

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™ Microprocessor  
Common Hardware Reference Platform (CHRP™)

System binding to:

IEEE Std 1275-1994

Standard for Boot

(Initialization, Configuration)

Firmware

Revision: 1.7 [Approved Version]

**Date: September 23, 1996**

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

# Purpose of this PowerPC Microprocessor CHRP™ System binding

This document specifies the application of Open Firmware to a PowerPC Microprocessor CHRP System, including requirements and practices to support unique hardware and firmware specific to the platform implementation. The core requirements and practices specified by Open Firmware must be augmented by system-specific requirements to form a complete specification for the firmware implementation of a PowerPC Microprocessor CHRP System. This document establishes such additional requirements pertaining to the platform and the support required by Open Firmware.

## Task Group Members

The following individuals were members of the Task Group that produced this document:

Rob Baxter, Motorola®

Mitch Bradley, FirmWorks

Steve Bunch, Motorola

Bob Coffin, IBM® (editor)

Ron Hochsprung, Apple®

John Kingman, IBM

Dr. Luan Nguyen, IBM

John O'Quin, IBM

Andy Rawson, IBM

Mike Segapeli, IBM

## Trademarks

The following terms, denoted by a registration symbol (®) or trademark symbol(™) on the first occurrence in this publication, are registered trademarks or trademarks of the companies as shown in the list below:

<b>Trademark</b>	<b>Company</b>
AIX™	International Business Machines Corporation
Apple	Apple Computer, Inc.
CHRP	Apple Computer, Inc.
Ethernet™	Xerox
IBM	International Business Machines Corporation
Microsoft®	Microsoft Corporation
Macintosh™	Apple Computer, Inc.

---

1	Mac™ OS	Apple Computer, Inc.
2	Motorola	Motorola, Inc.
3		
4	PowerPC	International Business Machines Corporation
5		
6	Windows NT™	Microsoft Corporation
7		

## Revision History

11	Revision 0.1	07/27/95	Initial draft; limited distribution to CHRP AIM(Apple, IBM & Motorola) Team
12			
13	Revision 0.2	08/15/95	Initial draft with distribution to CHRP AIM Team for review.
14			
15	Revision 0.3	09/08/95	Distribution to CHRP AIM Team for review. Added more function to Open
16			Firmware(properties & methods) and RTAS.
17			
18	Revision 0.4	10/20/95	Distribution to CHRP AIM Team for review. Changed basic format and
19			structure of document with appendices for proposals. Added
20			function/expanded function to: 'bootp', Power Management, System Nodes.
21			Expanded review process to include FlrePower, FirmWorks and SunSoft.
22			
23	Revision 0.5	11/11/95	Added description material to Appendix A for (proposed)interrupt support.
24			Expanded review process to include Bull, Canon & Toshiba.
25			
26	Revision 0.6	12/07/95	Added Appendix A for Future Direction. Differed certain function to next
27			version of CHRP System binding. CHanged old Appendix A to Appendix B
28			for an interrupt proposal as a recommended practice format. Numerous
29			editorial changes.
30			
31	Revision 0.7	02/07/96	Removed Appendix A; an example of CHRP Device Tree. Appendix A became
32			Interrupt Structure. Changed multi-boot in Section 3.0. Moved Sections of
33			PowerPC Processor binding Document to system binding; ELF, endianness,
34			properties, Client Callbacks. Numerous editorial changes. Internal IBM
35			Version.
36			
37	Revision 0.8	02/08/96	Minor changes to version 0.7 for review with AIM Members in Tempe, AZ.
38			
39	Revision 0.9	02/23/96	Changed Multi-boot and ELF sections to clarify. Made numerous editorial
40			changes for accuracy and clarification.
41			
42	Revision 1.0	04/10/96	Changes to reflect AIM working Cupertino Meeting on 03/7&8/96.
43			Added references(CHRP I/O device bindings) and edited references. Edited
44			boot flow(Section 3) to add path names and SGML tag format for bootinfo.
45			Combined ROM Node(5.2.1.1) to include OS-ROM function. Numerous
46			editorial changes to ISA and "disk-label" Support Package Sections.
47			
48	Revision 1.1	04/16/96	Made changes for AIM Meeting, 04/15/96 in Tempe, AZ. Edited references
49			15, 21, 23 thru 32. Edited Boot Overview and the /options Node. Edited
50			Table 1 end eliminated Note 2. Changed the tags in the 'bootinfo.txt' file.
51			Redefined the ICON string definition. Changed Boot Script Description to
52			<b>reboot-command</b> Description. Abbreviated Section 5.7.1.1 by referencing
53			[14]. Edited the ISA Device Support Section 8.0. Added Load Method,
54			

---

1			Section 10.1.3 and modified <code>"tape"</code> package(10.2).
2			
3	Revision 1.2	04/22/96	Made changes from AIM Meeting, 04/16/96 in Tempe, Az. Changed
4			'boot Overview(Section 3.0) and made 'bootinfo.txt' format use SGML
5			Tags. Added <code>reboot-command</code> . Added material to ISA Device Support
6			(Section 8.0). Eliminated Mac OS file support for the <code>"disk-label"</code>
7			package. Added/changed material to OF NVRAM SECTION 11.0.
8			
9	Revision 1.3	04/29/96	Removed PE Section 9.4.2. Made changes to CHRP Boot Flow Section 3.0.
10			Changed 'bootinfo.txt' tags in Section 3.1.3. Changed System Control Area
11			section 5.2.1 and Removed 5.2.1.2. Modified Section 12.0, OF NVRAM to
12			change partition definition. Changed Event Sources Node Section 5.9.1.
13			Changed ISA Device Support Section 8.0. Changed NVRAM Partition Section
14			12.0.
15			
16	Revision 1.4	05/02/96	Numerous editorial changes regarding Open Firmware practices about names
17			of properties and methods. Corrected the Table of Contents and added a List
18			of Tables. Added/changed material in Section 9.1, ISA Configuration Utility
19			(ICU). Modified NVRAM Partition Section 12.0.
20			
21	Revision 1.5	05/08/96	Numerous editorial changes regarding document to make consistent with
22			Open Firmware convention and improve readability. Changed Section 3.0 to
23			have a boot overview. Renamed most of section to 'Multi-boot'. Made
24			changes to <code>root</code> node Section 5. Changed Section 12.0; modified OS naming
25			convention and classified System Partition Configuration Variables as to
26			category(boolean, string or number). Rewrote all of Section 12.4, System
27			Partition.
28			
29	Revision 1.6	07/27/96	Some editorial changes. Added Interposition reference. Changed Table 1
30			to just reflect standard pathnames for install media. Added #size-cell property
31			in Section 5.1.1. Added RTAS property and RTAS Function Property Names.
32			Added interrupt-controller property to Section 5.5.1. Removed Interrupt
33			Properties. Changed built-in property, Section 5.6.2. Added Section 6.2.1,
34			Processor Node (slot-names-index property). Clarified Section 11.1,
35			disk-label Support package. Changed Table 12. Changed multi-initiator
36			script, Section 12.4.4.
37			
38	Revision 1.7	09/23/96	Added Reference [5] and version numbers/url's to other references. Removed
39			<code>"reserved"</code> property and reference from Section 5.3.1, PCI Host Bridge
40			Properties. Added reference to ROM Node methods, added <code>decode-unit</code>
41			and <code>encode-unit</code> methods and added arguments for the <code>set-args</code> method
42			for Section 5.1.3.2. Added clarification to "reg" property for the ROM Child
43			Node, Section 5.1.4.1. Added epow-events property, Section 5.8.1.2,
44			<code>"power-on-triggers"</code> . Added Sections 5.9, <code>reserved</code> Node and 5.10,
45			<code>/chosen</code> Node. Made Section 9.0, ICU optional. Added Section 11.4.1
46			Windows NT PE Support. Corrected Multi-initiator script example in Section
47			12.4.4.1.
48			
49			
50			
51			
52			
53			
54			

---

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

---

## Table of Contents:

1		
2		
3		
4		
5		
6		
7	1. Overview.....	13
8	1.1 General Requirements for Open Firmware.....	13
9	2. References and Terms.....	13
10	2.1 References.....	13
11	2.2 Terms.....	15
12	3. CHRP Boot Flow.....	16
13	3.1 CHRP Boot Overview.....	16
14	3.1.1 Additional Requirements for probe-all Method.....	17
15	3.1.2 CHRP Multiboot.....	17
16	3.1.3 Bootinfo Configuration Variables.....	18
17	3.1.4 Bootinfo Properties.....	18
18	3.1.5 Standard Locations for Bootinfo Objects.....	18
19	3.1.6 Bootinfo Objects.....	19
20	3.1.6.1 Bootinfo Entities.....	20
21	3.1.6.2 Bootinfo Character Sets.....	20
22	3.1.6.3 Element Tag Descriptions.....	21
23	3.1.6.4 CHRP-BOOT Element.....	21
24	3.1.6.5 OS-NAME element.....	21
25	3.1.6.6 BOOT-SCRIPT element.....	21
26	3.1.6.7 ICON element.....	21
27	3.1.6.7.1 BITMAP element.....	21
28	3.1.7 Multiboot Menu.....	22
29	3.2 Reboot-Command Variable Description.....	22
30	4. CHRP PowerPC Processor.....	23
31	4.1 Processor Endian-ness Support.....	23
32	4.1.1 Bi-Endian Booting.....	23
33	5. Open Firmware Platform Extensions.....	23
34	5.1 Open Firmware Root Node.....	24
35	5.1.1 Root Node Properties.....	24
36	5.1.2 Root Node Methods.....	25
37	5.1.3 ROM Node(s).....	25
38	5.1.3.1 ROM Node Properties.....	25
39	5.1.3.2 ROM Node Methods.....	26
40	5.1.4 ROM Child Node(s).....	27
41	5.1.4.1 ROM Child Node Properties.....	27
42	5.1.4.2 ROM Child Node Methods.....	27
43	5.2 Run Time Abstraction Services (RTAS) Node.....	28
44	5.2.1 RTAS Node Properties.....	28
45	5.2.2 RTAS Function Property Names.....	29
46	5.2.3 RTAS Node Methods.....	30
47		
48		
49		
50		
51		
52		
53		
54		

1	5.3 PCI Host Bridge Nodes . . . . .	30
2	5.3.1 PCI Host Bridge Properties . . . . .	31
3	5.3.2 Methods for PCI Host Bridge Nodes . . . . .	32
4	5.4 Memory Controller Nodes . . . . .	33
5	5.4.1 Memory Controller Node Properties . . . . .	33
6	5.5 Interrupt Controller Nodes . . . . .	34
7	5.5.1 Open PIC Interrupt Controller Node Properties . . . . .	34
8	5.6 Additional Node Properties. . . . .	35
9	5.6.1 Interrupt Properties . . . . .	35
10	5.6.2 Miscellaneous Node Properties . . . . .	35
11	5.7 Aliases Node . . . . .	36
12	5.8 Event Sources Node . . . . .	37
13	5.8.1 Child nodes of the Event Sources Node . . . . .	37
14	5.8.1.1 internal-errors . . . . .	37
15	5.8.1.2 epow-events . . . . .	37
16	5.8.1.3 power-management-events . . . . .	37
17	5.9 Reserved Node . . . . .	38
18	5.10 '/chosen' Node . . . . .	38
19	6. Symmetric Multi-Processors(SMP) . . . . .	38
20	6.1 SMP Platform Device Tree Structure . . . . .	38
21	6.2 SMP Properties . . . . .	39
22	6.2.1 Processor Node . . . . .	39
23	7. Device Power Management Properties/Methods . . . . .	39
24	7.1 System Node Properties . . . . .	40
25	7.1.1 Properties assigned to the RTAS node. . . . .	40
26	7.1.2 Properties of the power-management-events node . . . . .	41
27	7.2 Device Properties . . . . .	41
28	7.2.1 Properties for Power Domain Control Points . . . . .	43
29	7.3 Power Management Related Methods . . . . .	43
30	7.4 Power Management NVRAM Partition . . . . .	43
31	8. ISA Device Support . . . . .	43
32	8.1 Additional ISA Requirements. . . . .	44
33	8.1.1 ISA Addresses . . . . .	45
34	8.1.2 ISA Interrupts and DMA Channels . . . . .	45
35	8.1.3 ISA Unit Address Format . . . . .	45
36	8.1.4 ISA Additional Properties . . . . .	45
37	9. Configuration of Platform Resources . . . . .	45
38	9.1 ISA Resource Configuration. . . . .	46
39	9.1.1 ISA Resource Configuration Description . . . . .	46
40	9.1.1.1 ISA Configuration Handle Assignment . . . . .	46
41	9.1.2 ISA Legacy Configuration Data Format Requirements. . . . .	47
42	9.2 Power Management Resource Configuration . . . . .	49
43	9.2.1 Power Management Information Utility . . . . .	49
44	9.2.2 PM Configuration Process . . . . .	49
45	9.2.3 PM Configuration Format . . . . .	49
46	10. Client Program Requirements. . . . .	51
47		
48		
49		
50		
51		
52		
53		
54		



1	10.1 Load Address . . . . .	51
2	10.2 Initial Register Values. . . . .	52
3	10.3 I/O Devices State . . . . .	52
4	10.4 Client Program Format . . . . .	52
5	10.4.1 ELF-Format . . . . .	52
6	10.4.1.1 ELF Note Section . . . . .	52
7	10.4.1.2 Recognizing ELF-Format Programs . . . . .	53
8	10.4.1.3 Preparing ELF-Format Programs for Execution . . . . .	54
9	10.5 Additional Client Interface Requirements . . . . .	55
10	10.5.1 Client Interface Callbacks . . . . .	55
11	10.5.1.1 Real-Mode Memory Management Assist Callbacks . . . . .	55
12	10.5.1.2 Virtual Address Translation Assist Callbacks . . . . .	55
13	10.5.2 Client Interface Services . . . . .	55
14	11. Support Packages . . . . .	56
15	11.1 "disk-label" Support Package . . . . .	56
16	11.1.1 Media Layout Format . . . . .	56
17	11.1.1.1 FDISK Partition Types . . . . .	56
18	11.1.2 Open Method Algorithm . . . . .	57
19	11.2 "tape-label" Support Package . . . . .	63
20	11.2.1 Tape Format. . . . .	63
21	11.2.2 Tape bootinfo.txt File. . . . .	63
22	11.3 "network" Support Package . . . . .	64
23	11.4 Program-image formats. . . . .	64
24	11.4.1 CHRP PE Program-image Format Support . . . . .	65
25	11.4.1.1 Windows NT Veneer Implementation . . . . .	65
26	12. Open Firmware Related NVRAM Usage . . . . .	65
27	12.1 Open Firmware Partition. . . . .	66
28	12.2 Firmware Partition . . . . .	66
29	12.3 Hardware Partition . . . . .	66
30	12.4 System Partition . . . . .	66
31	12.4.1 Name . . . . .	67
32	12.4.2 Value . . . . .	67
33	12.4.3 Open Firmware Configuration Variables. . . . .	67
34	12.4.3.1 Boolean Configuration Variables . . . . .	67
35	12.4.3.2 Integer Configuration Variables . . . . .	68
36	12.4.3.3 String Configuration Variables . . . . .	68
37	12.4.3.4 Byte Configuration Variables . . . . .	68
38	12.4.4 nvramrc Script . . . . .	68
39	12.4.4.1 Multi-initiator Script . . . . .	69
40	12.5 Configuration Partition . . . . .	70
41	12.5.1 ISA PnP Configuration Resources. . . . .	70
42	12.5.2 Device Power Management Configuration Resources . . . . .	70
43	12.5.3 Error Log Partition. . . . .	70
44	12.6 Vendor-Defined Partition . . . . .	70
45	12.7 Free Space Partition . . . . .	70
46		
47		
48		
49		
50		
51		
52		
53		
54		

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

---

## List of Tables:

1		
2		
3		
4		
5		
6		
7	Table 1.....	Standard Pathnames for bootinfo.txt file 18
8	Table 2.....	Semantics of device state values 42
9	Table 3.....	Combinations of Device Power State/Domain Power Level 42
10	Table 4.....	ISA Open Firmware Resource Information 44
11	Table 5.....	ISA Card Configuration Directory 47
12	Table 6.....	ISA Card PnP-like Small Data Item Tag Format 47
13	Table 7.....	ISA Card PnP-like Large Data Item Tag Format 48
14	Table 8.....	Power Management Configuration Data Header 50
15	Table 9.....	Data Block Format 50
16	Table 10.....	Node Data Block Format 50
17	Table 11.....	Property Data Block Format 51
18	Table 12.....	Open Firmware Related NVRAM Partitions 66
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		

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

## 1. Overview

This document specifies the application of *IEEE Std 1275-1994 Standard for Boot (Initialization, Configuration) Firmware, Core Practices and Requirements, Core Errata, IEEE P1275.7* and appropriate Open Firmware Standards for PowerPC Microprocessor CHRP Computer Systems[3], including practices for client program interface and data formats.

### 1.1. General Requirements for Open Firmware

An Open Firmware implementation for a PowerPC Microprocessor CHRP platform *shall* implement the core requirements as defined in [1], core errata[2], the PowerPC Processor-specific extensions described in [6], other appropriate bindings and/or recommended practices contained in the references (Section 2.1), and the PowerPC Microprocessor CHRP binding specific extensions described in this document.

In addition, an Open Firmware implementation for a PowerPC Microprocessor CHRP platform *shall* implement the *Device Interface, Client Interface* and *User Interface* as defined in Open Firmware core document, reference [1].

The Open Firmware Documentation can be found on servers, ‘[playground.sun.com](http://playground.sun.com)’ and ‘[chrp.apple.com](http://chrp.apple.com)’. The URL for this documentation is the following:

<http://playground.sun.com/1275>

or

<http://chrp.apple.com/1275>

Open Firmware Documentation can also be found on an ftp server. The ftp server has the Open Firmware Documentation plus more archived documents. The URL for this documentation is the following:

<ftp://playground.sun.com/pub/1275>

Versions of the Open Firmware Documentation are maintained in sub-directories under the *1275* directory on a topic basis. For example, under the *bindings/* sub-directory, the *isa/*, *pci/* and *ppc/* sub-directories have versions of the *ISA/EISA/ISA-PnP binding*, *PCI Bus binding* and *PowerPC Processor binding* respectively. The sub-directory *postscript/* contains the latest postscript versions of the Open Firmware Documentation.

For ISA support, the CHRP System binding has additional restrictions and/or requirements imposed on the *ISA/EISA/ISA-PnP binding*(See “Additional ISA Requirements” on page 44.). Also, for completeness, the ISA I/O Device bindings ([24],[25],[26],[27],[28],[29],[30],[31] and [32]) are listed in the references. The I/O device bindings support the CHRP Standard I/O specified in the I/O Device Reference Document[4].

## 2. References and Terms

References and terms for the CHRP System binding are listed in this section.

### 2.1. References

This Open Firmware System binding 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] *Core Errata, IEEE P1275.7/D4.*

1 [3] *PowerPC Microprocessor Common Hardware Reference Platform: A System Architecture*. This document  
2 describes the platform specifications; hardware attributes, characteristics and I/O Sub-system content. The URL  
3 for the document is: <http://www.austin.ibm.com/tech/chrp/index.html>.

4 [4] *PowerPC Microprocessor Common Hardware Reference Platform: I/O Device Reference*. This document  
5 describes the platform Standard I/O Devices; hardware registers, register locations, and hardware attributes. The  
6 URL for the document is: <http://www.austin.ibm.com/tech/chrpio/index.html>.

7 [5] *CHRP Architecture Version 1.0 Change Notice*, 08/01/96. This document contains architectures changes to  
8 the base CHRP Architecture. The URL for this document is:  
9 [ftp://www.austin.ibm.com/pub/technology/spec/chrp/CHNG\\_C10.PS.Z](ftp://www.austin.ibm.com/pub/technology/spec/chrp/CHNG_C10.PS.Z)

10 [6] *PowerPC Processor binding to: IEEE Std 1275-1994, Standard for Boot (Initialization, Configuration)*  
11 *Firmware*, Version 2.0.

12 [7] *PCI Bus binding to: IEEE Std 1275-1994, Standard for Boot (Initialization, Configuration) Firmware*,  
13 Version 2.1.

14 [8] *ISA/EISA/ISA-PnP binding to: IEEE Std 1275-1994, Standard for Boot (Initialization, Configuration)*  
15 *Firmware*, Version 0.4.

16 [9] *PC Card binding to: IEEE Std 1275-1994, Standard for Boot (Initialization, Configuration) Firmware*,  
17 Version 1.2.

18 [10] *Open Firmware: Recommended Practice - Device Support Extensions*, Version 0.8.

19 [11] *Open Firmware: Recommended Practice - 16-color Text Extensions*, Version 1.2.

20 [12] *Open Firmware: Recommended Practice - 8-bit Graphics Extensions*, Version 1.2.

21 [13] *Open Firmware: Recommended Practice - Forth Source and FCode Image Support*, Version 1.0. This  
22 document describes program-image formats for FCode and Forth Source Code.

23 [14] *Open Firmware: Recommended Practice - Generic Names*, Version 1.3. This document describes a naming  
24 convention for Open Firmware device nodes.

25 [15] *Open Firmware: Recommended Practice - Interrupt Mapping*, Version 1.0. This document describes an  
26 interrupt structure for Open Firmware.

27 [16] *Open Firmware: Recommended Practice - TFTP Booting Extensions*, Version 0.8. This document describes  
28 network extensions to the OBP-TFTP Support Package.

29 [17] *Open Firmware: Recommended Practice - Interposition*, Version 0.2. This document describes a standard  
30 technique which allows file system services to be layered (interposed) on top of unmodified FCode device  
31 drivers. The recommended practice creates a new FCode and a new client interface service.

32 [18] *MS-DOS Programmer's Reference*, published by Microsoft. This document describes the MS-DOS  
33 partition, directory and FAT formats used by disk-label support package.

34 [19] *Peering Inside the PE: A Tour of the Win32 Portable Executable File Format*, found in the March, 1994  
35 issue of *Microsoft Systems Journal*.

36 [20] *ISO-9660, Information processing -- Volume and file structure of CD-ROM for information interchange*,  
37 published by International Organization for Standardization.

38 [21] *System V Application Binary Interface, PowerPC Processor Supplement*, Sunsoft. This document defines the  
39 PowerPC ELF format.

40 [22] *Inside Macintosh, SCSI Disk Partition Definition, CD-ROM ISBN 0-201-94591-6*.

41 [23] *Open PIC Multiprocessor Interrupt Controller Register Interface Specification*, Revision 1.2. Refer to URL  
42 of <http://www.amd.com/html/products/pcd/openpic/19725c.pdf> for the document.

- 1 [24] *CHRP ISA Keyboard/Mouse Controller Device binding*, Version 1.0.  
2 [25] *CHRP ISA DMA Controller Device binding*, Version 1.0.  
3 [26] *CHRP ISA Parallel Port Device binding*, Version 1.0.  
4 [27] *CHRP ISA Audio Device binding*, Version 1.0.  
5 [28] *CHRP ISA Interrupt Controller Device binding*, Version 1.1.  
6 [29] *CHRP ISA Serial Port Device binding*, Version 1.0.  
7 [30] *CHRP ISA Floppy Controller Device binding*, Version 1.0.  
8 [31] *CHRP Linear Frame Buffer Display Device binding*, Version 1.0.  
9 [32] *CHRP VGA Display Device binding*, Version 1.0.  
10 [33] *ISO Standard 8879:1986, Information Processing -- Text and Office Systems -- Standard Generalized  
11 Markup Language(SGML).*  
12 [34] *Plug and Play ISA Specification*, Version 1.0a, 05/05/1994.  
13 [35] *Clarification to the Plug and Play ISA Specification*, Version 1.0a, 12/10/94.

## 2.2. Terms

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

**ARP:** Address Resolution Protocol

**BOOTP:** Bootstrap Protocol

**CHRP:** Common Hardware Reference Platform, also PowerPC Platform.

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

**core errata:** Refers to *Core Errata, IEEE P1275.7*

**CPU:** Central Processing Unit

**ELF:** Executable and Linking Format. A binary object file format defined by [21] 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 [18].

**gateway:** Network connecting device

**host:** A computer. In particular a source or destination of messages from the point of view of the communication network.

**ICMP:** Internet Control Message Protocol

**ICU:** ISA Configuration Utility; Refers to CHRP platform program to assist a user to configure ISA resources.

**IETF:** Internet Engineering Task Force

**IP:** Internet Protocol

**IO:** Input/Output

**LAN:** Local Area Network

1 **NVRAM:** Non-volatile memory that is the repository for various platform, Open Firmware and operating system  
2 information that remains persistent across reboots, power management activities and/or cycles.

3 **Open Firmware:** The firmware architecture defined by the core specification[1] and errata[2], or, when used as  
4 an adjective, a software component compliant with the core specification and errata.  
5

6 **OpenPIC:** Is a register-level architectural definition of an interrupt controller that supports up to 32  
7 processors(Refer to [23] for more information).

8 **PCU:** Power Configuration Utility; Refers to CHRP platform program to assist a user to manage device power.  
9

10 **PE:** Portable Executable. A binary object file format defined by [19]; this format is used by Windows NT™  
11 operating system.

12 **PROM:** programmable read only memory

13 **real-mode:** The mode in which Open Firmware and its client are running with translation disabled; all addresses  
14 passed between the client and Open Firmware are real (i.e., hardware) addresses.  
15

16 **RFC:** Internet Request For Comments; part of the technical process of establishing a standard.  
17

18 **ROM:** Read Only Memory

19 **suspend:** A form of Power Management characterized by a fast recovery to full operation. Typically, system  
20 memory will not be powered off while in the suspend state.  
21

22 **TFTP:** Trivial File Transfer Protocol

23 **UDP:** User Datagram Protocol

24 **virtual-mode:** The mode in which Open Firmware and its client share a single virtual address space, and  
25 address translation is enabled; all addresses passed between the client and Open Firmware are virtual (translated)  
26 addresses.  
27

### 30 **3. CHRP Boot Flow**

31 This section gives a system boot process overview and defines the enhancements to the standard Open Firmware  
32 boot process that are present in the boot process for a CHRP platform.  
33

#### 34 **3.1. CHRP Boot Overview**

35 The CHRP platform performs a normal Open Firmware boot(See [1], *IEEE Std 1275-1994 Standard for Boot*  
36 *(Initialization, Configuration) Firmware, Core Practices and Requirements*) as stated in the Core Practice  
37 Document, Section 4.2.3, Start-up script evaluation. CHRP platforms provide an additional capability to assist  
38 the user in choosing which of several operating systems to boot. A key sequence can be used to interrupt the  
39 normal boot flow and present the user with a *multiboot menu*, which can be either graphical or text-based at the  
40 discretion of the platform's firmware, from which the user can choose one of the installed or installable  
41 operating systems. Presenting the user with this choice can also be made the default mode of operation at  
42 platform boot time, by means of the **auto-boot?** and **menu?** configuration variables.  
43

44 An overview of a CHRP platform boot sequence and the additions of the *multiboot menu* are given below:  
45  
46  
47  
48  
49  
50  
51  
52  
53  
54



### CHRP Start-up Script Sequence Evaluation:

- a) Power On Self Test(POST)
- b) System Initialization
- c) Evaluate the *script* (if **use-nvramrc?** is true)
- d) **probe-all** (evaluate FCode)
- e) **install-console**
- f) **banner**

-----

↓ (key chord for *multiboot menu* will be recognized and acted upon after this point) ↓

- g) Secondary Diagnostics and other system-dependent initialization
- h) Default boot (if **auto-boot?** is true)
- i) Entry to *multiboot menu* (if **menu?** is true)
- j) Invoke the command interpreter (if preceding step returns)

The CHRP boot flow described above occurs after all of the devices have been probed (i.e., by the execution of **probe-all**); See section 3.1.1. additional requirements for **probe-all** method.

The CHRP boot sequence defaults to a normal boot if the boolean variable **auto-boot?** is true and **diagnostic-mode?** is false. In this situation, the system *shall* then boot from information contained in the configuration variables **boot-device** and **boot-file**.

From the boot sequence above, entry to the *multiboot menu* may occur anywhere after step ‘f’, **banner**, if the platform key sequence(*multiboot menu*) has been depressed or in step ‘i’ if the boolean variable **menu?** is true.

#### 3.1.1. Additional Requirements for **probe-all** Method

Before probing for plug-in devices, Open Firmware *shall* execute the **probe** method, as with **execute-device-method**, of any built-in device nodes. The order of evaluation *shall* ensure that the **probe** method of a parent device node is executed before the **probe** method of any of its children.

**Note: During this built-in probing, /rom nodes will locate ROM based OSs. The FCode for these devices can publish their “bootinfo” properties that are used during the multiboot scenario as described below.**

#### 3.1.2. CHRP Multiboot

The boot choices identified to the user are defined by *bootinfo objects* which are located on various system media. Each *bootinfo object* contains information about one operating system, such as its name and description, an icon depicting it, and an Open Firmware command sequence to load and execute it. The locations where *bootinfo objects* can be found are specified by Open Firmware *device-specifiers* that are the values of configuration variables, the names of which are of the form “**bootinfo-*nnnnn*”**, where “*nnnnn*” is OS-specific. These *configuration variables* are stored in the System Partition in NVRAM and are published in the *device tree* as properties under the **/options** node. The *multiboot menu* will use these *configuration variables* to locate and parse *bootinfo* to obtain the OS icon, description, etc.

In addition to the `bootinfo-nnnnn` configuration variables, the *multiboot menu* will search the device tree for nodes containing `"bootinfo"` properties, which specify that the node can supply a *bootinfo object*. This is particularly useful for operating systems contained in ROMs.

**Note: The order prescribed by `probe-all` guarantees that these properties be created before the *multiboot menu* has been invoked.**

Different versions of the same operating system may each have their own *bootinfo* and associated *configuration variables*. Although it is possible to put *bootinfo* in any media location that Open Firmware can read, this specification defines standard locations for various types of media, to allow the firmware to establish the *bootinfo configuration variables* automatically in many cases.

### 3.1.3. Bootinfo Configuration Variables

A *bootinfo configuration variable* is any *configuration variable* that meets the following requirements:

- Its name is of the form `"bootinfo-nnnnn"`, where *nnnnn* is a string of at most 22 characters from the set of valid characters for Open Firmware configuration variable names. The exact value of `"nnnnn"` for a particular operating system may be chosen by that operating system. The naming convention for the operating system should be chosen to avoid possible naming conflicts between operating system vendors (See Section 12.4.1. on page 67.)
- Its value is an Open Firmware *device-specifier* that identifies an object (e.g. disk file, tape file, disk partition or `/rom` child node) whose contents are a *"bootinfo object"* as defined below.

### 3.1.4. Bootinfo Properties

Any node in the device tree can have a `"bootinfo"` property whose value specifies the arguments to use in opening that device in order to access its *bootinfo object*.

#### `"bootinfo"`

S

*prop-name*, locates the node's *bootinfo object*

*prop-encoded-array*: A string, encoded as with `encode-string`

The presence of this property signifies that the device has an associated *bootinfo object*. The value is a text string such that when this device's node `open` method is called, the value of text string that is passed to the device's node `open` method is `my-args`. When so opened, subsequent calls to the node's `"read"` method will yield the contents of the node's *bootinfo object*.

### 3.1.5. Standard Locations for *Bootinfo Objects*

The standard locations for *bootinfo objects* on various CHRP media and partition types is shown in the table below. An operating system must put its *bootinfo object* in the standard location in order to guarantee interoperability with the CHRP *multiboot menu* mechanism.

Table 1. Standard Pathnames for *bootinfo.txt* file

Name	Device/Partition	Notes
<b>Installation Media:</b>		
Any block device:	<i>device:partition</i> , <code>\ppc\bootinfo.txt</code>	Any file system format

Name	Device/Partition	Notes
Tape:	<i>device:0</i> (Note 1)	Presence of <i>bootinfo.txt</i> is optional
ROM:	<i>device:bootinfo</i>	<i>bootinfo</i> is the value of the " <b>bootinfo</b> " property in a <i>/rom</i> child node
Network:	Could specify <i>bootinfo.txt</i> or some other file from the Bootp server	Specifying <i>bootinfo.txt</i> from the Bootp server is optional

Note 1: If *bootinfo.txt* file is not present, file 0 should contain a program image file for a bootable tape.

### Example of installed("**bootinfo-nnnnn**") block device(disk):

#### ALIAS EXAMPLE:

*bootinfo-aix-4.3=disk:2* (The contents of partition 2, which is probably a "0x41" partition, on the default disk, is the *bootinfo.txt* file for a version of the AIX Operating System.)

*bootinfo-nt-4.0=disk:\os\winnt\bootinfo.txt* (The file *\os\winnt\bootinfo.txt* on the default partition of the default disk is the *bootinfo.txt* file for a version of the Windows NT Operating System.)

#### NON-ALIAS EXAMPLE:

*bootinfo-aix-4.4=/pci@ff500000/pci3,1000@10/sd0,0:3* (The contents of partition 3, which is probably a "0x41" partition, on the SCSI disk at target 0 unit 0, is the *bootinfo.txt* file for a version of the AIX Operating System.)

### 3.1.6. Bootinfo Objects

The information used by Open Firmware to display information in the *multiboot menu* and to locate and process an OS load image is contained within a sequence of text that is called a *bootinfo object*. The text comprising the *bootinfo object* uses SGML[33] syntax, with tags identifying the subordinate elements.

The following outline is a summary of the organization of the *bootinfo object*. Elements at the same level do not have any required order. The tags are illustrated in upper case, but *shall* be processed in a case-insensitive manner.

```
<CHRP-BOOT>
  <DESCRIPTION>
    . . . .
  </DESCRIPTION>
  <OS-NAME>
    . . . .
  </OS-NAME>
```

```

1      <BOOT-SCRIPT>
2
3      . . . .
4      </BOOT-SCRIPT>
5
6      <ICON
7          SIZE=ww, hh
8
9          COLOR-SPACE=r, g, b
10
11         >
12             <BITMAP>
13                 hh hh hh hh . . . .
14             </BITMAP>
15         </ICON>
16
17 </CHRP-BOOT>

```

**Note 1:** If ‘SIZE’ is not present, assume default of 64,64. If ‘COLOR-SPACE’ is not present, assume default of 3,3,2.

**Note 2:** Another <chrp-boot> tag sequence could define a different boot selection

**Note 3:** CHRP platforms will recognize only the tags between the beginning <chrp-boot> tag until the end </chrp-boot> tag. If a tag is unrecognized, the material will be ignored until the end tag. Other non-<chrp-boot> tags may be supported in the future. These additional selections would also be presented to the user as boot options.

### 3.1.6.1. Bootinfo Entities

SGML provides “entities” that provide symbolic names for text. When the entity names are contained within & and ‘;’, the entity is replaced with text as defined by the entity; i.e., entities provide a “macro” substitution capability. The *bootinfo object* may use entities to supply pathname components that depend upon the location of the file. Also, entities have been defined for the standard SMGL Tags for the presence of the ‘<’, ‘&’ and ‘>’ characters in the text as &lt; ; , & ; and &gt; ;. Within the <BOOT-SCRIPT> element, the following entities are defined with respect to the fully qualified pathname of the *bootinfo object*:

device	the device component.
partition	the partition component.
directory	the directory component.
filename	the filename component.
full-path	the entire fully qualified pathname.

The fully qualified pathname could be represented by the following text:

&device;[&partition;][,]&directory;\&filename;

**Note:** underlined portions illustrate where entities are positioned within the full pathname

### 3.1.6.2. Bootinfo Character Sets

The character set used by the bootinfo.txt file is ISO-8859-1 (Latin-1). Element tags and entity names are not case sensitive; all other text is case sensitive.



```

1     <icon size=64,64 color-space=3,3,2><bitmap>hh ... hh1</bitmap></icon>
2
3 </chrp-boot>

```

#### AIX Diagnostics Bootinfo Object Example:

```

4 <chrp-boot>
5
6     <description>AIX 4.2.D.0 Diagnostics</description>
7
8     <os-name>AIX 4.2.D.0 Diagnostics</os-name>
9
10    <boot-script>boot &device;:2 diag</boot-script>
11
12    <icon size=64,64 color-space=3,3,2><bitmap>hh ... hh1</bitmap></icon>
13
14 </chrp-boot>

```

Note 1: 64x64 icon size would have 8192 hex string characters

### 3.1.7. Multiboot Menu

If the boot sequence is interrupted by the multiboot key sequence, then the firmware *shall* present a *multiboot menu* that provides at least the functions listed below. The form of the menu (e.g. graphical or text- oriented) and the selection mechanism (e.g. numbered choices, arrow keys, or mouse) are platform-dependent.

#### Multiboot Required Functions:

- Locate all bootinfo objects specified by bootinfo configuration variables and device node **"bootinfo"** properties. For each *bootinfo object*, present a choice corresponding to each valid <chrp-boot> section contained therein. For each such choice, allow the user to either:
  - Execute the contents of that *bootinfo object's* <boot-script> element.
  - Set the **"boot-command"** configuration variable to the contents of that *bootinfo object's* <boot-script> element.
- Present a choice corresponding to each install device, which, when invoked, will attempt to locate a bootinfo object at the device's standard location (see Table 1).
- Allow the user to manage configuration variables
- Allow the user to invoke the Open Firmware user interface

Additional options that could be implemented would be to provide a means to get to diagnostics or specific platform options.

There *shall* be at least one key sequence to enter the multi-boot platform function for a CHRP platform.

**Note: Operating System have the responsibility to update the NVRAM System Partition Variable to reflect a change where the *bootinfo.txt* file is located; e.g., moving to a different disk device. Also, the OS is responsible for maintaining the contents of the *bootinfo.txt* file.**

## 3.2. Reboot-Command Variable Description

The OS can cause Open Firmware to execute a specified sequence of commands at the next boot by setting the value of the **"reboot-command"** configuration variable. CHRP firmware implementations *shall* implement the following configuration variable.

```
reboot-command( -- addr len ) N
```

One time or temporary reboot command.

The value of this configuration variable is a string consisting of zero or more lines of text, with lines separated by either <return>, <linefeed>, or <return><linefeed>.

1 During firmware start-up, just prior to checking the `auto-boot?` configuration variable for automatic  
2 booting, the firmware *shall* check the value of `reboot-command`. If the value is not the empty string,  
3 the firmware *shall* save the value to a temporary location, set `reboot-command` to the empty string, and  
4 evaluate the saved value as though it were a series of user interface command lines.

5 If the evaluation of `reboot-command` returns without executing, the firmware *shall* proceed with its  
6 normal start-up sequence. In typical usage, however, the value of `reboot-command` will include a `boot`  
7 command that starts a client program and does not return.

## 8 9 **4. CHRP PowerPC Processor**

10 Open Firmware defines a minimum cell size of 32 bits; therefore, only one cell is necessary to represent  
11 addresses up to 4GB (32 bits). Two cells are necessary to represent addresses above 4GB and within 64 bits.  
12 Also, two cells are necessary to represent sizes greater than 4GB.

### 13 14 15 **4.1. Processor Endian-ness Support**

16 An implementation of the PowerPC Open Firmware *shall* support both Big- and Little-Endian system  
17 implementations. This section describes added features to the core Open Firmware architecture features to  
18 support bi-endian booting.

#### 19 20 **4.1.1. Bi-Endian Booting**

21 The Configuration Variable `little-endian?` must be implemented. The basic concept of bi-endian support is  
22 to keep in the `little-endian?` variable a "cached" indication of the desired endian-ness of client programs  
23 (for example operating systems or their loaders). This variable indicates the expected endian-mode of a boot  
24 target; false (0) indicates Big-Endian, true (-1) indicates Little-Endian; the default value of `little-endian?`  
25 is implementation dependent.

26  
27 The client program must describe its endian-mode in the header section of its image as described in  
28 Section 10.4. on page 52. When Open Firmware is started, Open Firmware *shall* use the value of `little-`  
29 `endian?` to establish the endian-mode of a system. After Open Firmware locates and loads a client program,  
30 the correct endian-mode must be verified with the description in the header section of the client program image.  
31 If the cached value is correct, Open Firmware proceeds, with the system in its current endian-mode.

32  
33 If, however, Open Firmware determines that the endian-mode of the client program is different from what it had  
34 assumed, it must set `little-endian?` appropriately and reconfigure the system (hardware and firmware) for  
35 the new endian-mode, possibly by resetting the system as with `reset-all`.

36 **Note: The endian-mode applies to the entire hardware platform, including the processor(s).**  
37 **Open Firmware *shall* perform whatever steps are required to enable the platform to run**  
38 **in the specified mode.**

39 Client programs can use `setprop` to alter the value of `little-endian?`; users can alter it via the `setenv`  
40 command from the user interface (if present). The alteration of `little-endian?` *shall* not cause the platform  
41 to be reconfigured until the platform is re-booted.

42  
43 **Note: This mechanism introduces an extra configuration pass. However, this occurs only**  
44 **when switching the endian-mode from that which was last used. For most boots, Open**  
45 **Firmware will be appropriately configured, so that no additional overhead will occur.**

## 46 47 **5. Open Firmware Platform Extensions**

48 This section defines Open Firmware properties, methods, device tree structure and Client Interface Service  
49 requirements for a CHRP platform.

## 5.1. Open Firmware Root Node

This section defines additional properties and methods associated with CHRP platforms that CHRP operating systems expect to find in the root node. The unit address of all non-system nodes that are children of the root node *shall* have the same value each time the platform is booted; i.e., *shall* be invariant for each boot process.

**Note:** This requirement ensures that the PHB would have a stable unit address. Violations of this rule may require reinstallation of an OS.

**Note:** Typically, the root node corresponds to the common processor bus in a PowerPC platform.

### 5.1.1. Root Node Properties

This section defines the additional properties which *shall* be present in the root node unless otherwise specified.

**"#address-cells"** S

Standard property, encoded as with `encode-int`, that specifies the number of cells required to represent physical addresses on the processor bus. The value of "#address-cells" for the processor bus shall be 1 or 2 depending on whether there is any memory addressable at or above 4GB's.

**"#size-cells"** S

Standard property, encoded as with `encode-int`, that specifies the size of cells required to represent physical addresses on the processor bus. The value of "#size-cells" for the processor bus shall be 1 or 2 depending on whether there is any memory addressable at or above 4GB's.

**"clock-frequency"** S

Standard *property*, encoded as with `encode-int`, that represents the primary system bus speed (in hertz).

**"system-id"** S

Standard *property*, encoded as with `encode-string`, that contains the identification of the computer system (Reference the **"name"** property in the core specification[1]). This string should be unique across all systems and all manufacturers. An example of an address of this form is "0nnnnnnnnmmmmmmmm" where nnnnnnn is a sequence of 6 uppercase hexadecimal digits representing a 24-bit value that identifies manufacturer and mmmmmmm is a sequence of 6 uppercase hexadecimal digits representing a 24-bit binary number assigned by the manufacturer to assure uniqueness.

**Note:** For platforms with built-in ethernet or other IEEE 802-style interfaces, the 6-byte MAC address assigned to that interface meets the requirements and could be used as the system-id.

**"model"** S

*prop-name* is a printable string identifying the manufacturer and model number of the platform.

*prop-encoded-array*: Text string, encoded as with `encode-string`.

The value of this property is a vendor dependent string which identifies this platform via its manufacturer and model number.

**"device\_type"** S

*prop-name* is a printable string identifying the platform as CHRP Compliant.

*prop-encoded-array*: Text string, encoded as with `encode-string`.

The value of this property is a string, "chrp" which identifies the platform is CHRP Compliant.

**"pc-emulation"** S

**Note:** The recommended procedure to determine if the PC Emulation Function is present is to do a "test-method" (Refer to Section 10.5.2. on page 55) for the `set-pc-emulation` method. This property could be deleted in future CHRP Open Firmware Architecture Versions.

*prop-name* indicates the presence of the PC Emulation option for PCI Host Bridge 0 and the Memory Controller(s).



1 *prop-encoded-array*: <none>

2 This property *shall* be present if PCI Host Bridge 0 and the Memory Controller(s) support the PC emulation  
3 option, otherwise this property *shall* be absent. (The option may be enabled via the **set-pc-emulation**  
4 method, if the PC Emulation option is present, see Section 5.1.2. on page 25.)

5 **"exception-relocation-size"** **S**

6 *prop-name* specifies the fixed amount of address space above physical address 0xFFFF0000 which is  
7 relocated down into System Memory when the PC Emulation option is enabled.

8 *prop-encoded-int*: Integer, encoded as with **encode-int**.

9 The granularity of **"exception-relocation-size"** properties value is 4 KB and the minimum size  
10 is 12 KB. This property *shall* be present if the PC Emulation capability is available, otherwise this property  
11 *shall* be absent.

12 **"processor-hole"** **S**

13 **Note: The recommended procedure to determine if the Processor Hole Function is present**  
14 **is to do a "test-method" (Refer to Section 10.5.2. on page 55) for the "set-processor-hole"**  
15 **Method. This property could be deleted in future CHRP Open Firmware Architecture**  
16 **Versions.**

17 *prop-name* indicates that this PCI Host and Memory Controller has the optional processor-hole capability.

18 *prop-encoded-array*: <none>

19 This property *shall* be present if PCI Host Bridge 0 and the Memory Controller(s) support the processor-  
20 hole, otherwise this property *shall* be absent. (The processor hole may be enabled via the **set-**  
21 **processor-hole** method if the capability is supported, see Section 5.1.2. on page 25.)

22 **"platform-open-pic"** **S**

23 *prop-name* indicates the system Open Firmware Interrupt Controller physical addresses.

24 *prop-encoded-array*: List of system *phys-addr* values

25 This property value is a list of system physical addresses corresponding to **"reg"** property of the system  
26 OpenPIC Interrupt Controller

## 31 5.1.2. Root Node Methods

32 This section defines methods associated with the platform via "/" (the *root* node).

33 **set-processor-hole** ( true | false -- ) **M**

34 Enable(true) or disable(false) the processor-hole on PCI Host Bridge 0 and the Memory Controller(s). This  
35 method *shall* be present if the processor-hole capability is supported.

36 **set-pc-emulation** ( err-value tem-value enable? -- ) **M**

37 If the flag *enable?* is true, enable the PC Emulation option on PCI Host Bridge 0 and the Memory  
38 Controller(s). Also disable the system-memory-alias on PCI Host Bridge 0 and the Memory Controller(s).  
39 The input value *err-value* is loaded into the Exception Relocation Register (ERR) in the PCI Host Bridge(s)  
40 and Memory Controller(s). The input value *tem-value* is loaded into the Top of Emulated Memory (TEM)  
41 boundary in the PCI Host Bridge(s) and Memory Controller(s). If the flag *enable?* is false, the PC  
42 Emulation option on PCI Host Bridge 0 and the Memory Controller(s) is disabled. The state of PCI Host  
43 Bridge 0's initial memory aliases is undefined after it is disabled. This method *shall* be present if the PC  
44 Emulation option is present. The units for *err-value* and *tem-value* are in bytes.

## 47 5.1.3. ROM Node(s)

48 The ROM Node(s) *shall* be a child or children of the root node.

### 51 5.1.3.1. ROM Node Properties

52 Each ROM Node *shall* have the following properties:

1	<b>"name"</b>	<b>S</b>
2	Standard <i>prop-name</i> : The value of this property <i>shall</i> be "rom".	
3		
4	<b>"reg"</b>	<b>S</b>
5	Standard <i>prop-name</i> to define a <i>unit-address</i> for the node.	
6	<i>prop-encoded-array</i> : One ( <i>phys-addr</i> ; <i>size</i> ) pair.	
7		
8	The <i>phys-addr</i> of this property <i>shall</i> be the starting physical address of this ROM and the <i>size</i> value <i>shall</i>	
9	be 0. The <i>size=0</i> prevents a conflict with the <b>"reg"</b> of this node's children.	
10	<b>"#address-cells"</b>	<b>S</b>
11	Standard <i>prop-name</i> to define the address space representation of child nodes.	
12		
13	<i>prop-encoded-array</i> : an integer, encoded as with <b>encode-int</b> . Its value <i>shall</i> be identical to that of this	
14	node's parent's <b>"#address-cells"</b> value.	
15	<b>"ranges"</b>	<b>S</b>
16	Standard <i>prop-name</i> to define the address range that is decoded by this <b>/rom</b> node.	
17		
18	<i>prop-encoded-array</i> : One ( <i>child-phys</i> , <i>parent-phys</i> , <i>size</i> ) triple, where <i>child-phys</i> equals <i>parent-phys</i> and the	
19	number of cells of each corresponds to the parent's <b>"#address-cells"</b> value.	
20	<b>"available"</b>	<b>S</b>
21	Standard <i>prop-name</i> , to define available ROM resources.	
22		
23	<i>prop-encoded-array</i> : Arbitrary number of <i>phys-addr</i> ; <i>size</i> pairs. <i>Phys-addr</i> is a <i>phys.hi...phys.lo</i> list of	
24	integers, each integer encoded as with <b>encode-int</b> . <i>Size</i> is one or more integers, each encoded as with	
25	<b>encode-int</b> .	
26	The value of this property defines resources, managed by this package, that are currently available for use	
27	by a client program.	
28	<b>"write-characteristic"</b>	<b>S</b>
29	Standard <i>prop-name</i> , defines the ROM Technology.	
30		
31	<i>prop-encoded-array</i> : a string, encoded as with <b>encode-string</b> , where the value could equal "flash",	
32	"eeprom", "rom" or "nvram".	
33	<b>"cacheable"</b>	<b>S</b>
34	Open Firmware standard property indicating that the ROM is cacheable.	
35		
36	<i>prop-encoded-array</i> : <none>.	
37		
38	The presence of this property indicates that the ROM is cacheable.	

### 5.1.3.2. ROM Node Methods

If one or more ROM nodes are present, they *shall* each implement the following standard methods [1], Section 3.6.1. The **"reg"** property is used to determine which ROM the standard methods apply to for multiple ROM's.

The following methods must be defined by **/rom** node.

46	<b>open</b>	( -- true )	<b>M</b>
47	Standard method to prepare the ROM Node for subsequent use.		
48			
49	<b>close</b>	( -- )	<b>M</b>
50	Standard method to close the previously opened ROM Node.		
51			
52	<b>decode-unit</b>	( addr len -- phys.lo...phys.hi	<b>M</b>
53	Standard method to convert text unit-string to physical address.		
54			

1     **encode-unit**     ( phys.lo...phys.hi -- unit-str unit-len )     **M**

2     Standard method to convert physical address to text unit-string.

3  
4     **probe**             ( -- )                                     **M**

5     Open Firmware method used at boot time to probe all ROM's.

6     The **probe** method for ROM Nodes *shall* probe for FCode images within the address space defined by its  
7     **"reg"** property as defined herein. For each page within its address space, look for a valid FCode image. A  
8     valid FCode image is defined to start with an FCode-header(See section 5.2.2.5 in [1]) where the first byte  
9     is **start1**, the format byte is 0x08, the length field indicates that the FCode program is contained within  
10    the address space of the **/rom** node, and where the checksum is correct. (This probing must take into  
11    account the possibility that the ROM image is in the opposite endian-ness from which Open Firmware is  
12    currently running.)

13    If such an FCode image is found, a new child node *shall* be created by executing **new-device** and **set-**  
14    **args**, the FCode image copied to memory (taking into account the endian-ness) and the copy evaluated  
15    with **byte-load**. (The FCode program can use **my-unit** to create its **"reg"** property.). The arguments  
16    used by **set-args** are defined to be 0,0,unit-str,unit-len where unit-str is a text string  
17    representation of the physical address location for the FCode Image and unit-len is the length of the  
18    FCode Image.

#### 19    5.1.4. ROM Child Node(s)

20    This section describes the properties and methods for a ROM Child Node.

##### 21    5.1.4.1. ROM Child Node Properties

22    The following properties must be created by **/rom** child nodes.

23    **"name"**   **S**

24    Standard *prop-name* that names the child node.

25    Some physical ROM implementations may not fully decode their entire address range. This could lead to  
26    multiple images of the ROM to appear at different addresses, due to the "aliasing" of the ROM image. To  
27    prevent multiple device nodes from appearing in the device tree, the FCode for such ROMs should look for  
28    an already existing peer node that represents their image. This could be done, for example, by checking  
29    that any of the **peer** of the **child** of its **parent** node has a **"name"** property value that is the same as  
30    this node's FCode would create.  
31    This could be done, for example, by checking  
32    that any of the **peer** of the **child** of its **parent** node has a **"name"** property value that is the same as  
33    this node's FCode would create.

34    If such a node is found, the FCode should "abort" the evaluation of its FCode (e.g., by executing an **end0**)  
35    before creating its **"name"** property. Open Firmware *shall* remove a node when the FCode evaluation for  
36    the node does not result in a **"name"** property being defined.

37    **"reg"**   **S**

38    Standard *prop-name* that defines the child node address range for a ROM image(s).

39    *prop-encoded-array*: List of (*phys-addr*, *size*) specifications.

40    *Phys-addr* is encoded as with **encode-phys**, and *size* is encoded as with **encode-int**. The *phys-addr*  
41    is a base address of the ROM image and *size* is the length of the ROM image.  
42    The *phys-addr*  
43    is a base address of the ROM image and *size* is the length of the ROM image.

##### 44    5.1.4.2. ROM Child Node Methods

45    The following methods must be defined by **/rom** child nodes.

46    **open**                 ( -- true )                             **M**

47    Standard method to prepare this device for subsequent use.

48    The **open** method must be prepared to parse **my-args** for the case(s) when the node is being opened in  
49    order to access "files"; e.g., when the **bootinfo.txt** file is being accessed during the *multiboot menu*.

50    **close**             ( -- )                                     **M**

51    Standard method to close the previously opened device.  
52    Standard method to close the previously opened device.  
53    Standard method to close the previously opened device.  
54



The value of this property *shall* be the amount of contiguous real system memory required by RTAS, in bytes.

**"rtas-display-device"** S

*prop-name*, identifies RTAS Display Device

*prop-encoded-array*: An integer, encoded as with **encode-int**.

The value of this property *shall* be the *phandle* of the device node used by the RTAS display-character function.

**"rtas-error-log-max"** S

*prop-name*, identifies maximum size of an extended error log entry.

*prop-encoded-array*: An integer, encoded as with **encode-int**.

The value of this property *shall* be the maximum size of an extended error log entry, in bytes.

**"power-on-max-latency"** S

*prop-name*, specifies a future power on time capability.

*prop-encoded-array*: An integer, encoded as with **encode-int**.

The value of this property specifies the capability of the hardware to control the delay of system power on in days. If the property is present, the value *shall* indicate the maximum delay or latency in days. If the property is not present, the maximum delay or latency is 28 days.

### 5.2.2. RTAS Function Property Names

This section defines the property names associated with the various RTAS functions defined by Reference [5], Table 12, *RTAS Tokens for Functions*. Table 12 of [5] should be used as the reference for RTAS Functions currently implemented.

28	<b>"restart-rtas"</b>	S
29	<b>"nvram-fetch"</b>	S
30	<b>"nvram-store"</b>	S
31	<b>"get-time-of-day"</b>	S
32	<b>"set-time-of-day"</b>	S
33	<b>"set-time-for-power-on"</b>	S
34	<b>"event-scan"</b>	S
35	<b>"check-exception"</b>	S
36	<b>"read-pci-config"</b>	S
37	<b>"write-pci-config"</b>	S
38	<b>"display-character"</b>	S
39	<b>"set-indicator"</b>	S
40	<b>"get-sensor-state"</b>	S
41	<b>"set-power-level"</b>	S
42	<b>"get-power-level"</b>	S
43	<b>"assume-power-management"</b>	S
44	<b>"relinquish-power-management"</b>	S
45	<b>"power-off"</b>	S
46	<b>"hibernate"</b>	S
47	<b>"suspend"</b>	S
48	<b>"system-reboot"</b>	S

---

1	<code>"cache-control"</code>	S
2	<code>"freeze-time-base"</code>	S
3	<code>"thaw-time-base"</code>	S
4	<code>"stop-self"</code>	S
5	<code>"start-cpu"</code>	S
6	<code>"update-flash-and-reboot"</code>	S
7	<code>"update-flash"</code>	S
8		
9		

10  
11 For each of the indicated *prop-names*, the value returned for the property is a *token*.

12 *prop-encoded-array*: The value, *token*, is an integer encoded as with `encode-int`.

13  
14 If an RTAS function is implemented, there is a property name which corresponds to its function name. The  
15 value of this property is a *token*. This *token*, when passed to RTAS via its *rtas-call* interface (see below),  
16 invokes the named RTAS function. If a RTAS function is not implemented, there will not be a property  
17 corresponding to that function name. See the RTAS chapter of the CHRP Specification [3] for more  
18 information about RTAS functions.

### 19 5.2.3. RTAS Node Methods

20 The `instantiate-rtas` method is invoked by the operating system to instantiate the RTAS functionality.  
21 This is accomplished via the call-method *Client Interface Service*.

22  
23 `instantiate-rtas` ( *rtas-base-address* -- *rtas-call* ) M

24 Invoking the `instantiate-rtas` method binds the RTAS environment to a given location in System  
25 Memory and initializes the RTAS environment. The in parameter, *rtas-base-address*, is the physical address  
26 to which the RTAS environment is to be bound. This call indicates that RTAS is instantiated in a 32-bit  
27 mode. The amount of contiguous real memory that should be allocated for the RTAS environment is given  
28 by the value of the `"rtas-size"` property.

29 Upon completion of the `instantiate-rtas` method, an entry point address, *rtas-call*, is returned. The  
30 value of *rtas-call* specifies the physical address of the entry point into RTAS for future RTAS function calls.

31  
32 `instantiate-rtas-64` ( *rtas-base-address* -- *rtas-call* ) M

33 Invoking the optional `instantiate-rtas-64` method binds the RTAS environment to a given location  
34 in System Memory and initializes the RTAS environment. The in parameter, *rtas-base-address*, is the  
35 physical address to which the RTAS environment is to be bound. This call indicates that RTAS is  
36 instantiated in a 64-bit mode. The amount of contiguous real memory that should be allocated for the  
37 RTAS environment is given by the value of the `"rtas-size"` property.

38 Upon completion of the `instantiate-rtas-64` method, an entry point address, *rtas-call*, is returned.  
39 The value of *rtas-call* specifies the physical address of the entry point into RTAS for future RTAS function  
40 calls.

## 41 5.3. PCI Host Bridge Nodes

42 This section describes the PCI Host Bridge properties which are added or modified for a PowerPC Microproces-  
43 sor CHRP implementation. Refer to the *PCI Bus binding* [7] for the base PCI properties and methods. For each  
44 platform PCI Host Bridge, a `"reg"` property *shall* be present in the respective PCI Node.

45  
46  
47 **Note:** Since the RTAS PCI configuration access services do not have separate arguments  
48 identifying the PCI host bridge to which a service applies, platforms with multiple PCI  
49 host bridges must assign them unique bus numbers. An operating system must not reas-  
50 sign those bus numbers if it expects to make subsequent use of the RTAS PCI configura-  
51 tion access services..

### 5.3.1. PCI Host Bridge Properties

For each PCI Host Bridge (PHB) in the platform (called a PCI Bus Controller in the PCI Bus binding), a PCI Host Bridge Node *shall* be defined as a child node of the system bus, in accordance with the *PCI Bus binding*[7]. Each PCI PHB Node *shall* have a Unit Address defined in the **"reg"** property that is unique and persistent from each boot-to-boot. The following properties are modified or added by the PowerPC Microprocessor CHRP architecture and *shall* apply to each of these nodes.

Each PHB *shall* also have the **"used-by-rtas"** property, since rtas is used for PCI Configuration.

**"ranges"** S

Standard *prop-name*, defines this PHB's physical address ranges.

*prop-encoded-array*: From two to four (*child-phys*, *parent-phys*, *size*) specifications.

This property is mandatory for PCI Host Bridges in PowerPC Microprocessor CHRP implementations. The property value consists of four (*child-phys*, *parent-phys*, *size*) specifications, as described in the core document[1].

The first specification *shall* specify the configured address and size of this PHB's I/O Space. (I/O Space is shown as "BIO" to "TIO" in the Common Hardware Reference Platform Architecture [3].) The second specification *shall* specify the configured address and size of this PHB's Memory Space. (Memory Space is shown as "BPM" to "TPM" except for PHB0, when aliases are enabled, maps BPM0 to BIM, in the Common Hardware Reference Platform Architecture.) If present, the PHB method **change-address-map** can be used to reallocate these areas.

The third and fourth(components) specifications only apply to PHB 0. The third specification *shall* specify the configured address and size of the initial memory aliases. (The initial memory aliases are shown as "BIM" to "TPM0" in the PowerPC Microprocessor CHRP Architecture.) The third specification *shall* have a size of zero if the aliases are not enabled and *shall* be omitted if the aliases and the processor-hole are not enabled. The PHB method **set-initial-aliases** can be used to enable or disable the initial aliases. The fourth specification *shall* specify the configured address and size of the processor-hole. The fourth specification *shall* have a size of zero if the processor-hole is not enabled and *shall* be omitted if the processor-hole is not present. The root node method **set-processor-hole** can be used to enable or disable the processor-hole, if present.

**"model"** S

Standard *prop-name*, indicating this PHB's manufacturer, part number, and revision level. This property *shall* be present if this PHB does not supply the following standard PCI configuration properties which represent the values of standard PCI configuration registers: **"vendor-id"**, **"device-id"**, and **"revision-id"**.

*prop-encoded-array*: Text string, encoded as with **encode-string**.

The value of this property is a vendor dependent string which uniquely identifies this PHB and is correlated to its manufacturer, part number, and revision level. (see *Core* document[1] for more information.) The string value is device dependent, but *shall* supply information sufficient to identify the part to a level equivalent to the level achievable via the standard PCI configuration registers: **"vendor-id"**, **"device-id"**, and **"revision-id"**.

**"initial-memory-aliases"** S

*prop-name* indicates the status of the 16 MB initial memory aliases.

*prop-encoded-array*: <none>

This property *shall* be present if the initial memory aliases are enabled, otherwise this property *shall* be absent. (The aliases can be enabled or disabled via the **set-initial-aliases** method, see Section 5.3.2. on page 32.)

**"64-bit-addressing"** S

*prop-name* indicates this PHB's capability to address more than 4 Gigabytes of memory.

*prop-encoded-array*: <none>

This property *shall* be present if this PHB supports addressing more than 4 Gigabytes of memory, otherwise the property *shall* be absent. (If supported, the capability can be enabled via the **set-64-bit-addressing** method, see Section 5.3.2. on page 32.)

**"external-control"** S

*prop-name* indicates this PHB's ability to support the PowerPC external control facility.

*prop-encoded-int*: List of integers, each encoded as with **encode-int**.

The property value, if present, is a list of Resource ID's the version of the PowerPC external control facility supports. This property *shall* be present if this PHB supports the PowerPC external control facility, otherwise the property *shall* be absent.

**"64-bit-dma"** S

*prop-name* indicates this PHB's ability to support Dual Address Cycle for DMA.

*prop-encoded-array*: <none>

This property *shall* be present if this PHB supports Dual Address Cycle for DMA, otherwise the property *shall* be absent.

**"io-hole"** S

*prop-name* indicates that this PHB has enabled the io-hole capability, and the size of the hole.

*prop-encoded-int*: Integer, encoded as with **encode-int**.

This property *shall* be present if this PHB (phb,0 only) supports the io-hole, otherwise the property *shall* be absent. (If supported, the hole can be enabled via the **set-io-hole** method, see Section 5.3.2. on page 32.). The value of this property *shall* be the size of the io-hole in bytes. If the value is 0, the io-hole is disabled.

**"8259-interrupt-acknowledge"** S

*prop-name* indicates the address of the 8259 interrupt acknowledge facility.

*prop-encoded-array*: One *phys-addr* value

*Phys-addr* is encoded as with **encode-phys**.

The presence of this property indicates the presence of an 8259 interrupt acknowledge facility. A processor *load-byte* to this address will generate a Interrupt Acknowledge on the 8259 interrupt acknowledge facility.

### 5.3.2. Methods for PCI Host Bridge Nodes

The following *methods* are provided for the purposes of enabling, initializing, and/or modifying PHB features. If these features are present, the node *shall* implement the following methods. Operating systems can call these methods via the *client interface service* call-method function.

**change-address-map** ( new-mem-addr size new-io-addr size -- M  
okay? )

Relocate a PHB's peripheral spaces in the system address map, if possible.

**Note:** The **"ranges"** property may no longer reflect address spaces after execution of **change-address-map** and Open Firmware may no longer be capable of performing I/O.

**set-discontiguous-io** ( true | false -- ) M

Enable (true) or disable (false) a PHB's discontiguous I/O map.

**Note:** Once discontiguous I/O mode is enabled, the client should not call Open Firmware for I/O functions, since Open Firmware is not required to support the discontiguous I/O map.

**set-64-bit-addressing** ( location size enable? -- ) M



If the flag *enable?* is true, enable this PHB's 64-bit addressing capabilities and change the size and location of the PHB's TCE Table. If the flag *enable?* is false, disable this PHB's 64-bit addressing capabilities. For the case where all devices are 64-bit capable, both location and size can be specified as zero, in which case the 64-bit addressing capabilities will be enabled, but no TCE Table will be allocated. The Client Program *shall* initialize and maintain the Open Firmware and RTAS Areas to be mapped 1:1 as long as these areas are in use.

**Note: The `set-64-bit-addressing` method should not be executed if the PHB Device is not open.**

**`set-initial-aliases` ( true | false -- ) M**

Enable (true) or disable (false) the initial memory aliases for this PHB (PHB,0 only).

**`set-io-hole` ( true | false -- ) M**

Enable (true) or disable (false) the io-hole for PHB 0. This method *shall* be present if the io-hole capability is supported.

## 5.4. Memory Controller Nodes

This section describes *memory-controller* nodes and their properties.

### 5.4.1. Memory Controller Node Properties

For each Memory Controller in the platform, a *memory-controller* node *shall* be defined as a child of "/" (the root) and *shall* not have a "**ranges**" property. The following properties *shall* apply to each of these nodes.

A Memory Controller can also have the **used-by-rtas** property (See "Miscellaneous Node Properties" on page 35.), if it has functions abstracted by RTAS.

**"device\_type" S**

Open Firmware standard property.

*prop-encoded-array*: Text string, encoded as with **encode-string**.

The value of this property *shall* be "memory-controller".

**"reg" S**

Standard *prop-name*, defines the base physical address and size of this Memory Controller's addressable register space.

*prop-encoded-array*: One (*phys-address*, *size*) pair where *phys-address* is encoded as with **encode-phys**, and *size* is encoded as with **encode-int**.

The property value *shall* be the base physical address and size of this Memory Controller's register space.

**"model" S**

Standard *prop-name*, indicating this Memory Controller's manufacturer, part number and revision level.

*prop-encoded-array*: Text string, encoded as with **encode-string**.

The value of this property is a vendor dependent string which uniquely identifies this Memory Controller and *shall* be correlated to its manufacturer, part number, and revision level. (see Core document for more information.)

**"external-control" S**

*prop-name* indicates this Memory Controller's ability to support the PowerPC external control facility.

*prop-encoded-int*: List of integers, each encoded as with **encode-int**.

The property value, if present, is a list of Resource ID's the version of the PowerPC external control facility supports. This property *shall* be present if this Memory Controller supports the PowerPC external control facility, otherwise the property *shall* be absent.

**"error-checking" S**

Standard *prop-name*, defines the error checking capability of the node.

*prop-encoded-array*: a string, encoded as with **encode-string**, where the value could equal "none", "ecc", or "parity".

## 5.5. Interrupt Controller Nodes

This section describes the properties for the PowerPC Microprocessor CHRP interrupt controller node. An interrupt controller node *shall* not have the **"used-by-rtas"** property.

The CHRP Platform *shall* adhere to the *Open Firmware: Recommended Practice - Interrupt Mapping* (Refer to [15] for an interrupt structure Open Firmware representation) definition of the interrupt structure.

### 5.5.1. Open PIC Interrupt Controller Node Properties

An *open-pic* node *shall* be defined as a child of the bus to which the Open PIC is attached and through which registers are addressed. The following properties apply to this node.

**"name"** S

Standard *prop-name*: The value of this property *shall* be "interrupt-controller".

**"device\_type"** S

Open Firmware standard property.

*prop-encoded-array*: Text string, encoded as with **encode-string**.

The value of this property *shall* be "open-pic".

**"reg"** S

Standard *prop-name*, defines the base physical address(s) and size(s) of this Open PIC's addressable register space.

*prop-encoded-array*: List of (*phys-addr*; *size*) specifications.

*Phys-addr* is encoded as with **encode-phys**, and *size* is encoded as with **encode-int**.

The first entry in this list *shall* be the physical address and size decoded by the base Open PIC Interrupt Delivery Unit (IDU). Successive entries in this list *shall* be the physical addresses and sizes decoded by any additional Open PIC Interrupt Source Units (ISU).

**"compatible"** S

Standard *prop-name*, to define alternate **"name"** property values.

*prop-encoded-array*: The concatenation, with **encode+**, of an arbitrary number of text strings, each encoded as with **encode-string**.

The property value *shall* include "chrp,open-pic".

**"interrupt-ranges"** S

Standard *prop-name*, defines the interrupt number(s) and range(s) handled by the base Open PIC IDU and each additional Open PIC ISU, if any.

*prop-encoded-array*: List of (*int-number*, *range*) specifications.

*Int-number* is encoded as with **encode-int**.

*Range* is encoded as with **encode-int**.

The first entry in this list *shall* be zero (0) and the number of interrupts handled by the base Open PIC IDU. Successive entries in this list *shall* be the lowest interrupt number and the number of interrupts handled by each additional Open PIC ISU, if any. Successive entries *shall* include an equal or greater interrupt number. It is assumed that if *range* is zero in any entry, that unit does not handle any interrupts in that range. Note that there is a one-to-one correspondence between **"reg"** entries and **"interrupt-ranges"** entries.



## 5.7. Aliases Node

A *device alias*, or simply *alias*, is a shorthand representation of a *device-path*. *Aliases* are properties of the *aliases* node, encoded as with **encode-string**. *Aliases* are typically used by a user to facilitate not specifying a long path name at the User Interface ‘ok’ prompt.

An implementation of Open Firmware for a CHRP platform *shall* provide the following aliases as properties of the *aliases* node, if the corresponding device exists:

<b>"disk"</b>	<b>S</b>
<i>prop-name</i> , indicating the device path of the factory default disk.	
<b>"tape"</b>	<b>S</b>
<i>prop-name</i> , indicating the device path of the factory default tape.	
<b>"cdrom"</b>	<b>S</b>
<i>prop-name</i> , indicating the device path of the factory default CDROM.	
<b>"keyboard"</b>	<b>S</b>
<i>prop-name</i> , indicating the device path to the keyboard to be used for the User Interface.	
<b>"mouse"</b>	<b>S</b>
<i>prop-name</i> , indicating the device path to the mouse to be used for the User Interface.	
<b>"screen"</b>	<b>S</b>
<i>prop-name</i> , indicating the device path to the screen to be used for the User Interface.	
<b>"pc-keyboard"</b>	<b>S</b>
<i>prop-name</i> , indicating the device path of the factory default PC-style keyboard.	
<b>"pc-mouse"</b>	<b>S</b>
<i>prop-name</i> , indicating the device path of the factory default PC-style mouse.	
<b>"adb-keyboard"</b>	<b>S</b>
<i>prop-name</i> , indicating the device path of the factory default ADB-style keyboard.	
<b>"adb-mouse"</b>	<b>S</b>
<i>prop-name</i> , indicating the device path of the factory default ADB-style mouse.	
<b>"scsi"</b>	<b>S</b>
<i>prop-name</i> , indicating the device path of the factory default built-in SCSI device.	
<b>"com1"</b>	<b>S</b>
<i>prop-name</i> , indicating the device path of the factory default 16550-style serial port known as "com1."	
<b>"com2"</b>	<b>S</b>
<i>prop-name</i> , indicating the device path of the factory default 16550-style serial port known as "com2."	
<b>"scca"</b>	<b>S</b>
<i>prop-name</i> , indicating the device path of the factory default SCC-style serial port known as "SCCA."	
<b>"sccb"</b>	<b>S</b>
<i>prop-name</i> , indicating the device path of the factory default SCC-style serial port known as "SCCB."	
<b>"floppy"</b>	<b>S</b>
<i>prop-name</i> , indicating the device path of the factory default floppy drive.	
<b>"net"</b>	<b>S</b>
<i>prop-name</i> , indicating the device path of the factory default built-in network interface controller.	

1	<b>"rtc"</b>	<b>S</b>
2	<i>prop-name</i> , indicating the device path of the factory default real-time-clock chip.	
3		
4	<b>"nvram"</b>	<b>S</b>
5	<i>prop-name</i> , indicating the device path of the factory default NVRAM.	
6		

## 5.8. Event Sources Node

The Event-Sources Node describes the possible RTAS Error and Event Classes for Open PIC Interrupts (Refer to [3], Table 50). The Event-Sources Node *shall* be defined to be a child of the root device tree node. The following properties *shall* be defined for this node:

12	<b>"name"</b>	<b>S</b>
13	Open Firmware standard property. The value of this property <i>shall</i> be "event-sources".	

### 5.8.1. Child nodes of the Event Sources Node

The following specify standard child nodes of the /**event-sources** node. These nodes could be present in a CHRP platform.

#### 5.8.1.1. internal-errors

The presence of the node indicates that all or some of the function has been implemented using an OpenPIC interrupt.

23	<b>"name"</b>	<b>S</b>
24	Open Firmware standard property. The value of this property <i>shall</i> be "internal-errors".	
25		
26	<b>"open-pic-interrupt"</b>	<b>S</b>
27	The value of this property <i>shall</i> be a <i>prop-encoded-array</i> of integers encoded as with <b>encode-int</b> .	
28		
29	The value indicates one or more Open PIC interrupt source numbers that are used to report the internal-errors class of error conditions.	

#### 5.8.1.2. epow-events

The presence of the node indicates that all or some of the function has been implemented.

34	<b>"name"</b>	<b>S</b>
35	Open Firmware standard property. The value of this property <i>shall</i> be "epow-events".	
36		
37	<b>"open-pic-interrupt"</b>	<b>S</b>
38	The value of this property <i>shall</i> be a <i>prop-encoded-array</i> of integers encoded as with <b>encode-int</b> .	
39		
40	The value indicates one or more Open PIC interrupt source numbers that are used to report events of the epow-events class.	
41		
42	<b>"power-on-triggers"</b>	<b>S</b>
43	Standard <i>prop-name</i> defining the power-on trigger types implemented on a specific platform.	
44		
45	<i>prop-encoded-array</i> : an array of integers, each encoded as with <b>encode-int</b> . Each entry in this array is one of the values defined by Reference [5], Table 62b, <i>Defined Power-on Triggers</i> , for the power-on trigger types corresponding to the power-on trigger types available on the platform.	

#### 5.8.1.3. power-management-events

The presence of the node indicates that all or some of the function has been implemented.

51	<b>"name"</b>	<b>S</b>
52	Open Firmware standard property. The value of this property <i>shall</i> be "power-management-events".	

54



not be directly attached to the main physical bus(root node (“/”). Instead, cpu devices *shall* be children of the /cpus node.

The /cpus node *shall* have one child node of device type cpu for each processor. The value of the “reg” property of each cpu device node *shall* be an integer encoded as with **encode-int**, which *shall* be the index of the output of the Open PIC interrupt controller which is attached to this processor; i.e., the bit number corresponding to this processor in the Open PIC Destination Register. The *ihandle* of the “executing” processor *shall* be published in the “cpu” property of the /chosen node.

**Note:** The properties of a cpu device are already defined in the PowerPC Processor binding document[6]. The only change for symmetric multiprocessor (SMP) systems is that there will be a cpu device node under the /cpus node for each individual processor. Other properties of the cpu devices *shall* conform with the requirements stated in the PowerPC Processor binding document.

## 6.2. SMP Properties

The following properties are for a CHRP PowerPC SMP environment. These SMP properties will be under the /cpus Node.

**“slot-names”** S

*prop-name:* Describes platform labeling of add-in cpu/processor card slots.

*prop-encoded-array:* An integer, encoded as with **encode-int**, followed by a list of strings, each encoded as with **encode-string**.

The integer portion of the property value is a bitmask of possible processors; for each add-in slot on the bus, the bit corresponding to that slot’s ID number is set from least-significant to most-significant ID number. The number of following strings is the same as the number of slots; the first string gives the platform nomenclature for the slot with the smallest ID number, and so on. The CPU’s “slot-names-index” property can be used as an index into the bitmask integer of this property. The absence of this property indicates that no slots are present.

**“smp-enabled”** S

The presence of this property signifies that the platform is SMP enabled, even if it only has one processor.

*prop-encoded-array:* <null>

### 6.2.1. Processor Node

The following properties are for a CHRP PowerPC SMP environment. This SMP property will be under each /cpu Node.

**“slot-names-index”** S

*prop-name:* Identifies each cpu with a unique number.

*prop-encoded-array:* An integer, encoded as with **encode-int**.

The value of this integer is a platform unique number with a range from 0 to  $n-1$  for each CPU where  $n$  is the number of slots. This number is used to index into the “slot-names” property to identify the value of the string associated with the slot name.

## 7. Device Power Management Properties/Methods

This section defines standard platform node properties, device node properties, and methods related to power management. The properties and methods of this section *shall* be implemented on any platform which supports power management except where noted. However, it is still being enhanced. Operating system providers who want to ensure that the data needed for their power management policies is included should contact the authors of this document.

## 7.1. System Node Properties

The following defines properties are to be associated with the rtas and the power-management-events nodes of the device tree.

### 7.1.1. Properties assigned to the RTAS node

Power domains are a feature of CHRP platforms which support power management. Within the Open Firmware device tree, power domains are represented by a power domain identifier which is defined to be an integer in the range 0..n-1, where n is the number of power domains on the platform.

**"power-domains-tree"** S

Standard *prop-name* which defines the power domain hierarchy for this platform.

*prop-encoded-array*: An array of integers, each encoded as with **encode-int**, that is a flattened representation of the power domain dependency tree.

The array consists of a number of tuples, one for each power domain defined on the platform. Each tuple consists of the power domain identifier domain#, followed by the number of power levels #levels supported by the domain, followed by an array of tuples, one for each level. These tuples consist of a level identifier level, followed by the number of power sources from which the domain draws power, followed by an array of tuples (power-source-id, power). The power domain tuple is terminated by the number of children #children followed by a list of the domain identifiers of each child. The power values are expressed in milliamperes and include only the power consumed by support logic not represented as devices in the device tree including any RTAS abstracted devices within the particular power domain.

**"power-domains-controllers"** S

Standard *prop-name* which defines the power domain controllers (The CHRP architecture calls them power domain control points) present on this platform.

*prop-encoded-array*: an array of integers, each encoded as with **encode-int**. Each integer is the *phandle* of the device tree node that functions as the power domain controller for a domain. A single controller may serve as the control point for multiple domains. Each device which serves as a controller encodes the controls-power-domain property.

**"power-domains-names"** S

Standard *prop-name* used to define the user readable names for the power domains.

*prop-encoded-array*: an array of strings, each encoded as with **encode-string**, that are the user readable names for the domains.

The number of strings matches the number of domains and there is a one-to-one correspondence between the entries in the **"power-domain-controllers"** property and the entries in this array.

**"platform-power-sources"** S

Standard *prop-name* defining the platform power sources.

*prop-encoded-array*: an array of integers, each encoded as with **encode-int**. The array is structured as a number of tuples. Each of these tuples consists of the values source-voltage, (given in millivolts), peak-power, continuous-use-power (both expressed in milliamperes supplied at the stated voltage), and conversion-efficiency (expressed in percent).

**"power-sources-names"** S

Standard *prop-name* defining the platform power source names.

*prop-encoded-array*: an array of strings, each encoded as with **encode-string**, that are the user readable names for the power sources.



The number of strings match the number of power sources and is in one-to-one correspondence to the entries in the **"platform-power-sources"** property.

**"platform-battery-sources"** S

Standard *prop-name* defining the batteries utilized by a platform.

*prop-encoded-array*: an array of integers, each encoded as with **encode-int**. Each value in this array is the manufacturer's rated capacity of the battery expressed in milliwatt-hours.

**"battery-sources-names"** S

Standard *prop-name* defining the human-readable identifier of the batteries utilized by a platform.

*prop-encoded-array*: an array of strings, each encoded as with **encode-string**. Each entry in this array corresponds one-for-one with the batteries defined in the **"platform-battery-sources"** property.

### 7.1.2. Properties of the power-management-events node

**"power-type"** S

Standard *prop-name* defining the power management event types implemented on a specific platform.

*prop-encoded-array*: an array of integers, each encoded as with **encode-int**. Each entry in this array is one of the values defined by Reference [3] and extended by Reference [5], Table 62, for the power management event types corresponding to the event types available on the platform.

## 7.2. Device Properties

**"power-domains"** S

Standard *prop-name*: Indicates the power domains of which this device is a member. If the device is a member of the default power domain 0 alone, this property does not need to be provided.

*prop-encoded-array*: List of one or more domain-ids to which this device belongs. domain-ids is encoded as with **encode-int**.

The power-domains property should only list the domain-ids of the lowest power domain tree nodes in which this device has membership.

**"device-power-states"** S

Standard *prop-name* which describes the power states this device supports. This property *shall* be provided for each physical device which has multiple power states, if platform firmware provides device power state information.

*prop-encoded-array*: An array of integers, each encoded as with **encode-int** that defines the supported power states for this device.

The array consists of an integer representing the initial device power state after reset, followed by the number of power sources from which the device draws power, followed by an arbitrary number of tuples, one for each supported power state of the device. Each tuple consists of the state, followed by an array of tuples (power-source-id, power) giving the average power consumption from each power source during active use. This is followed by another array of tuples (power-source-id, power) giving the idle power consumption for each power source. Each power state tuple is terminated by the maximum expected power usage lifetime in seconds for the device if it were to remain in that state. The value power is stated in the millamperes consumed at the voltage supplied by the power source.

The value state *shall* be further constrained to have the following semantics:

Table 2. Semantics of device state values

Value	Semantics
100	This is the device's most responsive state.
20-99	The device is functional. The range represents a range of performance.
11-19	Reserved
10	Device is not operational, but retains its internal functional parameters.
1-9	Reserved
0	Device not functional, may lose internal functional parameters.

The semantics of device power states may be further defined by device type specific bindings.

The interaction of the defined semantics of device power state and domain power level is defined in Table 3. Combinations of Device Power State/Domain Power Level. Those combinations not marked are disallowed.

Table 3. Combinations of Device Power State/Domain Power Level

		Device Power State			
		100	99-20	10	0
Domain Power Level	Full On	Allowed	Allowed	Allowed	Allowed
	Reduced		Allowed	Allowed	Allowed
	Freeze			Allowed	Allowed
	Off				Allowed

#### "device-state-transitions"

S

Standard *prop-name* that describes the legal power state transitions supported by the device. This property *shall* be provided if platform firmware provides device power state information.

*prop-encoded-array*: an array of integers, each encoded as with **encode-int** that defines the legal power state transitions for this device.

The array is structured as a number of tuples, one for each possible transition. Each tuple consists of the starting state, followed by the destination state, followed by an array of tuples (power-source-id, power), one for each power source, followed by the time required to make the transition in microseconds, followed by the maximum count allowed for this transition. The starting state and destination state are values defined in the "**device-power-states**" property. The value power is stated in the milliamperes consumed.

#### "power-sources"

S

Standard *prop-name* which designates this device as a consumer of power sourced from a defined power source. This property *shall* be provided if platform firmware provides device power state information.

*prop-encoded-array*: an array of integers, each encoded as with **encode-int** that gives the list of power sources to which this device is connected. The values are indices into the platform-power-sources data structure

#### "power-management-mapping"

S

Standard *prop-name*. This optional property provides a device dependent mapping between device power state and commands which the device driver sends to its device. Also provides information concerning which device

power states are supported for each of the four domain power levels. See the device type binding for a definition of the property value.

### 7.2.1. Properties for Power Domain Control Points

The following are specific to devices which can act as power domain control points.

**"controls-power-domains"** **S**

Standard *prop-name* which designates the domains over which this device exercises control.

*prop-encoded-array*: an array of integers, each encoded as with **encode-int** that defines the domains for which this device can act a power domain control point.

A single device may serve as multiple logical control points.

## 7.3. Power Management Related Methods

This section defines methods associated with device tree nodes which serve as power domain controllers (CHRP architecture calls them control points).

**set-power-level** (domain# level -- actual-level) **M**

This method is only present for power domain controllers. The domain# is the power domain whose power level is altered, and level is the desired level. actual-level reports the level to which the domain was actually set.

**get-power-level** (domain# -- level) **M**

This method is only present for power domain controllers. The domain# is the power domain that is being queried. level is the current level at which the domain is now operating.

**system-off** ( -- ) **M**

Method to turn the system off. This method is attached to the root node of the device tree and is only present in a platform with software control over system power.

## 7.4. Power Management NVRAM Partition

The standard format for data located in the Power Management Configuration partition of the NVRAM is under development(see "Device Power Management Configuration Resources" on page 70).

## 8. ISA Device Support

The CHRP System binding identifies ISA platform support areas that are unique. These areas have to do with the DMA and Interrupt Controller attributes. In addition, it is desirable to further define certain choices for the CHRP standard ISA I/O and/or ISA pluggable resources where an OS expects to find information defining the attributes. These additional ISA dependencies for a CHRP Platform are in the area of DMA Modes, interrupt request states, address modes(I/O or memory), state of Interrupt and DMA Request Lines (tri-stateable: can be disabled) and the Unit Address Format. Open Firmware is responsible for programming ISA modes that are configured appropriately with the stated level of CHRP Platform support.

The following sections describe further restrictions and/or requirements for the CHRP Platform ISA Devices and/or Adapters:

- ISA Resources
  - ISA address mode
  - Interrupt triggering/level modes
  - DMA modes and width
- ISA Properties
- ISA Unit Address Format
- ISA Configuration Requirements

## 8.1. Additional ISA Requirements

Additional ISA Adapter/Device resource restrictions and/or formats are stated in this section for a CHRP Platform. Refer to *ISA/EISA/ISA-PnP binding*[8] for the various resources defined in this section. Open Firmware information for ISA is defined in the following Table 4. for a set of standard ISA I/O platform devices or ISA plug-in adapters/devices.

The meaning of the table headings below are defined to be:

- Device Tree - The Open Firmware device tree is expected to define the entries within property values. The operating system is expected to get the properties from the Open Firmware Device Tree.
- Device Binding - The Open Firmware device binding for that device is expected to define the entries within a property value, if necessary. The operating system is expected to get the properties from the Open Firmware device binding for that device.
- CHRP platform - The capabilities of the CHRP legacy interrupt controller or DMA controller create restrictions on the entries within property values, if necessary.

The ISA terms or nomenclature used in the table below relate to the *ISA/EISA/ISA PnP binding*[8] in the following manner:

interrupt level	irq#
interrupt triggering	type
dma channel #	dma#
DMA master or slave mode	busmaster
DMA width	width and countwidth
DMA timing	mode

Table 4. ISA Open Firmware Resource Information

Device Tree	Device Binding	CHRP Platform Restrictions on ISA Adapters/Cards
Interrupts - 'level' specified	Interrupts - 'triggering' specified	Interrupts - NO support for 'active high level sensitive' or 'high to low edge' triggering
DMA - 'channel #' specified	DMA - master or slave Mode specified DMA - channel and width count	DMA - No support for "C" Mode, 32-bit channel or 32-bit countwidth
Address - physical or aliased address specified	Address Mode - I/O or memory address specified	
DMA: ISA "Compatibility" Mode <i>shall</i> be supported(see note 1)	DMA: ISA "Compatibility" Mode <i>shall</i> be supported(see note 1)	DMA: ISA "Compatibility" Mode <i>shall</i> be supported(see note 1)

Note 1: ISA DMA slaves and masters must support the ISA compatible timing mode. ISA DMA slaves and masters are allowed to support additional timing modes that may be reported in the Open Firmware device tree. An operating system may restrict its mode operation to just the ISA compatible timing mode or may use the timing mode specified in the Open Firmware device tree node for that device.

### 8.1.1. ISA Addresses

All ISA Adapters *shall* use an I/O Address or memory address for operation with the ISA Bus per Table 4.

### 8.1.2. ISA Interrupts and DMA Channels

A CHRP Platform *shall* ensure that all I/O device states are consistent with the interrupt legacy controller and DMA controller specified in Table 4.

### 8.1.3. ISA Unit Address Format

A CHRP platform may use a ‘legacy’ or a ‘PnP’ Unit Address Format defined in Reference[8] as the first entry in a “**reg**” property. The ‘legacy’ Unit Address value used by a CHRP platform for I/O ISA devices *shall* remain invariant across platform boots.

This requirement ensures that ISA adapters or functions would have a stable unit address across boots. Violations of this rule may require reinstallation of an OS.

### 8.1.4. ISA Additional Properties

The CHRP Platform *shall* support two additional properties that identify when an adapter cannot disable interrupt lines or DMA Channels.

**“dma-enabled”** S

*prop-name* indicates that a DMA Channel is active

*prop-encoded-array*: a list of integers, each one encoded as with **encode-int**.

The property value is a list of integers for index indicating DMA channel numbers used by the adapter/device. The presence of this property signifies that the specified DMA Channels for the ISA adapter are enabled and thus not tri-stated.

**“interrupt-enabled”** S

*prop-name* indicates that interrupt request line is active

*prop-encoded-array*: a list of integers, each one encoded as with **encode-int**.

The property value is a list of integers for index indicating interrupt request lines used by the adapter/device. The presence of this property signifies that the specified interrupt request lines for the ISA adapter are enabled and thus not tri-stated.

## 9. Configuration of Platform Resources

Any computer platform is composed of standard components which are invariant (platform ‘built-in’ standard I/O and power management), optional components which are detectable (a second processor, for example), and configurable components which are self-identifying (system memory, for example). Most computer platforms also provide one or more industry standard I/O buses which allow the insertion of specialized functional adapter cards. These buses generally support a method for automatic identification, interrogation, and option selection of installed adapter cards. Some older adapter cards, however, may not support these methods or may supply incomplete information. Thus it is necessary to provide interactive utilities to aid in the setup of ‘legacy’ pluggable ISA adapter cards and to provide information describing the power management of platform resources.

A CHRP platform should be capable of configuring ISA resources, if required, and power management resources (if implemented). A CHRP platform with ISA slots should be capable of using both ‘legacy’ ISA adapter cards

1 and ‘Plug and Play’(PnP) ISA adapter cards (see the *ISA/EISA/ISA PnP binding*[8] and *Device Support Extensions binding*[10]). For the ISA PnP Adapter Cards, Open Firmware *shall* be capable of reading data resource records from the adapters to determine properties and device tree nodes. For ‘legacy’ ISA adapter cards, ISA Configuration Utility (ICU) support is optional and should be present to interact with the user to provide assignment of ISA resources; i.e., the I/O address, interrupt number, DMA channel, system memory and slot-names. The platform firmware should provide an ISA Configuration Utility(ICU) as defined in section 9.1., if the platform supports ISA plug-in adapters.

8 A CHRP Platform *shall* also have the capability of configuring power management resources, if power management is implemented by the platform, as defined in section 9.2.

## 11 9.1. ISA Resource Configuration

13 CHRP platforms should provide an ICU program which ‘legacy’ vendors and users of the platform can use to generate the Open Firmware configuration information for ISA devices, as described in this section. The information the ICU produces, if present, should be stored in the NVRAM Configuration Partition (NVRAM partition 0x71), and the information should have the format described in this section. In addition, the platform firmware should provide a user interface method (keyboard sequence, icon, or other means) for the user to enter the ICU at system boot time in order to enter configuration data before an operating system is executed.

### 20 9.1.1. ISA Resource Configuration Description

22 The CHRP Platform ISA Configuration Utility should contain appropriate functions which will determine the information needed to completely fill in the required data structures. This may require obtaining information from a human user, as well as interacting with and controlling hardware interface features of some ISA devices. A recommended set of interactions is:

- 26 • Interact with CHRP platform user to assign fixed ISA resources, if necessary
  - 27 • Determine status of hardware defined ISA adapter resources (switches)
  - 28 • Allocate ISA resources per adapter
  - 29 • Verify that resources don’t conflict with platform base I/O devices
- 31 • Interact with CHRP platform user to assign programmable ISA resources, if necessary
  - 32 • Set the state of software controlled ISA adapter resources (soft switches)
  - 33 • Allocate ISA resources per adapter
  - 34 • Verify that resources don’t conflict with platform base I/O devices and any fixed ISA adapter resources
- 36 • Assign ISA-PnP adapter and/or base device relocatable resources
  - 37 • Determine relocatable ISA resources
  - 38 • Allocate ISA resources per adapter
  - 39 • Verify that the assigned ISA-PnP adapter and/or base device resources do not conflict with ISA (base I/O devices and adapters) assigned resources
- 42 • Verify that ISA and ISA-PnP resources do not conflict with the PCI-to-ISA bridge/PCI bus resources
- 43 • Interact with CHRP platform user to provide the capability to display platform assigned ISA/ISA-PnP resources

#### 45 9.1.1.1. ISA Configuration Handle Assignment

47 Logical device handle numbers will be first assigned, starting with number “*l*” through “*m*”, to ISA legacy card devices. Resource assignment information gained from user input interactively will be tabulated and conflict checked by the ISA Configuration Utility (ICU). Conflict information will be displayed to the user and it will be the user’s obligation to re-adjust legacy card jumpers/switches to resolve conflicts and to provide updated resource information to the ICU. After completion of the assignment of legacy card resources, PnP device resource assignments will be made for logical devices with handles corresponding to PnP CSN assignments “*m+l*” through “*m+n*” via standard PnP isolation and auto-configuration methods. Card slot/connector

1 information will also be solicited interactively by the ICU for both legacy and PnP cards to support POST  
 2 diagnostic fault location requirements.

### 4 9.1.2. ISA Legacy Configuration Data Format Requirements

5 The ISA Data Format *shall* be in a 'PnP' TAG Format (Refer to [34] and [35] for detailed tag information)  
 6 and represent allocated resources only. The PnP Vendor extension tag will be used to represent the *slot*  
 7 *names/connector* information.

8  
 9 The ISA Configuration Data Format will follow the format conventions which have been defined for ISA Plug  
 10 and Play (PnP) attach cards. The PnP-like format given below will be used to characterize ISA legacy cards  
 11 which may also be present in a CHRP platform. The Configuration Manager will solicit additional information  
 12 via an interactive setup routine to obtain resource information for ISA legacy cards as well as physical attribute  
 13 information (such as card slot location and connector attributes) for all cards in general.

14 The following information will be stored in NVRAM to maintain configuration status and to characterize  
 15 resource requirements for ISA legacy cards. It has been assumed here that handles 1 through m+n have been  
 16 assigned by the Configuration Manager to logical devices in the attach card set.

18 Table 5. ISA Card Configuration Directory

# Bytes	Data Field
1	Number of Assigned Legacy Card Handles (0 thru m-1)
1	Number of Assigned PnP Card Handles (0 thru n-1)
6	Reserved
32	Card Slot List (Slot location by handle index)
32	Pointers to Legacy Card Resource Data (handle indexed)
64	Reserved for additional physical attributes

20  
 21  
 22  
 23  
 24  
 25  
 26  
 27  
 28  
 29  
 30  
 31  
 32 Data stored in NVRAM to characterize ISA legacy cards will make use of the following PnP-like format system  
 33 of tags for small data items:

34  
 35 Table 6. ISA Card PnP-like Small Data Item Tag Format

Tag	Data Item Type
0x15	Logical Device ID (Assigned Handle)
0x1C	Compatible Resource Requirement
0x23	IRQ Resource Requirement
0x2A	DMA Resource Requirement
0x30	Start Dependent Function(optional)
0x38	End Dependent Function(optional)
0x4B	I/O Resource Requirement(10 bit)
0x4C	I/O Resource Requirement(11 bit)
0x4D	I/O Resource Requirement(16 bit)

Table 6. ISA Card PnP-like Small Data Item Tag Format

Tag	Data Item Type
0x72	PhysicalAttributes(VendorExtension)-'slot names'/'connectors'
0x78	End Tag

Currently, the compatible device ID set for the legacy devices is restricted to that defined for equivalent PnP devices. Unlike PnP cards, however, the compatible device ID for legacy cards cannot be used to identify driver requirements.

Large data items will make use of the PnP-like format tags:

Table 7. ISA Card PnP-like Large Data Item Tag Format

Tag	Data Item Type
0x81	Memory Range Descriptor
0x82	ANSI String ID (optional)
0x86	32-bit Fixed Location Memory Range Description

The minimum base address and maximum base address fields of a tag type 0x81 Memory Range descriptor data item must be defined to be the same value. This is required to assure unambiguous interpretation of the memory range attributes as a resource assignment.

**An example of an ISA legacy card resource requirement description structure in NVRAM is as follows:**

TAG (0x82)	ANSI String Description (optional)
TAG (0x15)	Logical Device ID (Assigned Handle)
TAG (0x72)	Vendor Extension (Slot Location)
TAG (0x1C)	Compatible Device ID
TAG (0x23)	IRQ Resource Requirement
TAG (0x2A)	DMA Resource Requirement
TAG (0x4B)	I/O Port Resource Requirement (May be more than one)
TAG (0x81)	Memory Range Resource Requirement (May be more than one)
TAG (0x86)	32 Bit Fixed Memory Range Resource Requirement (May be more than one)
TAG (0x78)	End Tag.

The absence of a particular data item type indicates that no resource of that type is required or assigned. The presence of an empty (all zeroes) mask in IRQ/DMA resource data items also indicates that no resource of that type is required.

Because ISA devices can conflict in their use of DMA channels, memory addresses, I/O addresses, and interrupts, a configuration utility may detect conflicts which cannot be resolved by software. In such cases, the ICU needs to interact with a human user to suggest configuration changes such as changing jumper settings. It is recommended that the utility suggest non-conflicting ISA resources which the user could set the device to utilize.



---

## 9.2. Power Management Resource Configuration

For a platform which supports device power management, all platform power management related information *shall* be resident in the Open Firmware device tree prior to the transfer phase of software operation (see the definition of transfer phase in [3], Chapter 2). Dummy devices *shall* be placed in the device tree for all standard I/O bus connectors which are not in use to provide a node to assign the slot-names, power-domains, and power-sources properties.

Ultimately, the goal is that pluggable devices would not only identify themselves to platform firmware but would also provide all applicable power management related information. As an interim solution, a utility *shall* be provided either in the platform firmware ROM or supplied as a loadable Open Firmware utility on external media. This utility interacts with a person to obtain power management information concerning plug-in adapters and peripherals.

### 9.2.1. Power Management Information Utility

Any platform capable of being expanded via the addition of power-managed devices *shall* provide a device power management information utility. The purpose of the utility is to allow a person (end-user or system developer) to enter power management related device properties of plug-in adapters and peripherals which have no mechanism to automatically report this information to firmware or system software. The need for this utility will disappear as standard protocols are developed for interrogating pluggable adapters and devices to provide power management related information.

In the most general case, the devices to be added to a node representing a standard bus or I/O port are in the form of multilevel subtrees. The root of this subtree specifies the path to the node in the device tree where the subtree is to be grafted.

The utility determines the path to the node at which to graft the new devices by interacting with a person to receive the information. The utility uses the "**slot-names**" property to identify the location of the device for which it needs information. For example, the utility might prompt the user with, "Enter the name of the first device attached to the external scsi connector labeled 'SCSI1'."

A data structure describing the subtree is stored in NVRAM. The root node of this subtree contains an "**in-graft-node**" property which specifies the path to the parent node where the subtree is to be grafted into the Open Firmware device tree.

As adapters and devices are enhanced to support the automatic reporting of power management information the parent node would supply a method query-power-management-attributes which can be used by firmware to obtain this information without the need for this utility. Any information obtained by direct device interrogation may update that supplied via the PM NVRAM partition.

### 9.2.2. PM Configuration Process

When the platform is booted after a configuration change and the newly inserted adapter does not support the automatic reporting of power management information, firmware should prompt the user asking if he wishes to supply this information or potentially forfeit some or all of the power management capabilities of the device.

The utility records the information it obtains in the NVRAM Power Management Configuration Partition (NVRAM Signature of 0x71 and name *pm-config*). On a subsequent reboot, platform firmware uses the information saved in NVRAM to fill out the device tree adding new nodes and their properties, as well as adding properties and updating the values of properties of existing device tree nodes.

### 9.2.3. PM Configuration Format

The NVRAM power management configuration partition is designed to be accessed primarily by firmware, but the partition is designated global and the format is specified to allow a third party to write a power management information utility which runs on the booted operating system.

The data field of the power management NVRAM partition *shall* be defined as follows:

The data field is composed of a header, followed by a number of fixed length data blocks, and finally a variable length property list area. The length of the header and each data block is 8 bytes. The data blocks use 16-bit integer offsets into the partition as pointers to the data blocks and into the property list area. The base of this offset is the beginning of the partition. This effectively limits the size of the PM configuration area to 64K bytes. If more space is required, additional PM configuration partitions may be provided. Each pointer into the property list area locates the start of a null-terminated string which represents a list of property name/value pairs.

The following table specifies the format of the header:

Table 8. Power Management Configuration Data Header

Field	Size	Description
Version	1 byte	Designates the version of the PM Partition data area format
Subtree_ptr	2 bytes	Pointer to the first data block which describes a device subtree
Property_ptr	2 bytes	Pointer to first data block which describes a property list to be added to the base platform device tree
Reserved	3 bytes	Reserved

The PM Partition data area format value *shall* be 1.

The following table specifies the format of the data blocks:

Table 9. Data Block Format

Field	Size	Description
Block_type	1 byte	Designates the data block type
Data Block Data	7 bytes	Remainder of data block, format specific to data block type

Two data blocks are defined: one defining a device node and a second defining properties to be added to the base platform device tree.

The data block type field *shall* have the value 1 for a data block which describes a device node. The data block type field *shall* have a value 2 for a data block which describes a property.

Table 10. Node Data Block Format

Field	Size	Description
Block_type	1 byte	This field <i>shall</i> contain the value 0x01
Prop_list_ptr	2 bytes	Pointer to a null terminated string containing the property list for this node

Table 10. Node Data Block Format

Field	Size	Description
Child_ptr	2 bytes	Pointer to a data block defining a child node of this node. This pointer will be equal to 0x0000 if this node has no children.
Sibling_ptr	2 bytes	Pointer to a data block defining a sibling node of this node. This pointer will be equal to 0x0000 if this node has no siblings.
Reserved	1 byte	Reserved

Table 11. Property Data Block Format

Field	Size	Description
Block_type	1 byte	This field <i>shall</i> contain the value 0x02
Node_path	2 bytes	Pointer to a null terminated string giving the path name of the node to which the designated property list belongs.
Property_list_ptr	2 bytes	Pointer to a null terminated string containing the property list to be assigned to the designated node.
Reserved	3 byte	Reserved

The first node of a subtree *shall* have a **"name"** property equal to "/" and *shall* specify the **"in-graft-node"** property. The child\_ptr of this data block points to the first in a list of data blocks which describe the nodes which make up the subtree to be grafted onto the system tree.

The final area of the partition is a set of null terminated strings which represent property name/value pair lists. The last string in this area will be terminated by at least two null bytes. The property list for each node *shall* provide all the required PM properties and their values. These include **"power-domains"**, **"device-power-states"**, **"device-state-transitions"**, **"power-sources"**, **"power-management-mapping"**, and **"controls-power-domains"**.

## 10. Client Program Requirements

For CHRP platforms, the client program requirements are defined in Section 9 of Reference [6], with the following modifications. Open Firmware Client Programs for a CHRP platform *shall* execute in 32-bit mode with an Open Firmware cell size of 1.

### 10.1. Load Address

The client's load address is specified by the value of the **load-base** Configuration Variable. The value of **load-base** defines the default load address for *client programs* when using the **load** method. **Load-base** *shall* be a real address in real mode or a virtual address in virtual mode. Note that this address represents the area into which the client program file will be read by **load**; it does not correspond to the addresses at which the program will be executed. All of physical memory from **load-base** to either the start of Open Firmware physical memory or the end of physical memory, whichever comes first, *shall* be available for loading the client program.

**Note:** The `load-base` address represents the area into which the client program will be read by load and does not correspond to the address at which the program will be executed.

## 10.2. Initial Register Values

Table 2, "Initial Register Values", of Reference [6] is modified as follows:

- r3 -- *shall* be 0 on client program entry
- r4 -- *shall* be 0 on client program entry

## 10.3. I/O Devices State

With the exception of the stdin and stdout devices, Open Firmware *shall* close all devices with the following conditions true:

All Devices - no DMA and not interrupting

Normal I/O Devices - not responding to access PCI and ISA PnP Adapter/Devices

HOST Bridges - responding to config cycles and passing through config cycles to children

RTAS Devices - contract with Open Firmware to leave in state to perform intended function

## 10.4. Client Program Format

The data format of a client program compliant with this specification *shall* be either ELF (Executable and Linkage Format) as defined by [21], and extended by section 10.4.1.1., or PE (Portable Executable) as defined by [19]. The standard ELF format contains explicit indication as to the program's execution modes (e.g., 32- or 64-bit, Big- or Little-Endian). CHRP only supports the 32-bit version (i.e., ELFCLASS32) for 32 and 64 bit platforms.

**Note:** other client program formats may be supported, in an implementation specific manner, by an Open Firmware implementation.

A standard client program *shall* be statically linked, requiring no relocation of the image. The program's entry point (`e_entry`) *shall* contain the address of the first PowerPC instruction of the client program. It is the responsibility of the client program to establish the appropriate value of the TOC (`r2`), if necessary.

**Note:** the entry point is the address of the first instruction of the client program, not that of a procedure descriptor.

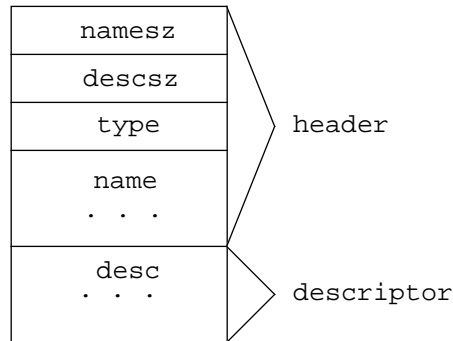
### 10.4.1. ELF-Format

This section defines how Open Firmware recognizes and prepares to execute an ELF-Format Program.

#### 10.4.1.1. ELF Note Section

Part of the process of loading a client program involves verifying its environmental requirements (e.g., endianness and address translation mode) against the current firmware configuration. The client's endianness can be directly determined by examining the ELF EI-DATA value; ELFDATA2LSB (1) implies Little-Endian while ELFDATA2MSB (2) implies Big-Endian. However, the other client requirements (e.g., address translation mode) are defined by means of an ELF Note Section (PT\_NOTE), pointed to by the program header. The following describes the format of the Note Section for a client program file.

As defined by [21], an ELF file can be "annotated" by means of Note Sections within the executable file. A Note Section contains a "header" followed by a (possibly null) "descriptor", as follows:



**Note: the endian format of the values corresponds to the endian-ness specified by the EI-DATA field of the file.**

The format of a Note Section header can be described by an CHRP Open Firmware struct as:

```

18 struct    \ Note Section header for Open Firmware
19 /L field   ns.namesz      \ length of ns.name, including NULL
20 /L field   ns.descrsz
21 /L field   ns.type
22 0 field   ns.name        \ NULL-terminated, /L padded

```

The ns.name field of the PowerPC Open Firmware Note Section shall be "PowerPC"; the ns.type field n shall be 0x1275.

Following the Note Section header is a descriptor (desc); the length (in bytes) of the descriptor is specified by a word in the Note Section's header (descsz). The interpretation of the descriptor depends upon the kind of Note Section in which it is contained. For Open Firmware, the format of the Note Section's descriptor can be described by an Open Firmware struct, as follows:

```

31 struct    \ Note Section descriptor for CHRP Open Firmware
32 /L field   ns.real-mode
33 /L field   ns.real-base
34 /L field   ns.real-size
35 /L field   ns.virt-base
36 /L field   ns.virt-size
37 /L field   ns.load-base

```

If the **ns.load-base** value is not -1, then that value is compared against the current value of the **load-base** configuration variable. If they are equal no further action is taken. If they are not equal then the **load-base** configuration variable is set to the value of **ns.load-base** and the system is rebooted.

#### 10.4.1.2. Recognizing ELF-Format Programs

The **init-program** shall recognize client program images that conform to all the requirements listed below as "ELF-format" programs.

In the description below, field names refer to fields within the ELF "file header" structure, which is assumed to begin at load-base, and offsets are relative to the beginning of that structure. Multi-byte numerical fields are interpreted according to the endianness specified by the "data" field at offset 5.

- a) The "e\_ident" field (at offset 0) contains the string "\7fELF", where "\7f" is a byte whose value is (hex) 7f. This indicates the beginning of an ELF file header.
- b) The "EI\_CLASS" field (at offset 4) contains the value 1. This indicates the 32-bit variant of the ELF format.
- c) The "e\_type" field (at offset 16) contains the value 2. This indicates that the ELF image is executable.

- 1 • d) The "e\_machine" field (at offset 18) contains the value 20. This indicates that the ELF image is for the Pow-
- 2 erPC instruction set.
- 3 • e) The "e\_version" field (at offset 20) contains the value 1.
- 4 • f) The "e\_flags" field (at offset 36) contains the value 0.

### 10.4.1.3. Preparing ELF-Format Programs for Execution

7 Upon recognition of the client program image at load-base as an ELF-format program, init-program *shall* pre-  
 8 pare the program for execution by performing the following sequence of steps.

9 In the description below, the fields mentioned by name are within ELF "program header" structures, unless  
 10 specified otherwise.

- 11 • a) Search for an ELF "note" section of type "1275" as defined in the section "ELF Note Section". If one is found,  
 12 and the values specified by its descriptor do not match the firmware's current operating mode, set the appropriate  
 13 configuration variables to the values specified in the note section descriptor, and restart the firmware so that it  
 14 will re-execute the **boot** command that resulted in the execution of **init-program**.
- 15 • b) Set the p\_paddr field for each PT\_LOAD segment equal to its p\_vaddr field value if **real-mode?** is false  
 16 and p\_paddr is -1. This effectively maps these segments v=r.
- 17 • c) Allocate and map, if required, sufficient physical memory for all program segments of type PT\_LOAD (i.e.  
 18 whose "p\_type" field contains the value 1) listed in the ELF image's program headers. Note that all PT\_LOAD  
 19 program segments that have a p\_paddr value that matches their location in physical memory need not be moved,  
 20 but the memory that they occupy must be claimed. This special case is added to allow large program images to  
 21 be loaded without the 2x memory required to move the segments.
- 22 • d) Copy the program headers to a "safe" location to guard against the possibility of them being overwritten by the  
 23 following steps.
- 24 • e) For each program segment of type "PT\_LOAD":  
 25
  - 26 • 1) Copy, if required, the initialized portion of the program segment from its current loca-  
 27 tion in the loaded image to the location given by the section's "p\_paddr" field.
  - 28 • 2) Fill the rest of the segment with zero bytes (i.e., fill "p\_filez - p\_memsz" bytes begin-  
 29 ning at the address "p\_paddr + p\_filez").
  - 30 • 3) If **real-mode?** is false, then map the program segment to the virtual address speci-  
 31 fied by p\_vaddr.
- 32 • f) Set the saved program state so that subsequent execution of **"go"** will begin execution at the address given by  
 33 the "e\_entry" field in the ELF file header. The e\_entry field is a physical address if **real-mode?** is *true* and is  
 34 a virtual address if **real-mode?** is *false*.

35 The implementation need not take precautions to ensure that the process of copying and zeroing program segments  
 36 does not overwrite the portions of the load image that have not yet been copied. In order to guarantee correct  
 37 copying, the value of the **load-base** configuration variable and the destination addresses of the various sections  
 38 must be such that such overwriting does not occur. One sufficient condition is that the region of memory beginning  
 39 at **load-base**, of size equal to the size of the loaded image, be disjoint from the regions of memory to which the  
 40 program segments are copied and zero-filled. Another sufficient condition is to specify a **load-base** in the  
 41 Notes Section (PT\_NOTE) that ensures that the PT\_LOAD segments are loaded at the address required by their  
 42 program headers and thus are not moved. There are other less-stringent sufficient conditions, especially for simple  
 43 ELF images with a small number of program segments that are to be copied to contiguous regions.

44 An implementation *shall* permit the ELF image to contain other program segments in addition to those described  
 45 above, but need not take any action beyond that defined above as a result of the presence of such other program  
 46 segments.

47 An implementation *shall* ignore all ELF sections. ELF sections are intended for binders, not loaders. Note that  
 48 the CHRP ELF Note Section is actual an ELF segment of type PT\_NOTE and thus the above does not apply to  
 49 it.

## 10.5. Additional Client Interface Requirements

This section describes processor assist callbacks for real and virtual memory management and a service.

### 10.5.1. Client Interface Callbacks

This section describes callbacks for memory management. These callbacks are provided by the client.

#### 10.5.1.1. Real-Mode Memory Management Assist Callbacks

claim\_mem

IN: [address] min\_addr, [address] max\_addr, size, align  
 OUT: throw-code, error, [address] real\_addr

Allocate contiguous physical memory between min\_addr and max\_addr of size bytes (128KB max for an area in the 0 to 16MB address range), with align alignment. The alignment boundary is the smallest power of two greater than or equal to the value of align; an align value of 1 signifies one-byte alignment. A non-zero error code *shall* be returned if the mapping can not be performed. If error code is zero (i.e. allocation succeeded) the routine returns the real address (real\_addr) of the physical memory block which was allocated for Open Firmware.

release\_mem

IN: [address] phys, size  
 OUT: throw-code

Free size bytes of physical memory starting at real address phys, making that physical memory available for later use. That memory must have been previously allocated by claim\_mem.

#### 10.5.1.2. Virtual Address Translation Assist Callbacks

alloc\_virt\_mem

IN: size  
 OUT: throw-code, error, [address] virt\_addr

Return the virtual address of a virtual memory area of size bytes aligned to a doubleword (8-byte) boundary. A non-zero error code *shall* be returned if the allocation can not be performed. If error code is zero (i.e. allocation succeeded) the routine returns the virtual address (virt\_addr) of the memory block which was allocated.

free\_virt\_mem

IN: [address] virt\_addr, size  
 OUT: throw-code

Free memory allocated by alloc\_virt\_mem. The values virt\_addr and size must correspond with memory previously allocated by alloc\_virt\_mem.

claim\_virt

IN: size, align  
 OUT: throw-code, error, [address] virt\_addr

Allocate a memory area of size bytes and alignment align. The alignment boundary is the smallest power of two greater than or equal to the value of align; an align value of 1 signifies one-byte alignment. A non-zero error code *shall* be returned if the allocation can not be performed. If error code is zero (i.e. allocation succeeded) the routine returns the virtual address (virt\_addr) of the memory block which was allocated.

release\_virt

IN: [address] virt, size  
 OUT: throw-code

Free size bytes of virtual memory starting at virtual address virt, making that physical memory and the corresponding ranges of virtual address space available for later use. That memory must have been previously allocated by claim\_virt.

### 10.5.2. Client Interface Services

Open Firmware *shall* provide the following *Client Interface Service*:

```

1 test-method
2     IN:  phandle, [string] method
3     OUT: missing-flag?
4     Tests whether the package method named method exists in the package phandle. missing-flag? is FALSE (0)
5     if the method exists or TRUE (-1) if the method does not exist.
6

```

## 11. Support Packages

This section describes the CHRP System binding specific requirements of Open Firmware support packages. These support packages are "disk-label" and "tape-label". For "network" and/or "obp-tftp" extensions, refer to [16]. These packages support the loading and executing of a client program. Another means of executing a Client Program is provided when an operating system ROM is a "bootable device"(Refer to Section 5.1.3. on page 25, as an example).

### 11.1. "disk-label" Support Package

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

The name of the partition (and, a file contained within the partition) can be explicitly specified by the user via the **load** or **boot** command, or can be implicitly specified by the value of the "**boot-device**" property of the "/options" node. The partition and filename are the ARGUMENTS portion of the final COMPONENT of the PATH\_NAME, as described in section 4.3.1 of [1].

The syntax for explicit partition/filename specification is given in section 11.1.2. below where partition identifies the partition to be used and filename is the name of a file within that partition. If partition is omitted, the default partition (as determined by the partition format) is used. If filename is omitted, the default filename (i.e., the filename component of the **boot-device** path-name) is used.

#### 11.1.1. Media Layout Format

This section describes the media layout formats of Client Program Images that the disk-label support package for a CHRP platform *shall* support; an implementation *may* support additional mechanisms, in an implementation-specific manner. The "disk-label" package for a CHRP platform *shall* support at least four(4) media layout types:

- FAT (FAT12 and FAT16 File System)
- FDISK (Partitions 4, 5, 6, 0x41 & 0x96)
- ISO-9660 (9660 File System)
- Mac OS (MAC Binary Image)

##### 11.1.1.1. FDISK Partition Types

The following FDISK partition types *shall* be supported:

Partition Type 4: FAT 12 or FAT 16 File System

Partition Type 5: Extended Chained Partitions

Partition Type 6: Extended Partitions

Partition Type 0x41: Single program image

Partition Type 0x96: ISO 9660 File System



1 FDISK partition type 0 is a free partition. Partition type 0x82 is reserved and should not be used by CHRP  
2 Architecture.

### 3 4 **11.1.2. Open Method Algorithm**

5 The **open** method of the "disk-label" support package *shall* implement a disk partition recognition  
6 algorithm that supports at least the set of disk formats that are supported by the following algorithm. The  
7 following algorithm is intended to support raw (uninterpreted) disks, raw partitions of disks beginning with an  
8 FDISK partition map, and files on FAT and ISO-9660 file systems both within FDISK partitions and by  
9 themselves on disks without a partition map.

10  
11 That **open** method *shall* accept an argument string (as returned by "**my-args**") with the following syntax  
12 (according to the algorithm below), where brackets denote an optional component:

13 [partition],[filename]  
14

15 If the argument string contains a comma, or if the argument string begins with a decimal digit, the partition  
16 component is deemed to be present. Note that the arguments above are not the client arguments with the boot  
17 command.

18 If the partition component is present, it selects the desired partition, where partition 0 refers to the entire disk,  
19 partition 1 refers to the first partition, partition 2 to the second, and so forth. If the partition component is  
20 absent and the disk has an FDISK or Mac partition map, the first "bootable" partition is used. If a "bootable"  
21 partition is not found, then fail in an implementation specific manner with an error.

22  
23 If the filename component is present, it selects a particular file within the file system on the disk or partition  
24 thereof.  
25

#### 26 27 **11.1.2.1. OPEN algorithm**

28 11.1.2.1.1. Set D.SIZE to -1

29 11.1.2.1.2. If ARGUMENT\$ is not the null string and the first  
30 character of ARGUMENT\$ is in the range '0'..'9' or is  
31 equal to ','  
32

33 11.1.2.1.2.1. LEFT-PARSE ( ARGUMENT\$ ) -> PARTITION\$, FILENAME\$

34 11.1.2.1.3. Else

35 11.1.2.1.3.1. Set PARTITION\$ to the null string

36 11.1.2.1.3.1.0.1. Set FILENAME\$ to ARGUMENT\$

37 11.1.2.1.4. If PARTITION\$ is not the null string

38 11.1.2.1.4.1. If PARTITION\$ is not a decimal number

39 11.1.2.1.4.1.1. Return FALSE

40 11.1.2.1.4.2. DECIMAL\_STRING\_TO\_NUMBER( PARTITION\$ ) -> PARTITION

41 11.1.2.1.4.3. If PARTITION is 0

42 11.1.2.1.4.3.1. GET\_DISK\_SIZE

43 11.1.2.1.4.4. Else

44 11.1.2.1.4.4.1. Read the first 512 bytes of the device into a buffer

45 11.1.2.1.4.4.2. SELECT\_EXPLICIT\_PARTITION( PARTITION )

46 11.1.2.1.4.4.3. If SELECT\_EXPLICIT\_PARTITION returned an error  
47  
48  
49  
50  
51  
52  
53  
54

```
1             indication
2 11.1.2.1.4.4.3.1. Return FALSE
3 11.1.2.1.5. Else \ PARTITION$ is NULL
4 11.1.2.1.5.1. Read the first 512 bytes of the device into a buffer
5 11.1.2.1.5.2. SELECT_ACTIVE_PARTITION
6 11.1.2.1.5.3. If SELECT_ACTIVE_PARTITION returned an error
7             indication
8 11.1.2.1.5.3.1. Return FALSE
9 11.1.2.1.6. \ (At this point, D.OFFSET is set to the beginning of
10 the selected partition and D.SIZE is set to the size
11 of that partition. If the entire disk was selected,
12 D.OFFSET is 0 and D.SIZE is the size of the disk.)
13 11.1.2.1.7. Call parent's "seek" method with an argument of 0,0.
14 11.1.2.1.8. Return TRUE
15
16 11.1.2.2. CHECK_FOR_BPB procedure
17
18 11.1.2.2.1. If the first four(4) bytes are EBCDIC 'IBMA'(hex
19 character string C9C2D4C1), then the sector does not
20 contain a BPB.
21
22 11.1.2.2.2. If the 16-bit little-endian quantity beginning at
23 buffer offset 510 is 0xAA55, and the 16-bit little-
24 endian quantity beginning at buffer offset 11 (which
25 is the BPB "bytes per sector" field) is either 256,
26 512, or 1024, and the byte at offset 16 (the BPB
27 "number of FATs" field is either 1 or 2, the sector is
28 deemed to contain a BPB. Otherwise, the sector does
29 not contain a BPB.
30
31 11.1.2.3. CHECK_FOR_ISO_9660 procedure
32
33 11.1.2.3.1. Read 512-byte sector 64 (the beginning of logical 2048-
34 byte sector 16) into a buffer.
35
36 11.1.2.3.2. If the byte at offset 0 contains the binary number "1",
37 and the 5 bytes beginning at offset 1 contains the
38 text string "CD001", the partition or raw disk is
39 deemed to contain an ISO 9660 file system. Otherwise,
40 the partition or raw disk is deemed not to contain an
41 ISO 9660 file system.
42
43 11.1.2.4. CHECK_FOR_FDISK procedure
44
45 11.1.2.4.1. If the buffer does not contain an FDisk partition map
46 signature of "AA55" as a 16-bit little-endian number
47 beginning at buffer offset 510, the buffer is deemed
48 not to contain an FDISK partition map.
49
50 11.1.2.4.2. If none of the partition type code field (the bytes at
51 buffer offsets 0x1C2,0x1D2, 0x1E2, and 0x1F2)
52 contains a recognizable partition type code (4,5, 6,
53 0x41, 0x96, or other types that may be recognized by
54 the implementation), the buffer is deemed not to
55 contain an FDISK partition map.
56
57 11.1.2.4.3. Otherwise, the buffer is deemed to contain an FDISK
```

---

1                   partition map.

2 11.1.2.4.4. The implementation may, at its option, apply additional

3                   validity tests to the partition map information.

4

5 **11.1.2.5. CHECK\_FOR\_MAC\_DISK procedure**

6 11.1.2.5.1. If the first(i.e., at the lowest offset) two bytes in

7                   the buffer contains the 16-bit big-endian signature

8                   0x4552, then the disk is deemed to be a Mac

9                   partitioned disk. Otherwise, the partition or raw

10                  disk is deemed not to be a Mac partitioned disk.

11                  **Note: Subsequent 512 byte sectors will contain Mac partition map entries, each of which**

12                  **begins with the 16-bit big-endian signature 0x504D. Each such partition map entry con-**

13                  **tains a field (V) indicating the total number of partition entries in the map.**

14

15 **11.1.2.6. INTERPOSE\_BY\_TYPE procedure**

16 11.1.2.6.1. If FILENAME\$ is not the null string

17 11.1.2.6.1.1. If PARTITION-TYPE is 0x96

18 11.1.2.6.1.1.1. INTERPOSE[17](ISO-9660 File System package,

19                                   FILENAME\$)

20 11.1.2.6.1.2. Else

21 11.1.2.6.1.2.1. If PARTITION-TYPE is FAT,

22 11.1.2.6.1.2.1.1. INTERPOSE (FAT File System package, FILENAME\$)

23

24 **11.1.2.7. SELECT\_ACTIVE\_PARTITION( PARTITION ) procedure**

25 11.1.2.7.1. CHECK\_FOR\_BPB

26 11.1.2.7.1.1. If the buffer contains a BPB

27 11.1.2.7.1.1.1. Set OFFSET to 0

28 11.1.2.7.1.1.2. Set D.SIZE to the maximum size of the disk in bytes,

29                                   as indicated by the information in the BIOS

30                                   Parameter Block

31

32 11.1.2.7.1.1.3. If FILENAME\$ is not the null string

33 11.1.2.7.1.1.3.1. INTERPOSE (FAT File System package, FILENAME\$)

34 11.1.2.7.1.1.4. Return OKAY

35 11.1.2.7.2. CHECK\_FOR\_FDISK

36 11.1.2.7.3. If the buffer contains an FDISK partition map

37 11.1.2.7.3.1. Search the Fdisk partition map, reading new 512-byte

38                                   sectors into the buffer if necessary to "chain" to

39                                   extended partition entries (i.e. ones whose type

40                                   byte at offset 4 contains "5") for the first(i.e.,

41                                   at the lowest offset)partition entry whose

42                                   "bootable" field (at offset 0 in the partition

43                                   entry) contains 0x80.

44

45 11.1.2.7.3.2. If a "bootable" partition was found:

46 11.1.2.7.3.2.1. Set PARTITION-TYPE to that entry's "type" field

47                                   (the byte at offset 4)

48 11.1.2.7.3.2.2. Set D.OFFSET to the byte offset from the beginning

49                                   of the disk of the beginning of the partition

50                                   denoted by that entry.

51

52 11.1.2.7.3.2.3. Set D.SIZE to the size of the partition denoted by

53

54

---

```
1           that entry.
2 11.1.2.7.3.2.4. INTERPOSE_BY_TYPE
3
4 11.1.2.7.3.2.5. Return OKAY
5     \ (If this point is reached, no partition was marked "bootable")
6 11.1.2.7.3.3. Search the FDisk partition map beginning in 512-byte
7           sector 0, reading new 512-byte sectors into the
8           buffer if necessary to "chain" to extended
9           partition entries, for the first partition(i.e.,
10          at the lowest offset) entry whose "type" byte is
11          neither 0 nor 5 (5 is the type code that indicates
12          a "chained" extended partition entry).
13 11.1.2.7.3.4. If one is found:
14 11.1.2.7.3.4.1. Set PARTITION-TYPE to that entry's "type" field
15                 (the byte at offset 4)
16 11.1.2.7.3.4.2. Set D.OFFSET to the byte offset from the beginning
17                 of the disk of the beginning of the partition
18                 denoted by that entry.
19 11.1.2.7.3.4.3. Set D.SIZE to the size of the partition in bytes
20                 denoted by that entry.
21 11.1.2.7.3.4.4. INTERPOSE_BY_TYPE
22 11.1.2.7.3.4.5. Return OKAY
23 11.1.2.7.3.5. Else \ (If this point is reached, the partition map
24                 did not contain any valid partition entries)
25 11.1.2.7.3.5.1. Return ERROR
26 11.1.2.7.4. CHECK_FOR_ISO_9660
27 11.1.2.7.5. If it is an ISO 9660 disk
28 11.1.2.7.5.1. GET_DISK_SIZE
29 11.1.2.7.5.2. If FILENAME$ is not the null string
30 11.1.2.7.5.2.1. INTERPOSE (ISO-9660 File System package, FILENAME$)
31 11.1.2.7.5.3. Return OKAY
32 11.1.2.7.6. CHECK_FOR_MAC_DISK
33 11.1.2.7.7. If this is a Mac partitioned disk
34 11.1.2.7.7.1. Search the Mac partition table for the first
35                 "bootable" partition. A partition is "bootable"
36                 when the pmPartStatus flags indicate that this is
37                 a valid, allocated, readable and bootable
38                 partition and the pmProcessor field contains
39                 "powerpc" (using case-insensitive matching).
40 11.1.2.7.7.2. If a Mac "bootable" partition is found
41 11.1.2.7.7.2.1. If FILENAME$ is "%BOOT"
42 11.1.2.7.7.2.1.1. If the Nth partition is marked bootable
43 11.1.2.7.7.2.1.1.1. Set D.OFFSET to the byte offset from the
44                    beginning of the disk to the beginning of
45                    the boot area, as given by the
46                    pmLgBootStart field.
47 11.1.2.7.7.2.1.1.2. Set D.SIZE to the size of the partition in bytes
48
49
50
51
52
53
54
```

---

1 denoted by pmBootSize.  
2 11.1.2.7.7.2.1.1.3. Return OKAY  
3 11.1.2.7.7.2.2. Else  
4 11.1.2.7.7.2.2.1. If the FILENAME\$ is the null string  
5 11.1.2.7.7.2.2.1.1. Set D.OFFSET to the byte offset of the "real"  
6 partition data  
7 11.1.2.7.7.2.2.1.2. Set D.SIZE to the size of the "real" partition  
8 data  
9 11.1.2.7.7.2.2.2. Else  
10 11.1.2.7.7.2.2.2.1. INTERPOSE\_BY\_TYPE  
11 11.1.2.7.7.2.2.3. Return OKAY  
12 11.1.2.7.7.3. Else  
13 11.1.2.7.7.3.1. Return ERROR  
14 11.1.2.7.7.4. (If this point is reached, no "bootable" partition was  
15 found)  
16 11.1.2.7.7.5. Return ERROR  
17  
18  
19  
20 **11.1.2.8. GET-DISK-SIZE procedure**  
21 11.1.2.8.1. Set OFFSET to 0  
22 11.1.2.8.2. If the parent has a "#blocks" method  
23 11.1.2.8.2.1. Execute the parent's "#blocks" and "block-size"  
24 methods  
25 11.1.2.8.2.2. Set D.SIZE to the product of the numbers they returned  
26 11.1.2.8.3. Else  
27 11.1.2.8.3.1. Set D.SIZE to -1  
28  
29  
30 **11.1.2.9. SELECT\_EXPLICIT\_PARTITION procedure**  
31 11.1.2.9.1. CHECK\_FOR\_BPB  
32 11.1.2.9.2. If the buffer contains a BPB  
33 11.1.2.9.2.1. If PARTITION is 1  
34 11.1.2.9.2.1.1. Set OFFSET to 0  
35 11.1.2.9.2.1.2. Set D.SIZE to the maximum size of the disk in bytes,  
36 as indicated by the information in the BIOS  
37 Parameter Block  
38 11.1.2.9.2.1.3. If FILENAME\$ is not the null string  
39 11.1.2.9.2.1.3.1. INTERPOSE (FAT File System package, FILENAME\$)  
40 11.1.2.9.2.1.4. Return OKAY  
41 11.1.2.9.2.2. Else \ Have a BPB, but PARTITION <> 1  
42 11.1.2.9.2.2.1. Return ERROR  
43 11.1.2.9.3. CHECK\_FOR\_FDISK  
44 11.1.2.9.4. If an Fdisk partition map is found  
45 11.1.2.9.4.1. Search the Fdisk partition map beginning in 512-byte  
46 sector 0, reading new 512-byte sectors into the  
47 buffer if necessary to "chain" to extended  
48 partition entries, for the N<sup>th</sup>, where N is the  
49 value of PARTITION, partition entry whose "type"  
50 byte is neither 0 nor 5 (5 is the type code that  
51  
52  
53  
54

---

```

1
2
3
4         indicates a "chained" extended partition entry).
5 11.1.2.9.4.2. If the Nth partition is found:
6 11.1.2.9.4.2.1. Set PARTITION-TYPE to that entry's "type" field
7                 (the byte at offset 4)
8 11.1.2.9.4.2.2. Set D.OFFSET to the byte offset from the beginning
9                 of the disk to the beginning of the partition
10                denoted by that entry.
11 11.1.2.9.4.2.3. Set D.SIZE to the size of the partition in bytes
12                denoted by that entry.
13 11.1.2.9.4.2.4. INTERPOSE_BY_TYPE
14 11.1.2.9.4.2.5. Return OKAY
15 11.1.2.9.4.3. Else \Nth partition does not exist
16 11.1.2.9.4.3.1. Return ERROR
17 11.1.2.9.5. CHECK_FOR_MAC_DISK
18 11.1.2.9.6. If this is a Mac partitioned disk
19 11.1.2.9.6.1. Search the Mac partition map for the Nth partition,
20                where N is the value of PARTITION.
21 11.1.2.9.6.2. If the Nth partition is valid, allocated, and readable
22 11.1.2.9.6.2.1. If FILENAME$ is %BOOT
23 11.1.2.9.6.2.1.1. If the Nth partition is marked bootable
24 11.1.2.9.6.2.1.1.1. Set D.OFFSET to the byte offset from the
25                    beginning of the disk to the beginning of
26                    the boot area, as given by the
27                    pmLgBootStart field.
28 11.1.2.9.6.2.1.1.2. Set D.SIZE to the size of the partition in bytes
29                    denoted by pmBootSize.
30 11.1.2.9.6.2.1.1.3. Return OKAY
31 11.1.2.9.6.2.1.2. Else \Nth partition not "bootable"
32 11.1.2.9.6.2.1.2.1. Return ERROR
33 11.1.2.9.6.2.2. Else
34 11.1.2.9.6.2.2.1. If FILENAME$ is not the null string
35 11.1.2.9.6.2.2.1.1. INTERPOSE_BY_TYPE
36 11.1.2.9.6.2.2.2. Return OKAY
37 11.1.2.9.6.3. Else \ (If this point is reached, the partition is
38                invalid)
39 11.1.2.9.6.3.1. Return ERROR
40 11.1.2.9.7. Else \ (If this point is reached, the partition map is
41                not regconized)
42 11.1.2.9.7.1. Return ERROR
43
44
45
46
47
48
49
50

```

This algorithm can be used to locate the correct partition and/or 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-specifier* (i.e., the field following the ':' of the last path component). Other formats may be recognized in an implementation-specific manner.

## 11.2. "tape-label" Support Package

The "tape-label" Support Package *shall* support tape as a standard "byte" device with the set of methods specified in [1], Section 3.7.3. Presence of the `bootinfo.txt` file is optional.

The `open` method *shall* accept an argument string, where brackets denote an optional component:


[file number]

where the first file on the tape media is located at file number 0.

### 11.2.1. Tape Format

The CHRP tape format *shall* consist of files ending with a file mark(FM). The first block of data will be identified as file 0. The `bootinfo.txt` file, if present, *shall* be located on the tape as file 0 (the first file). There *shall* be only one `bootinfo.txt` file on the tape media. Refer to Figure 1 below for the CHRP Tape format.

#### Optional *bootinfo.txt* File present



#### *bootinfo.txt* File not present




FIGURE 1 Tape Boot Format

### 11.2.2. Tape *bootinfo.txt* File

The `bootinfo.txt` file *shall* have included for each set of `<chrp-boot>` tags a set of `<boot-script>` tags that contains a pointer to the program image to be loaded (Refer to Section 3.1.6. on page 19). The form for this tape pointer will be:

device specifier = *device:file number*

EXAMPLE: device specifier = `tape:2` (For the specified set of `<chrp-boot>` tags, load the tape program image from file 2).

A `bootinfo.txt` file may contain a multiple set of `<chrp-boot>` tags where each one can point to a different tape file number. If a `bootinfo.txt` file is not present, file 0 should be a bootable file. Only file 0 will be loaded as a bootable image. No other files will be searched if a `bootinfo.txt` file is not present unless the file number to load is specified by an argument.

### 11.3. "network" Support Package

The "network" Support Package *shall* adhere to the *Recommended Practice - TFTP Booting Extension*[16] documentation functions and conventions.

### 11.4. Program-image formats.

Open Firmware must recognize a client program that is formatted as ELF [21] and PE [19]. Other formats *may* be handled in an implementation-specific manner. Reference [13] defines using FCode and Forth Program-Image Formats.

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.**

```

set restart? false

if the image is in ELF format(Refer to Section 10.4.1.2. on page 53)
    if the EI_DATA field does not match little-endian?
        set little-endian? appropriately.
        set restart? true

    locate the PowerPC Note Section

    if the Note Section's values do not match
        set Configuration Variables appropriately
        set restart? true

    if restart?
        restart the system, possibly by executing reset-all

else
    move and/or relocate the ELF image
    (Refer to Section 10.4.1.3. on page 54).

    set GO's context with initial register values

else if the image is in PE format
    if little-endian? is false
        set little-endian? to true.
```



```

1         restart the system, possibly by executing reset-all
2
3     else
4
5         move and/or relocate the PE image.
6
7         set GO's context with initial register values
8
9     else if the image is FCode Image(hex characters F1,08)
10
11         setup system and do subsequent go and perform a byte load of the FCode
12         image
13
14     else if the image is Forth Code Source Image(ASCII characters "<space>")
15
16         setup system to evaluate Forth Source Image
17
18     else if the image is a bootinfo.txt file (i.e., begins with "<CHRP-BOOT>")
19
20         setup system to parse the bootinfo.txt file
21
22     else
23
24         FAIL, in an implementation-specific manner.

```

### 11.4.1. CHRP PE Program-image Format Support

When the PE File Format is recognized, the Windows NT Veneer is loaded with the following environment:

- Open Firmware *shall* ensure that virtual address space from 0x80000000 to 0x81000000 (2GB to 2GB+16MB) is available and that any other virtual address space has not been mapped into the physical address space from 0 to 0x1000000 (0 to 16MB). If necessary, Open Firmware *shall* also modify the **real-base** configuration variable in such a way as to ensure that the Open Firmware image itself is not loaded into any part of the 0 to 0x1000000 physical address space.
  - For CHRP Systems with greater than 256MB's of physical memory, a suggested Open Firmware Image loading address is at the top of the first 256MB bank.
  - For CHRP Systems with equal to or less than 256MB's, a suggested Open Firmware Image loading address is at the top of physical memory.
- Open Firmware *shall* zero the first 256 bytes (0 to 0x100) of physical memory.
- Open Firmware *shall* set the **little-endian?** configuration variable to little endian mode.
- Open Firmware *shall* set the **real-mode?** configuration variable to virtual address mode.

#### 11.4.1.1. Windows NT Veneer Implementation

The current NT Veneer implementation calls the MMU node's standard **map** method through the Client Interface to map virtual address 0x80000000 to physical address 0. CHRP System's intention is to continue using NT calling the standard **map** method to map virtual memory to physical memory.

## 12. Open Firmware Related NVRAM Usage

This section describes the NVRAM partitions for a CHRP platform. The platform Open Firmware defines the currently-defined NVRAM partitions and tags(see table below). Other partitions whose tags and usage are defined in the CHRP specification [3], Chapter 8, should be treated as RESERVED FOR FUTURE USE.

Table 12. Open Firmware Related NVRAM Partitions

Signature	Type	Ownership <sup>1</sup>	Number of partitions	Description <sup>2</sup>
0x50	Open Firmware	Firmware	1	For general OF usage
0x51	Firmware	Firmware	0 to n	General FW usage
0x52	Hardware	Firmware	1 to n	For system information
0x70	System	Global	1	For configuration variables
0x71	Configuration	Global	0 to n	For ISA and Power Management utilities data resource
0x72	Error log	Global	0 to n	For errors during OF/RTAS
0x7E	Vendor-defined	Global	0 to n	“name” prefix
0x7F	Free Space	Global	0 to n	Marks free space

**Note 1: Ownership**

Firmware - Read/Write of partition done only by firmware

Global - Read/Write of partition can be done by firmware and/or operating system

**Note 2: Definitions**

OF - Open Firmware

FW - Firmware

RTAS - Run Time Abstraction Services (OS environment)

Note 3: OS's each have their own private partition, using one of the signatures 0xA0 through 0xAF, Signature Type OS.

**12.1. Open Firmware Partition**

This partition is for Open Firmware usage

**12.2. Firmware Partition**

This partition is for firmware usage.

**12.3. Hardware Partition**

This partition is used for platform specific operations; i.e., like storage of a vendor Vital Product Data(VPD).

**12.4. System Partition**

The System Partition, with name = '**common**', contains information that is accessible to both Open Firmware and operating systems. The contents of this partition are represented in the Open Firmware device tree as properties (i.e., (*name*, *value*) pairs) in the `/options` node. While Open Firmware is available, the operating system can alter the contents of these properties by using the `setprop` client interface service. When Open Firmware is no longer available, the operating system can alter the contents of the System Partition itself,

1 following the rules below for the formats of the *name* and *value*. Information is stored in the System Partition  
 2 as a sequence of (*name*, *value*) pairs in the following format:

3       *name* = *value*  
 4

5 where *name* follows the rules defined in Section 12.4.1 and *value* follows the rules defined in Section 12.4.2.  
 6 The end of the sequence of pairs is denoted by a null (0x00) byte.

### 8 **12.4.1. Name**

9 Since the data in the System Partition is an external representation of properties of the `/option` node, the  
 10 name component must follow the rules for *property names* as defined by Section 3.2.2.1.1 Property names of  
 11 [1]; i.e., a string of 1-31 printable characters containing no uppercase characters or the characters `'/'`, `'\'`,  
 12 `':'`, `'['`, `']'` or `'@'`. In addition to these rules, a naming convention is required for operating system  
 13 specific names to avoid name conflicts. Each such *name* must begin with the OS vendor's OUI followed by a  
 14 `'.'`; e.g., `aapl,xxx` or `ibm,xxx`. This introduces separate name spaces for each vendor in which it  
 15 manages its own naming conventions.

### 17 **12.4.2. Value**

18 The value component of System Partition data can contain an arbitrary number of bytes in the range 0x01-  
 19 0xFF, terminated by a null (0x00) byte. Bytes in the range 0x01-0xFE represent themselves. In order to  
 20 allow arbitrary byte data to be represented, an encoding is used to represent strings of 0x00 or 0xFF bytes.  
 21 This encoding uses the 0xFF byte as an escape, indicating that the following byte is encoded as:

22 **bnnnnnnn**

23 where **b**, the most-significant bit, is **0** to represent a sequence of 0x00 bytes or **1** to represent a sequence of  
 24 0xFF bytes. **nnnnnnn**, the least-significant 7 bits, is a binary number (in the range 0x01-0x7F) that  
 25 represents the number of repetitions of 0x00 or 0xFF.

### 29 **12.4.3. Open Firmware Configuration Variables**

30 Open Firmware configuration variables control the operation of Open Firmware. In addition to the standard  
 31 configuration variables defined in [1], other configuration variables are defined by the PowerPC Processor  
 32 binding [6] and by this binding. While such variables are stored in the System Partition as described above,  
 33 they have additional rules placed on the format of the value component. Each configuration variable is also  
 34 represented by a user interface word (of the same name) that returns stack value(s) when that word is evaluated.  
 35 Each also has a platform defined default value; the absence of a configuration variable in the System Partition  
 36 indicates that the value is set to its default value. The format of the external representation of configuration  
 37 variables, and their stack representation, is defined by Section 7.4.4.1 Configuration Variables of [1]; the format  
 38 depends upon the data type of the configuration variable. Whereas the internal storage format is not defined by  
 39 [1], this binding specifies them as described below. The names of configuration variables are followed by the  
 40 reference in which they are defined (where [\*] refers to this binding).

#### 43 **12.4.3.1. Boolean Configuration Variables**

44 The value of a boolean configuration variable is represented in the System Partition as the string `"true"` or  
 45 `"false"`. The following configuration variables are of type boolean:

46       **auto-boot?** [1]  
 47       **diag-switch?** [1]  
 48       **fcode-debug?** [1]  
 49       **oem-banner?** [1]  
 50       **oem-logo?** [1]  
 51       **use-nvramrc?** [1]  
 52  
 53  
 54

---

```

1   little-endian? [5]
2   real-mode? [5]
3   menu? [*]
4

```

### 12.4.3.2. Integer Configuration Variables

The value of an integer configuration variable is represented in the System Partition as a decimal number or a hexadecimal number preceded by "0x". The following configuration variables are of type integer:

```

9   screen-#columns [1]
10  screen-#rows [1]
11  security-#badlogins [1]
12  security-mode [1]
13  selftest-#megs [1]
14  real-base [5]
15  real-size [5]
16  virt-base [5]
17  virt-size [5]
18  load-base [5]
19
20
21

```

### 12.4.3.3. String Configuration Variables

The value of a string configuration variable is represented in the System Partition as the characters of the string. Where multiple "lines" of text are represented (e.g., in the nvramrc script), each line is terminated by a carriage-return (0x0D), a linefeed (0x0A), or carriage-return, linefeed sequence (0x0D,0x0A). The following configuration variables are of type string:

```

22
23
24
25
26
27
28
29  boot-command [1]
30  boot-device [1]
31  boot-file [1]
32  diag-device [1]
33  diag-file [1]
34  input-device [1]
35  nvramrc [1]
36  oem-banner [1]
37  output-device [1]
38  security-password [1]
39  bootinfo-nnnnn [*]
40  reboot-command [*]
41
42
43

```

### 12.4.3.4. Byte Configuration Variables

The value of a bytes configuration variable is represented by an arbitrary number of bytes, using the encoding escape for values of 0x00 and 0xFF. The following configuration variables are of type bytes:

```

44
45
46
47
48  oem-logo [1]
49

```

## 12.4.4. nvramrc Script

A CHRP Platform will implement the function `nvramrc` script as defined in [1] to allow an OS to change variable settings or initialize needed environment variables.

In addition, it is recommended for a CHRP Platform that has a hardware configuration with multiple SCSI Controllers on the same SCSI Bus to assign SCSI ID's to avoid possible bus address conflicts. The mechanism to do this is a **nvrामrc** script using a *multi-initiator* script defined below. The *multi-initiator nvrामrc* script shall use **banner** or **surpress-banner** executed within the script to modify or suppress the automatic execution of the Open Firmware start-up sequence; **probe-all install-console banner** (Reference [1], Section 4.2.3). The *multi-initiator* script should execute **probe-all** to construct the device nodes before the *multi-initiator* is assigned.

#### 12.4.4.1. Multi-initiator Script

A CHRP Platform recommends the use of a **nvrामrc** script to set the SCSI ID or *multi-initiator* for the next boot operation for a SCSI hardware configuration that is shared by two systems. That is, on SCSI buses where there are two or more SCSI Controller Adapters present on the same SCSI Bus. Below is a recommended script by which an OS can set the alternate SCSI Adapter ID to not conflict with the primary SCSI Controller.

**nvrामrc** script for SCSI multi-initiator (This script is represented using new line characters <nl> to parse into a readable format):

```
dev device-specifier<nl>
nnn encode-int " scsi-initiator-id" property<nl>
device-end<nl><null>
```

where '**nnn**' is an integer specifying the SCSI ID. The <nl> is the ASCII new line character and <null> is the ASCII null character. Note that the formatting is for visual purposes.

Ground Rules:

- The "**scsi-initiator-id**" property defines the SCSI ID for that node (typically adapter). Absence of this property implies that SCSI ID 7 is in use by the adapter.
- This property can be assigned to any SCSI adapter node via the **nvrामrc** script.
- Open Firmware processes the **nvrामrc** script prior to performing any operations on the SCSI bus.
  - **banner** or **surpress-banner** shall be executed within the script to suppress Open Firmware start-up sequence
  - **probe-all** shall be executed within the script before setting the *multi-initiator* or SCSI ID
  - **install-console** shall be executed within the script after the setting of the SCSI ID to complete the normal Open Firmware start-up sequence
- The **nvrामrc** script is contained in the NVRAM System Partition(0x70).

**An example of a multi-initiator script follows:**

Assume that the following are two SCSI adapter paths; /pci@ff500000/pci3,1000@10 and /pci@ff500000/pci3,1000@18.

Then the following example sets the SCSI ID for /pci@ff500000/pci3,1000@10 to 6 and SCSI ID for /pci@ff500000/pci3,1000@8 to 7 at each boot before any SCSI operations.

This sequence of code should be inserted in the **nvrामrc** script:

```
probe-all dev /pci@ff500000/pci3,1000@10<nl>
6 encode-int " scsi-initiator-id" property<nl>
device-end<nl>
dev /pci@ff500000/pci3,1000@8<nl>
7 encode-int " scsi-initiator-id" property<nl>
device-end<nl>
install-console banner1<nl>
<null>
```

Note 1: `suppress-banner` could also be used to suppress the normal Open Firmware start-up sequence of operations

## 12.5. Configuration Partition

The Configuration Partition (0x52) has two separate partitions with names of *isa-config* and *pm-config* (when power management is implemented by the platform). The ICU stores the ISA Resource information in the *isa-config* partition.

The device power management resources are stored in the *pm-config* partition by a CHRP platform Power Configuration Utility (PCU).

### 12.5.1. ISA PnP Configuration Resources

ISA legacy card resource data is stored in the *isa-config* partition of NVRAM as shown in Section 9.1.2. on page 47.

### 12.5.2. Device Power Management Configuration Resources

The format of power management information provided in the optional *pm-config* partition is defined in Section 9.2.3. on page 49.

### 12.5.3. Error Log Partition

This partition data format is defined in the CHRP Architecture for the possible error logs resulting from RTAS calls.

## 12.6. Vendor-Defined Partition

This partition is used as part of the naming convention to identify a partition within a signature group.

## 12.7. Free Space Partition

This partition signature is used to mark free NVRAM in a CHRP Platform

-- END OF DOCUMENT --