DOS Programmer's Reference*, published by Microsoft. This document describes the MS-DOS partition, directory and FAT formats used by `disk-label` support package.

[5] *Peering Inside the PE: A Tour of the Win32 Portable Executable File Format*, found in the March, 1994 issue of *Microsoft Systems Journal*.

[6] *Bootstrap Protocol*, Internet RFC 951; see also RFC 1532.

[7] *ISO-9660, Information processing -- Volume and file structure of CD-ROM for information interchange*, published by International Organization for Standardization.

[8] *System V Application Binary Interface, PowerPC Processor Supplement*, Sunsoft. This document defines the PowerPC specific ABI for System V and also gives details on the PowerPC ELF format.

[9] *Device Support Extensions to IEEE 1275-1994 Standard for Boot (Initialization, Configuration) Firmware*.

[10] *Open Firmware Recommended Practice: 16-color Text Extension*.

[11] *Open Firmware Recommended Practice: Graphics Extension*.

2.2. Terms

This standard uses technical terms as they are defined in the documents cited in "References", plus the following terms:

core, core specification: refers to *IEEE Std 1275-1994 Standard for Boot (Initialization, Configuration) Firmware, Core Practices and Requirements*

ELF: Executable and Linking Format. A binary object file format defined by [3][8] that is used to represent *client programs* in Open Firmware for PowerPC.

FDISK: Refers to the boot-record and partition table format used by MS-DOS, as defined in [4].

Open Firmware: The firmware architecture defined by the core specification or, when used as an adjective, a software component compliant with the core specification.

1 **PE:** Portable Executable. A binary object file format defined by [5]; this format is used by Microsoft's NT
2 operating system.

3 **Real-Mode:** The mode in which Open Firmware and its client are running with translation disabled; all
4 addresses passed between the client and Open Firmware are real (i.e., hardware) addresses.
5

6 **Suspend:** A form of Power Management characterized by a fast recovery to full operation. Typically, system
7 memory will not be powered off while in the suspend state.

8 **Virtual-Mode:** The mode in which Open Firmware and its client share a single virtual address space, and
9 address translation is enabled; all addresses passed between the client and Open Firmware are virtual (translated)
10 addresses.
11

12 13 3. Packages

14 This section describes the PowerPC Reference Platform-specific requirements of Open Firmware packages.
15
16

17 3.1. "disk-label" Support Package

18 This section describes the partition formats and the formats of client program images that the `disk-label`
19 support package for PowerPC *shall* support; an implementation *may* support additional mechanisms, in an
20 implementation-specific manner
21

22 The process of loading and executing a client program is described in two stages. The first stage determines
23 what file to read into memory. This is done by locating a file from the boot device, usually by means of a name
24 lookup within a directory contained within a disk "partition". The second stage examines the front portion
25 (header) of the image for "well-known" program formats. When the format of the image has been determined,
26 the loading is completed in a manner determined by that format.
27

28 The name of the file (and, the partition in which it is contained) can be explicitly specified by the user via the
29 `load` or `boot` command, or can be implicitly specified by the value of the "boot-device" property of the
30 `"/options"` node. The filename is the ARGUMENTS portion of the final COMPONENT of the PATH_NAME, as
31 described in section 4.3.1 of [1].

32 The syntax for explicit filename specification is as follows:
33

```
34 [n][,filename]
```

35 where `n` is the partition number to be used and `filename` is the name of a file within that partition. If `n`
36 is omitted, the default partition (as determined by the partition format) is used. If `filename` is omitted, the
37 default filename (i.e., the filename component of the "boot-device" path-name) is used. Partition number 0
38 refers to the entire device; this definition is independent of the existence of the device partition information.
39

40 3.1.1. Partition formats

41 Open Firmware shall implement the algorithm described as follow:
42

43 Algorithm locating boot file

```
44 read sector 0 (bootsector)  
45 if last 2 bytes of sector are 0AA55h (little-endian)  
46     if bsMedia == 0F8h \ FDISK partition on hard drive  
47         if an explicit partition has been requested  
48             select partition number n  
49         else  
50             select bootable partition (80h in peBootable field)  
51             use directory of the selected partition to locate file  
52     else (non-partitioned)  
53         use FAT-12/FAT-16 directory to locate file  
54     else
```

```
1 read sector 16.
2 if a valid ISO-9660 directory is found
3     locate the file, using the ISO-9660 directory.
4 else
5     FAIL, in an implementation-specific manner.
```

This algorithm can be used to locate the correct file and/or load image from the specified device. The boot device is selected as described in 7.4.3.2 of [1]. A filename can be explicitly given as the arguments field of the *device-specific* *er* (i.e., the field following the ':' of the last path component). Other formats may be recognized in an implementation-specific manner.

Although the above algorithm works independent of the device type, the following formats are strongly recommended for devices for portability.

Floppy Disk

1.44/2.88 MB, MFM floppy disks should be in FAT-12 format, as described in [4].

Hard Disk

Hard Disks should have an FDISK partition map, as described in [4].

Note: since the bootsector is used to contain boot program for floppies and the FDISK partition map for hard disks, the "disk-label" package must use the value of the bsMedia byte (located at offset 15h) to determine whether a partition map is present. If the value is 0F8h, it indicates a hard disk and a partition map should be present in the bootsector; any other value indicates a floppy disk.

CD-ROM

CD-ROMs should be formatted according to ISO-9660, as described in [7].

3.1.2. Program-image formats.

Open Firmware must recognize a client program that is formatted as either ELF [3][8] or PE [5]. PE format support is provided only for booting NT; all other clients *shall* use ELF. Other formats *may* be handled in an implementation-specific manner.

After locating the file, Open Firmware reads the image into memory at the location specified by the **load-base** Configuration Variable. Then, Open Firmware must perform the following procedure to prepare the image for execution.

init-program.

```
37 examine the header of the image.
38 set restart? false
39 if the image is in ELF format
40     if the EI_DATA field does not match little-endian?
41         set little-endian? appropriately.
42         set restart? true
43     locate the PowerPC Note Section
44     if the Note Section's descriptor is not correct
45         set Configuration Variables appropriately
46         set restart? true
47     if restart?
48         restart the system, possibly by executing reset-all
49     else
50         move and/or relocate the ELF image.
51 else
52     if the file is in PE format
53         if little-endian? is false
54             set little-endian? to true.
55             restart the system, possibly by executing reset-all
```

```
1     else
2         move and/or relocate the PE image.
3     else
4         FAIL, in an implementation -specific manner.
```

3.2. "obp-tftp" Support Package

The "obp-tftp" Support Package of an Open Firmware implementation (used to **load** from "network" devices) *shall* support the BOOTP protocol, as described in [6] and *shall* not require the server to support any vendor extensions.

4. Properties

This section describes the standard properties of PowerPC Reference Platform Open Firmware implementation.

4.1. "/openprom" node property

Open Firmware must implement the "halt-address" property under the "/openprom" node. The property value describes the location of physical memory where the firmware saves the halt procedure image. Halt procedure is the self-contained firmware callback service for system poweroff or reboot. Halt procedure is discussed further in Section 5.1.

"halt-address"

prop-encoded-array: address, length

address, encoded as with **encode-int**, is the real mode address of the halt procedure, and length, encoded as with **encode-int**, is the physical memory size for the halt procedure. If the halt procedure is located within read-write memory, length is the number of bytes of memory, encoded as with **encode-int**, occupied by the halt procedure. If the halt procedure is located within read-only memory, length is 0.

Note: "read-only memory" refers to memory that cannot be written in normal operation; ROM or FLASH ROM is considered to be "read-only"; while RAM, whether or not it is write-protected via an MMU, is considered to be "read-write".

The read-write memory, if any, that the halt procedure occupies must not be included within the "available" property of the "/memory" node.

5. Client Program Requirements

5.1. Halt callback

The halt callback allows a client program to invoke limited firmware services for turning-off or rebooting the machine. The halt procedure shall not return control to the client program that invoked it. Consequently, the halt procedure is not required to preserve any machine state on behalf of its caller.

PowerPC Reference Platform Open Firmware must implement the "halt-address" property within the "/openprom" node to support the halt callback. Detailed discussion about this property is presented in Section 4.1.

5.1.1. Halt procedure calling conventions

When control is transferred to the halt procedure, the system must be in the endian mode that was in effect when the firmware last had control of the system. The client program must establish the real mode address translation for the firmware.

5.1.2. Halt procedure argument

The client program must set GPR3 to point to the halt argument. The halt argument is the address of a null-terminated text string. In the following discussion, the string is called the "halt string".

5.1.2.1. Power-off Service

If the halt string is "power-off", the fi rmware must turn of the system power to the extent possible. If the fi rmware cannot turn of the system power, perhaps due to lack of hardware capability, the result is undefi ned, and the fi rmware must not try to return control to the client program or reboot the system. The suggested behavior in this case is to enter a fi rmware interactive mode, if available.

5.1.2.2. Suspend Service

If the halt string is "suspend", the fi rmware shall preserve the RAM image across a hardware suspend operation in the following manner:

1. The Firmware shall not alter RAM that is not devoted to fi rmware use (see System State at Call below).
2. If the fi rmware encounters an error or the hardware does not support suspend, the halt procedure shall set the appropriate return code (see System State at Return below) and return to the caller.
3. The fi rmware shall record suspend restoration parameters in a system-dependent manner
4. The fi rmware shall set the system power state to suspend in a system-dependent manner (this usually means powering down everything except the memory sub-system and its refresh).
5. When the system is subsequently powered up, the fi rmware performs the following steps:
 - Reads the suspend restoration parameters.
 - Initializes the system hardware as on IPL without disturbing the preserved RAM image
 - Does not load the standard client program (so as not to disturb the preserved RAM image).
 - Resets the suspend restoration parameters.
 - Returns to the caller of suspend (see system State at return below).

The System State at the Call to "suspend" shall be as follows:

Memory:

The memory range defi ned by the **real-base** and **real-size** confi guration variables is undefi ned but available for use by Open Firmware.

Master Processor:

GPR3 shall hold the starting address of the null-terminated text string "suspend".

Slave Processors:

In an SMP system only one processor shall call the Halt procedure with a halt string of "suspend", all other processors shall call the halt procedure with the halt string "slave". In this case, the halt procedure shall execute a program which does not change the state of memory.

I/O Sub-system:

All elements of the I/O sub-system shall be in a safe state. In a safe state, they shall not be transferring data to or from memory nor shall they cause an interrupt to any processor.

Interrupt Sub-system:

The processor's External Interrupt Enable bit (MSR[EE]) shall be 0 prior to calling the halt procedure.

The System State at a Return from "suspend" shall be as follows:

Memory:

1 The state of the memory locations, exclusive of the range defined by the "real-base" and "real-size"
2 configuration variables, shall be the same as the state they were in at the time of the call to the Halt
3 procedure.

4 Master Processor:

5
6 The state of the master processor shall be as defined for the "initial program state" as specified in [2] with
7 the following exceptions: The processor will start executing at the address which was in the master
8 processor's Link Register at the time the Halt procedure was called. GPR3 shall contain the return code
9 (see below).

10 Slave Processors:

11
12 The state of the slave processors shall be as follows: The MSR shall be set to the state specified for "initial
13 program state" as specified in [2]. All non-architected processor dependent registers shall be set to the
14 system default value. All other architected registers are undefined. The processor shall start executing at
15 the address which was in the slave processor's link register when the Halt procedure was called. GPR3
16 shall contain the return code 0.

17 I/O Sub-system:

18
19 All the elements of the I/O sub-system shall be initialized to a safe state. The device tree shall be re-
20 established as on IPL.

21 Interrupt Sub-system:

22
23 The processor's External Interrupt Enable bit (MSR[EE]) shall be 0.

24 Return Codes:

- 25
26 0: Restored from suspend.
27
28 -1: Hardware does not support suspend power state.
29

30 5.1.2.3. Other Halt Services

31
32 The firmware may recognize strings other than "poweroff" or "suspend" in a system-dependent manner. If the
33 halt string is not one of the recognized commands, the firmware must reboot the system. During the system reset
34 and the firmware restart, the firmware must preserve the halt string. Then, the firmware must evaluate the string
35 as if "auto-boot?" were true and "boot-command" were set to that string, without altering the values of the
36 "auto-boot?" and "boot-command" configuration variables. If the firmware includes an Open Firmware user
37 interface, the string can be any valid user interface command string. Otherwise, the firmware shall interpret the
38 string as shown in Table 1. Halt arguments.

39
40 **Table 1. Halt arguments**

41 String value	42 Meaning
43 boot	Load and execute the default client program.
44 <empty string>*	45 If the firmware has an interactive mode, enter that mode. 46 Otherwise the result is undefined, except that the firmware 47 shall not perform the "boot" action.
48 <any other string>	49 Behavior is system-dependent

50 * Note: <empty string> means a valid address that points to a null character. An address value of 0 is not
51 an empty string.
52
53
54

6. Extensions for PowerPC Reference Platform systems

This section describes the properties, methods, and device subtrees that are applicable to devices required by the PowerPC Reference Platform architecture. It is strongly recommended that other platforms follow these definitions for the corresponding devices.

6.1. Display devices

Display device packages (i.e., `device_type = "display"`) for PowerPC Reference Platform systems should include in their implementation all the properties and methods called for in [10] and [11].

6.2. Device support

Open Firmware implementations for PowerPC Reference Platform systems must implement those device types from the IEEE 1275 Device Support Extensions document [9] that are appropriate to the hardware present.

6.3. Conventions for devices on ISA and SCSI buses

This section defines the naming and device type conventions for typical devices on ISA and SCSI buses. The following lists are the values of the "name" and "device-type" properties of the devices on an ISA bus:

name	device_type
8042	
kbd	"keyboard"
mouse	"mouse"
floppy	"block"
com	"serial"
timer	"timer"
lpt	"parallel"
ide	"block"
nvrnm	"nvrnm"
rtc	"rtc"

Note: The "kbd" and "mouse" names are indented to show that they are the child nodes of the 8042 node.

Some systems use an I/O controller, often called a super I/O chip, which provides control functions of multiple I/O devices. When a system uses a super I/O chip, a device node representing the super I/O chip itself need not exist. Instead, the device nodes of the devices attached to the super I/O chip may be direct children of the bus node representing the bus to which the super I/O chip is attached.

The following are the values of the name and device-type properties of the devices on SCSI bus:

name	device_type
scsi	"scsi"
disk	"block"
tape	"byte"

Note: SCSI controller is considered a bus device in the device tree for PowerPC Open Firmware. The "disk" and "tape" names are indented to show that they are the child nodes of the scsi node.

It is strongly recommended that the "compatible" property be implemented for ISA and SCSI bus devices to help operating systems find appropriate device drivers for these devices.

6.4. /aliases node properties

An implementation of Open Firmware for the PowerPC Reference Platform shall provide the following aliases under "/aliases" node if an applicable device exists:

1 disk
2 tape
3 cdrom
4 keyboard
5 screen
6 scsi
7 com1
8 com2
9 floppy
10 net
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54