



Solaris™ OS Guide for New System Administrators

Rolf Kersten

Sun Microsystems GmbH, Germany

Revised October 30, 2003

Copyright 2003 Sun Microsystems, Inc. All Rights Reserved.

Sun, Sun Microsystems, the Sun logo, Ultra, Sun Blade, Sun OpenGL for Solaris, Solaris, Solaris Volume Manager, Solstice DiskSuite, PGX32, XIL, XGL, Sun Enterprise, OpenWindows, and SunSolve are trademarks or registered trademarks of Sun Microsystems, Inc. in the United States and other countries.

All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. in the United States and other countries. Products bearing SPARC trademarks are based upon an architecture developed by Sun Microsystems, Inc.

OpenGL is a registered trademark of Silicon Graphics, Inc.

UNIX is a registered trademark in the U.S. and other countries, exclusively licensed through X/Open Company, Ltd.

Change History

Date	Version	Comment
25.05.2002	V 1.8e	First English release
12.06.2002	V 1.9e	First feedback (thanks to Daniel Schmidberger!)
19.07.2002	V 1.10e	Added iostat chapter
04.09.2002	V 1.11e	Edits from Matt Ruetz
05.12.2002	V 1.12e	Miscellaneous corrections and additions
17.01.2003	V 1.13e	Additions and corrections from Halvard Halvorsen (www.sundot.com) and Sven Fink
11.04.2003	V 1.14e	Shells and Environment Variables chapter, more from Sven and Andreas Fatum
16.04.2003	V 1.15e	xcpustate, OGL&DGA
27.05.2003	V 1.16e	Handling CD-Rs written on other platforms
02.07.2003	V 1.17e	Mounting USB Memory Sticks
30.10.2003	V 1.18e	Lots of editorial changes, /var/sadm/install/contents file, Cloning systems with FC-AL boot disks

Table of contents

1 System Booting and Testing.....	2
1.1 The OK-Prompt.....	2
1.2 Testing at a Console.....	2
1.3 Rebooting a Sun Workstation.....	2
1.4 Testing the System.....	3
1.4.1 Banner.....	3
1.4.2 probe.....	3
1.4.3 test all.....	3
1.4.4 printenv and devalias.....	4
1.4.5 Disabling the Network-Interface Error Message.....	4
2 Administering Disks.....	5
2.1 Logical Device Names.....	5
2.2 Partitions.....	5
2.3 Crash Dumps.....	6
2.4 Example.....	7
2.5 Removable Devices.....	8
2.5.1 Volume Management.....	8
2.5.2 Floppy.....	8
2.5.3 CD-ROM.....	8
2.5.4 CD-R and CD-RW.....	9
2.5.5 USB Memory Sticks.....	9
2.5.6 Tapes.....	9
2.6 Cloning a Disk.....	10
2.6.1 Cloning Using Internal Disks.....	10
2.6.2 Cloning Using a Crossover Cable.....	11
2.7 Mirroring the Boot Disk with SDS.....	11
3 Swap Space.....	14
3.1 Persistent Configuration in /etc/vfstab.....	14
3.2 "Raw Partition" Versus Plain File.....	14
3.3 Temporary Configuration.....	14
4 Installation and Maintenance.....	15
4.1 Installation From CD.....	15
4.2 Packages.....	15
4.3 Patches.....	16
5 X-Server and CDE User Interface.....	17
5.1 General Observations.....	17
5.2 Customizing the Login Screen.....	17
5.3 Startup with Multiple Screens, Xinerama, 24-Bit Color Depth.....	17
5.3.1 Multi-Screen.....	17
5.3.2 Xinerama.....	18
5.3.3 24-Bit Color Depth.....	18
5.4 A User-Defined CDE Front Panel.....	18
5.4.1 Actions.....	18
5.4.2 Icons.....	19
5.4.3 Front Panels.....	19
6 Frame Buffers and OpenGL.....	20
6.1 fbconfig.....	20
6.2 OpenGL.....	21
6.3 Patch Matrix for Graphics.....	22
7 Administration.....	23
7.1 Hosts.....	23
7.2 Serial Ports.....	23
7.3 Printers.....	23
7.4 User Creation.....	23
7.5 Shells and Environment Variables.....	24

8	Network Connections.....	26
8.1	/etc-files.....	26
8.2	Virtual Interfaces.....	27
8.3	ifconfig.....	27
8.4	Routing.....	27
8.5	NIS-Client.....	28
8.6	DNS-Client.....	28
8.7	DHCP-Client.....	29
8.8	sys-unconfig.....	29
8.9	10 or 100 Mbit, Half- or Full-Duplex?.....	29
8.10	Network Settings: Temporary vs. Persistent	30
9	Network File System.....	31
9.1	NFS-Server.....	31
9.2	NFS-Client.....	31
9.3	Differences From Other Popular UNIX Systems.....	31
9.4	Automounter.....	32
10	Tools.....	33
10.1	Compiler.....	33
10.2	GNU-Tools.....	33
10.3	SunPCi™ and SunPCi-II.....	33
10.4	StarOffice™ Software.....	33
10.5	Netscape Navigator™ Software	34
10.6	Miscellaneous.....	34
11	Hint Collection.....	35
11.1	Top Three Hints.....	35
11.2	Analysis of the Hardware Configuration.....	35
11.2.1	prtdiag.....	35
11.2.2	prtconf / path_to_inst.....	35
11.2.3	Sun Explorer Software.....	35
11.3	Analysis of Runtime Data.....	36
11.3.1	time/ptime.....	36
11.3.2	top / prstat / ps.....	36
11.3.3	vmstat / mpstat.....	36
11.3.4	iostat.....	36
11.3.5	sdtpfometer, xpcpustate, and sdtprocess.....	36
11.3.6	32 or 64 Bit?.....	37
11.4	Debugging.....	37
11.4.1	Truss.....	37
11.4.2	pstack, pmap, and So On.....	37
11.4.3	snoop.....	37
11.5	Common Pitfalls.....	37
12	Books and Web Sites.....	38
12.1	Books.....	38
12.2	Web Sites.....	38

Introduction

Within this document you can find the most relevant 20 percent of technical information you will need to know when working with Solaris™ and Sun™ platforms. It should answer 80 percent of your questions about the Solaris Operating System (OS).

This document is not intended to replace detailed books or comprehensive FAQs written about the Solaris OS (books and links are listed in the Appendix). Rather, this document is intended to serve as a "how to" collection for those who are new to Solaris OS system administration or those who are not administering Solaris systems on a daily basis. This content focuses on graphical workstations used for technical applications and should also be helpful to those who need to set up Sun systems for demonstrations or benchmarks.

Two special notes:

- Always keep track of changes made to your system by using `achangelog`.
- Keep in mind that this document contains system-level procedures that can provoke system unavailability. Do not apply these procedures to productive systems unless you have made a careful analysis of the consequences.

1 System Booting and Testing

1.1 The OK-Prompt

Each Sun system has a Boot-PROM mode that can be accessed prior to starting the Solaris OS or while the OS is running. You may enter this mode by pressing STOP-A (Hold down the STOP-key while pressing "a"). If the STOP-key is not available (PC keyboard), try `Ctrl-BREAK`.

Once you are in the Boot-PROM mode, an OK-Prompt is displayed. If the Solaris OS has been running, it is frozen now, but can be "reanimated" immediately by typing the command "go" at the OK-Prompt.

1.2 Testing at a Console

The Boot-PROM mode can be used either by means of a directly connected keyboard and monitor or by using a separately connected terminal.

If the system does not detect a keyboard while booting, input and output are automatically redirected to the first serial interface (`/dev/ttya`, often named "A"). You may attach a terminal or another computer (using a crossover serial cable) to this port.

If this other computer happens to be a Sun workstation, you can use `tip hardware` to monitor the boot console of the observed machine. After issuing the `tip` command, all boot messages will appear in the corresponding terminal window. To issue a STOP-A command on the observed host, you have to change `tip` to `telnet-mode` first with `Strg-]`. After that, enter a BREAK: `telnet> send brk`

1.3 Rebooting a Sun Workstation

Prior to rebooting, a Sun workstation should always be shut down in an appropriate manner. For this task, several alternative commands exist:

The best and cleanest methods are:

- # `init 5` Shut down all processes, switch off system afterwards
- # `init 6` Shut down all processes, reboot system afterwards
- # `init 0` Shut down all processes, display OK prompt afterwards

In all three cases, the `K*` scripts in `/etc/rc0.d` are executed before the system shuts down or reboots. This is important for databases that have to be shut down, for web servers that are cleaning process ID files, and so on.

System Booting and Testing

If there is no shell window available, try the preceding commands using a remote login. If that is not possible, the next step should be:

- STOP-A
- OK sync Sync file systems. Panic the OS afterwards, then reboot.

Attention: Watch out for big OS core dumps in `/var/crash/<hostname>` afterwards.

If STOP-A is not possible anymore, the computer has to be switched off and on.

1.4 Testing the System

Boot the system. Press STOP-A while the memory is being tested.

The following paragraphs will describe the most important commands only.

A detailed description can be found in "Hardware Diagnostics for Sun[™] Systems: A Toolkit for System Administrators" (SunSolve Infodoc Number 23476 at <http://sunsolve.sun.com/pub-cgi/show.pl?target=content/content9>).

1.4.1 Banner

`banner` displays the following pieces of information:

- number, type, and frequency of CPUs
- memory
- MAC-Address and Host-ID
- Type of Console-Frame buffer
 - Sun-Globe: Creator 3D
 - Sun with face: Elite 3D
 - Self-describing Logo: PGX, Raptor GFX (=PGX32), Expert3D(lite) XVR-500, XVR-1000, and so on.

1.4.2 probe

`probe-scsi-all` lists all internal and external SCSI-attached devices. For IDE disks, use:
`probe-ide-all`.

1.4.3 test all

`test all` runs all self-test procedures and can be used to find out details such as whether a network cable is working properly.

1.4.4 `printenv` and `devalias`

`devalias` lists all aliases, while `printenv` lists all variables. The variable `boot-device`, for example, defines from which path a machine shall boot. In most cases the value of this variable is `disk`. `disk` itself is an alias pointing to a particular device name.

Hint: To reset all variables to their default values, use `set-defaults` at the OK-prompt or press `STOP-N` while powering-on the system. Warning: By setting all variables to their default values, you could lose your boot device definition, and the system may not boot afterwards.

1.4.5 Disabling the Network-Interface Error Message

When network interfaces are present with no cable attached during boot time, "Interface: No link" error messages are printed multiple times for each interface. These can be disabled by setting `setenv tpe-link-test? false` at the OK prompt.

Note: This fix speeds up boot times when network cables are detached temporarily. For permanently unconnected interfaces, shut down the interface with `ifconfig unplu Mbyte <interface>` (using a root shell).

2 Administering Disks

2.1 Logical Device Names

All disks are numbered following a common naming scheme. This system consists of four parts:

- `c` Controller number
- `t` Target-ID (for SCSI-devices this number corresponds to the SCSI-ID)
- `d` Drive Number or LUN for RAID-devices (one ID, multiple disks)
- `s` Slice, Partition number (see below)

These logical device names are stored as links in the directories `/dev/dsk` and `/dev/rdsk`. These links point to physical device names, but they are only needed for tasks like changing the boot device. Typical Solaris OS commands for administering drives expect logical device names as their parameters.

Example: `c0t0d0s0` typically is the first drive in a workstation, `c0t1d0s0` is the second one, `c0t6d0s0` is the internal CD-ROM-drive (often with SCSI-ID 6), and so on.

A refresh of the `/dev/dsk` and `/dev/rdsk` directories occurs either while booting (`boot -r` at the OK-Prompt) or during Solaris OS operation with `devfsadm (drvconfig;disks;devlinks` prior to the Solaris 8 OS). After this refresh, logical links should exist for all devices found upon boot time with `probe-scsi-all`.

2.2 Partitions

On a workstation with two internal disk drives, the command `format` generates the following output:

```
c0t0d0
c0t1d0
```

This listing displays disks, not single slices or partitions. The purpose of `format` is to create and administer these slices (or partitions). A disk may have up to eight slices. These may be displayed using the `format` sub-command `print - partition`.

If a default installation was used, the Solaris installer tool divides the OS disk into several slices: `/` (root), `/usr`, `/var` and `/export/home`. A Solaris 8 OS default installation creates three slices:

(i) approximately 1 Gbyte root, (ii) swap (depending on memory capacity) and (iii) the remaining disk space for `/export/home`.

For workstations, this partitioning scheme is not optimal. The capacity of a partition cannot be expanded without data loss. In the case of a partition running out of space, symbolic links to other partitions have to be created. It is useful, especially for servers, to reserve one partition for the `/var` directory, so that growing log files or big files stored in `/var/tmp` by users cannot fill up the root partition, thus causing OS problems.

A workstation OS disk should be partitioned this way:

```
slice 0: /          root-Partition, incl. /usr, /var, /opt plus
                /export/home
slice 1: swap      swap-Partition, as large as physical memory
slice 2: backup   Do not change this partition!
                It is always as big as the whole disk
slice 7: sds      A small space reserved for Solstice DiskSuite(TM)
                metadata
```

To be prepared for the use of the Solstice DiskSuite logical volume manager (SDS), it is recommended to pre-serve some space (30 Mbyte, approximately 60 cylinders) for copies of the "state replica databases." (More on that later. See also the chapter on Mirroring the Boot Disk with SDS.)

Starting with the Solaris 8 02/02 OS, Live Upgrade automatically will be installed. Using Live Upgrade, you can update a copy of the current root environment while the system is still running. To use this feature (for example, for testing a new OS version without changing the currently running system), there has to be enough space for creating an alternate boot environment (ABE). Modern disks typically allow sufficient capacity for this partitioning:

```
slice 0: /          root-partition, incl. /usr, /var, /opt
slice 1: swap      swap-Partition, as large as physical memory
slice 2: backup   Do not change this partition!
                It is always the size of the whole disk
slice 3: /export  OS-independent data, home-directories
slice 4: <not mounted>, to be used as ABE with Live Upgrade
                ( about 150% of the expected size of / )
slice 7: sds      A small space reserved for Solstice DiskSuite metadata
```

The important partitions in this example ("/" and "/export") should be mounted "logging." This avoids lengthy and manual file system checks during reboot after a machine was not shut down properly. This option is not enabled by default. After installation, you need to replace "-" with "logging" in the last column of "/etc/vfstab" manually.

2.3 Crash Dumps

Starting with the Solaris 7 OS, crash dumps are written to the swap partition by default. This is another reason why swap should have the size of the physical memory.

During the next reboot, this file is read back from swap, compressed, and written to /var/crash/<host-name>, to allow for later analysis. Therefore, /var (or /, if no special partition for /var was defined) should also not be too small.

If a system with lots of gigabytes uses the swap partition as a dump device, the boot process can take some time because of the copying procedure mentioned above. In these cases, it is best to define a dedicated dump partition using dumpadm. Copying from a dedicated dump device (rather than from swap) to /var/crash is performed in the background.

Hint: Check and clean up the content of `/var/crash/<hostname>` frequently, so you have enough space in this directory for upcoming crash dumps.

2.4 Example

Suppose you have a Sun Enterprise™ 450 system and a bootable disk with a fully configured demo installed on it, but you do not know two things:

1. Where to put the disk (the system has 20 slots, and the system only boots correctly when the disk is placed in the same slot as it has been at installation time)
2. The root password

Start with putting the disk into an arbitrary slot, then follow this procedure:

Step 1: Insert a bootable CD (a Solaris Installation CD) and reboot the system. Interrupt with `STOP-A`, type `boot cdrom -s`. `cdrom` is an alias to the internal CD-ROM drive. If you are using an external drive, find out what the device path is (using `probe-scsi-all`) and boot from this path: `boot <long, cryptic device-path> -s`

Once the system has booted, you have a root-shell without entering a password.

Step 2: Now we are going to mount the disk. First let the Solaris OS rebuild the device tree with `devfsadm`. There should be eight new files in `/dev/dsk` afterwards.

Now use format subcommands `disk-print-partition` to verify that the new disk is correctly attached and to find out on which slice the root-Partition has been defined. Normally, slice 0 holds the root-partition. Let us assume you inserted the disk into slot 2 of the system. Format should have recognized the disk as `c1t2d0` (SCSI-ID 2 attached to the first controller), with slice `c1t2d0s0` as the root partition. So you should be able to mount the disk: `fsck -Y /dev/rdisk/c1t2d0s0; mkdir /tmp/disk; mount /dev/dsk/c1t2d0s0 /tmp/dsk`

Step 3: First, reset the root password to "empty" by deleting the crypted characters between the colons in `/tmp/dsk/etc/shadow`. Then look into `/tmp/dsk/etc/vfstab` to find out if this disk was located in slot 2 in the original system. If that was the case, the root file system "/" should be mounted on `/dev/dsk/c1t2d0s0`. If the mount point you find in `/tmp/dsk/etc/vfstab` is different, switch the slot of the disk accordingly.

Step 4: Let's assume we really found `/dev/dsk/c1t2d0s0` in `/tmp/dsk/etc/vfstab`. This means the disk occupied slot 2 in the original system. Typically, a Sun Enterprise 450 system would boot from the first slot (`c1t0d0s0`). Use the `/dev/dsk` directory to find out which physical device name `c1t2d0s0` points to, and write it down.

Step 5: Reboot, `STOP-A` and type: `boot <this physical device-Name>`

Step 6: You could define an alias `nvalias mydisk <this physical device-Name>` so you only need to type `boot mydisk` from now on. If you then define `setenv boot-device mydisk`, the system will always boot from `mydisk`.

2.5 Removable Devices

2.5.1 Volume Management

Removable devices are handled by Solaris Volume Manager software. This Volume Manager detects CD-ROMs and automatically mounts them to `/cdrom/cdrom0`.

Solaris Volume Manager will stumble when manually circumvented, for example, when a CD is removed using a paper clip instead of the appropriate command: `eject cdrom`. In such a case the Volume Manager can be restarted with `/etc/init.d/volmgt stop; /etc/init.d/volmgt start`.

2.5.2 Floppy

Insertion of a floppy disk (UNIX® or MS-DOS format) has to be communicated to the Volume Manager manually: `volcheck`

If `volcheck` was successful, the disk is mounted on `/floppy/floppy0`. Do not eject a floppy disk manually; always use `eject floppy`.

2.5.3 CD-ROM

CD-ROMs do not need `volcheck`. They are automatically detected and mounted to `/cdrom/cdrom0`. Ejecting the CD using `eject cdrom` is only possible while no shell or command occupies the `cdrom` path. `fuser -c /cdrom/cdrom0` prints out the process ID of that shell or command (`fuser -k` kills that process immediately).

Without volume management, CD-ROMs must be explicitly mounted and ejected (example for SCSI-ID 6):

```
- mkdir /tmp/cdrom
- mount -F hfs /dev/dsk/c0t6d0s0 /tmp/cdrom
- (Work with CD mounted on /tmp/cdrom)
- umount /tmp/cdrom
- eject /dev/dsk/c0t6d0s0
```

Note: The Solaris OS detects file names with more than 8.3 characters on a CD only if that CD has been created with Standard ISO-9660 Rock Ridge extensions. If only Microsoft Joliet extensions (default for most Microsoft

Windows CD Burning tools) are present, all file names are truncated to 8.3 (that is, “long-f.txt” instead of “LongFilename.txt”). Always create both Rock Ridge and Joliet extensions when burning CDs to make sure they are readable on all platforms.

Workaround: Create a zip file with a name no longer than 8.3 characters and burn that file onto the CD. Unzipping the archive on the target platform will properly restore all long file names.

2.5.4 CD-R and CD-RW

Bundled with Solaris OS 8 software is a tool, `cdrw`, to burn CD-R or CD-RW media. This tool burns an ISO-Image, which can be created from a Solaris OS file system or directory using `mkisofs`.

Look at this example for the exact parameters:

```
#!/bin/sh
# ./nfsd <PARAMETER1> <PARAMETER2>
#     PARAMETER1 = Application ID
#     PARAMETER2 = from where
#
# ./nfsd SunNetManager /net/center2000/export/pkg/SUNWsnm

# set Umask
umask 022

# Remove old image
rm /data/disk.img

# Create image
mkisofs -A $1 -d -l -L -o /data/disk.img -P "Sun Microsystems" -p
"Solaris SPARC" -r -R -J -V "$1" -v $2

# Burn
cdrw -i /data/disk.img
```

2.5.5 USB Memory Sticks

Most newer Sun systems have USB ports. So that the memory stick will be recognized and mounted by the Volume Manager, create a directory `/rmdisk` with the same ownership and permission attributes as `/cdrom`. After a FAT file system Memory Stick is inserted into the USB slot, its content should be readable as `/rmdisk/<name>`. If this is not the case, look at freshly created links in `/dev/dsk/. . .` pointing to “usb” paths and try a manual mount:

```
mkdir /tmp/usbstick
mount -F pcfs /dev/dsk/c1t0d0s2:c /tmp/usbstick
```

2.5.6 Tapes

Tape drives are not under the control of Volume Management, and instead must be handled with commands like `tar`. Typically tape drives are connected to the system as external SCSI devices. Connection and configuration can be performed while the Solaris OS is running. After running `devfsadm` (Solaris OS 8) or `drvcon-`

Administering Disks

`fig;tapes;devlinks` with previous Solaris OS versions, you will find some logical links in the `/dev/rmt` directory. 0 describes the first tape drive found, 1 describes the second tape drive, and so on.

After that, you may access a specific tape drive using `tar tvf /dev/rmt/0`.

Note: Tapes written with SGI systems often can only be read using a different blocksize:

```
tar tvfb 512 /dev/rmt/0
```

In some rare cases an additional byte swap has to be performed:

```
dd if=/dev/rmt/0 conv=swab | tar tvfb 512 -
```

2.6 Cloning a Disk

Using this recipe you may "clone" one system to another. Both systems must belong to the same hardware family. For example, an Ultra™ 10 workstation can be cloned to another Ultra 10 machine, but it cannot be cloned to an Ultra 60 workstation.

If the hardware configuration of the target system does have minor differences to the root machine (different frame buffer, NICs), the target system should be reconfigured after the cloning process using `boot -r`.

2.6.1 Cloning Using Internal Disks

This method is appropriate for systems with easy-to-change internal disks:

- Ultra 30, Ultra 60, Ultra 80, or Sun Blade™ 100 workstations or E220R/E420R/E250/E450.
- Sun Enterprise™ 220R, 420R, 250, or 450 servers.

The following is an example using an Ultra 60 workstation:

- Put the target disk into the second internal disk bay.
- Boot with single-user mode (`STOP-A; boot -s`).
- Configure all disks using `devfsadm`.
- Type `format` to look up the names of the disks. In an Ultra 60 workstation, you will find `c0t0d0` and `c0t1d0`. `c0t1d0` is the topmost internal disk.
- Determine which disk is the source disk and which one is the target disk (and do not forget to back up your data!). In our case `c0t0d0` (the bottom disk) acts as the source and `c0t1d0` (the upper disk) is the target.
- `dd if=/dev/rdisk/c0t0d0s2 of=/dev/dsk/c0t1d0s2 Bs=128k` (This command does a 1:1 copy of the whole disk, including the partition table. If the source disk is a 9 Gbyte disk, while the target disk is of 18 Gbyte capacity, 9 Gbyte would be wasted.) Warning: Double-check this step, so you do not accidentally destroy your source disk!
- `fsck -Y /dev/rdisk/c0t1d0s0` (File system check of the root partition, typically slice 0.)
- Remove the source disk, and put the clone disk into the previous source slot. Reboot your newly created clone system.
- Systems with internal FC-AL disks (such as on Sun Blade™ 2000 workstations and Sun Fire™ 280R or V880 servers) need additional work: Boot the system using a bootable Solaris CD-ROM, mount the root

filesystem and remove the files `/etc/path_to_inst` and `/etc/path_to_inst.old`. Then reboot the system interactively with `OK> boot -ar` and select "recreate `/etc/path_to_inst`" during boot.

2.6.2 Cloning Using a Crossover Cable

This method is not as fast as the internal disk copying process described in the preceding section, but it is suitable for systems where the internal disks cannot be changed easily (as in an Ultra 10 or Ultra 5 workstation.)

Warning: Remember to use the `/.rhosts` file only for this task. Delete it after completion to avoid a permanent security risk.

Target	Source
<pre>boot use "Command line login" at login screen Stop as many processes causing disk I/O as possible TERM=vt100 vi /etc/inet/hosts -> Insert target and source ifconfig hme0 plumb ifconfig hme0 <target> netmask 255... up -> Test with rsh source ls rsh source "dd if=/dev/rdisk/c0t0d0s2 ibs=128k obs=8k" dd of=/dev/rdisk/c0t0d0s2 ibs=8k obs=128k (For 9 Gbyte wait about 40 minutes.) STOP-A boot cdrom -s (with bootable CD) fsck -Y /dev/rdisk/c0t0d0s0</pre>	<pre>Insert target-IP in /.rhosts sync; sync</pre>

2.7 Mirroring the Boot Disk with SDS

Solaris Volume Manager (SVM), formerly known as Solstice DiskSuite (SDS) software, can configure disk subsystems for RAID 0, 1, 0+1 or 5. The most interesting feature for workstation installations is RAID 1, which allows you to mirror the disk that holds the root file system to be prepared for disk failures.

SVM is part of the Solaris 9 OS. For the Solaris 8 OS, you can find the SDS software on the second CD in the directory "EA." For the Solaris OS 7 or 2.6, Server Edition, look on the "Easy Access Server" CD.

Let's assume a workstation has two internal disks, `c0t0d0` and `c0t1d0`. Both disks should be partitioned in exactly the same way: Slice 0 for / (root), slice 1 for swap and a small slice 7 for the State Replica Databases. Slice 2 is always defined as Backup-Partition (see chapter "Partitioning").

Administering Disks

To copy the partition information from the first disk to the second disk, use this command:

```
# prtvtoc /dev/rdisk/c0t0d0s2 | fmthard -s - /dev/rdisk/c0t1d0s2
```

State Replica Databases are small files holding state information about the mirror. Each disk in a two-way mirror should hold at least three State Replica Databases:

```
# metadb -a -f -c 3 c0t0d0s3 c0t1d0s3
```

In the next step, create the first submirror on the first disk, holding the root file system:

```
# metainit -f d50 1 1 c0t0d0s0
```

Then create the second submirror on the second disk:

```
# metainit d52 1 1 c0t1d0s0
```

Build the first half of the mirror, using the first submirror:

```
# metainit d54 -m d50
```

Use `metaroot` to modify `/etc/vfstab` and `/etc/system` in a way that the system is able to boot from the metadevice `d54`:

```
# metaroot d54
```

Now reboot the system, then attach the second submirror `d54`. This action forces a "Mirror resync," copying the contents of the first disk onto the second one. This can impact system performance for a few minutes.

```
# metattach d54 d52
```

To avoid error messages during an "emergency boot" from the second disk, the swap partition should be mirrored in the same fashion. Otherwise, the swap partition on the broken disk, `c0t0d0s1` would still be referenced in `/etc/vfstab`.

Now the system may be booted either from the first disk (`boot disk`) or from the second disk (`boot disk1`).

If one of the disks is broken, Solstice DiskSuite software stops the boot process. This reason for this is that, to successfully boot after a failure, the majority of State Replica Databases must "survive." On our two-disk mirror only three out of the six replicas would have survived.

Servers with more than two disks should always have a third disk configured with State Database Replicas. Two-disk workstations have to be fixed manually, as described in the following example:

- State Database Replicas are configured on `c0t0d0s7` and `c0t1d0s7`, disk `c0t0d0` fails.
- The boot process stops with a Solstice DiskSuite error message and switches to Single User mode.
- Now the State Database Replicas on the failed disk have to be unconfigured:

```
metadb -d -f c0t0d0s7
```
- After the next reboot, partition the replacement disk exactly as in the surviving disk.
- Having done that, you now can create new State Database Replicas on the replacement disk:

```
metadb -a -c 3 c0t0d0s7
```
- Upon the next reboot, the mirror will be resynchronized.

Administering Disks

Hint: If you have a two-disk system, create three State Database Replicas on one disk, and four on the other one. Statistically there is a 50 percent chance to successfully reboot after a failure without manually using `metadb`. If you can tolerate unattended reboots after a disk failure in a two-disk configuration, create the same number of State Database Replicas on both disks and create this `/etc/system` entry: `set md:mirrored_root_flag=1`. More information on this topic can be found in the Sun BluePrints publication "Configuring Boot Disks With Solaris™ Volume Manager Software," (see: <http://www.sun.com/blueprints/1002/817-0407-10.pdf>).

3 Swap Space

3.1 Persistent Configuration in /etc/vfstab

```
#device          device          mount          FS          fsck          mount
mount
#to mount        to fsck          point          type        pass         at boot
options
#
#/dev/dsk/c1d0s2 /dev/rdisk/c1d0s2 /usr          ufs         1            yes
-
/dev/dsk/c0t0d0s1 -            -            swap        -           no           -
/usr/local/swapfile.1 -            -            swap        -           no           -
...
swap           -            /tmp        tmpfs       -           yes         -
```

A swap file system can use a raw partition like `/dev/dsk/c0t0d0s1` or plain files like `/usr/local/swapfile.1`.

You may define multiple swap file systems. All of these are logically combined by mounting them to the `/tmp` mount point. This means as long as physical memory is available, writing files to `/tmp` will write to and use up memory.

This example:

```
/dev/dsk/c0t0d0s1 -            -            swap        -           no           -
/dev/dsk/c0t1d0s1 -            -            swap        -           no           -
/dev/dsk/c0t2d0s1 -            -            swap        -           no           -
...
swap           -            /tmp        tmpfs       -           yes         -
```

creates an additional 6 Gbyte of virtual memory (given that all partitions have a size of 2 Gbyte). The Solaris OS automatically stripes access to all disks.

3.2 "Raw Partition" Versus Plain File

The fastest swapping is no swapping at all. If physical swapping is unavoidable, use multiple partitions striped over a set of disks. Only use plain files if you need memory and no extra partitions or disks are available.

3.3 Temporary Configuration

Plain files can be used to add swap space during standard Solaris OS operations:

```
# mkfile 100m /usr/local/swapfile.1
# swap -a /usr/local/swapfile.1
# swap -l
```

This addition is temporary until the next reboot. Until then, `swap -d /usr/local/swapfile.1` unconfigures the additional swap space.

4 Installation and Maintenance

4.1 Installation From CD

During installation, you may follow all the instructions printed on the screen. There are only a few pitfalls possible:

- Name service: Choose NONE. If you choose something like NIS at this moment, and the network is not yet properly set up and no NIS server is available, the system may "hang" waiting for a response from the NIS server. It is better to configure NIS or DNS after completing the basic installation of the network (see one of the following chapters).
- Installation cluster: For workstations, you should always choose the "Entire Plus OEM Software"-cluster to make sure everything gets installed. For dedicated servers, unnecessary packages should be removed afterwards to minimize security risks. The Sun BluePrints publication "Minimizing the Solaris^(TM) Operating Environment for Security: Updated for Solaris 9 Operating Environment" provides more information (see: <http://www.sun.com/blueprints/1102/816-5241.pdf>).
- Disk partitioning: To obtain the chance to partition the disk according to the recommendations given in the chapter "Partitioning" (root, swap and one small Solstice DiskSuite slice), choose "custom," not "default," when prompted for an installation method.

Starting with the Solaris OS 7, the Installation CD comes with a pre-patched Solaris OS image. To upgrade previous quarterly updates to this image, use the "Maintenance Update."

For example: Installing the "Maintenance Update 7" (which comes with the Solaris OS 8 02/02 version) onto the Solaris OS 8 07/01 version, will upgrade the Solaris OS patches (and only those -- that is, there are no functionality enhancements) to the status found on the 02/02 image. The upgrade installation writes a mark into the `/etc/release` file: "Solaris 8 07/01 Maintenance Update 7."

The "Recommended Patch Cluster" is a different mechanism, since you can have a different combination of patches and patch versions in the "Recommended Patch Cluster" every month. This cluster can be downloaded from the BigAdmin Portal (<http://www.sun.com/bigadmin>) and should be installed after every Solaris OS installation.

After an installation, `root` can only log in at the system console, that is, a terminal connected directly to the system or a keyboard. To enable a Remote Login via `telnet`, `rsh`, and so on, one file has to be modified first: Comment out the line `CONSOLE=/dev/console` in `/etc/default/login`. Warning: This is a security risk!

4.2 Packages

Basically, packages are subdirectories or `tar/compress/zip` compressed subdirectories, which follow a common name scheme: The first four uppercase letters correspond to the U.S. stock ticker symbol (or a fantasy symbol) of the origin company and the following lowercase letters describe the function of the package.

Example: The package containing the drivers for the SunPCi-card is named `SUNWspci2`. "SUNW" for the ticker symbol, "spci2" for the function.

You will install a package by using `pkgadd` with the directory storing the package contents:

```
pkgadd -d /cdrom/cdrom0
```

`pkginfo` prints information on installed packages, and `pkgrm <Package-Name>` removes a package.

Example: You have just changed the frame buffer card in your workstation, removed Creator3D (ffb) and installed Elite3D (afb). Now, you encounter an error message while booting, telling you that "no console framebuffer" is available. This indicates that no Elite3D drivers are installed.

First, verify this speculation with: `pkginfo | grep -i afb`

This list would typically be empty if you had done a plain default Solaris OS installation when the Creator3D, not the Elite3D, was present in the system. Drivers for the Elite3D can be found on the Solaris OS CD, subdirectory `s0/Solaris/Product`.

Change to this directory, and enter the following command: `pkgadd -d . *afb*`

After a reboot, all Elite3D drivers are now installed and available.

To find out which files belong to which installed packages, consult the file `/var/sadm/install/contents`.

To find out which files are present in packages on a Solaris CD-ROM, for example, look into the "pkgmap" files: `grep /cdrom/cdrom0/s0/Solaris_8/Product/*/pkgmap <file>`

4.3 Patches

A patch number contains two parts: the six-digit patch number, followed by a two-digit version number, for example: 123456-78.

Prior to the Solaris 8 OS, patches were supplied in `*.tar.Z` format. Solaris 8 OS patches are packaged as `*.zip` archives. For installing a patch, copy the archives to `/tmp` and extract them. Do not forget to remove the archives from `/tmp` after the installation has been completed:

```
zcat 123456-78.tar.Z | tar xvf -
unzip 123456-78.zip
```

After that, `patchadd <directory (= Patch number)>` installs the patch. In our example, `patchadd 123456-78`.

If no special parameters are given, `patchadd` installs patches in a way that they always can be removed with `patchrm <patch number>`. In order to be able to use this function, enough space in `/var/sadm/patch` is mandatory.

With `showrev -p | grep <patch number>` you can find out if a patch has already been installed.

5 X-Server and CDE User Interface

5.1 General Observations

The search sequence for X-Server or CDE-settings is always:

1. \$HOME/.dt/...
2. /etc/dt/...
3. /usr/dt/...

System-wide configuration files should never be changed in `/usr/dt`. Always use `/etc/dt` to maintain a fallback position in case something goes wrong. User-specific settings like menu extensions or changes to the front panel should always be stored in `$HOME/.dt`. In the following sections, `<DT>` will be used as a placeholder for all of the above.

5.2 Customizing the Login Screen

File: `<DT>/config/$LANG/Xresources`

Logo (Pixmap xpm or Bitmap xbm, beware of 24bit, extensive loading time!):

```
Dtlogin*logo*bitmapFile: /usr/local/lib/X11/dt/bitmaps/Mylogo.bm
```

General Welcome:

```
Dtlogin*greeting*labelString: Here's %LocalHost%
```

Welcome after user name is entered:

```
Dtlogin*greeting*persLabelString: Hello %s
```

5.3 Startup with Multiple Screens, Xinerama, 24-Bit Color Depth

First look into `/dev/fbs` and find out how many frame buffers are available and how they are named. Let's assume we found `/dev/fbs/ifb0` and `/dev/fbs/ifb1`, two Expert3D frame buffers.

Now we copy the file `Xservers` from `/usr/dt/config` to `/etc/dt/config` and edit them. The last line should now look like this:

```
:0Local local_uid@console root /usr/openwin/bin/Xsun :0 -nobanner
```

5.3.1 Multi-Screen

Multi-Screen means that CDE runs on every screen with a dedicated front panel. The mouse pointer can be moved between screens, but windows cannot. The display running on `ifb0` is called

```
:0.0, the display on ifb1 :0.1
```

```
:0Local local_uid@console root /usr/openwin/bin/Xsun :0 -dev /
dev/fbs/ifb0 -dev /dev/fbs/ifb1 -nobanner
```

5.3.2 Xinerama

Xinerama (available since Solaris OS 7 11/99) is one virtual screen covering all physical screens. The mouse pointer and windows can be moved among screens, and CDE draws only to one front panel. There is only one display number :0.0. Xinerama is only possible with frame buffers of the same type.

```
:0Local local_uid@console root /usr/openwin/bin/Xsun :0 +xinerama
-dev /dev/fbs/ifb0 -dev /dev/fbs/ifb1 -nobanner
```

Xinerama can define an overlap in X- and Y-direction. This can be used for edge blending with overlapping projectors or to "double" an image on two monitors (`xoverlap=<horizontal resolution>`). Example for 20% Edge Blending with 1280 pixel horizontal resolution:

```
:0Local local_uid@console root /usr/openwin/bin/Xsun :0 +xinerama
-xoverlap 256 -dev /dev/fbs/ifb0 -dev /dev/fbs/ifb1 -nobanner
```

5.3.3 24-Bit Color Depth

CDE itself uses 8-bit color depth (256 colors) by default. This is acceptable, as each window can define its own individual color depth. Some frame buffers only have one color table. This may cause color map flashing with some applications. A workaround for this is to start CDE itself with 24-bit color depth. On PGX, you have to switch the frame buffer to 24-bit with `fbconfig -depth 24`.

Note: Some legacy applications expect the default visual to be 8-bit, not 24-bit. This can cause problems.

```
:0Local local_uid@console root /usr/openwin/bin/Xsun :0 -dev /
dev/fbs/ifb0 defdepth 24 -dev /dev/fbs/ifb1 defdepth 24 -nobanner
```

5.4 A User-Defined CDE Front Panel

The application start mechanism in CDE is based on so-called "actions." Their behavior defines what happens when the user clicks on a CDE-Icon. Once defined, actions are referenced in various environments: menus, file manager, MIME-types, and definition files for the front panel.

5.4.1 Actions

CDE looks for *.dt-files in the following order:

```
$HOME/.dt/types
/etc/dt/types/$LANG/
/usr/dt/types/$LANG/
```

Example `StarOffice52.dt` (after installation of StarOffice in `$HOME/.dt/types`):

```
ACTION StarOffice52
{
  LABEL      StarOffice 5.2
  TYPE       COMMAND
  EXEC_STRING „/opt/Office52/program/soffice" „,(File)Args%"
  ICON       So52
  WINDOW_TYPE NO_STDIO
```

```

    DESCRIPTION StarOffice
}

```

5.4.2 Icons

The definition of an action also contains an icon name. An icon name is resolved in this order:

```

$HOME/.dt/icons
/etc/dt/icons/$LANG/
/usr/dt/icons/$LANG/
naming scheme: <name>.<size=m|t|l>.<format=bm|pm>

```

So52, for example, is the icon file in `$HOME/.dt./icons/So52.m.bm`

5.4.3 Front Panels

The front panel can be extended easily by adding *.fp-files into `$HOME/.dt/types`. The default definitions are stored in:

- `/etc/dt/types/$LANG/`
- `/usr/dt/types/$LANG/`

Example (`$HOME/.dt/types/tools.fp`):

```

CONTROL Tools          (Always define the main control before sub panel)
{
    TYPE                icon
    CONTAINER_TYPE     BOX
    CONTAINER_NAME     Top
    POSITION_HINTS      first
    ICON               Dtagen    (This is a folder icon, Icon must
                                exist!)
    LABEL              Tools
}

SUBPANEL Toolspanel
{
    CONTAINER_NAME     Tools
    TITLE              My Tools
}

CONTROL StarOffice52
{
    TYPE                icon
    CONTAINER_TYPE     SUBPANEL
    CONTAINER_NAME     Toolpanel
    ICON               So52
    LABEL              Star Office 5.2
    PUSH_ACTION        StarOffice52    (compare ACTION StarOffice52)
    DROP_ACTION        StarOffice52
}

```

You may build a whole menu bar by adding additional CONTROLS in the SUBPANEL.

6 Frame Buffers and OpenGL

6.1 fbconfig

With the Solaris 8 OS, `fbconfig` is the unified interface for configuring frame buffers. Prior to the Solaris 8 OS software, each frame buffer driver had its own configuration utility, with a naming scheme corresponding to the logical device-names in `/dev/fbs`:

- `ffbconfig` Creator, Creator3D
- `afbconfig` Elite3D
- `ifbconfig` Expert3D, Expert3Dlite
- `m64config` PGX24 (onboard U5, U10), PGX64
- `gfxconfig` PGX32

Because the parameters are exactly the same, we will only mention `fbconfig` in the following paragraphs. If you do not specify a device for the frame buffer, `fbconfig` operates on the default (console) frame buffer.

You can find out which frame buffers are installed in a system by using `ls /dev/fbs`. Afterwards, they can be addressed directly using their device names:

```
fbconfig -dev /dev/fbs/afb1 -propt -prconf
```

prints configuration information (which frame buffer model, which resolution, and so on) for the second Elite3D frame buffer configured in the system. The first Elite3D would have the device name `afb0`.

```
fbconfig -res \?
```

shows a list with possible resolutions. A given resolution could be copied to the command:

```
fbconfig -res „1280x1024x86" -try -now
```

using "cut&paste".

`-try` and `-now` are options. `-try` changes the resolution for 10 seconds only. This way you can determine if the monitor is able to synchronize to the video signal. `-now` switches the resolution immediately, without the need to restart the X-Server. In some cases, however, the geometry of the X-Server may look distorted after `-now`. An X-Server restart using `Exit` and `Login` solves this problem.

```
fbconfig -g 1.7
```

sets the gamma factor to 1.7. The default value is 2.2, which often is perceived as too bright. To be activated, the new gamma factor does not need an X-Server restart. After a reboot, the newly set value is still valid.

6.2 OpenGL

OpenGL® will not be installed automatically during the Solaris OS installation process, although it has been packaged with the Solaris OS CDs since Solaris 2.5.1 OS.

The command `/usr/openwin/demo/GL/ogl_install_check` checks if OpenGL is installed correctly and which version it is. Note: The rotating ring also comes in handy when judging gamma correction values (see above).

It is recommended to always use the latest version of OpenGL. Xinerama, for example, is only supported with OpenGL 1.2.1 or later. OpenGL can be downloaded from <http://www.sun.com/solaris/opengl>.

Look at the `ogl_install-check` output: If it prints `"GLX: context is direct"`, OpenGL uses DGA and applications may run at full speed. If it prints `"GLX: context is indirect"`, OpenGL does not use DGA, probably because the application was called from a `"su -"` environment.

If you log into the window system using one user ID and then switch to another (using `"su"`), OpenGL will switch to "indirect mode." Due to a security feature in the Solaris OS, `"su"`ing to a user other than the one who started the window system disables DGA (Direct Graphics Access) to the frame buffer, thus making OpenGL graphics slower than necessary.

If you wish, you can remove this security feature and thus make graphics fast for all users by doing the following:

- Become root
- Edit the permission in `/etc/logindevperm` from 0600 to 0666:

```
/dev/console 0666 /dev/fbs/* # frame buffers
```
- Reboot

Note that any user will now have read/write access to the contents of your screen.

DGA will also be switched off when the `$DISPLAY` variable is set to `"hostname:display"`, e.g. `"dumbledore:0.0"` or `"dumbledore.sun.com:0.0"`. It's better to stick with the default (`":0.0"`) setting after locally logging into the machine.

6.3 Patch Matrix for Graphics

There is a small set of patches that are important for each installation where a frame buffer is involved. The collection depends on the Solaris OS version, OpenGL version, and type of frame buffer.

The following table only lists the patch numbers. Please download the most recent versions of the corresponding patches from <http://sunsolve.sun.com> and install them in the following order:

1. Frame buffer patch
2. X-Server patch
3. OpenGL
4. OpenGL patches

Graphics Patch Matrix			
GRAPHICS CARDS	Solaris 9 OS	Solaris 8 OS	Solaris 7 OS
Expert3D (lite)	112540	108576	108787
XVR-500	112540	108576	not supported
XVR-1000	112565	112564	not supported
Elite3D	111620	108604	106144
- and XFB			106148
- and VIS		109872	106147
Creator3D	112621	108605	106145
- and XFB			106148
- and VIS		109872	106147
OpenGL 1.2.2	32 Bit: 111993	32 Bit: 111993	32 Bit: 111993
	64 Bit: 111994	64 Bit: 111994	64 Bit: 111994
OpenGL 1.2.3	32 Bit: 112628	32 Bit: 112628	32 Bit: 112628
	64 Bit: 112629	64 Bit: 112629	64 Bit: 112629
XSun	112785	108652	108376
PGX32 /64	112622	109154	107716

7 Administration

7.1 Hosts

You may use `admintool` or `smc` to create hosts, but it is also convenient to directly edit the file `/etc/inet/hosts`. Note: The file `/etc/hosts` is a link to `/etc/inet/hosts`.

A copy like `cp /etc/hosts.old /etc/hosts` would not overwrite the correct file.

A default Solaris OS installation starts the `sendmail-daemon`. This daemon writes error messages if no alias with a dot in its name is defined for the primary host name. To suppress the error message, define the primary host in the file `/etc/inet/hosts` like this:

```
192.168.0.2    myhost    myhost.domain.com    loghost
```

7.2 Serial Ports

Again, the best tool for infrequent use is `admintool` or, with the Solaris 9 OS, use the "Solaris Management Console" `smc`.

7.3 Printers

Again, the best tool for infrequent use is `admintool`. The Solaris OS 9 comes with a "Printer Setup Wizard" `/usr/sadm/admin/bin/printmgr`.

7.4 User Creation

The best tool for infrequent creation of user accounts is `admintool`. Beginning with the Solaris 9 OS, use the "Solaris Management Console" `smc`. Please do not specify `/home/<username>` as the home directory when you create a new user. It's better to use:

```
/export/home/<username>.
```

By default `/home/<username>` is a target of the Solaris automounter and is not writeable because of this. Some third-party software tries to use `/home/xyz`. Change the defaults or shut down the automounter to avoid any problems.

7.5 Shells and Environment Variables

The correct setting of environment variables is one of the most error-prone tasks during system setup. Unfortunately, the syntax for how to set those variables and how to set them permanently for one particular user differs greatly depending on which shell is used (or is defined as default for the user).

The most common shells are `sh`, `ksh` and `csch`. Beginning with the Solaris 8 OS, the popular `bash` shell is also available. `sh` is the default shell for `root` after installation. The default shell for a particular user is defined in the `/etc/passwd` file, for example: `demo:x:2003:10:Demo`

User: `/export/home/demo:/bin/csh`

This table describes how to set variables using the different shells:

Shell	Setting	Appending	Retrieving	Comments
<code>sh</code>	<code>NAME=value1</code> <code>export NAME</code>	<code>NAME=value2:\${value1}</code>	<code>set (internal)</code> <code>env (exported)</code>	Use "export" if subprocesses should see the value of NAME
<code>ksh</code>	<code>NAME=value1</code> <code>export NAME</code>	<code>NAME=value2:\${value1}</code>	<code>set (internal)</code> <code>env (exported)</code>	Use "export" if subprocesses should see the value of NAME
<code>csh</code>	<code>setenv NAME value1</code>	<code>setenv NAME value2:\${value1}</code>	<code>printenv</code>	No "export" needed
<code>bash</code>	<code>NAME=value1</code>	<code>NAME=value2:\${value1}</code>	<code>set env printenv</code>	No "export" needed

Every shell reads a set of files during startup. A login shell (typically started after the user has logged into a system using `telnet`, `ssh`, and so on.) reads a different set of files than a subshell. The following table describes which files are read:

Command	Shell	----- Files Read in This Order ----->			
Subshell (calling <code>sh</code> , <code>ksh</code> , <code>csh</code> or <code>bash</code> explicitly) <code>su</code>	<code>sh</code>				
	<code>ksh</code>				
	<code>csh</code>	<code>~/.cshrc</code>			
	<code>bash</code>	<code>~/.bashrc</code>	<code>~/.inputrc</code>	<code>/etc/inputrc</code>	
<code>su -</code>	<code>sh</code>	<code>/etc/profile</code>	<code>\${HOME}/.profile</code>		
	<code>ksh</code>	<code>/etc/profile</code>	<code>~/.profile</code>		
	<code>csh</code>	<code>~/.cshrc</code>	<code>~/.login</code>		
	<code>bash</code>	<code>~/.bash_profile</code>	<code>~/.bash_login</code>	<code>/etc/profile</code>	<code>~/.profile</code>

Administration

Command	Shell	----- Files Read in This Order ----->			
telnet, ssh, rlogin	sh	/etc/profile	\${HOME}/.profile		
	ksh	/etc/profile	~/.profile		
	csch	~/.cshrc	~/.login		
	bash	~/.bash_profile	~/.bash_login	/etc/profile	~/.profile

Hint: ksh, csch and bash use "~" as a shortcut to the user's home directory

8 Network Connections

8.1 /etc-files

Solaris software stores network configuration files in the directory `/etc`. All configuration data, such as IP addresses, gateways, and so on, can be defined and changed at runtime level. However, only those definitions stored in `/etc` files will survive the next reboot.

The host name can be found in these files:

- `/etc/nodename`
- `/etc/hostname.<interface-name>`
- `/etc/inet/hosts`
- `/etc/net/{ticlts,ticots,ticotsord}/hosts`

`/etc/nodename` defines the name of the system. Looking from a network perspective, a system with multiple network interfaces and multiple IP addresses can have multiple names, each one defined in `/etc/inet/hosts`. One of these names may be equal to the `nodename`, but this is not necessarily the case.

Certain files create the connection between IP-addresses and interface names. These are named `/etc/hostname.<interface-name><number>`.

The most common interface names are:

<code>le</code>	Ethernet on older SPARC™ systems
<code>hme</code>	FastEthernet on most UltraSPARC™ systems
<code>eri</code>	FastEthernet on most UltraSPARC-III systems
<code>qfe</code>	FastEthernet on QuadFastEthernet-extension cards

Example: An Ultra 60 workstation with QFE (QuadFastEthernet) has five network interfaces: `hme0` (on board), `qfe0`, `qfe1`, `qfe2` and `qfe3` (on the QFE-card). `Nodename` is `hermione`, and the `/etc/inet/hosts` looks like this:

```
127.0.0.1          localhost
157.168.34.12     hermione  hermione.domain.com  loghost
191.133.23.10     dumbledore
```

The identity of the system in the corporate LAN using the address `157.168.34.12` should be `hermione`. Accordingly, the file `hostname.hme0` must contain the name `hermione`.

`dumbledore` is an address assigned by an Internet Service Provider (ISP): `191.133.23.10`. That cable is connected to the first QFE port, `qfe0`. So the file `hostname.qfe0` contains the string `dumbledore`.

`qfe1-3` are currently not used. Do not create `hostname.*` files for unused interfaces to avoid Solaris OS error messages complaining about interfaces with no network cables attached.

A system with more than one network interface will automatically route between the subnets (in our example `157.168...` and `191.133...`). During boot, the message "machine is a router" is printed. If this behavior is not desired, create an empty file `/etc/notrouter` and reboot.

The netmasks for the IP-addresses are defined in the file `/etc/netmasks`.

Do not change the nodename while the X-Server is running! It is better to log out from the current session and login again using "Command Line Login" before changing the nodename or host name.

8.2 Virtual Interfaces

Occasionally, it may be useful to define two IP addresses for a system with only one physical network interface. For example, a demo workstation for use at a trade show should keep its corporate IP address while getting an additional address to communicate with the other hosts at the booth.

This can be done with "virtual interfaces." The procedure is exactly the same as it was described for two physical interfaces in the last chapter. The only difference is that we do not use `hme0` and `qfe0` this time, but `hme0` and `hme0:1`. So if you write `dumbledore` into `/etc/hostname.hme0:1`, the cable provided by the ISP at the trade show can be connected to `hme0`, a QFE card will not be necessary.

8.3 ifconfig

`ifconfig -a` lists all interface definitions, physical as well as virtual ones, with IP-addresses and netmasks. Additionally, you can use `ifconfig` to define IP addresses and netmasks during runtime. Let's go back to our Ultra 60 workstation example again, but assume that this time `dumbledore` has not been defined in `/etc/inet/hosts` so that `qfe0` will be unused after reboot.

The first step now is to activate `qfe0` (load the driver): `ifconfig qfe0 plumb`

To startup the interface type:

```
ifconfig qfe0 191.133.23.10 netmask 255.255.255.0 up
ifconfig -a shows the new configuration.
```

It is the same procedure for a virtual interface:

```
ifconfig hme0:1 plumb
ifconfig hme0:1 191.133.23.10 netmask 255.255.255.0 up
```

8.4 Routing

The default router (or standard-gateway) can be entered with its IP address into the file: `/etc/defaultrouter`. If you specify a host name in that file, make sure the host name is defined in `/etc/inet/hosts`.

The equivalent runtime commands are:

- Delete all current routes: `route flush`
- Define 1.2.3.4 as default router: `route add default 1.2.3.4`
- Check with `netstat -r`.

If more than one network interface is available in your workstation, Solaris OS automatically turns on ip-forwarding.

- To check if ip-forwarding is switched on, use: `ndd -get /dev/ip_forwarding`
- To turn it on, use: `ndd -set /dev/ip_forwarding 1`
- To turn ip-forwarding off, use: `ndd -set /dev/ip_forwarding 0`

Another way of turning ip-forwarding off is to create the file `/etc/notrouter`, for example, by issuing the command `touch /etc/notrouter`.

8.5 NIS-Client

1. First enter the name of the NIS domain into `/etc/defaultdomain` (case-sensitive)
2. Then define the NIS server in `/etc/inet/hosts`
3. Start `ypinit -c` and type in the name of the NIS server
4. `/usr/lib/netsvc/yp/ypstart` or reboot

Now you have to tell the Solaris OS the tables (host names, users, and so on) for which it should use NIS. This is done with the file `/etc/nsswitch.conf`. To make Solaris software ask the NIS server if it cannot find a host name in the local `/etc/inet/hosts` file, append `nis` after `files` in the line starting with `hosts`.

There is a predefined file called `nsswitch.nis` in the `/etc`-directory, which contains NIS definitions for all available tables.

8.6 DNS-Client

1. Enter your DNS server into the file `/etc/resolv.conf`
`nameserver 192.168.0.1`
Now test the DNS server using `nslookup www.sun.com`
2. Tell Solaris OS software to use DNS when resolving host names. Edit `/etc/nsswitch.conf`, and add `dns` after `files` in the line starting with `hosts`.

Now, not only `nslookup www.sun.com`, but also `ping www.sun.com` should be successful.

8.7 DHCP-Client

An interface can be configured to use DHCP if you create two empty files: `/etc/hostname.<interface>` and `/etc/dhcp.<interface>`. The Solaris OS software waits 30 seconds for a DHCP server to answer. This can be modified in `/etc/dhcp.<interface>: WAIT <time in seconds>`.

It is possible to start DHCP on an interface at runtime (in this example for interface `qfe0`):

```
ifconfig qfe0 plumb
ifconfig qfe0 dhcp start
```

8.8 sys-unconfig

If you want to change the network configuration and you are unsure which files are involved, there is an easy solution: `sys-unconfig` shuts down the machine. After the next reboot (enter `boot` at the OK prompt), all network-related questions you know from the Solaris OS installation tool are asked again (IP address, gateway, netmask, and so on).

8.9 10 or 100 Mbit, Half- or Full-Duplex?

You can query the state of a network interface with the command `ndd /dev/<interface-driver>.<interface-driver>` typically is `le`, `hme`, `eri` or `qfe`:

```
- ndd /dev/eri link_status: 0 = down, 1 = up
- ndd /dev/eri link_speed: 0 = 10 Mbit, 1 = 100 Mbit
- ndd /dev/eri link_mode: 0 = half-duplex, 1 = full-duplex
```

The Solaris OS default value for all network interface parameters is "autonegotiation." If the attached network equipment is not capable of autonegotiation, all parameters can be set manually. This is done by setting all unwanted properties to zero.

An example for 10 Mbit half-duplex:

```
ndd -set /dev/eri adv_100fdx_cap 0
ndd -set /dev/eri adv_100hdx_cap 0
ndd -set /dev/eri adv_10fdx_cap 0
ndd -set /dev/eri adv_10hdx_cap 1
ndd -set /dev/eri adv_autoneg_cap 0
```

Shut down the interface with `ifconfig unplu Mbyte` and restart it with `ifconfig plu Mbyte` afterwards. The settings can also be defined in the `/etc/system` file (`set eri:eri_adv_10hdx_cap = 1`).

8.10 Network Settings: Temporary vs. Persistent

What	Valid Only for Runtime	Persistent After Reboot
Hostname	hostname <new name> Do not use within CDE!	/etc/nodename /etc/hostname.<Interface> /etc/inet/hosts
IP-Address	ifconfig <Interface> plumb ifconfig <Interface> <IP> up	/etc/inet/hosts
Netmask	ifconfig <Interface> <netmask> up	/etc/netmasks
Gateway	route add default <IP Gateway>	/etc/defaultrouter
DHCP-Client	ifconfig <Interface> plumb ifconfig <Interface> dhcp start	touch /etc/hostname.<Interface> touch /etc/dhcp.<Interface>

9 Network File System

9.1 NFS-Server

Solaris OS software stores information about directories, which are shared and can be mounted from other systems in the network, in the file `/etc/dfs/dfstab`. This file is a script containing `share` commands.

If this file does not contain any `share` commands, the Solaris OS does not start the NFS server daemon during reboot. Even when explicitly called, `/etc/init.d/nfs.server start` does nothing as long as `dfstab` is empty.

Example: Sharing `/export/home`

1. Insert `share /export/home` into `/etc/dfs/dfstab`.
2. If this has been the first entry, execute `/etc/init.d/nfs.server start`.
3. Enter `shareall`.
4. Check with `share if /export/home` has been correctly exported.

9.2 NFS-Client

A NFS client does not need any special daemons, but it is important that the NFS server is known by name. An IP-address (`mount 192.168.0.1:/export/home`) is not enough. All NFS servers must be defined in `/etc/inet/hosts`, NIS or DNS.

You can mount a network directory with `mount server:/export/home /tmp/home`. If this mount should be reestablished automatically at the next reboot, insert the mount into the file `/etc/vfstab`.

Note: When copying files from network directories as user "root," keep in mind that the local "root" does not act as "root" on the network directory, but as "nobody" (unless shared on the server with `anon=0`). Files in this directory that do not have read permission for "other" are not copied. To copy all files from a network directory as "root", do **not** use:

```
root# cd server:/verz; tar cvf - | (cd verz; tar xvf -)
```

Better use "rsh" to get root read permissions on the server:

```
root# rsh server "cd /verz; tar cvf -" | (cd verz; tar xvf -)
```

9.3 Differences From Other Popular UNIX Systems

The file `/etc/dfs/dfstab` is a Solaris OS specialty. Other popular UNIX systems (Linux, AIX, and so on) use `/etc/exports` which has a different syntax. `shareall` is known there as `exportfs -a`.

9.4 Automounter

The automounter allows automatic mounts of file systems when the user accesses the respective directory. In large networks with lots of hosts and users, it will be very helpful to always have directory paths like `/net/<hostname>` or `/home/<username>`.

Two files define the basic automounter operation: `/etc/auto_master` and `/etc/auto_home`.

`auto_master` comes predefined for the behavior described above. `/net` is automatically extended with the proper host name when accessed. How `/home` should be handled is defined in `auto_home`.

```
/net    -hosts    -nosuid,nobrowse
/home   auto_home -nobrowse
```

If `auto_home` contains the line

```
*      /export/home/&
```

the automounter becomes active when somebody enters `/home/joe` and mounts `/export/home/joe` to this mount point.

10 Tools

10.1 Compiler

The C, C++, and Fortran compilers typically are installed in the `/opt/SUNWsprow/bin` directory, while common tools like `make` are in `/usr/ccs/bin`.

10.2 GNU-Tools

Almost all GNU-Tools are available as precompiled packages for the Solaris Operating System. They can be downloaded from <http://www.sunfreeware.com>. Beginning with the Solaris 8 OS, some important tools like Perl, gzip, and apache are automatically installed. Most of the remaining GNU-Tools come as precompiled packages on the Solaris OS Companion CD.

10.3 SunPCi™ and SunPCi-II

Most workstations can be equipped with a SunPCi card. This is a complete PC on a PCi card, running Windows (98, ME, NT, 2000, XP). Here are some hints for the operation of this card:

- Encapsulate `/opt/SUNWspci2/bin/sunpci` into a small script that sets some environment variables:
`setenv NVL_INTERFACE=hme0` (If SunPCi should use the `hme0` network interface)
`setenv LANG=C`
`setenv KBCP=850` (Example: German keyboard and codepage)
`setenv KBTYP=GR` (Remember MS-DOS ?)
- Make the C: image writeable not only for root, but for each user on the workstation: `chmod 666`. When a user starts `sunpci` for the first time, a subdirectory `pc` is created in the HOME directory of the user. To make the user work with the global C: image, interrupt the following creation of a user-specific image and edit the `sunpci.ini` file to point to the global image, as in `pc/C.diskimage`, for example.

10.4 StarOffice™ Software

If StarOffice software is installed as "root" with `setup /net`, all users can choose a "workstation installation" during their individual install. A workstation installation only copies 1 Mbyte of data into the home directory of the user and otherwise uses the central installation. If StarOffice software is installed in `/opt/Office60`, a user can start his or her workstation installation by calling `/opt/Office60/program/soffice`.

10.5 Netscape Navigator™ Software

Beginning with Solaris 8 OS, Netscape Navigator software is used as the standard browser. It is installed in `/usr/dt/bin/netscape`.

10.6 Miscellaneous

Hint: The registration reminder when starting the desktop can be disabled with "More Information – Never Register."

To disable the warning for all users, edit `/etc/default/solregis` and put in the following line: `DISABLE=1`

11 Hint Collection

11.1 Top Three Hints

- C-Shell: Define `set filec` and `set history=100 savehist=50` in `~/.cshrc`. The shell will complete file names and directories once you hit the ESC key, and `history` lists the last 100 commands. With `!13` command `#13` can be re-issued. `!$` can be used as a variable for the last word of the former command, like in: `mkdir /export/home/demo/test`, then `cd !$`.
- Beginning with the Solaris 8 OS, the `tcsh` is available, which allows editing the command line and browsing the history using the cursor keys.
- Process control: `ps -ef | grep <search text>` and then `kill <PID>` to search and kill a process can be replaced with `pkill <search text>` since the Solaris 8 OS `kill -9 -1` terminates all processes belonging to the current user.
- File editing: Those who do not like `vi` can use the CDE-Editor `dtpad`.

11.2 Analysis of the Hardware Configuration

11.2.1 prtdiag

The command `/usr/platform/`uname -i`/sbin/prtdiag -v` prints information on the number and type of CPUs, RAM, extension cards, and so on.

11.2.2 prtconf / path_to_inst

`prtconf` lists the hierarchical device tree. The mapping of physical devices (names in the `/devices`-directory) with logical devices (drivers in the `/dev`-directory) is described in the file `/etc/path_to_inst`. This file is re-written during a `boot -r` or `devfsadm`. Editing this file is necessary after having moved an adapter card from one PCI slot to another one. `/dev/hme1` is connected to a specific PCI-Slot in `/etc/path_to_inst`. A simple reboot would not redefine this relationship. Warning: Editing `/etc/path_to_inst` manually requires high-level experience with device path handling. Editing errors can lead to unbootable systems!

11.2.3 Sun Explorer Software

This tool can be downloaded from SunSolve (<http://sunsolve.Sun.COM/pub-cgi/show.pl?target=explorer/explorer>). It collects all important system information. It can be used to prepare service calls, backup configuration data, snapshot configurations, and so on.

11.3 Analysis of Runtime Data

11.3.1 time/ptime

One of the simpler runtime analysis tools is `time` (0,1 sec. resolution) or `ptime` (1 msec resolution), which can be added to a command line:

```
time gzip j2sdk.tar.gz
```

The result is divided into `user` (=application) time, `system` (=OS) time, and elapsed time.

11.3.2 top / prstat / ps

A tool known as `top` on other platforms is `prstat` for Solaris OS software. `/usr/ucb/ps -aux` shows similar results, but they are unsorted and without auto-update.

```
ps -L -p <process-ID> lists whether a process has threads and, if so, the number of threads.
```

11.3.3 vmstat / mpstat

`vmstat <frequency in seconds>` lists details of the system load. The following rows are the most interesting ones:

<code>free:</code>	Free memory (near zero through Solaris 7 OS, because of disk cache used. Beginning with Solaris 8 OS, the correct size of free memory)
<code>page pi/po:</code>	Paging activity. The system does not swap until I/O traffic on the swap device is reported!
<code>cpu us:</code>	User-time. This percentage is consumed by applications
<code>cpu sy:</code>	System-time. This percentage is consumed by the Solaris OS itself.
<code>cpu id:</code>	Idle-time.

`mpstat` prints similar information, one line per processor.

11.3.4 iostat

`iostat -xtcP <frequency in seconds>` lists the I/O load of a system. One line is printed for each partition or NFS-mount. The rows "`kr/s`" and "`kw/s`" show read and write throughput as Kbyte/sec. A "`svc_t`" bigger than 100 indicates that the disk is spending more time with head positioning than data transfer. Beginning with the Solaris 9 OS, "`-txciP`" can be used to print the partition names in "`c0t0d0s0`" format.

11.3.5 sdtperfmer, xcpustate, and sdtprocess

`sdtperfmer` is a graphical tool, which draws the `vmstat`-data as a bar or line chart. A minimized version of this tool is displayed in the CDE front panel.

Unfortunately `sdtperfmer` only prints the cumulated CPU and I/O load. To get a bar chart for every CPU and disk installed in the system, download `xcpustate` from <ftp://ftp.sunmde.com/pub/gregns/xcpustate>. Start the tool with "`xcpustate &`" or "`xpustate -disk &`".

`sdtprocess` lists all processes in a sorted way. It is possible to look further into process properties, terminate processes, and so on.

11.3.6 32 or 64 Bit?

`isainfo -kv` displays if the kernel is a 32-bit or 64-bit kernel. The 64-bit kernel is used automatically, starting with UltraSPARC II systems. With UltraSPARC II systems, the 32-bit kernel can be loaded explicitly with `boot kernel/unix` instead of `boot kernel/sparcv9/unix` at the OK prompt. With UltraSPARC III systems, only 64-bit kernels are available.

11.4 Debugging

11.4.1 Truss

The command `truss <command>` lists all system calls that an application calls during its execution. Because all parameters of these system calls and the return values are listed, one can see, for example, where an application looks for shared libraries or configuration files, and at which location they are finally found.

11.4.2 pstack, pmap, and So On

The "p*-command" prints out detailed information belonging to processes:

```
pstack <PID>    Prints the last function calls of a process.
pldd <PID>      Prints a list of shared libraries.
pmap -x <PID>   Prints the memory usage of the process and its modules.
```

11.4.3 snoop

`snoop` prints all packets that are sent through a network interface.

11.5 Common Pitfalls

Observation	Hint
Boot process takes very long	OBP settings: Is <code>diag-switch? = true ??</code>
Cannot login with ftp (but telnet, ssh is fine)	User disallowed to use ftp in <code>/etc/ftpusers</code>
<code>vi</code> does not display correctly	Variable <code>TERM</code> defined? If not, set <code>TERM=vt100</code>

12 Books and Web Sites

12.1 Books

- A valuable collection of hints and tricks for UNIX:
Jerry Peek, Tim O'Reilly, Mike Loukides, *UNIX Power Tools*
(ISBN 1565922603)
- A detailed view on all UNIX variants, including Solaris OS and Linux.
Evi Nemeth, Garth Snyder, Scott Seebass, Trent R. Hein, *UNIX System Administration Handbook*
(ISBN 0130206016)

12.2 Web Sites

<http://www.sun.com>

Home page of Sun Microsystems

<http://www.sun.com/bigadmin>

Everything a system administrator may need

<http://www.sun.com/blueprints>

BluePrints publications on various topics, with monthly additions

<http://docs.sun.com>

All documentation for Sun products online

<http://sunsolve.sun.com>

Technical FAQs, infodocs, symptom/resolution database,
patch downloads, and so on

<http://www.sunfreeware.com>

Freeware tools for Solaris OS, precompiled

<http://www.sunhelp.org>

FAQ portal, other information

<http://bhami.com/rosetta.html>

"A Sysadmin's Universal Translator (ROSETTA STONE)

OR What do they call that (command) in this world (Unix)?"