We make the net work.

Sun microsystems

Solaris™ OE Guide for New System Administrators

The 20% of Solaris knowledge that solves 80% of your needs

Rolf Kersten

Sun Microsystems GmbH, Germany

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

Table of contents

Introduction

This document contains the most relevant 20% of technical information you will need to know when working with Solaris™ and Sun platforms. It should answer 80% of your Solaris questions.

This is not intended to replace detailed books on the Solaris Operating Environment (OE) or comprehensive FAQs (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 system administration or those that are not administering Solaris systems on a daily basis. The focus of this document is on graphical workstations used for technical applications and should also be helpful to those that need to set up Sun systems for demos or benchmarks.

1 System Booting and Testing

1.1 The OK-Prompt

Each Sun system has a Boot-PROM mode which can be accessed prior to starting the Solaris OE or while it 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 OE 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 via a directly connected keyboard and monitor or 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 hardwire` 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 do exist:

The best and cleanest method 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 of the 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 or web servers cleaning process ID files etc.

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

- `STOP-A`
- `OK sync`   Sync filesystems. Panic the OS afterwards, then reboot.

   Attention: Watch out for big OS coredumps 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 only describe the most important commands.

A detailed description can be found in "Hardware Diagnostics for Sun TM Systems: A Toolkit for System Administrators" (Infodoc Number 23476 on http://sunsolve.sun.com).

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-Framebuffer
    - Sun-Globe: Creator 3D
    - Sun with face: Elite 3D
    - Self-describing Logo: PGX, Raptor GFX (=PGX32), Expert3D(lite) XVR-500, XVR-1000 etc.

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. This can be used to find out id e.g. a network cable is working properly.

1.4.4 printenv and devalias

`devalias` lists all aliases, `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` during power-on of the system.

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. This can be disabled by setting `setenv tpe_link test false` at the OK prompt.

2 Administering disks

2.1 Logical Device-Names

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

- c   Controller number
- t   *T*arget-ID. For SCSI-devices this number corresponds to the SCSI-ID
- d   *D*rive Number or LUN for RAID-devices (one ID, multiple disks)
- s   *S*lice, Partition number (see next paragraph)

These logical device names are stored as links in the directories `/dev/dsk` and `/dev/rdsk`. These links point to physical device names. But those are only needed for tasks like changing the boot device. Typical Solaris OE commands for administering drives expect logical device names as their parameter.

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).

A refresh of the `/dev/dsk` and `/dev/rdsk` directories occurs either while booting (`boot -r` at the OK-Prompt) or during Solaris OE operation with `devfsadm` (`drvconfig;disks;devlinks` prior to the Solaris 8 OE). 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` subcommand `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 OE default installation creates three slices:
(i) approximately 1GB root, (ii) swap (depending on memory capacity) and (iii) the remaining diskspace for `/export/home`.

For workstations, this partitioning scheme is not the optimum. 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 logfiles or big files stored into `/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™ metadata
```

In order to be prepared for the use of the Solstice Disksuite logical volume manager (SDS), it is recommended to preserve some space (one cylinder) 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 OE, "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 (e.g.. 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 mountet>, to be used as ABE with Live Upgrade
                 ( about 150% of the expected size of / )
slice 7: sds    A small space reserved for Disksuite metadata
```

2.3 Crashdumps

Starting with the Solaris 7 OE, crashdumps 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/<hostname>`, in order to allow for later analysis. Therefore, /var (or /, if no special partition for /var was defined) should also not be sized too small.

If a system with lots of GBytes uses the swap partition as dump device, the boot process can take some time because of the copying procedure mentioned above. In these cases, it is better to define a dedicated dump partition using `dumpadm`. Copying from a dedicated dump device – and not from swap - to `/var/crash` is performed in the background.

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: (i) 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) and (ii) what the root password is. Start with putting the disk into an arbitray 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 OE 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/rdsk/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 filesystem "/" should be mounted on `/dev/dsk/c1t2d0s0`. If the mountpoint 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, 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 the Solaris Volume Manager. This Volume Manager detects CD-ROMs and automatically mounts them to `/cdrom/cdrom0`.

Solaris Voume 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 successfull, 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.

Without Volume Management, CD-ROMs must be explicitly mounted and ejected (example for SCSI-ID 6):
  - `mkdir /tmp/cdrom`
  - `mount -F hsfs /dev/dsk/c0t6d0s0 /tmp/cdrom`
  - (Work with CD mounted on `/tmp/cdrom`)
  - `umount /tmp/cdrom`
  - `eject /dev/dsk/c0t6d0s0`

2.5.4 CD-R und CD-RW

Bundled with Solaris 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 filesystem or directory using `mkisofs`.

Look at this example for the exact parameters:

```
#! /bin/sh
# ./nfscd <PARAMETER1> <PARAMETER2>
#       PARAMETER1 = Application ID
#       PARAMETER2 = from where
#
# ./nfscd 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 Tapes

Tape drives are not under the control of Volume Management. They 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 OE is running. After running `devfsadm` (Solaris 8 OE) or `drvconfig;tapes;` `devlinks` with previous Solaris versions you will find some logical links in the `/dev/rmt` directory. `0` describes the first tape drive found and `1` the second one etc.

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

Note: Tapes written with SGI systems often can only be red 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" a system to another one. Both systems must belong to the same hardware family. Cloning an Ultra™ 10 workstation to another Ultra 10 machine is OK, 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 framebuffer, 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.

Internal FC-AL disks (such as on Sun Blade 2000 workstations and Sun Fire™ 280R or V880 servers) may cause problems.

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 backup your data!) In our case `c0t0d0` (the bottom disk) acts as the source, `c0t1d0` (the upper disk) is the target.
- `dd if=/dev/rdsk/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 9GB disk, but the target disk is of 18GB capacity, 9GB would be wasted)
- `fsck -Y of=/dev/dsk/c0t1d0s0`   (Filesystem-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.

2.6.2 Cloning using a crossover cable

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

Target	Source

`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/rdsk/c0t0d0s2`
`            ibs=128k obs=8k" |`
`            dd of=/dev/rdsk/c0t0d0s2`
`            ibs=8k obs=128k`
(For 9GB wait about 40 minutes.)
`STOP-A`
`boot cdrom -s` (with bootable CD)
`fsck -Y /dev/rdsk/c0t0d0s0`
	Insert target-IP in `/.rhosts`
	`sync; sync`

2.7 Mirroring the boot disk with SDS

Solaris Logical Volume Manager, formerly known as Solstice DiskSuite software can configure disk subsystems for RAID 0, 1, 0+1 or 5. The most interesting feature for workstation installations is RAID 0: Mirroring the disk which holds the root filesystem to be prepared for disk failures.

You can find Solstice DiskSuite software for Solaris 8 platform on the second CD in the directory "EA". For the Solaris 7 OE or 2.6, Server Edition, look on the "Easy Access Server" CD.

Let's assume a workstation has two internal disks, `c0t0d0` und `c0t1d0`. Both disks should be partitioned exactly in 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").

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 filesystem:

```
  # 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 is because 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 expected on the failed disk have to be unconfigured:
  `metadb -d -f  c0t0d0s7`
• After the next reboot, partition the replacement disk exactly in the same way as 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.

Hint: If you have a two-disk system, create three State Database Replicas on one disk, four on the other one. Statistically there is a 50% chance to successfully reboot after a failure without using `metadb` manually.

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/rdsk/c1d0s2 /usr           ufs     1       yes     -
/dev/dsk/c0t0d0s1          -       -      swap    -       no      -
/usr/local/swapfile.1     -       -      swap    -       no      -
...
swap      -       /tmp    tmpfs    -       yes     -
```

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

You may define multiple swap filesystems. 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.

An example:

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

creates additional 6GB of virtual memory (given that all partitions have a size of 2GB). The Solaris OE 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 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 most happily follow all the instructions printed on the screen. There are only a few pit-
falls 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 NIS server to respond. It is
   better to configure NIS or DNS after the basic installation of the network (see one of the following chapters) has
   been completed.
– Installation cluster. You should always choose the "Entire Plus OEM Software"-cluster. Otherwise, only drivers
   for hardware present in the system would be installed. To activate hardware added later, additional drivers have
   to be subsequently installed from the CD.
– Disk partitioning: To obtain the chance to partition the disk according to the recommendations given in the chap-
   ter "Partitioning" (root, swap and one small Solstice DiskSuite slice), choose "custom", not "default" when
   prompted for an installation method.

Starting with the Solaris 7 OE, the Installation CD comes with a pre-patched Solaris 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 8 02/02 OE) onto the Solaris 8
07/01 OE upgrades the Solaris patches (and only those, 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
http://www.sun.com/bigadmin and should be installed after every Solaris installation.

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

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 US stock ticker symbol (or a fantasy symbol) of the
origin company, 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, `pkgrm <Package-Name>` removes a package.

Example: You have just changed the framebuffer 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 will typically be empty if you had done a plain default Solaris installation when the Creator3D, not the Elite3D was present in the system. Drivers for the Elite3D can be found on the Solaris CD, subdirectory `s0/Solaris/Product.`

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

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

4.3 Patches

A patch number contains of two parts. The six-digit patch number, followed by a two-digit version number, e.g.: 123456-78.

Prior to the Solaris 8 OE, patches were supplied in `*.tar.Z` format. Solaris 8 patches are packaged as `*.zip` archives. For installing the 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 fall-back 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, 24bit 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 it. The last line should now look like this:
```
  :0 Local  local_uid@concole   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, windows can not. The display running on `ifb0` is called
`:0.0,` the display on `ifb1 :0.1`

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

5.3.2 Xinerama

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

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

Xinerama allows to 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:

```
:0 Local  local_uid@concole   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 OK, since each window can define its own individual color depth. Some frame buffers only have one color table. This may cause colormap flashing with some applications. A workaround for this is to start CDE itself with 24bit color depth. On PGX, you have to switch the framebuffer to 24bit with `fbconfig -depth 24`

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

```
:0 Local  local_uid@concole   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 OE, `fbconfig` is the unified interface for configuring frame buffers. Prior to Solaris 8 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

The parameters are exactly the same. So 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 etc.) 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 ogl_install_check

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

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

6.3 Patch matrix for graphics

There is a small set of patches which are important for each installation where a frame buffer is involved. The collection depends on the Solaris 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 OE**	**Solaris 8 OE**	**Solaris 7 OE**
Expert3D (lite)	112540	108576	108787
XVR-500	112540	108576	not supported
XVR-1000			not supported
Elite3D		108604	106144
- and XFB			106148
- and VIS		109872	106147
Creator3D		108605	106145

Graphics Patch Matrix			
- and XFB			106148
- and VIS		109872	106147
OpenGL 1.2.2		32 Bit: 111993	32 Bit: 111993
		64 Bit: 111994	64 Bit: 111994
OpenGL 1.2.3		32 Bit: 112628	32 Bit: 112628
		64 Bit: 112629	64 Bit: 112629
XSun		108652	108376
PGX32 /64		109154	107716

7 Administration

7.1 User creation

The best tool for infrequent creation of user accounts is `admintool`. Beginning with Solaris 9 platform, 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 3^rd^-party software tries to use `/home/xyz`. Change the defaults or shut down the automounter to avoid any problems.

7.2 Serial Ports

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

7.3 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 installation starts the sendmail-Demon. This demon writes error messages if no alias with a dot in its name is defined for the primary hostname. 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
```

7.4 Printers

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

8 Network connections

8.1 /etc-files

Solaris software stores network configuration files in the directory `/etc`. All configuration data like IP-addresses, gateways etc. can be defined and changed at run-time level. However, only those definitions stored in `/etc` files will survive the next reboot.

The hostname can be found in three files:
- `/etc/nodename`
- `/etc/hostname.<interface-name>`
- `/etc/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 does not necessarily need to be 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:

`le`	Ethernet on older SPARC® systems
`hme`	FastEthernet on most UltraSPARC® systems
`eri`	FastEthernet on most UltraSPARC-III systems
`qfe`	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 `u60`, the `/etc/inet/hosts` looks like this:

```
127.0.0.1              localhost
157.168.34.12          u60   u60.domain.com    loghost
191.133.23.10          u60extern
```

The identity of the system in the corporate LAN using the address `157.168.34.12` should be `u60`. Accordingly, the file `hostname.hme0` must contain the name `u60`.

`u60extern` is an address assigned by an internet service provider: `191.133.23.10`. That cable is connected to the first QFE port, `qfe0`. So the file `hostname.qfe0` contains the string `u60extern`.

`qfe1-3` are currently not used. Do not create `hostname.*` files for unused interfaces to avoid Solaris 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 hostname.

8.2 Virtual Interfaces

Occasionally, it may be useful to define two IP addresses for a system with only one physical network interface. A demo workstation for use at a trade show, for example, that 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` und `hme0:1`. So if you write `u60extern` 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 run time. Let's go back to our Ultra 60 workstation example again, but assume that this time `u60extern` 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`.

The equivalent run time 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`.

8.5 NIS-Client

– First enter the name of the NIS domain into `/etc/defaultdomain` (case-sensitive)

– Then define the NIS server in `/etc/inet/hosts`

– Start `ypinit -c` and type in the name of the NIS server

– `/usr/lib/netsvc/yp/ypstart` or reboot

Now you have to tell the Solaris OE for which tables (host names, users etc.) 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 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>`. Solaris 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 run time (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 way: `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 installation tool are asked again (IP address, gateway, netmask etc.).

8.9 10 or 100MBit, 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 = 10MBit, 1 = 100MBit
- ndd /dev/eri link_mode:    0 = half-duplex, 1 = full-duplex
```

The Solaris 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 10MBit half-duplex:
```
ndd -set /dev/eri adv_autoneg_cap 0
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
```

Shut down the interface with `ifconfig unplumb` and restart it with `ifconfig plumb` 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 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 OE does not start the NFS server demon 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 demons. 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 to other popular UNIX systems

The file `/etc/dfs/dfstab` is a Solaris OE specialty. Other popular UNIX systems (Linux, AIX...) 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. It will be very helpful in large networks with lots of hosts and users 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 behaviour described above. `/net` is automatically extended with the proper hostname 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/SUNWspro/bin` directory, common tools like `make` in `/usr/ccs/bin`.

10.2 GNU-Tools

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

10.3 Sun PCi 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 & codepage)
    setenv KBTYPE=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, interupt the following creation of a user-specific image and edit the `sunpci.ini` file to point to the global image (`/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 1MB 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 workstation installation by calling `/opt/Office60/program/soffice`.

10.5 Netscape Navigator™ Software

Beginning with Solaris 8 platform, 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 3 Hints

- C-Shell: Define `set filec` and `set history=100` in `~/.cshrc`. The shell will complete filenames and directories once you hit the ESC key. `history` lists the last 100 commands. With `!13` command #13 can be re-issued. `!$` can be used as a variable for the last word of the former command, like in: `mkdir /ex-port/home/demo/test`, then `cd !$`.
  Beginning with the Solaris 8 platform, 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 OE. `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 etc.

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.

11.3 Analysis of runtime data

11.3.1 time/ptime

One of the simpler run time 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 software. `/usr/usb/ps -aux` shows similar results, but not sorted and without auto-update.

`ps -L -p <process-ID>` lists if a process has threads and 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:

free:	Free memory (until Solaris 7 near zero because of disk cache used. Beginning with Solaris 8 the correct size of free memory)
page pi/po:	Paging activity. The system does not swap until I/O traffic on the swap device is reported!
cpu us:	User-time. This percentage is consumed by applications
cpu sy:	System-time. This percentage is consumed by Solaris itself
cpu id:	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 Kbytes/s. A "svc_t" bigger than 100 indicates that the disk is spending more time with head positioning than data transfer.

11.3.5 sdtperfmeter, sdtprocess

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

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

11.3.6 32 or 64 bit?

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

11.4 Debugging

11.4.1 Truss

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

11.4.2 pstack, pmap etc.

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.

12 Books and Web Sites

12.1 Books

A valuable collection of Hints&Tricks for UNIX:

```
Jerry Peek, Tim O'Reilly, Mike Loukides "UNIX Power Tools"
ISBN 1565922603
```

A detailed view on all UNIX-variants incl. Solaris software 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	Homepage of Sun Microsystems
http://www.sun.com/bigadmin	Everything a sys admin needs
http://www.sun.com/blueprints	Blueprints on various topics, monthly additions
http://docs.sun.com	All documentation for Sun products online
http://www.sunfreeware.com	All Freeware-Tools for Solaris, precompiled
http://www.sunhelp.org	FAQ portal, other information