



Disk Layouts, Booting, and Installation for Solaris on PowerPC

Background

Open Firmware and Conventional Firmware

Solaris® on PowerPC™ targets systems based on the PowerPC Reference Platform specification¹ (PReP). PReP requires systems shipped after June, 1995, to support the IEEE 1275 specification for Open Firmware² (“OF”), with additions specified by the PowerPC supplement.³ However, many systems that will ship during 1995, including the IBM platform that is the reference target for the initial Solaris implementation, will not have OF, but instead will boot using an *ad hoc* mechanism developed by IBM and described in the PReP specification as “conventional firmware” (“CF”).

Open Firmware Booting

The PowerPC supplement calls for OF to boot a specified file from either a DOS (“FAT-12” or “FAT-16”) file system or from an ISO-9660 file system. Locating the file to boot depends only on the data, not the kind of device on which it resides. If the medium appears to contain a DOS-style partition table in its first 512 bytes, OF boots from a specified partition or from the active partition (which in either case is presumably a DOS file system); otherwise, OF looks for an ISO-9660 directory beginning at the ISO-9660 defined offset of 32 Kbytes from the beginning of the disk.

1. *PowerPC Reference Platform Specification Version 1.1*, IBM Corporation, October 14, 1994.

2. *IEEE Standard 1275-1994: Standard for Boot Firmware: Core Requirements and Practices*.

3. *PowerPC processor binding to IEEE Std 1275-1994 Standard for Boot (Initialization, Configuration) Firmware*, Draft Revision 1.0, February 16, 1995.



Conventional Firmware Booting

CF boots only from a special, type 0x41, partition on a partitioned DOS disk. The 0x41 partition contains a load image which must be entirely self-sufficient because CF provides no callbacks into the firmware. That is, the loaded program must contain drivers for all the devices used during the remainder of bootstrapping, or at least drivers for devices from which it may load any other needed drivers. The specification is not entirely clear with respect to diskettes, but seems to require a diskette to be partitioned or to contain an ISO-9660 directory. The only example of a bootable diskette that we have seen contains a partition table and a 0x41 partition.

Approach

In order to unify the design and to reduce the design, architecture review, and implementation effort required to support both OF- and CF-based platforms, CF-based platforms under Solaris boot a program that emulates the Open Firmware client interface. This program, called “Virtual Open Firmware” (VOF) provides most of the client interface of OF⁴, including:

1. Access to a complete and correct device tree
2. Drivers for any devices used during bootstrapping or installation, including:
 - CD
 - Hard drive
 - Diskette
 - Network
 - Keyboard
 - Display
 - Serial port
 - Mouse (optional)
3. OF *packages* and methods, including:
 - `obp-tftp`, which provides the protocols needed for network booting and installation.
 - Terminal emulation for graphical displays.
 - Support for DOS and ISO-9660 file systems.

4. The subset of the 1275 client interface provided by VOF is described in *Interim Bootstrapping and Configuration Strategy for Solaris on PowerPC*, Steven Zucker, SunSoft, August, 1994.



- Memory allocation, freeing, and mapping.

The Platform Support consolidation is providing the first implementation of VOF for the target IBM platforms. Although VOF is platform-specific because of the first two items above, most of the implementation consists of framework code that handles client calls, the mechanics of device tree construction and properties, and code in the third item above, all of which is common across platforms. Adapting the reference VOF implementation to new platforms, especially those using the same buses and devices as the reference platform, is not very onerous. SunSoft will use it as the basis for VOF for other platforms and will make it available in kit form to OEMs, who can produce VOF for similar platforms with relatively little effort.

The most common platform adaptation involves only the addition of support for new devices, some of which are boot devices. The reference VOF implementation requires relinking of new drivers with object files for the base VOF implementation to create a new VOF module.⁵

Assumptions

1. OEMs (or SunSoft, if it provides the VOF implementation) provide some means of identifying all the buses and devices on the system and of determining their properties (e.g., addresses, interrupts, etc.). For self-identifying and configurable buses such as PCI, and for devices known to be part of the system (e.g., on the motherboard), this is not very difficult. For other buses (e.g., ISA) it may require a means of recording the system configuration in nonvolatile memory for VOF to read at boot time.
2. The native firmware on non-OF platforms supports the CF model of booting from a 0x41 partition.
3. An OF platform provides FCODE drivers for all boot devices. To use I/O cards without FCODE drivers during booting, the OEM must provide a means for installing FCODE drivers (e.g., in FLASH) where it is accessible to and recognized by the platform's OF implementation.

5. For booting from diskette, it would be possible to load needed VOF device drivers from a DOS file system on the diskette. Because most systems (all that we know of) use the same diskette interface, this would make adapting to additional boot devices as easy as adding the required VOF driver to a DOS diskette. However, this approach is impractical for booting from hard drives because it would require VOF drivers for all the various hard drive (e.g., SCSI) adapters to be built into the base VOF implementation. Rather than have separate mechanisms for loading drivers from diskettes and hard drives, Solaris requires VOF to be statically linked.



Requirements

Installation

1. Solaris must be installable from the same CD on OF platforms and CF platforms.
2. CF platforms must be installable with VOF and Solaris drivers and packages on a bootable diskette provided by an OEM. Installation of Solaris on any CF platform requires a bootable VOF diskette.
3. OF platforms not supported by the CD should be installable with Solaris drivers and packages on an auxiliary (non-bootable) diskette.

Driver Updates

The installation requirements above include the ability to install Solaris on a system with boot devices not supported by the Solaris CD. Once a system has been installed, driver updates are only a matter of installing ordinary Solaris packages on a running Solaris system.

Booting

On an OF platform Solaris must be able to boot from any device, including the net, for which there is an FCODE driver. After installation on a CF platform, Solaris must be bootable:

- Directly from a hard drive, or
- From the net after the platform-specific VOF has been loaded.

Disk Layout

CD

The Solaris CD contains a partition table with three partitions as shown in the table below.. With respect to the ISO-9660 standard, the whole disk (partition 0) is an ISO-9660 disk. The other partitions reside in space not occupied by ISO-9660 files. All the partitions are sized as needed; their sizes are known when the CD is created, and the CD is read-only so there is no need for them to be any larger.



Table 1 CD Partitions

| Partition | Type | Contents |
|--------------|----------|------------------------------------|
| 1 | DOS | Bootstraps |
| 2 | Solaris | Solaris VTOC; slice 0 is miniroot |
| (Whole disk) | ISO-9660 | Bootstraps, Solaris /usr directory |

The CD layout and contents are similar to those on x86 except for the addition of the DOS partition. As currently specified in the PowerPC Supplement to IEEE 1275, OF first looks for a partition table on the boot device, and only if it fails to find one, boots a selected file from the ISO-9660 file system. Because Solaris install requires the disk to have both the ISO-9660 and the Solaris partitions, the CD must contain a partition table, and OF (unless the Supplement is amended) will boot from one of the partitions, not from the ISO-9660 file system. Therefore, the Solaris CD includes a DOS partition containing OF bootstraps. Systems that first check for the presence of an ISO-9660 directory will find copies of the bootstraps in the ISO-9660 partition of the Solaris CD. If the Supplement is changed to check first for the ISO-9660 file system, there will be no need for the DOS file system on the CD.

Diskette

CF can boot only from a 0x41 partition, so even a diskette must contain a partition table. However, neither Solaris nor DOS supports partitioned diskettes. Solaris on PowerPC works around this difficulty by placing a partition table in the first block of the diskette. The table specifies a 0x41 partition for use by CF. The “whole disk” is a DOS file system, as shown in the table below.

Table 2 Diskette Partitions

| Partition | Type | Contents |
|--------------|------|---------------------------------------|
| 1 | 0x41 | VOF |
| (Whole disk) | DOS | Bootstraps, drivers, Solaris packages |

The 0x41 partition is created only so that CF can find VOF. So far as all other software is concerned, including DOS and Solaris, the diskette can be treated as an ordinary DOS diskette.⁶



It appears that the current *fdformat* program with the **-B** flag specifying the VOF module will create the diskette, so long as the VOF module contains in its first 512 bytes the appropriate DOS diskette file system information (media type, etc.) along with the partition table. This has two advantages:

1. The file cannot be deleted and replaced with a noncontiguous file, which would not boot because the 0x41 partition must be contiguous, and
2. The *fdformat* program already supports it, whereas otherwise a program to create a contiguous file on the DOS partition and the partition table would be required.

It has the disadvantage that it requires a separate utility to access the VOF module to install it on the hard drive. A final decision has not yet been made on this issue.

Hard Drive

A hard drive contains a partition table specifying three partitions, as shown in the table below. The Solaris partition contains the standard VTOC specifying configurable

Table 3 Hard Drive Partitions

| Partition | Type | Contents |
|-----------|---------|-------------------|
| 1 | DOS | Bootstraps |
| 2 | Solaris | Solaris partition |
| 3 | 0x41 | VOF |

slices. As for diskettes, the 0x41 partition is used only by CF for loading VOF. The 0x41 partition should lie at the “far end” of the Solaris partition so that, if and when CF is no longer supported, the space occupied by the 0x41 partition can be more easily incorporated into the Solaris partition.

6. In the unlikely event that the firmware on any CF-based platform checks for and rejects partitioned diskettes, Solaris for that platform must use a separate diskette containing only a 0x41 partition for booting, with other files on another diskette.



Boot Sequences

Booting from Hard Drive

Unless specified otherwise in an explicit boot command given through the OF (or VOF) user interface, the partition and file to be booted by OF is specified in OF's `boot-device` configuration variable. For Solaris, this will normally be the DOS partition and a file containing the `ufsboot` secondary bootstrap⁷. The `boot-file` configuration variable specifies the file which the secondary bootstrap is to load.

Under CF, the firmware boots the VOF module from the 0x41 partition, and the remainder of the booting proceeds as above. Note that VOF supports the OF `setenv` user command for setting configuration variables.

Booting from Net

OF supports booting from a network using the `bootp` and `tftp` protocols. The firmware uses `bootp` to determine the host's network address and the network server and boot program for the host, then uses `tftp` to load the boot program (normally `inetboot` for Solaris).

The CF specification does not support network booting, so there is no defined mechanism for booting from a network using CF. However, one can boot VOF from a floppy, hard drive, or CD (or, using platform-specific mechanisms outside of the PReP specification, from the net). The reference VOF implementation provides full support for network booting so long as a suitable VOF network driver is available.

Booting from CD

Under the search path in the current PowerPC Supplement, OF will find a partition table in the first sector of the Solaris CD. Therefore, OF will boot a selected file from one of the partitions. Presumably (the specification is not clear in this regard) this must be a (the) DOS file system, as there is no mechanism for specifying an ISO-9660 partition in the DOS partition table.

7. SunSoft's ON consolidation may decide to introduce a level of indirection in booting, placing a "bootblk" program that loads the "real" bootstrap into the DOS partition. If so, the indirection will be transparent to the rest of Solaris and the bootstraps themselves.



SunSoft has proposed that the OF Supplement for PowerPC be changed to specify that OF look first for an ISO-9660 filesystem, and not a partition table. This would make it possible for to eliminate the DOS partition and to have the secondary bootstraps (the files loaded by OF) reside only in the ISO-9660 directory. If this proposal is not adopted, or is not adopted in time to affect all OF implementations, both the ISO-9660 and DOS file systems on the Solaris CD will contain the Solaris secondary boot programs (currently `ufsboot` for booting from the miniroot in the Solaris partition and `inetboot` for booting from a network). If OF checks first for the presence of an ISO-9660 file system, it will find the boot program in the ISO-9660 file system; if it uses the partition table, it will find the boot program in the DOS partition.

CF will always boot VOF from the 0x41 partition, and VOF, once booted, will then load the Solaris secondary boot from either the DOS or ISO-9660 filesystem.

Installation

Follows x86 Model

Solaris must be installable on PowerPC systems some or all of whose boot devices are not supported by the Solaris CD. In this respect, the PowerPC is similar to x86, and the same installation mechanisms apply. In particular, the following interfaces specified in PSARC case 1993/674, *3rd-Party Driver Boot/Install*, shall be used on PowerPC:

- The `SOLARIS.MAP` file and corresponding changes in `ufsboot` that allows files on the boot diskette to overlay or be logically appended to files on the CD, and
- The installation shell scripts executed by installation software at various points in the installation process.

These mechanisms are simple and general and have proved sufficient to meet the identical needs on x86. As described in the PSARC case materials, the ON consolidation will provide the `SOLARIS.MAP` file; the Admin consolidation (RMTC) will provide the installation shell scripts and assure that they are executed as on x86. The ON consolidation will also assure that the `ufsboot` changes for supporting the `SOLARIS.MAP` file are integrated into the PowerPC version. The platform consolidation (SunSoft South) is the “consumer” of these interfaces and will probably provide at least some of the resources for the ON deliverables.



Deviations from x86

Use of BOOTP Protocol

Earlier SPARC and x86 systems used the RARP (Reverse Address Resolution Protocol) for determining the various servers involved in network booting. PowerPC uses the BOOTP protocol. This will call for some changes in the `add_install_client`, `setup_install_server`, and related server scripts, as well as distribution of BOOTP daemons for SPARC, x86, and PowerPC.

Installation of VOF

An additional task that must be accomplished during installation on CF systems is the installation of VOF on the hard drive. The `installboot` program will be modified to do this on PowerPC.

Installation without Boot Diskette

OF supports booting directly from CD. OF systems with FCODE for their network adapters and possibly some CF systems also support booting directly from a network. Therefore, in contrast to x86, OF-based platforms supported by the Solaris CD can be installed from CD without a boot diskette, and some platforms can be installed via network without a boot diskette. The installation procedures must determine whether or not there is a diskette before applying operations that depend on it.