
# Operating Environments: Solaris™ 8 Operating Environment Installation and Boot Disk Layout

---

*By Richard Elling - Enterprise Engineering*

*Sun BluePrints™ OnLine - March 2000*

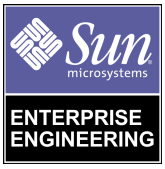


**http://www.sun.com/blueprints**

**Sun Microsystems, Inc.**
901 San Antonio Road
Palo Alto, CA 94303 USA
650 960-1300 fax 650 969-9131

Part No.: 806-4978-10
Revision 01, March 2000

# Operating Environments: Solaris™ 8 Operating Environment Installation and Boot Disk Layout

The Solaris™ Operating Environment Version 8 does not fit on a single compact disk (CD). A new Solaris™ Web Start installation procedure is available which uses Java™ technology-based graphical wizards to configure and install the operating environment. The Solaris Web Start installation procedure enables changing CDs during the installation process. Previous installation procedures using JumpStart™ software and Interactive Installation are available and have been updated to accommodate the multiple CDs required for a full Solaris 8 Operating Environment installation.

Solaris Web Start imposes some changes on the boot disk layout for standalone systems being installed from CD. This article discusses these changes and recommends boot disk layouts for desktop and workgroup server systems.

## Installation Procedures

The Solaris 8 Operating Environment can be installed by three different procedures: Solaris Web Start, JumpStart software, and Interactive Installation. Each of these procedures may be used with local media (CD) or from an install server over the network.

## Solaris™ Web Start

Solaris Web Start uses a Java technology-based graphical user interface (GUI) that guides the user through the installation tasks. Wizards guide the installer through the installation process. If the system does not have a keyboard, mouse, and graphics display then a command line interface is available. The command line interface offers the same configuration options as the GUI but is not nearly as interesting or user friendly.

Many Solaris Operating Environment software products are now delivered with the Solaris Web Start style wizards for installation, in addition to the traditional Solaris Operating Environment packages.

The disadvantage of using Solaris Web Start is that it takes longer to perform installations. The Solaris Web Start boot image is first loaded onto a disk slice. The system is rebooted using the Solaris Web Start boot image. This allows the Solaris Operating Environment Installation CDs to be changed during the installation. The implications of loading the Solaris Web Start boot image onto a disk slice are discussed below.

Solaris Web Start is recommended as the procedure for initial Solaris 8 Operating Environment installation by both novice and experienced Solaris Operating Environment systems administrators. Experienced Solaris Operating Environment system administrators may later choose to implement custom JumpStart software or Interactive Installation for their production systems.

## JumpStart™ and Custom JumpStart Software

*JumpStart* software provides a way to install groups of systems automatically and identically. A set of rules is used to determine the hardware and software configuration for the installation. Configuration parameters such as disk slice allocations are specified by a default profile that is selected based on the model of disk drive on the system. In a JumpStart software installation the installer is not given a choice of which software is to be installed.

*Custom JumpStart* software is a JumpStart software installation with a custom rule file. These rules may be accessed locally on a floppy diskette or remotely from a JumpStart server. The custom rule file allows detailed changes to be made in the installation system configuration, disk slice allocations, and software to be installed.

Custom JumpStart software is the most efficient method for installing systems in an enterprise. Custom JumpStart works especially well when unattended, centrally managed, and configuration controlled installation is desired. The time and cost saving gained by using custom JumpStart software are well worth the investment of building the custom JumpStart server and rules files.

## Interactive Installation

Interactive Installation is the method most familiar to Solaris Operating Environment systems administrators. The interactive installation remains virtually unchanged for the Solaris 8 Operating Environment release. The changes involve support for new Solaris 8 Operating Environment features such as DNS, DHCP, and IPv6 client support. Both command line and Motif GUI are available.

# System Types: Server, Client, and Standalone

A *server* is a system that provides services or file systems, such as home directories or mail files, for other networked systems.

An *install client* is a system which gets its operating system installation image from a server.

An *install server* is a server that provides the Solaris Operating Environment software for installation on install clients.

A *boot server* is a system used to boot an install client to be installed over the network.

A *JumpStart server* is a system providing the rules file which contains the hardware and software configuration for the install client. The rules file is a text file that contains a rule for each system or group of systems for installing the Solaris Operating Environment. Each rule distinguishes a group of systems based on one or more attributes and links each group to a profile, which is a text file that defines how the Solaris Operating Environment software is to be installed.

A boot, install, and JumpStart server are often the same system. However, if the system on which the Solaris 8 Operating Environment is to be installed is located in a different subnet than the install server, then a boot server is required on the install client's subnet.

A single boot server can provide Solaris 8 Operating Environment boot software for multiple releases, including the Solaris 8 Operating Environment boot software for different platforms. For example, a SPARCserver™ boot server could provide the Solaris 7 and 8 Operating Environments boot software for SPARCserver based systems. The same SPARCserver boot server could also provide the Solaris 8 Operating Environment boot software for Intel Architecture (IA) based systems.

A *standalone system* stores the Solaris Operating Environment software on its local disk and does not require installation services from an install server. Typically a standalone system will load the Solaris Operating Environment software from a locally attached CD.

# Web Start Boot Disk Slice Layout

The Solaris 7 Operating Environment default boot disk layout is representative of previous versions of the Solaris Operating Environment. The Solaris 8 Operating Environment, Solaris Web Start default boot disk layout is significantly different. The Solaris 8 Operating Environment standalone system installation procedure requires a slice to hold the miniroot. Typically this slice would be also used as the swap device and thus represents a suitable temporary holding place. This is reminiscent of the SunOS™ software 4.x installation procedure.

The problem with the SunOS software 4.x installation procedure is that a system administrator had to make a guess at the desired size of the root slice, /. This is because the swap space which held the miniroot was physically located at the end of the root slice. Not only did this complicate the installation procedure, but often led to poor decisions on the root slice size. It was common to have the root slice be too small since typical disk sizes were small, often less than 500MB. The small root slice would occasionally need to be adjusted to make room for additional software or SunOS software upgrades. The typical procedure to increase the size of the SunOS software 4.x root slice was to boot into the miniroot, resize the root slice, build a new filesystem, and either reinstall SunOS or recover from backup tapes.

The Solaris 8 Operating Environment standalone installation reintroduces the miniroot loaded into swap space. However, the location of the swap slice is physically moved to the beginning of the disk. The slice numbers remain the same, but the physical location of the swap slice 1 has been switched with the root slice 0. Table 1 shows the slice layouts for Solaris 7 and 8 Operating Environments

**Table 1: Default Boot Disk Slice Layouts**

| Solaris 7 Layout | | | | Solaris 8 Layout | | | |
|---|---|---|---|---|---|---|---|
| Slice | Blocks | | Use | Slice | Blocks | | Use |
| | Begin | Blocks | | | Begin | End | |
| 0 | 0 | n1 | / | 1 | 0 | m1 | swap |
| 1 | n1 + 1 | n2 | swap | 0 | m1 + 1 | m2 | / |

**Table 1: Default Boot Disk Slice Layouts**

| Solaris 7 Layout | | | | Solaris 8 Layout | | | |
|---|---|---|---|---|---|---|---|
| Slice | Blocks | | Use | Slice | Blocks | | Use |
| | Begin | Blocks | | | Begin | End | |
| 2 | 0 | end of disk | backup | 2 | 0 | end of disk | backup |
| 6 | n2 + 1 | n3 | /usr | | | | |
| 7 | n3 + 1 | end of disk | /export/ home | 7 | m2 + 1 | end of disk | /export/ home |

Figure 1 is a graphical representation of the disk layouts for Solaris 7 and 8 Operating Environments.

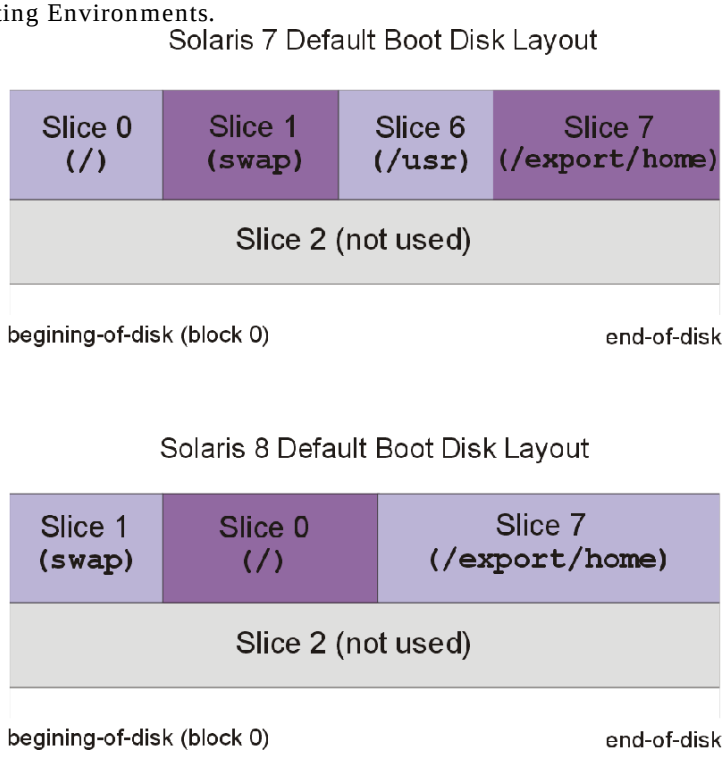


**FIGURE 1** Graphical View of Default Boot Disk Layout

The problem with resizing the root file system during system installation is now mitigated since the resizing procedure will not require the swap space containing the miniroot image to be moved.

## Reserving Space for Logical Volume Managers

Boot disks can be managed by Logical Volume Managers (LVM) such as Solstice DiskSuite™ (SDS) software and Veritas Volume Manager (VxVM). The Sun BluePrints™ book titled *Guide to High Availability: Configuring boot/root/swap*, published in July 1999 by Prentice Hall, ISBN#0-13-016306-6, goes into detail on how to encapsulate and mirror boot disks. The book is available through http://www.sun.com/books, amazon.com, fatbrain.com, or Barnes & Noble bookstores. It is a good practice to reserve a few megabytes of disk space and a slice for use by a LVM. For most servers using Solstice DiskSuite software or Veritas VxVM, a reserved slice of 2 MB should suffice.

## Swap Device Recommendations

32 bit Solaris Operating Environment systems are limited to using only the first 2GB ($2^{31}$-1 bytes) of the swap device. 64 bit Solaris Operating Environment systems are virtually unlimited and support the first $2^{63}$-1 bytes. $2^{63}$ bytes or more than 9223 Petabytes is a number much larger than any contemporary storage device.

The total amount of virtual memory is the sum of the physical memory (RAM) plus the sum of the size of the swap devices. The minimum virtual memory size is 32 MB. Systems with only 32 MB of RAM are almost impossible to purchase new as most systems have 64 MB or more.

It is common for 32 bit Solaris Operating Environment systems to use multiple swap devices. Swap devices can be added and deleted on the fly using the swap command. The kernel will write to the swap devices in a round-robin manner changing swap devices for every 1 MB written. This is similar in concept to RAID-0 stripe with an interlace of 1 MB. This allows the swap load to be balanced across multiple swap devices. However, the kernel will not actually write to the swap devices until physical memory is full. Thus the total swap device size is not typically required to be large.

The performance implications of multiple swap devices are more difficult to ascertain. The access time of a page on a swap device will be approximately four orders of magnitude greater than the access time of a page of memory. If a system is actively using the swap devices, then performance tends to suffer. The physical placement of the active swap devices may also impact the performance. However, a bad case of head contention on a modern disk drive will lead to at worst, a single order of magnitude difference in access time. This will be dwarfed by the four orders of magnitude cost of actively swapping. Thus it is reasonable to use multiple swap devices on a single physical disk, especially for 32 bit systems. If the system is continually actively using the swap devices, then adding RAM will help significantly improve performance, whereas adding or relocating swap devices will be much less effective.

For standalone systems being installed via Solaris Web Start the default swap device size is 512 MB. This swap device must accommodate the miniroot which is required to be at least 422 MB.

Interactive installation allows the swap device size to be set to any value. A recommended value is given based on the size of the system boot disk, but there is no requirement for the swap device size. Interactive installation allows multiple swap devices to be created.

JumpStart software will base the default swap space size on the amount of physical memory in the install client system. Custom JumpStart can be used to override these defaults. JumpStart software makes the size of swap no more than 20 percent of the disk where it is located, unless there is free space left on the disk after laying out the other file systems. If free space exists, JumpStart software allocates the free space to swap, and if possible, allocates the amount shown in Table .

**Table 2: Default JumpStart Swap Device Size**

| Physical Memory (MB) | JumpStart Default Swap Device Size (MB) |
|---:|---:|
| 16 - 64 | 32 |
| 64 - 128 | 64 |
| 128 - 512 | 128 |
| greater than 512 | 256 |

Additional swap device space may be required based on application needs. The swap command can be used to add swap devices on a system without causing an outage. The additional swap devices can be built as files in a file system. This flexibility allows the final decision on the swap device size to be deferred until demand dictates a change.

# The Backup Slice

Slice 2 has historically been specified as the entire disk. This slice is shown with a tag of "backup" (numerically 5) in the output of prtvtoc as shown in Listing 1. This slice

```
# prtvtoc -s /dev/dsk/c0t0d0s0
*                          First     Sector    Last
* Partition  Tag  Flags    Sector     Count    Sector  Mount
Directory
       0      2    00    1049328  15790319  16839647   /
       1      3    01          0   1049328   1049327
       2      5    00          0  16839648  16839647
```

is not normally used by the Solaris Operating Environment. However, there may be other utilities and systems management products which expect the backup slice to represent the entire disk. It is recommended to leave the configuration of the backup slice as is.

# One File System or Two?

Over the years there has been debate about whether to use one large file system for the entire Solaris Operating Environment installation or, multiple smaller file systems. Given modern hardware technology and software enhancements to the UNIX® Fast File system, the case for multiple file systems seems anachronistic

.

**Table 3: Multiple File System Issues**

| Issue | Advantages | Disadvantages |
|---|---|---|
| Backup and restore. | Can set policies on backup for each file system independently, eg. `/export/home` is backed up every night, `/usr` is backed up once per week. | Backup procedure is complex and requires multiple passes.<br><br>Restoration requires multiple passes. |
| File system check, fsck, time. | `fsck` is run concurrently on non-root file systems.<br><br>`fsck` on read-only and inactive file systems is quick.<br><br>UFS logging significantly improves `fsck` time but is not enabled by default. | Concurrent `fsck` on file systems using a single disk causes disk head contention. |
| File system size. | Solaris 8 Operating Environment has a 1 TB file system size limit for UFS. | SunOS 4 has a 2 GB file system size limit for UFS which is much less than the size of a modern disk drive. |
| Swap device size. | 32-bit systems have a 2GB swap device size limit.<br><br>64-bit systems have a swap device size limit of $2^{63}$-1 bytes. | Multiple swap devices consume slices. Only 8 slices are available per disk. |
| SunOS default core dump device. | Solaris 7-8 Operating Environments have the dumpadm command which allows customization of the dump process and core dump destination. | Previous releases of Solaris Operating Environment placed OS core dumps in the first swap device defined in the `/etc/vfstab file`. |

# Recommended Solaris 8 Operating Environment Boot Disk Layouts

Recommended Solaris 8 Operating Environment boot disk layouts are described in Table 4. The possible configurations for servers can be more complicated and use

**Table 4: Recommended Boot Disk Slice Layouts**

| Standalone Desktop Layout | | | | Server Layout | | | |
|---|---|---|---|---|---|---|---|
| Slice | Blocks | | Use | Slice | Blocks | | Use |
| | Begin | Blocks | | | Begin | End | |
| 1 | 0 | n1 | swap | 1 | 0 | m1 | swap |
| 0 | n1 + 1 | n2 | / | 0 | m1 + 1 | m2 | / |
| 2 | 0 | end of disk | backup | 2 | 0 | end of disk | backup |
| | | | | 4 | m3 + 1 | end of disk | reserve for LVM |
| | | | | 7 | m2 + 1 | m3 | /export |

slices for other, exported file systems.

An example of a disk layout for a standalone desktop with a single 9GB disk is shown in Table 5.

**Table 5: Standalone Desktop Example**

| Slice | Blocks | | Size (MB) | Use |
|---|---|---|---|---|
| | Begin | End | | |
| 1 | 0 | 1049327 | 537 | swap |
| 0 | 1049328 | 16839647 | 8084 | / |
| 2 | 0 | 16839647 | 8622 | backup |

An example of a disk layout for a workgroup server with a mirrored 9GB boot disk is shown in Table 6. A LVM is to be used for mirroring the root disk. Note that the

**Table 6: Workgroup Server Example**

| Slice | Blocks | | Size (MB) | Use |
|---|---|---|---|---|
| | Begin | End | | |
| 1 | 0 | 1049327 | 537 | swap |
| 0 | 1049328 | 4121712 | 1573 | / |
| 2 | 0 | 16839647 | 8622 | backup |
| 4 | 16834607 | 16839647 | 2 | LVM |
| 7 | 4121712 | 16834606 | 6509 | /export |

root and `/export` file systems are contiguous. This allows resizing of the file systems later without forcing a reinstallation of the operating environment or LVM.

# Conclusion

The Solaris 8 Operating Environment requires multiple CDs for installation. A new Java technology-based installation procedure, Solaris Web Start, simplifies installation but has a different boot disk layout than JumpStart software or interactive installation. These changes were discussed and a recommended boot disk layout for desktop and small workgroup servers given.

# References

Solaris Operating Environment installation manuals.

*Solaris 8 Operating Environment Advanced Installation Guide.*

Manual page on *swap(1m), prtvtoc(1m).*

Sun BluePrints *Guide to High Availability: Configuring boot/root/swap* by Jeannie Johnstone Kobert, Prentice Hall, July 1999, ISBN#0-13-016306-6. The book is available through http://www.sun.com/books, amazon.com, fatbrain.com, or Barnes & Noble bookstores.

### *Author's Bio: Richard Elling*

*Richard Elling is a Senior Staff Engineer in the Enterprise Engineering group at Sun Microsystems in San Diego, California. Prior to the Enterprise Engineering group, Richard had been a Field Systems Engineer at Sun for five years. Richard was the Sun Worldwide Field Systems Engineer of the year in 1996. Prior to Sun, Richard was the Manager of Network Support for the College of Engineering at Auburn University. Richard was also a design engineer for a startup microelectronics company, and he worked for NASA doing electronic design and experiments integration for Space Shuttle missions.*