


-
-
-
-
-

SCO® UNIX® Operating System Maintenance Supplement v4.2



-
-
-
-

© 1983-1993 The Santa Cruz Operation, Inc.
© 1980-1993 Microsoft Corporation.
© 1989-1993 UNIX System Laboratories, Inc.
All Rights Reserved.

No part of this publication may be reproduced, transmitted, stored in a retrieval system, nor translated into any human or computer language, in any form or by any means, electronic, mechanical, magnetic, optical, chemical, manual, or otherwise, without the prior written permission of the copyright owner, The Santa Cruz Operation, Inc., 400 Encinal, Santa Cruz, California, 95060, U.S.A. Copyright infringement is a serious matter under the United States and foreign Copyright Laws.

The copyrighted software that accompanies this manual is licensed to the End User only for use in strict accordance with the End User License Agreement, which should be read carefully before commencing use of the software. Information in this document is subject to change without notice and does not represent a commitment on the part of The Santa Cruz Operation, Inc.

SCO UNIX Software is commercial computer software and, together with any related documentation, is subject to the restrictions on U.S. Government use as set forth below.

If this procurement is for a DOD agency, the following DFAR Restricted Rights Legend applies:

RESTRICTED RIGHTS LEGEND: Use, duplication or disclosure by the Government is subject to restrictions as set forth in subparagraph (c)(1)(ii) of rights in Technical Data and Computer Software Clause at DFARS 252.227-7013. Contractor/Manufacturer is The Santa Cruz Operation, Inc., 400 Encinal Street, Santa Cruz, CA 95060.

If this procurement is for a civilian government agency, the following FAR Restricted Rights Legend applies:

RESTRICTED RIGHTS LEGEND: This computer software is submitted with restricted rights under Government Contract No. _____ (and Subcontract No. _____, if appropriate). It may not be used, reproduced, or disclosed by the Government except as provided in Paragraph (g)(3)(i) of FAR Clause 52.227-14 or as otherwise expressly stated in the contract. Contractor/Manufacturer is The Santa Cruz Operation, Inc., 400 Encinal Street, Santa Cruz, CA 95060.

SCO, Open Desktop, The Santa Cruz Operation, the Open Desktop logo, and the SCO logo are registered trademarks of The Santa Cruz Operation, Inc. in the USA and other countries.

All other brand and product names are or may be trademarks of, and are used to identify products or services of, their respective owners.

SCO NFS was developed by Lachman Technology, Inc. (LTI)
SCO NFS is derived from LTI SYSTEM V NFS, a joint development of LTI and Sun Microsystems.
SCO TCP/IP is derived from Lachman Technology, Inc.
SYSTEM V STREAMS TCP/IP, a joint development of LTI and Convergent Technologies.

Date: 1 February 1993
Document Version: 3.2.4B

Introduction **1**

How to use these Notes	1
What you need to know	2
Disk space requirement	3

Chapter 1 ***Installing MSv4.2*** **5**

Installation instructions	5
Installing another package after MSv4.2 has been installed	6
Installing the TCP/IP Development System on a SCO UNIX system with a partially installed Link Kit	6
Replacement /boot on MSv4.2 distribution	6
Link Kit installation for sar(ADM) usage	6
Removal instructions	7

Chapter 2 ***Updating to MSv4.2*** **9**

Operating system notes	9
Object reuse	9
New default value for ULIMIT	9
New MODE_SELECT kernel parameter for slow printing	9
Additional options to mkdev parallel	10
Additional option to mkdev pty	10
corex(C) converts new-style core files	10
Multi-volume CPIO backups with SCSI tapes	10
New ORTSFL flag for hardware flow control	10
getserno(C) prints binary file serial number	11
mail(C) now makes all non-printable characters visible	11
Updates to tar(C)	11

Reducing interrupt trigger level to stop data loss with FIFO buffers	12
Slow performance due to insufficient cache	12
Increased MAXSC value improves performance	14
New maximum value for NOFILES kernel parameter	14
Using the audit reduction program with a high NOFILES value	14
Additional protocol for UUCP	15
Additional "cache" option for boot	15
Documentation updates	15
User's Reference	15
System Administrator's Reference	16

Chapter 3

Updates in MSv4.1 17

Operating system notes	17
scsibadblk(ADM)	17
ttyhog	17
Additional option to uname(C)	17
New option gives extra 30% compression using compress(C)	17
Patching object modules inside archives	18
Using the dd(C) command with Exabyte tapes	18
SCO Xsight Release 2.2.0 Supplement diskette	19
Everex and Tandberg tape controllers (ISA and EISA only)	19
IBM external floppy drive support	19
/etc/shadow and VP/ix vpixadm (ADM)	19
Running vtp, mscreen and xterm simultaneously	20
Restoring user accounts from Version 2.0 to Version 4.0 and later of the operating system	20
Possible password corruption when upgrading from Version 2.0 to Version 4.0	21
New version of /shlib/protlib_s	23
Installation security defaults	23

Documentation updates	23
User's Reference	23
System Administrator's Reference	24
System Administrator's Guide	25
Release Notes (or Release Notes Addendum)	25

Introduction

These notes contain installation instructions for Maintenance Supplement v4.2 for SCO® UNIX® System V Release 3.2 Operating System Version 4.0 (referred to as MSv4.2). Also included is information on the operating system software updates in this supplement, and a list of documentation updates.

Install MSv4.2 on a system running either SCO UNIX Operating System Version 4.0 or Version 4.1 to increment the operating system to a full SCO UNIX Operating System Version 4.2.

Installing this Maintenance Supplement:

- Updates your operating system software from Version 4.0 or 4.1 to Version 4.2.

The updates include bug fixes in the software and new features which are discussed in these notes.

- Updates your operating system online documentation from Version 4.0 or 4.1 to Version 4.2.

The updates include revised and new manual pages which are provided in the software for online usage.

How to use these notes

These notes are distributed in this printed format and also are provided online in the software. To access the online version type the following:

```
man msv4.2notes
```

Notes sections to read:

- **If you have the SCO UNIX Operating System Version 4.0 running on your system**, then having installed this supplement, your resulting operating system will contain the updates from the Maintenance Supplement Version 4.1 in addition to the updates in this supplement, Version 4.2.

In this case, read both chapters 2 and 3 as they contain details of the software and documentation updates for your current system.

- **If you have the SCO UNIX Operating System Version 4.1 running on your system**, then installing this supplement will add only the Version 4.2 software and documentation updates to your operating system.

In this case, you only need to read chapter 2 as it contains details of the software and documentation updates for this release. Chapter 3 contains information which you already received with the Version 4.1 supplement.

As part of the documentation updates in this release, Reference Manual pages that have been modified are listed in these notes. These reference pages have been included in the software so your system will have the most recent version once you have installed the supplement.

What you need to know

To install and make use of this supplement, you must be running either of the following on your system:

- SCO UNIX System V Release 3.2 Operating System Version 4.0.
- SCO UNIX System V Release 3.2 Operating System Version 4.0 with Maintenance Supplement v4.1.

After installation, the version of the operating system will be incremented to Release 3.2 Version 4.2.

NOTE This supplement should not be used with SCO® Open Desktop® / SCO® Open Server™ products. If you are running SCO Open Desktop 2.0 on your system, contact your supplier for the relevant product.

When installing the Maintenance Supplement with other SCO UNIX products, please do so in the order described below. The Maintenance Supplement may be removed if necessary, so that the order of installation may be adhered to. (Note that it must be fully removed before it can be reinstalled.)

SCO® MPX™ Release 3.0

WARNING This Maintenance Supplement does not work with SCO MPX Release 2.0. If you are running this Release 2.0 on your system, contact your supplier for the relevant SCO MPX 3.0 product before proceeding with this installation.

If the installation detects SCO MPX Release 2.0 on your system, it will be removed and MSv4.2 installed. Having installed MSv4.2, you will be prompted to insert and install your new SCO MPX Release 3.0 distribution.

Alternatively, you could manually remove SCO MPX Release 2.0, install this MSv4.2 supplement and then install SCO MPX 3.0. For more information, see the documentation provided with SCO MPX.

SCO® NFS®

This Maintenance Supplement should be installed *before* SCO NFS Release 1.2.0 (if this is not already installed) or 1.2.1.

SCO® TCP/IP

This Maintenance Supplement should be installed *before* SCO TCP/IP Release 1.2.0 (if this not already installed) or 1.2.1.

IBM EFS

This Maintenance Supplement should be installed *after* the IBM Enhanced Feature Supplement Release 1.0.0.

SCO UNIX System V Release 3.2 Development System Version 4x

The Maintenance Supplement may be installed *before* or *after* the Development System. If you have not already installed your Development System, then we recommend you install this MSv4.2 product first as this ensures an error-free installation environment.

You should always remove products in reverse order to that in which they were installed. This is because the Maintenance Supplement restores operating system files, including configuration information, to the state they were in before the Maintenance Supplement was installed. This means that if products are not removed in reverse order, inconsistencies may occur.

Disk space requirement

Approximately 15 MBytes (30 512-KByte blocks) are required for the MSv4.2 product, including the file backup procedure performed during the installation.

Ensure that you have the necessary disk space on your system before proceeding to install this supplement.

For more information on archiving the backed-up files, see the section, "Installation instructions" of chapter 1. This makes more space available after the installation.

Chapter 1

Installing MSv4.2

This chapter provides instructions for installing and removing this Maintenance Supplement v4.2 on your SCO system.

Installation instructions

To install MSv4.2 for SCO UNIX Release 3.2 Operating System Version 4.0 or 4.1, follow these steps:

1. In single user mode, type **custom**(ADM), then press (Return).

NOTE You cannot use the **-m** option of the **custom** command to specify the floppy drive to be used. You must use the same drive that was used to install the original Version 4.0 system, which must be drive 0.

2. Select "Install," then "A New Product," then "Entire Product," pressing (Return) after each selection.
3. Now follow the prompts to continue with the installation. **custom** will select the packages to be installed, depending on the packages that are already present on your system. Appropriate installation messages will be displayed on the console as the process progresses.

If the Link Kit is partially installed, or not installed at all, you are notified of this during installation, at which point you have the option to fully install the Link Kit. You are recommended to answer "Yes" to this prompt because the Maintenance Supplement contains updates to the Link Kit. Also, you will not have the option to create a new kernel at the end of the installation procedure unless the Link Kit is fully installed.

During installation, the Version 4.0 or 4.1 kernel files are automatically saved before being updated by the Maintenance Supplement. You have the option to move these files to an archive medium, thus freeing up space on your hard disk. If you do not archive them, they are kept in the */usr/lib/custom/pre-v4.2* directory. If you want to remove the Maintenance Supplement at some later time, you will need these files, so keep them safe.

You can archive the files after installation by changing to the */usr/lib/custom/pre-v4.2* directory, and running the following command:

```
find ./ -print | cpio -ocBO dev -M "Insert next volume"
```

During installation, a check is made to see if a parallel port device driver is installed; if one is not, you have the option to install one.

Installing another package after MSv4.2 has been installed

If you install a **custom** package on a system that includes files updated by the Maintenance Supplement, you will be prompted to insert the Maintenance Supplement volumes again. Updates to the package are contained on the Maintenance Supplement volumes.

Installing the TCPIP Development System on a SCO UNIX system with a partially installed Link Kit

The Maintenance Supplement includes a new version of `/usr/include/sys/wait.h`. This is needed if you intend to use the TCP/IP Development System. If you complete the MSv4.2 installation without updating the UNIX Link Kit (the UNIX Link Kit was not already installed, and you choose not to install it during the Maintenance Supplement installation), then this file will need to be manually installed from volume UB04 or UB05 after the UNIX Development System is installed.

Replacement lboot on MSv4.2 distribution

For those machines that cannot boot the operating system, a replacement `/boot` is included on the first distribution volume of this Maintenance Supplement.

To use this version of `/boot`, reboot your system with UB1, the first Maintenance Supplement volume, in drive 0 (primary). When the "Boot" prompt is displayed, replace this floppy with the N1 floppy from the operating system distribution and continue with the normal OS installation. Your system should now boot normally.

Continue with the installation. When the "Add Additional Software" prompt is displayed, install the Maintenance Supplement.

NOTE You may need to use UB1 more than once, depending on whether the installation requires a reboot (that is, if you have non-standard tape/CD-ROM parameters).

Link Kit installation for sar(ADM) usage

If you need to use **sar** on your system, you need to have the Link Kit installed so the updates to **sar** in this release can be included in the rebuilt kernel.

If the Link Kit is not installed, on system startup, the following message will be displayed:

```
can't find MP-SYSINFO in this kernel
```

This means that `sar(ADM)` has been disabled until the kernel is rebuilt with the new `sar`.

Removal instructions

To remove the Maintenance Supplement, follow the instructions below.

1. In single user mode, type `custom`, then press (Return).
2. Select "Remove," then press (Return).
3. Use the <Down Arrow> key to select SCO UNIX Maintenance Supplement v4.2, then press (Return).
4. Select "ALL," then press (Return). (The Maintenance Supplement cannot be partially removed.)

The Maintenance Supplement is removed, and preserved pre-v4.2 system files are restored. The version of the operating system reverts to the pre-installation system which is either v4.0 or v4.1.

NOTE If your system has both MSv4.1 and MSv4.2 products installed and you wish to remove the maintenance products, then you must remove the MSv4.2 product first followed by the MSv4.1.

If the pre-v4.2 files were archived during installation of the Maintenance Supplement, you are prompted to insert the archive media.

When you are prompted to rebuild the kernel, you must answer "Yes" to return to the pre-v4.2 system. If you answer "No," you may have problems using any devices that were installed during or after installation of the Maintenance Supplement. For example, if you selected to install a parallel port during the installation, it will not work with your pre-v4.2 system unless you rebuild the kernel.

Chapter 2

Updating to MSv4.2

This chapter contains information relating to the operating system and documentation updates which are provided for Version 4.2 of the Maintenance Supplement.

Operating system notes

This section contains information about the updates to your operating system.

Object reuse

The Object Reuse requirements stated in the Department of Defense's *Trusted Computer System Evaluation Criteria* (also known as the "Orange Book") are satisfied in SCO UNIX by clearing unallocated disk and memory areas (object) before they are allocated to a new process (subject).

New default value for ULIMIT

ULIMIT specifies, in 512-byte blocks, the size of the largest file that an ordinary user may write. The default value for this parameter is now 2097151. Doubling this new default value does not exceed the absolute allowable limit.

New MODE_SELECT kernel parameter for slow printing

If you have a parallel printer which prints abnormally slowly, check that your configuration matches the information given in the "Installing printers" chapter in the *Hardware Configuration Guide*.

If your printer is still slow, that is, of the order of four seconds per line, your printer may be deselecting itself after receiving each line of text.

The parameter value of `MODE_SELECT` can be changed by executing `/etc/conf/cf.d/configure` and changing the appropriate option. `MODE_SELECT` is a new kernel parameter, which is defined in `mtune` file and has a default value of 1 which enables mode-select checking. Setting this value to 0 will disable mode-select checking, and hence remove the printing delay incurred by these checks.

Additional options to mkdev parallel

When using **mkdev parallel**, in addition to adding and removing parallel ports, you can now choose to view the current configuration or get more help for this **mkdev** session.

Also, the allocation of device nodes to driver ports is now set at configuration time and remains the same for the life cycle of the driver. In the event that a driver is removed, the device node remains unallocated until it is reallocated.

Additional option to mkdev pty

mkdev pty, which is used to create pseudo-ttys, has been updated to include a third option. This option provides information on the number of pseudo-ttys configured in the system.

corex(C) converts new-style core files

A new utility, **corex(C)**, converts the new-style *core* dump file to the old-style format for backwards compatibility. The new format enables debugging of the dump image by providing information on the process' address space (in particular, any attached shared memory areas), the state of the system at the time of the core dump and some additional information on the system.

This new format is fully documented in the revised **core(FP)** manual page in the Development System (Version 4.2).

Multi-volume CPIO backups with SCSI tapes

Having completed the first volume and inserted the next as prompted by the system, the system must be given time to initialize the tape drive before you press <Return> to continue the backup. This initialization can be indicated by the active drive light and /or by the resetting sound from the physical drive.

Appropriate error messages are displayed if you press <Return> before the initialization.

New ORTSFL flag for hardware flow control

In addition to **CTSFLOW** and **RTSFLOW** flags, **ORTSFL** is now available for **CTS** and **RTS** handshaking and provides hardware flow control. The available combinations are detailed in the **stty(C)** and **termio(M)** manual pages in this supplement.

getserno(C) prints binary file serial number

A new utility, `getserno(C)` prints the SCO serial number, if present, as branded on the named binary file. Appropriate error codes are returned depending on the status of the binary file.

mail(C) now makes all non-printable characters visible

When `mail(C)` either displays messages on a user's tty or when messages get piped through the `PAGER`, `mail(C)` makes all non-printable characters visible.

A new variable, `visible`, has been added to `mail(C)`. This variable is enabled by default. This can be disabled by executing the `unset visible` command either at the `"&"` prompt within the mail program for this session or by specifying this command in your `$HOME/.mailrc` file.

When the visible switch is on and the `mail(C)` output is going to a tty or through the `PAGER`, non-printable characters are transformed as follows:

Control characters (with the exception of tabs, newlines and formfeeds) are transformed into `^X` (`^` followed by `X`), where `X` is the key pressed with the `<Ctrl>` key (for example, `<Ctrl>A`, octal 0001, is transformed to `^A`). The `` character (octal 0177) is transformed to `^?`.

Non-ASCII characters between octal 0200 and 0237 are transformed into `M-^X`, where `X` is the control character specified by the seven low order bits (for example, octal 0201 would be transformed into `M-^A`). Octal 0200 is transformed into `M-^@`.

Characters between octal 0040 to 0176 and octal 0240 to 0377 inclusive, are considered printable and are not transformed.

Updates to tar(C)

This section lists the updates to the `tar(C)` utility:

- Using `tar(C)` with variable block-sized tape devices:
When using `tar(C)` on tape devices which support variable block sizes for reading and/or writing, the archives are created in 512-byte blocks by default. The 10 kilobytes value given in `/etc/default/tar` is not used.
- Using `tar(C)` on filesystems without long filename support:
You are reminded that the `-T` option of the `tar(C)` command may be used to truncate filenames on tar archives to 14 characters. This ensures that they may be read on old-style systems that do not support greater than 14 character-length filenames.

- Using `tar(C)` for archive extraction as *root*:

When extracting the contents of a tar archive as *root*, if there is a need to create directories, then the ownership and group identification of the directory is set to that of the first file on the archive which belongs within that directory.

Reducing interrupt trigger level to stop data loss with FIFO buffers

Serial adapters that use chips with FIFO buffers and allow 16-byte buffering of characters can lose characters if the interrupt routine dispatch time is slow. By default, the interrupt trigger level is set to 14. If this trigger level is not correctly set, data will be lost as incoming data will overwrite data which is still waiting to be transferred out.

You can detect if you are losing data by executing `sar(ADM)` with the `-g` option. `sar -g` detects chip overflows and displays error messages. A 0 error code means no overflow errors are detected. If overflow errors are detected, you should lower the trigger level to give the interrupt routine more time to respond. The lower the trigger level, the better the chance of copying out the characters before an overflow occurs in the adapter. This inherently reduces your buffering capacity.

This trigger level can be set in the file `/etc/conf/npack.d/sio/space.c`. One line appears for each minor device of the "sio" driver. Edit the appropriate line (determined by examining the minor number of the device node you wish to affect). Substitute the symbols `FCR_Rx1`, `FCR_Rx4`, `FCR_Rx8` or `FCR_Rx14` to set the trigger level respectively. A lower trigger level gives the kernel more time to handle the interrupt before further incoming characters are lost, but increases the overall interrupt burden on the CPU. Do not change the format or number of lines since the script `/usr/lib/mkdev/serial` depends on the format of this file.

Slow performance due to insufficient cache

Most motherboards which have external (motherboard) cache require a certain amount of cache to correspond to system memory. For instance, in the most typical design, 64Kbytes of cache are required for each 16Mbytes of RAM.

If the system has more RAM than its cache can handle, it generates a signal whenever memory outside the cachable area is accessed. This signal tells the motherboard cache not to pay attention to this particular memory access.

On most motherboards, this same signal causes the Intel 486 or Pentium internal cache to ignore that memory access. The result is that memory below a certain address is cached by both the internal and external caches, while memory above that address is not cached at all. This causes a striking difference in performance, up to a factor of eleven in certain tests. This problem is sometimes referred to as "anti-caching".

If you suspect this problem, enter the following shell script (save it as */tmp/times.sh*):

```
:
while :
do
    set `timex sh -c 'echo "for (i = 0; i < 10000; i++)" | nice -2 bc' 2>&1`
    echo "date` real=$2 user=$4 sys=$6"
done | tee /tmp/times.log
```

Run it for several hours while your system is being used for various normal tasks. If the longest "user" time is more than 1.5 times the shortest "user" time, your system probably has this problem. (It is normal for "real" time to vary a great deal.)

If running the script shows the problem, take the following steps:

1. Check BIOS setup. Many BIOS's can be set to cache only part of memory. Some are configured that way from the factory. The BIOS should normally be set to allow caching of all main memory.
2. Check system or motherboard documentation, or contact the hardware manufacturer. Make sure you have enough cache for your RAM. If you are unsure and cannot find the necessary information, upgrade the motherboard to the maximum amount of cache it supports. Having extra cache will not degrade performance.
3. While waiting for a cache upgrade to be installed, the problem may in some cases be partially alleviated by loading the operating system kernel entirely within the first 16MBytes of memory. To do this, add the string "mem=/L" to the end of the **DEFBOOTSTR** entry in */etc/default/boot*. You can also do this for a single boot session by entering it at the "Boot:" prompt as follows:

```
Boot
: defbootstr mem=/L
```

- Alternatively, you could temporarily restrict your system to only use memory which is cachable by the currently installed cache. For instance, if the system has 64MBytes RAM and 128Kbytes cache, it is likely that only the first 32MBytes is cachable. Until upgrade hardware arrives, you could set the system to use only the first 32MBytes of RAM:

Boot

: defbootstr mem=1m-32m

See **boot(HW)** for more information about “boot” keywords.

Increased MAXSC value improves performance

The default value for the **MAXSC** kernel tunable parameter has been increased from 1 to 8. This means that up to 8 pages are passed to the **swap()** routine for swapping, rather than just one.

The **swap()** routine has also been modified to determine if the disk adapter driver can handle scatter-gather requests, and if so, pass multiple pages to the driver in a scatter-gather buffer header. This improves performance on systems which swap heavily by using scatter-gather to swap multiple pages out at once, instead of swapping pages out synchronously one by one.

New maximum value for NOFILES kernel parameter

The **NOFILES** parameter represents the maximum number of files that can be held open by any one process at a time. In the previous release, the parameter values ranged from 60 (minimum and default) to 150 (allowable maximum). In this release, **NOFILES** can be tuned between 60 (minimum and default) and 11000 (allowable maximum). This allows large applications, like databases, that need to have a large number of files open simultaneously, to improve their performance. Increase this value if advised by your application vendor to do so.

If you need to modify this parameter, then you will most likely need to update a number of other kernel tunable parameters accordingly. For example, an increase in **NOFILES** requires a corresponding increase in the **NINODE** (number of allocated inode entries) and **NFILE** (number of allocated open file entries) parameters.

This parameter value can be changed by executing **/etc/conf/CF.d/configure** and selecting the appropriate option.

Using the audit reduction program with a high NOFILES value

If you configure your kernel with a high **NOFILES** value, the audit reduction program, **reduce(ADM)**, may be unable to allocate sufficient memory to process an audit session.

If this occurs, you will see a message similar to the following:

```
Error on malloc of proc/file table space
```

If you choose to execute this audit reduction program, ensure that the following memory requirement for **reduce** is satisfied:

```
memory required = (NPROC * NOFILES * 8) + 1 MByte
```

Additional protocol for UUCP

The *t* protocol has been added to the UUCP system. The UUCP system now supports *e*, *f*, *g*, *x* and *t* protocols.

The “*t*” protocol is for network connection to remote machines.

Additional “cache” option for boot

A new option, “*cache*”, has been added to **boot**(HW). This controls the i80486 and Pentium internal cache. The following flags are recognized:

- /n** The internal cache is off after the kernel is loaded. This may be necessary for some machines where there are problems with cache coherency (this occurs when DMA does not notify the internal cache that memory has been written to directly).
- /y** The internal cache is on after the kernel is loaded (this is the default). Machine performance will be enhanced if caching is enabled.
- /d** Flush cache code is disabled. Booting will take less time if flushing is disabled, however, this may cause some machines to fail to boot.
- /e** Flush cache code is enabled (this is the default).

Documentation updates

This section lists documentation updates which are either revisions to currently documented sections of the Versions 4.0 and 4.1 release or new items not described elsewhere in these notes.

User's Reference

The following manual pages have been updated and are supplied online, as part of the MSv4.2 distribution. The new manual pages to this release are marked with a dagger(†).

Table 2-1 Revised manual pages

Intro(C)	awk(C)	basename(C)	cd(C)
chgrp(C)	chmod(C)	clear(C)	compress(C)
corex(C)†	cpio(C)	cpset(C)†	crypt(C)
dd(C)	devnm(C)	df(C)	dirname(C)
disable(C)	du(C)	enable(C)	find(C)
getopts(C)	getsermo(C)†	join(C)	ksh(C)
lp(C)	mail(C)	man(C)	mcart(C)†
pcpio(C)	pg(C)	ps(C)	pstat(C)
sed(C)	sh(C)	sort(C)	stty(C)
tape(C)	tar(C)	test(C)	time(C)
tput(C)	uname(C)	uuencode(C)	console(M)
error(M)	idas(M)†	idld(M)†	messages(M)†
termio(M)	termios(M)		

System Administrator's Reference

The following manual pages have been updated and are supplied online, as part of the MSv4.2 distribution. The new pages to this release are marked with a dagger(+). The manual pages marked with a double-dagger(‡), **msv4.2notes(ADM)**, is the online version of these MSv4.2 notes and **relnotes(ADM)** is the online SCO UNIX Release 3.2 Operating System Version 4.2 Release Notes.

Table 2-2 Revised manual pages

bcheckrc(ADM)	cbackup(ADM)	custom(ADM)	deliver(ADM)
divvy(ADM)	dmesg(ADM)	fixperm(ADM)	fsck(ADM)
fstyp(ADM)	initcond(ADM)	lpadmin(ADM)	lpmove(ADM)†
memsize(ADM)†	mkfs(ADM)	mount(ADM)	msv4.2notes(ADM)‡
reduce(ADM)	relnotes(ADM)‡	repckman(ADM)†	sar(ADM)
sfmt(ADM)	shutdown(ADM)	submit(ADM)	sysadmsh(ADM)
timex(ADM)	authorize(F)†	btld(F)	default(F)†
devassign(F)†	files(F)†	group(F)	mapchan(F)
mdevice(F)	mmdftailor(F)	m SCSI(F)	null(F)
perms(F)†	prpw(F)†	purge(F)	sdevice(F)
systems(F)	ttys(F)†	boot(HW)	fd(HW)
hd(HW)	ramdisk(HW)	screen(HW)	scsi(HW)
tape(HW)			

Chapter 3

Updates in MSv4.1

This chapter contains information relating to the operating system and documentation updates which were provided for Release v4.1 of the Maintenance Supplement.

Operating system notes

This section contains information relating to the operating system and to certain items of supported hardware.

scsibadblk(ADM)

A new utility, **scsibadblk(ADM)**, is included with this Maintenance Supplement. The **scsibadblk** utility scans selected areas of a hard disk for “bad” blocks, which can then be reallocated.

ttyhog

Please note that **ttyhog** has been moved from *io.a/tty.o* to *kernel/space.c*, and takes its value from the tunable parameter **TTHOG**. **TTHOG** is the size of the raw queue of the tty driver; increasing its value may improve the serial I/O performance of a slow system. The value of **TTHOG** may be changed using **configure(ADM)**; it configures the current raw tty queue size.

Additional option to uname(C)

A new option, **-A**, has been added to **uname(C)**; it prints the activation state of the operating system. This value is the same as the license field displayed by **uname -X**.

New option gives extra 30% compression using compress(C)

compress(C) includes a new option, **-H**. Using this option can save approximately 30% more space than when the option is not used. **uncompress(C)** automatically detects when files have been compressed using the **-H** option, and processes them appropriately. Previous versions of **uncompress** do not recognize files compressed using the **-H** option.

Note that files compressed with this option are not portable to non-SCO systems or previous SCO OS releases.

file(C) reports regular compressed files as “compressed data,” and files that were compressed using **-H** as “LZH-compressed data.”

Patching object modules inside archives

You may encounter problems when using `_fst` to patch object modules, if the modules contain more than one section of the same type. Run `idld -r` (refer to `idld(M)`) on the object module before you attempt to patch it, so that multiple occurrences of any section of the file are combined into one.

For example, you may wish to change the kernel variable `noclistfreq`. This variable controls the kernel's behavior when it is running out of character processing blocks (clists). By default, `noclistfreq` is set to `-1`, which causes the kernel to complain at most once per minute. If `noclistfreq` is `0`, the kernel prints a warning message every time an attempt to allocate a clist fails. If `noclistfreq` is positive, it sets the number of clock ticks (normally 100Hz) between warning messages.

To make the warning print at most every 10 minutes, set `noclistfreq` to $(10 * 60 \text{ seconds/minute} * 100 \text{ ticks/second}) = 60000$. (Note that no warnings print unless the system is running out of clists, a condition that is normally corrected by increasing the kernel parameter `NCLIST`.)

```
cd /etc/conf/pack.d/kernel
idar xv io.a clist.o      # extract the object file
mv clist.o clist.o.orig # preserve original version
idld -r clist.o.orig -o clist.o # use idld to coalesce similar sections
_fst -w clist.o          # patch —
$d                        # set decimal radix
noclistfreq/W 60000     # patch in new value
$q                        # quit _fst
idar rv io.a clist.o     # insert updated object file into archive
```

Using the dd(C) command with Exabyte tapes

Do not use the `dd` command to put individual data files onto Exabyte tapes; extracting the files may cause extraneous characters to be appended to the original data. You can, however, use the `dd` command with Exabyte tapes to store and extract `tar(C)` or `cpio(C)` images.

SCO Xsight Release 2.2.0 Supplement diskette

Older versions of SCO Xsight Release 2.2.0 contain a Supplement diskette which was required for SCO UNIX System V Release 3.2 Operating System installations. You must not install this Supplement diskette on this release of SCO UNIX System V. This diskette contains older utilities and configuration files which, if installed, will lead to problems with these utilities as well as with other operating system utilities. The files that are overwritten by the Supplement diskette include:

```
/usr/bin/compress  
/usr/bin/tar  
/etc/termcap  
/usr/lib/mkdev/mouse
```

If this diskette is accidentally installed, you must first reinstall the **compress** utility. You can use **custom(ADM)** to do this. After reinstalling the **compress** utility, you must reinstall the remaining three files.

Everex and Tandberg tape controllers (ISA and EISA only)

Please note that specifications given in the *Release Notes* for either Everex or Tandberg tape controllers, apply to both types.

IBM external floppy drive support

The SCO UNIX Release 3.2v4 Operating System includes support for the IBM 1.2 megabyte 5.25 inch external floppy drive. To configure this device, edit */etc/conf/sdevice.d/f15* to change the "N" to a "Y", then relink the kernel. Refer to your hardware documentation for further information.

letcshadow and VP/ix vpixadm (ADM)

Removing the */etc/shadow* file when installing VP/ix (as recommended in the operating system *Release Notes*) causes the encrypted password strings to be moved from */etc/shadow* to */etc/passwd*. However, if you prefer passwords not to be stored in */etc/passwd* because of the associated reduction in security, you can do one of the following:

- Edit the *vpixadm* script, to remove the test for a single x character in the encrypted password field of */etc/passwd*. (There are two instances of this in the *vpixadm* script.) For example, change:

```
grep "$(_uname):x:" /etc/passwd /dev/null 2&1 && {
```

to:

```
grep "$(_uname):P:" /etc/passwd /dev/null 2&1 && {
```

- Whenever you add a user:
 1. Take the system into single user mode.
 2. Use **pwunconv**(ADM) to remove */etc/shadow*.
 3. Use **vpixadm** to add the user.
 4. Recreate */etc/shadow* using **pwconv**(ADM).
 5. Take the system into multiuser mode.

NOTE This problem is fixed with Support Level Supplement app342.

Running vtp, mscreen and xterm simultaneously

When running Microsoft LAN Manager for UNIX Systems Release 1.1.0 on Version 2.0 of the operating system, a conflict was observed between **vtp** and other utilities such as **mscreen**. To resolve this problem, it was necessary to add an **-e** flag to the **vtpsrv** entry in the file */usr/lib/lm/lmx.servers*; this flag sets the vtp server to expect enabled pseudo ttys.

On Release 3.2v4.0 of the operating system, the addition of the **-e** flag is no longer necessary. **vtp**, **mscreen** and **xterm** all use disabled pseudo ttys, so **vtpsrv** should not be run with the **-e** flag.

Restoring user accounts from Version 2.0 to Version 4.0 and later of the operating system

After performing a fresh installation of Release 3.2v4.0 of the operating system, (that is, if you choose not to perform an Upgrade Installation), user accounts may be restored from a previous Version 2.0 backup.

The **ap**(ADM) command has been extended to include two additional options, **-u**, and **-g**. The **-u** option allows account profiles from earlier SCO UNIX systems to be restored. The **-g** option causes **ap** to include group membership in the account profile information output. When restoring and updating, if group membership information is available, **ap** processes it automatically.

An amended version of **ap** is supplied online as part of this Maintenance Supplement. To restore users from a Version 2.0 backup, follow the instructions below.

NOTE This procedure can only be carried out by the super user, or by a user having the **auth** subsystem authorization, and both the **chown** and **execsuid** kernel authorizations.

1. Back up the home directories of the users on the earlier system, using **cpio** or **tar**. (Do not back up these files using absolute pathnames. For example, if your accounts are in */usr*, run your backup command from that directory, not from */*.)
2. Make a copy of */etc/passwd*, */etc/group*, and */tcb/files/auth* from the earlier system. For example, to archive the required files into the file *acctnt_file*, type:

```
cd /
tar cf acctnt_file etc/passwd etc/group tcb/files/auth
```

NOTE These files are not backed up with absolute pathnames.

3. On the new system, create a directory, change to the new directory and extract the saved files into it. Take care not to overwrite the */etc/passwd* and */etc/group* files on the new system. For example, to extract the files from *acctnt_file* type:

```
mkdir /tmp/old_accnts
cd /tmp/old_accnts
tar xf acctnt_file
```

The new directory should now contain the directories *etc* and *tcb*.

4. Edit the *passwd* file from the earlier system, to reflect the new location of the users' home directories (if these are different).
5. Run the **ap** command, as shown below, to update the system profile files with the new entries:

```
/tcb/bin/ap -vu 'pwd'
```

This will not overwrite existing entries; a warning message will be issued instead.

6. Restore the users' home directories to the desired location with their original permissions.

Possible password corruption when upgrading from Version 2.0 to Version 4.0

If your protected password database is not in a consistent state, you may encounter some password corruption when you upgrade from Version 2.0 to Version 4.0. The severity of the problem depends on the response you gave to the "Is your current system in relaxed security mode?" prompt.

- Answering "Yes" results in all accounts being given bogus passwords, and thus they cannot be logged into. This includes the root account (and hence system maintenance mode).

- Answering “No” results in encrypted passwords being held in */etc/passwd*, and */etc/shadow* not being created. In this case the system is still usable, but strict C2 security is compromised.

To avoid the problem, before upgrading from Version 2.0, run **tcback** to remove any lock files, and run **/etc/pwck** and fix any syntax errors reported. Then run **/tcb/bin/authck -p**, and use Accounts ⇄ User ⇄ Examine:Identity in **sysadmsh(ADM)** to set the type field to “individual” for any accounts reported as not having protected password database entries.

If the upgrade has already triggered the problem, but the system is not relaxed, the above procedure can be used to repair the *passwd* file and protected password database (note that **authck -p** will now offer to add missing protected password database entries for you, so running **sysadmsh(ADM)** is not necessary). **/etc/pwconv** should then be run to move the encrypted passwords from */etc/passwd* to */etc/shadow*.

If the upgrade has already triggered the problem and the system is relaxed, it is impossible to log in. To correct the problem, boot the system from the N1/N2 disks, paying attention to the following points:

1. Start a shell by typing **shell** at the prompt for keyboard type.
2. Run **fsck** to check the root filesystem.
3. Mount */dev/hd0root* on */mnt* by typing **mount /dev/hd0root /mnt**
4. Type the command:

```
/mnt/bin/chroot /mnt /bin/su root -c passwd root
to set a new root password.
```

5. Unmount */mnt* by typing **umount /mnt**
6. Reboot the system and enter system maintenance mode.

From this point, either new passwords can be set on all accounts, or */etc/passwd* and */tcb/files/auth/?.** can be restored from a backup of the pre-upgrade system. If you restore pre-upgrade passwords, you should run **/etc/pwck** and fix any syntax errors reported. You should then use Accounts ⇄ Default ⇄ System in **sysadmsh(ADM)** and change the following field to “Yes:”

‘Protected Password database is the master’

Finally, you should use System ⇄ Configure ⇄ Security in **sysadmsh(ADM)** and enter ‘traditional’ as the security defaults package.

Note that the root password reverts to its pre-upgrade setting.

New version of /shlib/protlib_s

A problem has been identified when either */etc/passwd* or */etc/group* is edited: ordinary users will appear to have no subsystem authorizations until one of the following takes place:

- The system is rebooted.
- A login is executed.
- A command requiring the **auth** subsystem authorization is run (for example, adding users, changing passwords, etc.).
- A command requiring any subsystem authorization is run by root (for example, printer administration).

This Maintenance Supplement includes a replacement version of */shlib/protlib_s*, which will correct the problem.

Installation security defaults

This Maintenance Supplement corrects the following two errors in the files that are read by **relax**(ADM) when security defaults are set:

- The variable **DIALUPPRINTER**, used to control dial-up printers, is not added to */etc/default/lpd*. This problem occurs with all security defaults.
- In the improved security defaults, the variable **REUSEUID** in */etc/default/login* is set to "Yes," allowing users to be removed and unretired. This is contrary to C2 security guidelines, which the improved defaults are designed to enforce.

This Maintenance Supplement includes new versions of the */tcblib/relax/*etc_def* files. To correct the problems with **DIALUPPRINTER** and **REUSEUID**, either run **relax** to reset the security defaults, or edit */etc/default/lpd* and */etc/default/login* to correct the entries manually.

Documentation updates

This section lists documentation updates which are either revisions to currently documented sections of the Version 4.0 release or new items not described elsewhere in these notes.

User's Reference

The following manual pages have been updated and are supplied online, as part of Release v4.1 of the Maintenance Supplement distribution:

- **chmod**(C)
- **cron**(C)

- **crontab**(C)
- **dos**(C)
- **echo**(C)
- **kill**(C)
- **tail**(C)
- **translate**(C)
- **vi**(C)
- **wait**(C)
- **getty**(M)
- **login**(M)
- **mapkey**(M)

System Administrator's Reference

The following manual pages have been updated and are supplied online, as part of Release v4.1 of the Maintenance Supplement distribution:

- **ap**(ADM)
- **custom**(ADM)
- **dparam**(ADM)
- **fixmog**(ADM)
- **fsck**(ADM)
- **fsname**(ADM)
- **mkfs**(ADM)
- **rmuser**(ADM)
- **unretire**(ADM)
- **gettydefs**(F)
- **issue**(F)
- **passwd**(F)
- **log**(HW)
- **screen**(HW)
- **tape**(HW)

There are also two new manual pages supplied online:

- **scsibadblk**(ADM)
- **shadow**(F)

System Administrator's Guide

- Chapter 4, "Default account configuration"

The section entitled "Subsystem authorizations" contains a table of authorizations and provinces. You are reminded that the **auth** authorization should only be assigned to the most trusted individuals, because they will be able to modify the root account.

The section entitled "Administering user accounts" contains a table of "System default security parameters." The C2 Features section of this table includes the security parameter "Users can be deleted." The entry against the improved security level should read "No" rather than "Yes"; users cannot be removed from a C2 system (although they can be retired).

- Chapter 6, "Adding multiport cards, memory and other bus cards"

The section entitled "Adding and configuring parallel ports" states that the system configures one parallel port automatically (*/dev/lp0*). This is incorrect; no parallel port is configured automatically on this version of the operating system. You must invoke **mkdev parallel** to configure a parallel port.

- Chapter 7, "Using printers"

The warning in the section entitled "Configuring a dialup printer" is no longer applicable, and should be ignored.

- Chapter 11, "Maintaining system security"

The section entitled "The sticky bit and directories" indicates that the sticky bit may not be placed on a directory by that directory's owner. This is incorrect; both the super user and the directory owner may set the sticky bit.

- Chapter 17, "Tuning system performance"

The section entitled "Tunable system parameter descriptions" contains a number of incorrect mtune values. (Refer to the file */etc/conf/cf.d/mtune* for correct tunable parameter ranges and values.)

Release Notes (or Release Notes Addendum)

- Chapter 3, "Before installing your software"

The section entitled "If you plan to install TCP/IP Release 1.1.3f" instructs you to use **cp(C)** to copy some smtp files. However, this will cause the files to lose their original ownerships. You are therefore advised to use the copy **-o** command instead, so that the original ownerships of these files are preserved. **copy(C)** is part of the Extended Utilities, which are installed as a **custom(ADM)** package.

The section entitled "Installing the SCO UNIX System V Release 3.2 Development System Version 4.0," describes a problem with the PERMS package of `custom(ADM)`. This problem has now been corrected.

Please note that the section entitled "Installing an older version of Xsight" applies both to SCO Xsight Release 2.2 and to SCO Xsight Release 2.3.

The number of xterms available with SCO Xsight Release 2.3 is initially restricted to 11. To increase this number, run `mkdev pty xnodes`. This command increases the maximum number of available xterms to equal the number of available pseudo ttys (subject to kernel parameters). Each xterm uses a master and a slave tty.

- Chapter 4, "New features and enhancements in this release"
Step 1. of the section entitled "Additional options for the TCP/IP protocol interface" contains an entry for the file `/etc/inetd.conf`. While this entry is correct for SCO TCP/IP Release 1.1.3, it should read as the following for Release 1.2:

```
uucp stream tcp nowait NOLUID /usr/lib/uucp/uucpd uucpd
```

The section entitled "Pseudo tty" implies that you may encounter problems running SCO MultiView Release 1.6.5 or SCO Portfolio Version 1.0 on Version 4.0 of the operating system. This is misleading; SCO Portfolio Version 2.0 does not function correctly, but SCO MultiView Release 1.6.5 and SCO Portfolio Version 1.0 will run as expected.

- Chapter 6, "Administering your system"
The section entitled "Sector remapping for bad sectors" states that several manufacturers do not do Automatic Write Remapping (AWRE). This is inaccurate; some drive models do not support AWRE, and in addition, some of those that do support AWRE must have the feature enabled. Note that `scsibadblk(ADM)` can turn on AWRE.



1 February 1993



AJ01302P002

62724