

AIX Version 7.2

*Commands Reference, Volume 6, v - z*

**IBM**



AIX Version 7.2

*Commands Reference, Volume 6, v - z*

**IBM**

**Note**

Before using this information and the product it supports, read the information in "Notices" on page 273.

This edition applies to AIX Version 7.2 and to all subsequent releases and modifications until otherwise indicated in new editions.

© **Copyright IBM Corporation 2015, 2017.**

US Government Users Restricted Rights – Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

---

# Contents

<b>About this document</b> . . . . .	<b>v</b>
Highlighting . . . . .	v
Case sensitivity in AIX . . . . .	v
ISO 9000. . . . .	v
Support for the single UNIX specification . . . . .	v

<b>V</b> . . . . .	<b>1</b>
vacation Command . . . . .	1
val Command (SCCS) . . . . .	3
varyoffvg Command. . . . .	4
varyonvg Command. . . . .	5
vc Command . . . . .	9
vgrind Command . . . . .	11
vi or vedit Command . . . . .	13
view Command . . . . .	32
vmh Command . . . . .	32
vmo Command . . . . .	34
vmstat Command . . . . .	46
vpdadd Command . . . . .	59
vpddel Command . . . . .	60
vsdata1st Command . . . . .	61
vsdchgserver Command . . . . .	63
vsdelnode Command . . . . .	65
vsdelvg Command . . . . .	66
vsdnode Command. . . . .	67
vsdsk1st Command . . . . .	69
vsdvg Command . . . . .	71
vsdvgts Command . . . . .	73

<b>W</b> . . . . .	<b>75</b>
w Command . . . . .	75
wait Command . . . . .	76
wall Command . . . . .	77
wallevent Command . . . . .	78
watch Command . . . . .	80
wc Command . . . . .	81
what Command . . . . .	83
whatis Command . . . . .	84
whatnow Command . . . . .	85
whereis Command . . . . .	89
which Command . . . . .	90
which_fileset Command . . . . .	91
who Command . . . . .	92
whoami Command . . . . .	95
whodo Command . . . . .	95
whois Command . . . . .	97
whom Command . . . . .	98
wlmassign command. . . . .	100
wlmcheck command . . . . .	103
wlmcntrl Command . . . . .	105
wlmstat Command . . . . .	107
wol command . . . . .	112
wparerr Command . . . . .	113
wparexec Command . . . . .	114
wpar_reg_script Command. . . . .	120

wparprnterr Command . . . . .	121
write Command . . . . .	122
writesrv Daemon . . . . .	126
wtmpfix Command . . . . .	127
wump Command . . . . .	128

<b>X</b> . . . . .	<b>131</b>
X Command. . . . .	131
x_add_fs_fpe Command. . . . .	141
x_add_nfs_fpe Command . . . . .	142
x_rm_fpe Command . . . . .	143
xargs Command . . . . .	144
xauth Command . . . . .	147
xclock Command . . . . .	150
xcmsdb Command . . . . .	152
xdm Command. . . . .	153
xfindproxy Command . . . . .	166
xfs Command . . . . .	167
xget Command. . . . .	169
xhost Command . . . . .	171
xinit Command. . . . .	172
xkbcomp Command . . . . .	174
xkbevd Daemon . . . . .	175
xkbprint Command . . . . .	176
xlock Command . . . . .	178
xlsfonts Command . . . . .	180
xmbind Command . . . . .	181
xmkmf Command . . . . .	182
xmwlm Command . . . . .	183
xmodem Command . . . . .	184
xmodmap Command. . . . .	186
xmpeek Command . . . . .	188
xmscheck Command . . . . .	189
xmtopas Command . . . . .	191
xntpd Daemon . . . . .	192
xntpdC Command . . . . .	195
xpr Command . . . . .	202
xpreview Command . . . . .	204
xprofiler Command . . . . .	207
xrdb Command . . . . .	208
xsend Command . . . . .	211
xset Command . . . . .	212
xsetroot Command . . . . .	215
xss Command . . . . .	216
xstr Command . . . . .	217
xterm Command . . . . .	219
xwd Command. . . . .	241
xwud Command . . . . .	242

<b>Y</b> . . . . .	<b>245</b>
yacc Command. . . . .	245
yes Command . . . . .	247
ypbind Daemon . . . . .	248
ypcat Command . . . . .	249
ypinit Command . . . . .	250

ypmatch Command . . . . .	252
yppasswd Command . . . . .	253
yppasswdd Daemon . . . . .	254
yppoll Command . . . . .	255
yppush Command. . . . .	256
ypserv Daemon . . . . .	257
ypset Command . . . . .	258
ypupdated Daemon . . . . .	259
ypwhich Command . . . . .	260
ypxfr Command . . . . .	261
<b>z . . . . .</b>	<b>265</b>
zcat Command . . . . .	265

zdump Command . . . . .	266
zic Command . . . . .	267

**Notices . . . . . 273**

Privacy policy considerations . . . . .	275
Trademarks . . . . .	275

**Index . . . . . 277**

---

## About this document

This document provides end users with complete detailed information about commands for the AIX<sup>®</sup> operating system. The commands are listed alphabetically and by category, and complete descriptions are given for commands and their available flags. If applicable, each command listing contains examples. This volume contains AIX commands that begin with the letters v through z. This publication is also available on the documentation CD that is shipped with the operating system.

---

## Highlighting

The following highlighting conventions are used in this document:

<b>Bold</b>	Identifies commands, subroutines, keywords, files, structures, directories, and other items whose names are predefined by the system. Bold highlighting also identifies graphical objects, such as buttons, labels, and icons that the you select.
<i>Italics</i>	Identifies parameters for actual names or values that you supply.
Monospace	Identifies examples of specific data values, examples of text similar to what you might see displayed, examples of portions of program code similar to what you might write as a programmer, messages from the system, or text that you must type.

---

## Case sensitivity in AIX

Everything in the AIX operating system is case sensitive, which means that it distinguishes between uppercase and lowercase letters. For example, you can use the **ls** command to list files. If you type **LS**, the system responds that the command is not found. Likewise, **FILEA**, **FiLea**, and **filea** are three distinct file names, even if they reside in the same directory. To avoid causing undesirable actions to be performed, always ensure that you use the correct case.

---

## ISO 9000

ISO 9000 registered quality systems were used in the development and manufacturing of this product.

---

## Support for the single UNIX specification

The AIX operating system is designed to support The Open Group's Single UNIX Specification Version 3 (UNIX 03) for portability of operating systems based on the UNIX operating system. Many new interfaces, and some current ones, have been added or enhanced to meet this specification. To determine the correct way to develop a UNIX 03 portable application, see The Open Group's UNIX 03 specification on The UNIX System website (<http://www.unix.org>).





---

## V

The following AIX commands begin with the with the letter *v*.

---

### vacation Command

#### Purpose

Returns a message to the sender that the mail recipient is on vacation.

#### Syntax

```
vacation [ { -I | User } ] | [ { -f Number [ Unit ] | User } ]
```

#### Description

The **vacation** command returns a message to the sender of a mail message to notify the sender that the recipient is on vacation. The intended use is in a **\$HOME/.forward** file that allows messages to come to you while also sending a message back to the sender.

**Note:** Sendmail version 8.9.3 and subsequent releases have a security enhancement that will ignore the **.forward** file if *either* of the following conditions exist:

- The **.forward** file has group or world writeable permissions
- Any of **.forward** file's parent directories have group or world writable permissions

If you think that the vacation program is not working because the **.forward** file is being ignored, check the permissions. If you must have group or world writeable permissions on any of the parent directories of the **.forward** file, then set the DontBlameSendmail option in the sendmail configuration file with the appropriate values.

The **vacation** command expects a **\$HOME/vacation.msg** file containing a message to be sent back to each sender. If this file does not exist, the **vacation** command looks for **/usr/share/lib/vacation.def**, a systemwide default vacation message file. It should be an entire message, including any desired headers, such as From or Subject. By default, this message is sent only once a week to each person who sends mail to you. Use the **-f** flag to change the frequency intervals at which the message is sent. The names of the people who send messages are kept in the files **\$HOME/vacation.pag** and **\$HOME/vacation.dir**. These files are created when the **vacation** command is initialized for your user ID using the **-I** (uppercase i) flag.

If the **-I** flag is not specified, the **vacation** command reads the first line from the standard input for a From line to determine the sender. If no text is available from standard input, the command returns an error message. All properly formatted incoming mail should have a From line. No message is sent if the From header line indicates that the message is from Postmaster, MAILER-DAEMON, or if the initial From line includes the string-REQUEST@ or if a Precedence: bulk or Precedence: junk line is included in the header.

#### Flags

Item	Description
<b>-I</b>	Initializes the <b>\$HOME/vacation.pag</b> and <b>\$HOME/vacation.dir</b> files. Execute the <b>vacation</b> command using this flag before you modify your <b>\$HOME/.forward</b> file.
<b>-fNumber [Unit ]</b>	Specifies the frequency interval at which the vacation message is sent. The <i>Number</i> parameter is an integer value and the <i>Unit</i> parameter specifies a time unit. The <i>Unit</i> parameter can be one of the following:
<b>s</b>	Seconds
<b>m</b>	Minutes
<b>h</b>	Hours
<b>d</b>	Days
<b>w</b>	Weeks

**Note:** The **-f** flag cannot be used with the **-I** flag.

## Examples

1. Before you use the **vacation** command to return a message to the sender saying that you are on vacation, you must initialize the **\$HOME/vacation.pag** and **\$HOME/vacation.dir** files. To initialize these files, type:

```
vacation -I
```

2. Modify the **.forward** file. For example, Mark types the following statement in the **.forward** file:  
mark,|"/usr/bin/vacation mark"

The sender receives the message that is in the **\$HOME/vacation.msg** file, or if the file does not exist, the default message found in the **/usr/share/lib/vacation.def** file. If neither of these files exist, no automatic replies are sent to the sender of the mail message and no error message is generated. If either of these files exist, the sender receives one vacation message from mark per week, regardless of how many messages are sent to mark from the sender.

3. If the following entry is contained in your **.forward** file,  
mark, |"/usr/bin/vacation -f10d mark"

The sender receives one vacation message from mark every ten days, regardless of how many messages are sent to mark from the sender.

4. To create a vacation message that is different from the default vacation message, create the file **\$HOME/vacation.msg** and add your message to this file. The following is an example of a vacation message:

```
From: mark@odin.valhalla (Mark Smith)
Subject: I am on vacation.
Delivered-By-The-Graces-Of: the Vacation program
I am on vacation until October 1. If you have something urgent,
please contact Jim Terry <terry@zeus.valhalla>.
--mark
```

5. To cancel the vacation message, remove the **.forward** file, **.vacation.dir** file, **.vacation.pag** file, and **.vacation.msg** file from your **\$HOME** (login) directory:

```
rm .forward .vacation.dir .vacation.pag .vacation.msg
```

## Files

Item	Description
<code>\$HOME/.forward</code>	Contains the names of people who you want your mail to be forwarded to.
<code>/usr/share/lib/vacation.def</code>	Contains the systemwide default <b>vacation</b> message.
<code>\$HOME/.vacation.dir</code>	Contains the names of people who have sent mail to you while the <b>vacation</b> command was being used.
<code>\$HOME/.vacation.msg</code>	Contains your personalized <b>vacation</b> message.
<code>\$HOME/.vacation.pag</code>	Contains the names of people who have sent mail to you while the <b>vacation</b> command was being used.
<code>/usr/bin/vacation</code>	Contains the <b>vacation</b> command.

#### Related information:

.forward file  
 Sending a vacation message notice

## val Command (SCCS)

### Purpose

Validates SCCS files.

### Syntax

```
val [ -s ] [ -rSID ] [ -mName ] [ -yType ] File ...
```

### Description

The **val** command reads the specified file to determine if it is a Source Code Control System (SCCS) file meeting the characteristics specified by the accompanying flags. If you specify a - (minus) for the *File* value, the **val** program reads standard input and interprets each line of standard input as **val** flags and the name of an SCCS file. An end-of-file character terminates input.

The **val** command displays messages to standard output for each file processed.

### Flags

Each flag or group of flags applies independently to each named file. The flags can appear in any order.

Item	Description
<code>-mName</code>	Compares the <i>Name</i> value with the SCCS <b>31</b> identification keyword in the specified file.
<code>-r SID</code>	Specifies the SID of the file to be validated. The SID must be valid and unambiguous.
<code>-s</code>	Suppresses the error message normally written to standard output.
<code>-yType</code>	Specifies a type to compare with the SCCS identification keyword in the specified file.

### Exit Status

The **val** command returns 0 if successful for all files; otherwise, it returns an 8-bit code that is a disjunction of the possible errors. It is interpreted as a bit string in which set bits (from left to right) are interpreted as follows:

Item	Description
0x80	Missing file argument.
0x40	Unknown or duplicate option.
0x20	Corrupted SCCS file.
0x10	Cannot open file or file not SCCS.
0x08	SID is invalid or ambiguous.
0x04	SID does not exist.
0x02	, y mismatch.
0x01	31, m mismatch.

**Note:** The **val** command can process two or more files on a given command line and can process multiple command lines (when reading standard input). In these cases, an aggregate code is returned; a logical OR of the codes generated for each command line and file processes.

## Example

To determine if file `s.test.c` is an SCCS text file, enter:

```
val -ytext s.test.c
```

### Related information:

List of SCCS commands

get command

## varyoffvg Command

### Purpose

Deactivates a volume group.

### Syntax

```
varyoffvg [ -s ] VolumeGroup
```

### Description

The **varyoffvg** command deactivates the volume group specified by the *VolumeGroup* parameter along with its associated logical volumes. The logical volumes first must be closed. For example, if the logical volume contains a file system, it must be unmounted.

To activate the volume group, use the **varyonvg** command.

**Note:** To use this command, you must either have root user authority or be a member of the **system** group.

You can use the System Management Interface Tool (SMIT) to run this command. To use SMIT, enter:

```
smit varyoffvg
```

### Note:

- A volume group that has a paging space volume on it cannot be varied off while the paging space is active. Before you deactivate a volume group with an active paging space volume, ensure that the paging space is not activated automatically at system initialization, and then reboot the system.
- The **varyoffvg** command discards any background space reclamation process that is running for the volume group. To identify whether a space reclamation is running, you can use the **lvmstat** command with **-r** option.

## Flag

Item	Description
------	-------------

-s	Puts the volume group into System Management mode, so that only logical volume commands can be used on the volume group. In this mode, no logical volume can be opened or accessed by users.
----	--

## Exit Status

This command returns the following exit values:

Item	Description
0	Successful completion.
>0	An error occurred.

## Security

**Attention RBAC users and Trusted AIX users:** This command can perform privileged operations. Only privileged users can run privileged operations. For more information about authorizations and privileges, see Privileged Command Database in *Security*. For a list of privileges and the authorizations associated with this command, see the **lssecattr** command or the **getcmdattr** subcommand.

## Examples

1. To deactivate volume group vg03, enter:  
varyoffvg vg03
2. To deactivate volume group vg02, but allow logical volume commands to continue to take effect, enter:

```
varyoffvg -s vg02
```

Logical volumes within the volume group cannot be opened, but logical volume commands continue to take effect.

## File

Item	Description
/usr/sbin/varyoffvg	Contains the <b>varyoffvg</b> command.

### Related reference:

“varyonvg Command”

### Related information:

Paging space concepts

RBAC

Trusted AIX

---

## varyonvg Command

### Purpose

Activates a volume group.

### Syntax

```
varyonvg [ -b ] [ -c ] [ -f ] [ -M ltgsize ] [ -n ] [ -p ] [ -r ] [ -s ] [ -t ] [ -u ] [ -k loc|rem ] [ -d ] [ -o ] [ -O ]  
volumegroup
```

## Description

The **varyonvg** command activates the volume group specified by the *volume group* parameter and all associated logical volumes. A volume group that is activated is available for use. When a volume group is activated, physical partitions are synchronized if they are not current. Physical volumes that are in the PVMISSING state and that have been replaced will be returned to the PVACTION state by the **varyonvg** command.

**Note:** If a physical volume is part of a dump device, the **varyonvg** command cannot return it to the PVACTION state. To run the command effectively, temporarily change the dump device.

A list of all physical volumes with their status is displayed to standard output whenever there is some discrepancy between the Device Configuration Database and the information stored in the Logical Volume Manager (LVM). The volume group may or may not be varied on. You must carefully examine the list and take proper action depending on each reported status to preserve your system integrity.

While varying on in concurrent mode, if the varyon process detects that there are logical volumes which are not previously known to the system, their definitions are imported. The permissions and ownership of the new device special files are duplicated to those of the volume group special file. If you have changed the permissions and/or ownership of the device special files of the logical volume on the node it was created, you will need to perform the same changes on this node.

**Restriction:** Classic Concurrent mode is not supported in AIX 5.3.

If the *volume group* cannot be varied on due to a loss of the majority of physical volumes, a list of all physical volumes with their status is displayed. To vary on the *volume group* in this situation, you will need to use the force option.

The **varyonvg** command fails to vary on the volume group if a majority of the physical volumes are not accessible (no Quorum). This condition is true even if the quorum checking is disabled. Disabling the quorum checking will only ensure that the volume group stays varied on even in the case of loss of quorum.

The *volume group* will not vary on if any physical volumes are in the PV\_MISSING state and the quorum checking is disabled. This condition is true even if a quorum of disks are available. To vary on in this situation either use the force option or set an environment variable MISSINGPV\_VARYON to TRUE (set this value in **/etc/environment** if the volume group needs to be varied with missing disks at the boot time).

In the above cases (using the force vary on option and using the MISSINGPV\_VARYON variable), you take full responsibility for the *volume group* integrity.

**Requirement:** To use this command, you must either have root user authority or be a member of the **system** group.

You can use the System Management Interface Tool (SMIT) to run this command. To use SMIT, enter:  
smit varyonvg

## Flags

Item	Description
-b	<p>Breaks disk reservations on disks locked as a result of a normal <b>varyonvg</b> command. Use this flag on a volume group that is already varied on.</p> <p><b>Notes:</b></p> <ul style="list-style-type: none"> <li>• This flag unlocks all disks in a given volume group.</li> <li>• The <b>-b</b> flag opens the disks in the volume group using <b>SC_FORCED_OPEN</b> flag. For SCSI and FC disks this forces open all LUNS on the target address that this disk resides on. Volume groups should therefore not share target addresses when using the <b>varyon -b</b> option.</li> <li>• The <b>-b</b> flag can cause a system to hang if used on a volume group that contains an active paging space.</li> </ul>
-c	<p>Varies the volume group on Enhanced Concurrent mode. This is only possible if the volume group is Concurrent Capable or Enhanced Concurrent Capable and the system has the PowerHA<sup>®</sup> SystemMirror<sup>®</sup> product loaded and available. If neither is true, the volume group fails the varyon operation.</p> <p><b>Requirement:</b> Enhanced Concurrent volume groups use Group Services. Group Services must be configured prior to activating a volume group in this mode.</p>
-d	<p>Allows data divergence. The <b>-d</b> flag only takes effect when you try to bring the volume group online while the cache at the opposite site might contain nonmirrored data updates and that cache is not accessible. If the <b>varyonvg</b> command detects that you might use back-level data and you do not specify the <b>-d</b> flag, the command fails with a meaningful error message.</p> <p>For more information about asynchronous mirroring of geographic LVM, see <i>Geographic Logical Volume Manager for PowerHA SystemMirror Enterprise Edition</i>.</p>
-f	<p>Allows a volume group to be made active that does not currently have a quorum of available disks. All disks that cannot be brought to an active state will be put in a removed state. At least one disk must be available for use in the volume group. The <b>-f</b> flag (used to override quorum loss) is ignored if the volume group has not lost quorum. If a disk is put into removed state, use the <b>chpv -v a PVname</b> command to bring the disk back to active state.</p>
-k loc   rem	<p>Keeps data from the local mirror copy or remote mirror copy. You can specify the following attributes with the <b>-k</b> flag:</p> <p><b>loc</b>      Retains the local mirror copy data. Local means local physical volumes and not primary site</p> <p><b>rem</b>      Retains the remote mirror copy data. Remote means remote physical volumes and not remote site.</p> <p>For more information about asynchronous mirroring of geographic LVM, see <i>Geographic Logical Volume Manager for PowerHA SystemMirror Enterprise Edition</i>.</p>
-M ltgsize	<p>Statically sets the <i>ltgsize</i> of the volume group. Valid values for <i>ltgsize</i> include 128K, 256K, 512K, 1M, 2M, 4M, 8M, 16M, 32M, and 128M. If any disk in the volume group is not configured with a maximum transfer of <i>ltgsize</i> or greater, the varyonvg command will fail.</p>
-n	<p>Disables the synchronization of the stale physical partitions within the <i>volume group</i> parameter.</p>
-o	<p>Allows using data from partitions that are stale in the copy you selected but fresh in the other copy. The <b>varyonvg</b> command fails if you specify the <b>-k</b> flag to preserve either local copy or remote copy in the data divergence case and the <b>varyonvg</b> command cannot preserve the complete copy because some partitions are not fresh in the local or remote copy that you selected. You can override the failure by specifying the <b>-o</b> flag to use data from partitions that are stale in the copy that you selected but fresh in the other copy. The <b>-o</b> flag is only valid with the <b>-k</b> flag.</p> <p>For more information about asynchronous mirroring of geographic LVM, see <i>Geographic Logical Volume Manager for PowerHA SystemMirror Enterprise Edition</i>.</p>
-p	<p>All physical volumes must be available to use the <b>varyonvg</b> command.</p>
-r	<p>Varies on the volume group in read-only mode. This mode prevents:</p> <ul style="list-style-type: none"> <li>• Write operations to logical volumes</li> <li>• LVM metadata updates</li> <li>• Stale partitions synchronization</li> </ul> <p><b>Restriction:</b> Mounting a JFS file system on a read-only logical volume is not supported.</p> <p><b>Restriction:</b> All LVM high-level commands that require the LVM metadata update will fail the request in this mode.</p>

Item	Description
-s	Makes the volume group available in System Management mode only. Logical volume commands can operate on the volume group, but no logical volumes can be opened for input or output. <b>Restriction:</b> Logical volume commands also cannot read or write to or from logical volumes in a volume group varied on with the <b>-s</b> flag. Logical volumes that attempt to write to a logical volume in a volume group varied on with the <b>-s</b> flag (such as <b>chvg</b> or <b>mklvcopy</b> ) may display error messages indicating that they were unable to write to and/or read from the logical volume.
-t	Checks the timestamps in the Device Configuration Database and the Logical Volume Manager. If there is a discrepancy in the timestamps, the <b>syncldvdm</b> command is issued to synchronize the database. <b>Tip:</b> This check is always done if the Volume Group is varied on in concurrent mode.
-u	Varies on a volume group, but leaves the disks that make up the volume group in an unlocked state. Use this flag as part of the initial varyon operation of a dormant volume group.
-O	Force a varyon operation of the volume group even if it is varied on in some other node. <b>Note:</b> In AIX 61 TL8 and later releases, if the volume group created is not allowed to varyon in non-concurrent mode in more than one node then the <b>varyonvg</b> command updates LVM metadata and ODM with the varyon state of the volume group. During varyon time, the <b>varyonvg</b> command reads this data and fails if the volume group is already varied on in another node. The <b>varyoffvg</b> command will reset the varyon state of the volume group during the varyoff time. If the system crashes before varying off the volume group or the volume group is forced off, the <b>varyonvg</b> command fails after you reboot the system. In this scenario, use <b>-O</b> flag to force a varyon operation of the volume group.

**Attention:** The base design of LVM assumes that only one initiator can access a volume group. The PowerHA SystemMirror product does work with LVM in order to synchronize multi-node accesses of a shared volume group. However, multi-initiator nodes can easily access a volume group with the **-b** and **-u** flags without the use of PowerHA SystemMirror. You must be aware that volume group status information might be compromised or inexplicably altered as a result of disk protect (locking) being bypassed with these two flags. If you use the **-b** and **-u** flags, data and status output cannot be guaranteed to be consistent.

## Exit Status

This command returns the following exit values:

Item	Description
0	Successful completion.
>0	An error occurred.

## Security

**Attention RBAC users and Trusted AIX users:** This command can perform privileged operations. Only privileged users can run privileged operations. For more information about authorizations and privileges, see Privileged Command Database in *Security*. For a list of privileges and the authorizations associated with this command, see the **lssecattr** command or the **getcmdattr** subcommand.

## Examples

- To activate volume group vg03, enter:  
varyonvg vg03
- To activate volume group vg03 without synchronizing partitions that are not current, enter:  
varyonvg -n vg03

## Files



Item	Description
/usr/sbin	Contains the <b>varyonvg</b> command directory.
/tmp	Stores the temporary files while the command is running.

**Related reference:**

“varyoffvg Command” on page 4

**Related information:**

chvg command

RBAC

Trusted AIX

## vc Command

### Purpose

Substitutes assigned values for identification keywords.

### Syntax

```
vc [ -a ] [ -t ] [ -s ] [ -cCharacter ] [ Keyword=Value ]...
```

### Description

The **vc** command copies lines from standard input to standard output. The flags and keywords on the command line and control statements in the input modify the resulting output. The **vc** command replaces user-declared keywords with the value assigned on the command line. Keywords can be replaced both in text and in control statements.

### Control Statements

A control statement is a single line beginning with a control character (the default control character is a : (colon)). Control statements provide conditional processing of the input. The allowable types of control statements are:

**:if** *Condition*

*Text*

Item	Description
<b>:end</b>	Writes all the lines between the <b>:if</b> statement and the matching <b>:end</b> to standard output only if the condition is true. You can nest <b>:if</b> and <b>:end</b> statements. However, once a condition is false, all remaining nested <b>:if</b> and <b>:end</b> statements are ignored. See the Condition Syntax section for the syntax of conditions and allowable operators.
<b>:dcl</b> <i>Keyword</i> , [ <i>Keyword</i> . . . ]	Declares specified keywords. All keywords must be declared.
<b>:asg</b> <i>Keyword=Value</i>	Assigns the specified value to the specified keyword. An <b>:asg</b> statement takes precedence over keyword assignment on the <b>vc</b> command line. A later <b>:asg</b> statement overrides all earlier assignments of the associated keyword. The keywords that are declared but not assigned <i>Values</i> , have null values.
<b>::</b> <i>Text</i>	Removes the two leading control characters, replaces keywords with their respective values, and then copies the line to standard output.
<b>:on</b> or <b>:off</b>	Turns on or off keyword replacement on all lines.
<b>:ctl</b> <i>Character</i>	Changes the control character to the <i>Character</i> value.
<b>:msg</b> <i>Message</i>	Writes a message to standard error output in the form: Message(n): message where n is number of the input line on which the message appeared.

<b>Item</b>	<b>Description</b>
<code>:err</code> <i>Message</i>	Writes an error message to standard error. The <code>vc</code> command stops processing and returns an exit value of 1. The error message is in the form:  ERROR: message ERROR: err statement on line n (vc15)

## Condition Syntax

The items and statements allowed are:

<b>condition</b>	::=OR statement ::=NOR statement
<b>OR statement</b>	::=AND statement ::=AND statement   OR statement
<b>AND statement</b>	::=expression ::=expression & AND statement
<b>expression</b>	::=( OR statement ) ::=value operator value
<b>operator value</b>	::= = or != or < or > ::= ASCII string ::= numeric string

The available condition operators and their meanings are:

<b>Item</b>	<b>Description</b>
=	Equal
!=	Not equal
&	AND
&	OR
>	Greater than
<	Less than
()	Used for logical groupings
NOT	May only occur immediately after the <i>if</i> , and when present, inverts the value of the entire condition.

The > and < (greater-than and less-than) operate only on unsigned integer values; for example, 012 > 12 is false. All other operators take strings as modifiers; for example, 012 != 12 is true. The precedence of the operators, from highest to lowest precedence, is as follows:

- = != > < (all of equal precedence)
- &
- &|

Parentheses can be used to alter the order of precedence.

Values must be separated from operators or parentheses by at least one blank or tab.

## Keyword Replacement

A keyword must begin and end with the same control character used in control statements. A keyword may be up to nine alphanumeric characters, where the first character must be alphabetic. Keyword values can be any ASCII string. A numeric keyword *Value* is an unsigned string of digits. Values cannot contain tabs or spaces.

## Flags

Item	Description
-a	Replaces keywords surrounded by control characters with their assigned value in all text lines (not just those beginning with two control characters).
-c <i>Character</i>	Uses the <i>Character</i> value as the control character. The <i>Character</i> parameter must specify an ASCII character.
-s	Does not display the warning messages normally displayed to standard error.
-t	Ignores all characters from the beginning of a line up to and including the first tab character for detecting a control statement. If the <b>vc</b> command finds a control character, it ignores all characters up to and including the tab.

## Exit Status

This command returns the following exit values:

Item	Description
0	Successful completion.
>0	An error occurred.

## Examples

- Examples of *Keyword=Value* assignments are:

```
numlines=4
prog=acctg
pass4=yes
```

The **vc** command removes all control characters and keywords from input text lines marked with two control characters as it writes the text to standard output.

- To prevent a control character from being interpreted, precede it with a backslash, as in the following example:

```
::the :prog: program includes several of the following\:
```

The **:prog:** keyword is replaced by its value, but the **\:** is passed to standard output as **:** (colon).

Input lines beginning with a **\** (backslash) followed by a control character are not control lines, and are copied to standard output without the backslash. However, the **vc** command writes lines beginning with a backslash and no following control character without any changes (including the initial backslash).

## File

Item	Description
<code>/usr/bin/vc</code>	Contains the <b>vc</b> command.

### Related information:

List of SCCS commands

---

## vgrind Command

### Purpose

Formats listings of programs that are easy to read.

### Syntax

```
vgrind [ -f ] [ -n ] [ -t ] [ -x ] [ -PPrintdev ] [ -TName ] [ - ] [ -dFile ] [ -h Header ] [ -lLanguage ] [ -sSize ] [ File ... ]
```

## Description

The **vgrind** command formats (grinds) the program sources specified by the *File* parameters in an easily readable style using the **troff** command. Comments are placed in italics, keywords in boldface, and the name of the current function is listed down the margin of each page as it is encountered.

The **vgrind** command runs in either filter mode or regular mode.

In filter mode, the **vgrind** command acts as a filter in a manner similar to the **tbl** command. Standard input is passed directly to standard output except for lines bracketed by the following **troff**-like macros:

Item	Description
<b>.vS</b>	Starts processing.
<b>.vE</b>	Ends processing.

The preceding lines are formatted according to the **vgrind** command conventions. The output from this filter can be passed to the **troff** command for output. There is no particular ordering with the **eqn** or **tbl** command.

In regular mode, the **vgrind** command accepts input files, processes them, and passes them in order to the **troff** command, the appropriate postprocessor, and then the printer.

In both modes, the **vgrind** command passes without converting lines, beginning with a decimal point.

The **vgrind** command supports only ASCII keywords defined in either the standard */usr/share/lib/vgrindefs* language definitions file or any alternately specified file by the **-d** flag.

## Flags

Item	Description
<b>-f</b>	Forces filter mode.
<b>-n</b>	Forces no keyword bolding.
<b>-t</b>	Causes formatted text to go to standard output.
<b>-x</b>	Outputs the index file in an easily readable format. The index file itself is produced whenever the <b>vgrind</b> command is run with the <b>index</b> file in the current directory. The index of function definitions can then be run off by running the <b>vgrind</b> command with the <b>-x</b> flag and the <i>File</i> parameter.
<b>-PPrintDev</b>	Sends the output to <i>Printdev</i> Printer using the <b>qprt</b> command. If this flag is not specified, the <b>PRINTER</b> environment variable is used. If the <b>PRINTER</b> environment variable is not set, the system default is used.
<b>-TName</b>	Creates output for a <b>troff</b> device as specified by the <i>Name</i> parameter. The output is sent through the appropriate postprocessor. The default is the <b>ibm3816</b> postprocessor.
<b>-</b>	Forces input to be taken from standard input (default if the <b>-f</b> flag is specified).
<b>-dFile</b>	Specifies an alternate language definitions file (default is the <i>/usr/share/lib/vgrindefs</i> file).
<b>-h Header</b>	Specifies a particular header to put on every output page (default is the file name). <b>Note:</b> A blank space is required after the <b>-h</b> flag before the <i>Header</i> variable.

Item	Description
<b>-lLanguage</b>	Specifies the language to use. Currently known languages are:
<b>c</b>	C (the default). Function names can be preceded on a line only by spaces, tabs, or an asterisk. The parenthetical options must also be on the same line.
<b>csH</b>	CSH.
<b>p</b>	PASCAL. Function names must be displayed on the same line as the <b>function</b> or <b>procedure</b> keywords.
<b>m</b>	MODEL. Function names must be displayed on the same line as the <b>isbeginproc</b> keyword phrase.
<b>sh</b>	SHELL.
<b>r</b>	RATFOR.
<b>mod2</b>	MODULA2.
<b>yacc</b>	YACC.
<b>isp</b>	ISP.
<b>I</b>	ICON.
<b>-s Size</b>	Specifies a point size to use on output (exactly the same as a <b>.ps</b> request).

## Files

Item	Description
<b>index</b>	Contains the file the where source for the index is created.
<b>/usr/bin/vgrind</b>	Contains the <b>vgrind</b> command.
<b>/usr/share/lib/tmac/tmac.vgrind</b>	Contains the macro package.
<b>/usr/share/lib/vfontedpr</b>	Contains the preprocessor.
<b>/usr/share/lib/vgrindefs</b>	Contains the language descriptions.

### Related information:

qprt command  
tbl command  
troff command  
vgrindefs file format

---

## vi or vedit Command

### Purpose

Edits files with a full-screen display.

### Syntax

```
{ vi | vedit } [ -l ] [ -R ] [ -tTag ] [ -v ] [ -wNumber ] [ -yNumber ] [ -r [ File ] ] [ { + | -c } { Subcommand } ] [ File ... ]
```

### Description

The **vi** command starts a full-screen editor based on the underlying **ex** editor. Therefore, **ex** subcommands can be used within the **vi** editor. The **vedit** command starts a version of the **vi** editor intended for beginners. In the **vedit** editor, the **report** option is set to 1, the **showmode** option is set, and the **novice** option is set, making it a line editor.

You start the **vi** editor by specifying the name of the file or files to be edited. If you supply more than one *File* parameter on the command line, the **vi** editor edits each file in the specified order. The **vi** editor

on an existing file displays the name of the file, the number of lines, and the number of characters at the bottom of the screen. In case of multibyte locales the number of characters need to be interpreted as the number of bytes.

Since the vi editor is a full-screen editor, you can edit text on a screen-by-screen basis. The vi editor makes a copy of the file you are editing in an edit buffer, and the contents of the file are not changed until you save the changes. The position of the cursor on the display screen indicates its position within the file, and the subcommands affect the file at the cursor position.

## vi Editor Limitations

The following list provides the maximum limits of the vi editor. These counts assume single-byte characters.

- 256 characters per global command list
- 2048 characters in a shell escape command
- 128 characters in a string-valued option
- 30 characters in a tag name
- 128 map macros with 2048 characters total
- 1,048,560 lines silently enforced
- The macro name and the macro text are limited to 100 characters.

**Note:** The vi editor supports a maximum of 2 GB edit buffer.

## vi Editing Modes

The vi editor operates in the following modes:

Item	Description
<b>command mode</b>	When you start the vi editor, it is in command mode. You can enter any subcommand except those designated for use only in the text input mode. The vi editor returns to command mode when subcommands and other modes end. Press the Esc key to cancel a subcommand.
<b>text-input mode</b>	You use the vi editor in this mode to add text. Enter text input mode with any of the following subcommands: the <b>a</b> subcommand, <b>A</b> subcommand, <b>i</b> subcommand, <b>I</b> subcommand, <b>o</b> subcommand, <b>O</b> subcommand, <b>cx</b> subcommands (where the <i>x</i> represents the scope of the subcommand), <b>C</b> subcommand, <b>s</b> subcommand, <b>S</b> subcommand, and <b>R</b> subcommand. After entering one of these subcommands, you can enter text into the editing buffer. To return to command mode, press the Esc key for normal exit or press Interrupt (the Ctrl-C key sequence) to end abnormally.
<b>last-line mode</b>	Subcommands with the prefix <b>:</b> (colon), <b>/</b> (slash), <b>?</b> (question mark), <b>!</b> (exclamation point), or <b>!!</b> (two exclamation points) read input on a line displayed at the bottom of the screen. When you enter the initial character, the vi editor places the cursor at the bottom of the screen, where you enter the remaining characters of the command. Press the Enter key to run the subcommand, or press Interrupt (the Ctrl-C key sequence) to cancel it. When the <b>!!</b> prefix is used, the cursor moves only after both exclamation points are entered. When you use the <b>:</b> prefix to enter the last-line mode, the vi editor gives special meaning to the following characters when they are used before commands that specify counts:  %        All lines regardless of cursor position  \$        Last line  .        Current line  <b>Note:</b> The history of last line mode subcommands can be navigated using the Up and Down Arrow keys.

## Customizing the vi Editor

You can customize the vi editor by:

- Setting vi editor options
- Defining macros
- Mapping keys
- Setting abbreviations

## Setting vi Editor Options

The following list describes the vi editor options you can change with the **set** command. The default setting for these options is **off**. If you turn on one of these toggle options, you can turn it off again by entering the word **no** before the option. If you want to discontinue the **autowrite** vi option, enter **noaw**, where **no** turns off the option and **aw** specifies the **autowrite** option.

**Note:** Do not include parentheses when entering vi options.

vi Option (Abbreviation)	Description
<b>autoindent (ai)</b>	Indents automatically in text input mode to the indentation of the previous line by using the spacing between tab stops specified by the <b>shiftwidth</b> option. The default is <b>noai</b> . To back the cursor up to the previous tab stop, press the Ctrl-D key sequence. This option is not in effect for global commands.
<b>autoprin (ap)</b>	Prints the current line after any command that changes the editing buffer. The default is <b>ap</b> . This option applies only to the last command in a sequence of commands on a single line and is not in effect for global commands.
<b>autowrite (aw)</b>	Writes the editing buffer to the file automatically before the <b>:n</b> subcommand, the <b>:ta</b> subcommand, the Ctrl-A, Ctrl -], and Ctrl -T key sequences, and the <b>!</b> subcommand if the editing buffer changed since the last <b>write</b> subcommand. The default is <b>noaw</b> .
<b>backtags (bt)</b>	Allows the Ctrl-T subcommand to return the file editing position to the location where the previous Ctrl-] subcommand was issued. If <b>nobacktags</b> is set, then Ctrl-T is the same as Ctrl-]. The default is <b>backtags</b> .
<b>beautifying text (bf)</b>	Prevents the user from entering control characters in the editing buffer during text entry (except for tab, new-line, and form-feed indicators). The default is <b>nobf</b> . This option applies to command input.
<b>closepunct (cp=)</b>	Handles a list of closing punctuation, especially when wrapping text ( <b>wraptyp</b> option). Precedes multicharacter punctuation with the number of characters; for example, <b>cp=3. .;}</b> . The <b>vi</b> command does not split closing punctuation when wrapping.
<b>directory (dir=)</b>	Displays the directory that contains the editing buffer. The default is <b>dir = /var/tmp</b> .
<b>edcompatible (ed)</b>	Retains <b>g</b> (global) and <b>c</b> (confirm) subcommand suffixes during multiple substitutions and causes the <b>r</b> (read) suffix to work like the <b>r</b> subcommand. The default is <b>noed</b> .
<b>exec (exec)</b>	If not set, ignores any <b>.exec</b> file in the current directory during initialization, unless the current directory is that named by the <b>HOME</b> environment variable. The default is <b>noexec</b> .
<b>hardtabs (ht=)</b>	Tells the vi editor the distance between the hardware tab stops on your display screen. (This option must match the tab setting of the underlying terminal or terminal emulator.) The default is <b>ht=8</b> .
<b>history (hist=)</b>	Sets the limit on last line mode history commands. The initial value is <b>hist=32</b> . The history size is zero ( <b>hist=0</b> ) for the <b>tvi</b> command.
<b>ignorecase (ic)</b>	Ignores distinction between uppercase and lowercase while searching for regular expressions. The default is <b>noic</b> .
<b>linelimit (ll=)</b>	Sets the maximum number of lines, as per the <b>-y</b> command-line option. This option only is effective if used with the <b>.exec</b> file or the <b>EXINIT</b> environment variable.
<b>lisp (lisp)</b>	Removes the special meaning of <b>(</b> , <b>{</b> , <b>[</b> , and <b>]</b> and enables the <b>=</b> (formatted print) operator for s-expressions, so you can edit list processing (LISP) programs. The default is <b>noisp</b> .
<b>list (list)</b>	Displays text with tabs (^I) and the marked end of lines (\$). The default is <b>noisp</b> .

vi Option (Abbreviation)	Description
magic (magic)	Treats the . (period), [ (left bracket), and * (asterisk) characters as special characters when searching for a pattern. In off mode, only the ( ) (parentheses) and \$ (dollar sign) retain special meanings. However, you can evoke special meaning in other characters by preceding them with a \ (backslash). The default is <b>magic</b> .
mesg (mesg)	Turns on write permission to the terminal if set while in visual mode. This option only is effective if used with the . <b>exrc</b> file or the <b>EXINIT</b> environment variable. The default is <b>on</b> .
modeline (modeline)	Runs a vi editor command line if found in the first five or the last five lines of the file. A vi editor command line can be anywhere in a line. For the vi editor to recognize a command line, the line must contain a space or a tab followed by the <b>ex:</b> or <b>vi:</b> string. The command line is ended by a second : (colon). The vi editor tries to interpret any data between the first and second colon as vi editor commands. The default is <b>nomodeline</b> .
novice	Indicates whether you are in <b>novice</b> mode. You cannot change the value by using the <b>set</b> command.
number (nu)	Displays lines prefixed with their line numbers. The default is <b>nonu</b> .
optimize (opt)	Speeds the operation of terminals that lack cursor addressing. The default is <b>noopt</b> .
paragraphs (para=)	Defines vi macro names that start paragraphs. The default is <b>para=IPLPPPQPP\ Lipplipnbp</b> . Single-letter <b>nroff</b> macros, such as the <b>.P</b> macro, must include the space as a quoted character if respecifying a paragraph.
partialchar (pc=)	Appears in the last display column where a double-wide character would not be displayed completely. The default character is - (minus sign).
prompt	Prompts for a new vi editor command when in command mode by printing a : (colon). The default is <b>on</b> .
readonly (ro)	Sets permanent read-only mode. The default is <b>noreadonly</b> .
redraw (redraw)	Simulates a smart workstation on a dumb workstation. The default is <b>nore</b> .
remap	Allows defining macros in terms of other macros. The default is <b>on</b> .
report (re=)	Sets the number of times you can repeat a command before a message is displayed. For subcommands that produce many messages, such as global subcommands, the messages are displayed when the command sequence completes. The default is <b>report=5</b> .
scroll (scr=)	Sets the number of lines to be scrolled when the user scrolls up or down. The default is 1/2 of the window size, rounded down.
sections (sect=)	Defines vi macro names that start sections. The default is <b>sect=NHSHHH\ HUuhsh+c</b> . Single-letter <b>nroff</b> macros, such as the <b>.P</b> macro, must include the space as a quoted character if respecifying a paragraph.
shell (sh=)	Defines the shell for the ! subcommand or the :! subcommand. The default is the login shell.
shiftwidth (sw=)	Sets the distance for the software tab stops used by the <b>autoindent</b> option, the shift commands ( > and < ), and the text input commands ( the Ctrl-D and Ctrl-T key sequences). This vi option only affects the indentation at the beginning of a line. The default is <b>sw=8</b> .
showmatch (sm)	Shows the ( (matching left parenthesis) or { (left bracket) as you type the ) (right parenthesis) or } (right bracket). The default is <b>nosm</b> .
showmode (smd)	Displays a message to indicate when the vi editor is in input mode. The default is <b>nosmd</b> .
slowopen (slow)	Postpones updating the display screen during inserts. The default is <b>noslow</b> .
tabstop (ts=)	Sets the distance between tab stops in a displayed file. The default is <b>ts=8</b> .
tags (tags =)	Defines the search path for the database file of function names created using the <b>ctags</b> command. The default is <b>tags=tags\ /usr/lib/tags</b> .
term (term=)	Sets the type of workstation you are using. The default is <b>term=\$TERM</b> , where <b>\$TERM</b> is the value of the <b>TERM</b> shell variable.
terse (terse)	Allows the vi editor to display the short form of messages. The default is <b>noterse</b> .
timeout (to)	Sets a time limit of two seconds on an entry of characters. This limit allows the characters in a macro to be entered and processed as separate characters when the <b>timeout</b> option is set. To resume use of the macro, set the <b>notimeout</b> option. The default is <b>to</b> .



vi Option (Abbreviation)	Description
<b>ttytype</b>	Indicates the tty type for the terminal being used. You cannot change this value from the vi editor.
<b>warn (warn)</b>	Displays a warning message before the ! subcommand executes a shell command if it is the first time you issued a shell command after changes were made in the editing buffer but not written to a file. The default is <b>warn</b> .
<b>window (wi=)</b>	Sets the number of lines displayed in one window of text. The default depends on the baud rate at which you are operating: 600 baud or less, 8 lines; 1200 baud, 16 lines; higher speeds, full screen minus 1 line.
<b>wrapmargin (wm=)</b>	Sets the margin for automatic word wrapping from one line to the next. The default is <b>wm=0</b> . A value of 0 turns off word wrapping.
<b>wrapscan (ws)</b>	Allows string searches to wrap from the end of the editing buffer to the beginning. The default is <b>ws</b> .
<b>wraptype (wt=)</b>	Indicates the method used to wrap words at the end of a line. The default value is <b>general</b> . You can specify one of the following four values: <ul style="list-style-type: none"> <li><b>general</b> Allows wraps on word breaks as white space between two characters. This setting is the default.</li> <li><b>word</b> Allows wraps on words.</li> <li><b>rigid</b> Allows wraps on columns and before closing punctuation.</li> <li><b>flexible</b> Allows wraps on columns, but one character of punctuation can extend past the margin.</li> </ul>
<b>writeany (wa)</b>	Turns off the checks usually made before a <b>write</b> subcommand. The default is <b>nowa</b> .

To see a list of the vi editor settings that have changed from the default settings, enter **set** and press the spacebar. Press the Enter key to return to the command mode.

To see a complete list of the vi editor settings, enter **set all**. Press the Enter key to return to the command mode.

To turn on a vi editor option, enter **set Option**. This command automatically returns you to the command mode.

To turn on multiple vi editor options, enter **set Option Option Option**. This command turns on the three designated vi editor options and returns you to the command mode.

To turn off a vi editor option, enter **set noOption**. This command automatically returns you to the command mode.

To change the value of a vi editor option, enter **set Option=Value**. This command automatically returns you to the command mode.

You can use the **:set** subcommand of the vi editor to set options for this editing session only, or to set options for this editing session and all future editing sessions.

To set or change vi editor options *for this editing session only*, enter the **:set** subcommand from the command line.

To set vi options *for all editing sessions*, put the **:set** subcommand in the **EXINIT** environment variable in the **.profile** file (read by the shell on login) or put the **set** subcommand into a **.exrc** file. The vi editor first looks for the **EXINIT** environment variable and runs its commands. If the **EXINIT** environment variable does not exist, the vi editor then looks for the **\$HOME/.exrc** file and runs its commands. Last, and regardless of any previous results, the vi editor looks for the local **.exrc** file and runs its commands.

**Note:** This process is true except with the **tv**i command (trusted vi). In this instance, the vi editor looks for and runs only the **/etc/.exrc** file.

For information about changing an option by setting the **EXINIT** environment variable, see the description of environment variables in the **environment** file.

The **.exrc** file can contain subcommands of the form **set Option=Value**; for example:

```
set cp=3 . . ;
```

To include a comment in the **.exrc** file, use a " (double quotation mark) as the first character in the line.

## Defining Macros

If you use a subcommand or sequence of subcommands frequently, you can use the vi editor to define a macro that issues that subcommand or sequence.

To define a macro, enter the sequence of subcommands into a buffer named with a letter of the alphabet. The lowercase letters a through z overlay the contents of the buffer, and the uppercase letters A through Z append text to the previous contents of the buffer, allowing you to build a macro piece by piece.

For example, to define a buffer macro named c that searches for the word corner and makes the third line after the word corner the current line, enter the following command:

```
o /corner/+3
```

Then press the Esc key and enter the following command:

```
"c
```

where c is the name of the buffer macro.

To add text to the previous contents of the defined buffer, enter the o viSubcommand, press the Esc key, and enter "CapitalLetter, where the *CapitalLetter* variable specifies an uppercase letter A through Z. For example, to build a buffer macro named T that searches for the word corner and allows you to add more commands, enter the following command:

```
o corner
```

Then press the Esc key and enter the following command:

```
"T
```

where T is the name of the buffer macro. You can repeat this process at any time to add more vi subcommands to the same buffer.

For example, to add commands that move the cursor to the previous line and delete that line, enter the following command:

```
o -dd
```

where - (minus sign) means to move the cursor up one line, and dd means to delete the current line. Press the Esc key and enter the following command:

```
"Tdd
```

To start the macro, enter @Letter, where the *Letter* variable specifies the letter name of the buffer macro you want to use. To use the same macro again, enter @@ (two at symbols). For example, enter @T to start the T buffer macro and run the **search**, **move cursor**, and **delete line** commands. Enter @@T to start the T buffer macro again.

The character set used by your system is defined by the collation table. This table affects the performance of vi macros.

## Mapping Keys

You can use the **:map**, **:map!**, and **:ab** subcommands to map a keystroke to a command or a sequence of commands. The **:map** subcommand is used in the command mode. The **:map!** and **:ab** subcommands are used in the text input mode. You can map keys for this editing session and all future editing sessions or only for the current editing session from either mode.

To map keys *for all future editing sessions*, put the subcommand into a **\$HOME/.exrc** file. Each time you start the vi editor, it reads this file. The mapping remains in effect for every editing session.

To map keys *for the current editing session only* from the *command mode*, start the subcommand during the vi editor session. To map keys for the current editing session only from the *text input mode*, enter the subcommand on the command line during the vi editor session. The mapping remains in effect only for the current editing session.

**Attention:** If you use an IBM® 3161 ASCII display station, IBM 3163 ASCII display station, or IBM 3101 ASCII display station, the default key-mapping of the vi editor can cause you to lose data. To see the default mapping, issue a **:map** subcommand. Specific problems arise with the Esc-J or Shift-J key sequence. These key sequences delete all information from the current position of the cursor to the end of the file. To avoid problems, change this key sequence using a **.exrc** file.

The **:map**, **:map!**, and **:ab** subcommands are defined and used as follows:

<b>Item</b>	<b>Description</b>
<b>:map</b>	Defines macros in the command mode. The <b>:map</b> subcommand allows you to run a specified command or sequence of commands by pressing a single key while in the vi editor.

To map keys in the command mode, start the vi editor with an empty editing buffer and do not name a vi file using the **vi** command or type anything into the buffer after the vi editor starts. You can use the **:map** subcommand to do the following:

- To map a character to a sequence of editing commands, enter:  
:map Letter viSubcommand
- To unmap a character previously mapped in command mode, enter:  
:unmap Letter
- To display a list of current mappings for the command mode, enter  
:map

The following keys are not used by the vi editor, but are available for use with the **:map** subcommand in the command mode:

- Letters g, K, q, V, and v
- Control key sequences Ctrl-A, Ctrl-K, Ctrl-O, Ctrl-W, and Ctrl-X
- Symbols \_ (underscore), \* (asterisk), \ (backslash), and = (equal sign)

Although you can map a key that is already used by the vi editor, the key's usual function is not available as long as the map is in effect. Some terminals allow you to map command sequences to function keys. If you are in LISP mode, the = (equal sign) cannot be used because it is used by the vi editor.

To map the letter v to the sequence of commands that would locate the next occurrence of the word map and change it to the word MAP, enter the following command:

```
:map v /map<Ctrl-V><Enter>cwMAP<Ctrl-V><Esc><Ctrl-V><Enter>
```

The previous example instructs the vi editor to locate the next occurrence of map (/map<Ctrl-V><Enter>), change map to MAP (cwMAP), end the change-word subcommand (<Ctrl-V><Esc>), and enter the command (<Ctrl-V><Enter>).

**Requirement:** To prevent the vi editor from interpreting the Enter key, it must be preceded by the Ctrl-V key sequence when being mapped. This condition is also true of the Esc, Backspace, and Delete keys.

To map the control characters Ctrl-A, Ctrl-K, and Ctrl-O, simultaneously press the Ctrl key and the letter. For example, to map the Ctrl-A key sequence to the sequence of commands that saves a file and edits the next one in a series, enter the following command:

```
:map <Ctrl-A> :w<Ctrl-V><Enter>;n<Ctrl-V><Enter>
```

To map the control characters Ctrl-T, Ctrl-W, and Ctrl-X, you must first escape them with the Ctrl-V key sequence.

**Item**      **Description**

To map the | (pipe symbol), you must first escape it with the two Ctrl-V key sequences, as illustrated by the following example that maps the character g to the sequence of commands that escapes to the shell, concatenates the file /etc/motd, and pipes the output to the wc command:

```
:map g :!cat /etc/motd <Ctrl-V><Ctrl-V>| wc<Ctrl-V><Enter>
```

If your terminal permits you to map function keys, you must reference them with the #number key sequence to designate the number of the function key that you want to map. In the following example, the F1 function key is mapped to the sequence of commands that deletes a word and moves the cursor three words down:

```
:map #1 dwww
```

In order for function key mapping to work, the output of the function key for your terminal type must match the output defined in the **terminfo** file. These definitions are denoted by the *kfnumber* entries, where kf1 represents the F1 function key, kf2 represents the F2 function key, and so on. If the output that you get when you press the function key does not match this entry, you must use the terminal's setup mode to correct the settings to match these terminal database entries before any mapping can occur.

You can also map certain keyboard special keys, such as the Home, End, Page Up, and Page Down keys. For most terminals, these keys are already mapped in the vi editor. You can verify this mapping by using the **:map** subcommand. If these keys are not already mapped, you can use the **:map** subcommand as follows:

```
:map <Ctrl-V><End> G
:map <Ctrl-V><Home> 1G
:map <Ctrl-V><PageUp> <Ctrl-F>
:map <Ctrl-V><PageDown> <Ctrl-B>
```

To get a listing of all current maps in the command mode, enter the **:map** subcommand. The preceding examples are then displayed as follows:

```
v          v          /map<Ctrl-M>cwMAP<Ctrl-[>Ctrl-M>
<Ctrl-A> <Ctrl-A>    :w<Ctrl-M>:n<Ctrl-M>
g          g          :!cat /etc/motd | wc <Ctrl-M>
```

**Tip:** The Ctrl-V and Enter key sequence is displayed as the Ctrl-M key sequence, and the Ctrl-V and Esc key sequence is displayed as the Ctrl-[ key sequence.

**:map!**

Maps character strings to single keys while in text input mode. To map keys in the text input mode, start the vi editor with an empty editing buffer and do not name a vi file using the **vi** command or type anything into the buffer after the vi editor starts. You can use the **:map!** subcommand to do the following:

- To map a letter to one or more vi strings in text input mode, enter:  
:map! Letter String
- To unmap a letter previously mapped in text input mode, enter:  
:unmap! Letter
- To display a list of existing strings that are mapped to specific keys in text input mode, enter:  
:map!

Typing the mapped key while in text input mode produces the specified string. The Ctrl-V and Esc key sequence puts you into command mode, backs up to the beginning of the current word (**bbw**), and starts the **cw** (change-word) subcommand. For example:

```
:map! % <Ctrl-V><Esc>bbcw
```

When typing text, if you realize that you have mistyped a word, you can change it by pressing the % (percent) key and retyping the word. You are automatically returned to insert mode.

**Important:** Be careful when choosing keys to be used for the **:map!** subcommand. Once keys have been mapped, they can no longer be input as text without first issuing the **:unmap!** subcommand.

<b>Item</b>	<b>Description</b>
<b>:ab</b>	<p>Maps a key or sequence of keys to a string of characters for use in the text input mode. The <b>:ab</b> subcommand is useful when inputting text that possesses several repetitive phrases, names, or titles.</p> <p>The following example replaces the word <code>city</code> with the phrase <code>Austin, Texas 78759</code> whenever it is typed in text input mode and followed by a white space, period, or comma:</p> <pre>:ab city Austin, Texas 78759</pre> <p>For example, if while inputting text, you type the following:</p> <pre>My current residence is city.</pre> <p>Pressing the Tab key expands the word <code>city</code> to read:</p> <pre>My current residence is Austin, Texas 78759.</pre> <p>The abbreviation is not expanded within a word. For example, if you type <code>My current residence iscity</code>, the word <code>iscity</code> is not expanded.</p> <p>If the <b>:map!</b> subcommand is used to map abbreviations for insert mode, then all occurrences of the abbreviations are expanded regardless of where it occurs. If you used the <b>:map!</b> subcommand for the preceding example (<code>:map! city Austin, Texas 78759</code>), then whenever you type the word <code>city</code>, regardless of what precedes or follows, the word will be expanded to <code>Austin, Texas 78759</code>. Therefore, the word <code>iscity</code> becomes <code>isAustin, Texas 78759</code>.</p> <p><b>Important:</b> Be careful when choosing the keys that are used for the <b>:ab</b> subcommand. Once keys are defined, they can no longer be input as text without first issuing the <b>:unab</b> subcommand.</p>

## Setting Abbreviations

The **set** command has behavior similar to the **map!** command except that the **set** command substitutes the string for the abbreviation only when the abbreviation is a separate word. You can use the **set** command of the vi editor to:

- List existing abbreviations
- Remove an abbreviation
- Set (define) an abbreviation

**Tip:** Start the vi editor with an empty editing buffer. Do not name a vi file using the **vi** command or type anything into the buffer after the vi editor starts. Press the Esc key to be sure you are in the command mode.

<b>Item</b>	<b>Description</b>
<b>To list abbreviations</b>	Enter the <b>:ab</b> command to list existing abbreviations. Press the Enter key to return to command mode.
<b>To remove abbreviations</b>	Enter the <b>:anab</b> <i>Abbreviation</i> command to remove an abbreviation, where the <i>Abbreviation</i> variable specifies the character string you do not want abbreviated any more.
<b>To set (define) an abbreviation</b>	Enter the <b>:ab</b> <i>Abbreviation String</i> command to set an abbreviation, where the <i>Abbreviation</i> variable specifies the character string being defined as an abbreviation and the <i>String</i> variable specifies the character string being abbreviated. The abbreviation can be substituted for the string only when the abbreviation is a separate word.
	For example, if you enter the <b>:ab kn upper</b> command and then type <code>acknowledge</code> while in the text input mode, the set abbreviation string is not started because the <code>kn</code> string in the word <code>acknowledge</code> is not a separate word.
	However, if you type the <b>:ab kn upper</b> command and then type <code>make the kn line all kn case</code> while in the text input mode, the result is <code>make the upper line all uppercase</code> .

## Flags

Item	Description
<code>-cSubcommand</code>	Carries out the ex editor subcommand before viewing with vi begins. The cursor moves to the line affected by the last subcommand to be carried out. When a null operand is entered, as in <code>-c'</code> , the vi editor places the cursor on the first line of the file. The <code>-c</code> flag is incompatible with the <code>+</code> flag. Do not specify both flags at the same time.
<code>-l</code>	Enters the vi editor in LISP mode. In this mode, the vi editor creates indents appropriate for LISP code, and the <code>(, ), {, }, [[, and ]]</code> subcommands are modified to act appropriately for LISP.
<code>-r[File]</code>	Recovers a file after a vi editor or system malfunction. If you do not specify the <i>File</i> variable, the vi editor displays a list of all saved files.
<code>-R</code>	Sets the <b>readonly</b> option to protect the file against overwriting.
<code>-tTag</code>	Edits the file containing the <i>Tag</i> variable and positions the vi editor at its definition. To use this flag, you must first create a database of function names and their locations using the <b>ctags</b> command.
<code>-v</code>	Enters the vi editor in the verbose mode.
<code>-wNumber</code>	Sets the default window size to the value specified by the <i>Number</i> variable. This flag is useful when you use the vi editor over a low-speed line.
<code>-yNumber</code>	Overrides the maximum line setting of 1,048,560 with any value greater than 1024. You should request twice the number of lines that you require because the vi editor uses the extra lines for buffer manipulation.
<code>+[Subcommand]</code>	Carries out the ex editor subcommand before editing begins. If you do not specify the <i>Subcommand</i> variable, the cursor is placed on the first line of the file. This <code>+</code> flag is incompatible with the <code>-c</code> flag. Do not specify both flags at the same time.

## vi General Subcommand Syntax

Use the following general syntax to enter subcommands:

`[Named_Buffer] [Operator] [Number] Object`

**Tip:** Square brackets indicate optional items.

Item	Description
<code>[Named_Buffer]</code>	Specifies a temporary text storage area.
<code>[Operator]</code>	Specifies the subcommand or action; instructs the vi editor.
<code>[Number]</code>	Specifies either the extent of the action or a line address as a whole number.
<code>Object</code>	Specifies what to act on, such as a text object (a character, word, sentence, paragraph, section, character string) or a text position (a line, position in the current line, screen position).

## Counts before Subcommands

You can put a number in front of many subcommands. The vi editor interprets this number in one of the following ways:

- Go to the line specified by the *Number* parameter:  
5G  
10Z
- Go to the column specified by the *Number* parameter:  
25|
- Scroll the number of lines up or down specified by the *Number* parameter:  
10Ctrl-U  
10Ctrl-D

## vi Editor Subcommands

Use the subcommands to perform these kinds of actions:

- Moving the cursor
- Editing text

- Manipulating files
- Other actions

## Moving the Cursor

Use subcommands to move the cursor within a file in these ways:

- Moving within a line
- Moving within a line by character position
- Moving to words
- Moving by line position
- Moving to sentences, paragraphs, or sections
- Moving by redrawing the screen
- Paging and scrolling
- Searching for patterns
- Marking a specific location in a file and returning

### Moving within a Line

Enter the following subcommands in command mode. You can cancel an incomplete command by pressing the Esc key. If you need information about the format of vi subcommands, see vi General Subcommand Syntax.

Item	Description
Left Arrow or h or Ctrl-H	Moves the cursor one character to the left.
Down Arrow or j or Ctrl-J or Ctrl-N	Moves the cursor down one line (it remains in the same column).
Up Arrow or k or Ctrl-P	Moves the cursor up one line (it remains in the same column).
Right Arrow or l	Moves the cursor one character to the right.

### Moving within a Line by Character Position

Enter the following subcommands in command mode. You can cancel an incomplete command by pressing the Esc key. If you need information about the format of vi subcommands, see vi General Subcommand Syntax.

Item	Description
^	Moves the cursor to the first nonblank character.
0	Moves the cursor to the beginning of the line.
\$	Moves the cursor to the end of the line.
fx	Moves the cursor to the next <i>x</i> character.
Fx	Moves the cursor to the last <i>x</i> character.
tx	Moves the cursor to one column before the next <i>x</i> character.
Tx	Moves the cursor to one column after the last <i>x</i> character.
;	Repeats the last <b>f</b> , <b>F</b> , <b>t</b> , or <b>T</b> subcommand.
,	Repeats the last <b>f</b> , <b>F</b> , <b>t</b> , or <b>T</b> subcommand in the opposite direction.
<i>Number</i> l	Moves the cursor to the specified column.

### Moving to Words

Enter the following subcommands in command mode. For more information about the format of vi subcommands, see vi General Subcommand Syntax.



Item	Description
w	Moves the cursor to the next small word.
b	Moves the cursor to the previous small word.
e	Moves the cursor to the next end of a small word.
W	Moves the cursor to the next big word.
B	Moves the cursor to the previous big word.
E	Moves the cursor to the next end of a big word.

## Moving by Line Position

Enter the following subcommands in command mode. If you need information about the format of vi subcommands, see vi General Subcommand Syntax.

Item	Description
H	Moves the cursor to the top line on the screen.
L	Moves the cursor to the last line on the screen.
M	Moves the cursor to the middle line on the screen.
+	Moves the cursor to the next line at its first nonblank character.
-	Moves the cursor to the previous line at its first nonblank character.
Enter	Moves the cursor to the next line at its first nonblank character.

## Moving to Sentences, Paragraphs, or Sections

Enter the following subcommands in command mode. You can cancel an incomplete subcommand by pressing the Esc key. If you need information about the format of vi subcommands, see vi General Subcommand Syntax.

Item	Description
(	Places the cursor at the beginning of the previous sentence, or the previous s-expression if you are in LISP mode.
)	Places the cursor at the beginning of the next sentence, or the next s-expression if you are in LISP mode.
{	Places the cursor at the beginning of the previous paragraph, or at the next list if you are in LISP mode.
}	Places the cursor at the beginning of the next paragraph, at the next section if you are in C mode, or at the next list if you are in LISP mode.
]]	Places the cursor at the next section, or function if you are in LISP mode.
[[	Places the cursor at the previous section, or function if you are in LISP mode.

## Moving by Redrawing the Screen

Enter the following subcommands in command mode. You can cancel an incomplete subcommand by pressing the Esc key. If you need information about the format of vi subcommands, see vi General Subcommand Syntax.

Item	Description
z	Redraws the screen with the current line at the top of the screen.
z-	Redraws the screen with the current line at the bottom of the screen.
z.	Redraws the screen with the current line at the center of the screen.
/Pattern/z-	Redraws the screen with the line containing the character string, specified by the <i>Pattern</i> parameter, at the bottom.

## Paging and Scrolling

Enter the following subcommands in command mode. You can cancel an incomplete subcommand by pressing the Esc key. If you need information about the format of vi subcommands, see vi General Subcommand Syntax.

Item	Description
Ctrl-U	Scrolls up one-half screen.
Ctrl-D	Scrolls down one-half screen.
Ctrl-F	Scrolls forward one screen.
Ctrl-B	Scrolls backward one screen.
Ctrl-E	Scrolls the window down one line.
Ctrl-Y	Scrolls the window up one line.
z+	Pages up.
z^	Pages down.

## Searching for Patterns

Enter the following subcommands in command mode. You can cancel an incomplete subcommand by pressing the Esc key. If you need information about the format of vi subcommands, see vi General Subcommand Syntax.

Item	Description
[ <i>Number</i> ]G	Places the cursor at the line number specified by the <i>Number</i> parameter or at the last line if the <i>Number</i> parameter is not specified.
/ <i>Pattern</i>	Places the cursor at the next line containing the character string specified by the <i>Pattern</i> parameter.
? <i>Pattern</i>	Places the cursor at the next previous line containing the character string specified by the <i>Pattern</i> parameter.
n	Repeats the last search for the text specified by the <i>Pattern</i> parameter in the same direction.
N	Repeats the last search for the text specified by the <i>Pattern</i> parameter in the opposite direction.
/ <i>Pattern</i> ! <i>Number</i>	Places the cursor the specified number of lines after the line matching the character string specified by the <i>Pattern</i> parameter.
? <i>Pattern</i> ?- <i>Number</i>	Places the cursor the specified number of lines before the line matching the character string specified by the <i>Pattern</i> parameter.
%	Finds the parenthesis or brace that matches the one at current cursor position.

## Editing Text

The subcommands for editing enable you to perform the following tasks:

- Marking a specific location in a file and returning
- Adding text to a file
- Changing text while in input mode
- Changing text from command mode
- Copying and moving text
- Restoring and repeating changes

### Marking a Specific Location in a File and Returning

Enter the following subcommands in command mode. You can cancel an incomplete subcommand by pressing the Esc key. If you need information about the format of vi subcommands, see vi General Subcommand Syntax.

Item	Description
"	Moves the cursor to the previous location of the current line.
"	Moves the cursor to the beginning of the line containing the previous location of the current line.
mx	Marks the current position with the letter specified by the <i>x</i> parameter.
`x	Moves the cursor to the mark specified by the <i>x</i> parameter.
'x	Moves the cursor to the beginning of the line containing the mark specified by the <i>x</i> parameter.

## Adding Text to a File (Text Input Mode)

Enter the following subcommands in command mode to change the vi editor into text input mode. If you need information about the format of vi subcommands, see vi General Subcommand Syntax.

Item	Description
<i>aText</i>	Inserts text specified by the <i>Text</i> parameter after the cursor. End text input mode by pressing the Esc key.
<i>AText</i>	Adds text specified by the <i>Text</i> parameter to the end of the line. End text input mode by pressing the Esc key.
<i>iText</i>	Inserts text specified by the <i>Text</i> parameter before the cursor. End text input mode by pressing the Esc key.
<i>IText</i>	Inserts text specified by the <i>Text</i> parameter before the first nonblank character in the line. End text input mode by pressing the Esc key.
o	Adds an empty line below the current line. End text input mode by pressing the Esc key.
O	Adds an empty line above the current line. End text input mode by pressing the Esc key.

## Changing Text While in Input Mode

Use the following subcommands only while in text input mode. These commands have different meanings in command mode. If you need information about the format of vi subcommands, see vi General Subcommand Syntax.

Item	Description
Ctrl-D	Goes back to previous autoindent stop.
^ Ctrl-D	Ends autoindent for this line only.
0 Ctrl-D	Moves cursor back to left margin.
Esc	Ends insertion and returns to command state.
Ctrl-H	Erases the last character.
Ctrl-Q	Enters any character if xon is disabled.
Ctrl-V	Enters any character.
Ctrl-W	Erases the last small word.
\	Quotes the erase and kill characters.
Ctrl-?	Interrupts and ends insert or the Ctrl-D key sequence.

## Changing Text from Command Mode

Use the following subcommands in command mode. An incomplete subcommand can be canceled by pressing the Esc key. If you need information about the format of vi subcommands, see vi General Subcommand Syntax.

Item	Description
<b>C</b>	Changes the rest of the line (same as <b>c\$</b> ).
<b>c</b>	Changes a line.
<b>cc</b>	Changes a word.
<b>cw</b>	Changes a word to the text specified by the <i>Text</i> parameter.
<b>cwText</b>	Changes a word to the text specified by the <i>Text</i> parameter.
<b>D</b>	Deletes the rest of the line (same as <b>d\$</b> ).
<b>dd</b>	Deletes a line.
<b>dw</b>	Deletes a word.
<b>J</b>	Joins lines.
<b>rx</b>	Replaces the current character with the character specified by <i>x</i> .
<b>RText</b>	Overwrites characters with the text specified by the <i>Text</i> parameter.
<b>s</b>	Substitutes characters (same as <b>cl</b> ).
<b>S</b>	Substitutes lines (same as <b>cc</b> ).
<b>u</b>	Undoes the previous change.
<b>x</b>	Deletes a character at the cursor.
<b>X</b>	Deletes a character before the cursor (same as <b>dh</b> ).
<b>&lt;&lt;</b>	Shifts one line to the left.
<b>&lt;L</b>	Shifts all lines from the cursor to the end of the screen to the left.
<b>&gt;&gt;</b>	Shifts one line to the right.
<b>&gt;L</b>	Shifts all lines from the cursor to the end of the screen to the right.
<b>~</b>	Changes letter at the cursor to the opposite case.
<b>!</b>	Indents for LISP.

## Copying and Moving Text

Use the following subcommands in command mode. An incomplete subcommand can be canceled by pressing the Esc key. If you need information about the format of vi subcommands, see vi General Subcommand Syntax.

Item	Description
<b>p</b>	Puts back text from the undo buffer after the cursor.
<b>P</b>	Puts back text from the undo buffer before the cursor.
<b>"xp</b>	Puts back text from the <i>x</i> buffer.
<b>"xd</b>	Deletes text into the <i>x</i> buffer.
<b>y</b>	Places the object that follows (for example, <b>w</b> for word) into the undo buffer.
<b>"xy</b>	Places the object that follows into the <i>x</i> buffer, where <i>x</i> is any letter.
<b>Y</b>	Places the line in the undo buffer.

## Restoring and Repeating Changes

Use the following subcommands in command mode. An incomplete subcommand can be canceled by pressing the Esc key. If you need information about the format of vi subcommands, see vi General Subcommand Syntax.

Item	Description
------	-------------

<b>u</b>	Undoes the last change.
----------	-------------------------

**Tip:** After an undo, the cursor moves to the first non-blank character on the updated current line.

<b>U</b>	Restores the current line if the cursor has not left the line since the last change.
----------	--

<b>.</b>	Repeats the last change or increments the " <b>np</b> " command.
----------	--

**Note:**

1. This subcommand will repeat the last change, including an undo. Therefore, after an undo, repeat performs an undo rather than repeat the last change.

2. This subcommand is not meant for use with a macro. Enter @@ (two at signs) to repeat a macro.

<b>"n p</b>	Retrieves the <i>n</i> th last delete of a complete line or block of lines.
-------------	---

## Manipulating Files

The subcommands for manipulating files allow you to do the tasks outlined in the following sections:

- Saving changes to a file
- Editing a second file
- Editing a list of files
- Finding file information

### Saving Changes to a File

Use the following subcommands in command mode. If you need information about the format of vi subcommands, see vi General Subcommand Syntax.

Item	Description
------	-------------

<b>:w</b>	Writes the edit buffer contents to the original file. If you are using this subcommand within the ex editor, you do not need to type the : (colon).
-----------	---

<b>:w File</b>	Writes the edit buffer contents to the file specified by the <i>File</i> parameter. If you are using this subcommand within the ex editor, you do not need to type the : (colon).
----------------	---

<b>:w! File</b>	Overwrites the file specified by the <i>File</i> parameter with the edit buffer contents. If you are using this subcommand within the ex editor, you do not need to type the : (colon).
-----------------	---

### Editing a Second File

Enter the following subcommands in command mode. If you need information about the format of vi subcommands, see vi General Subcommand Syntax.

Item	Description
------	-------------

<b>:e File</b>	Edits the specified file. If you are using this subcommand from the ex editor, you do not need to type the : (colon).
----------------	---

<b>:e!</b>	Re-edits the current file and discards all changes.
------------	---

<b>:e + File</b>	Edits the specified file starting at the end.
------------------	---

<b>:e + Number File</b>	Edits the specified file starting at the specified line number.
-------------------------	---

<b>:e #</b>	Edits the alternate file. The alternate file is usually the previous file name before accessing another file with a <b>:e</b> command. However, if changes are pending on the current file when a new file is called, the new file becomes the alternate file. This subcommand is the same as the <b>Ctrl-A</b> subcommand.
-------------	---

<b>:r File</b>	Reads the file into the editing buffer by adding new lines below the current line. If you are using this subcommand from the ex editor, you do not need to type the : (colon).
----------------	--

<b>:r !Command</b>	Runs the specified command and places its output into the file by adding new lines below the current cursor position.
--------------------	---

<b>:ta Tag</b>	Edits a file containing the <i>Tag</i> tag starting at the location of the tag. To use this subcommand, you must first create a database of function names and their locations using the <b>ctags</b> command. If you are using this subcommand from the ex editor, you do not need to type the : (colon).
----------------	--

Item	Description
Ctrl-]	Edits a file containing the tag associated with the current word starting at the location of the tag. To use this subcommand, you must first create a database of function names and their locations using the <b>ctags</b> command. Ctrl-T edits a file at the editing position where the previous Ctrl-] subcommand was issued. If multiple Ctrl-] subcommands have been issued, then multiple Ctrl-T subcommands can be used to return to previous editing positions where Ctrl-] subcommands were issued.
Ctrl-A	Edits the alternate file. The alternate file is usually the previous current file name. However, if changes are pending on the current file when a new file is called, the new file becomes the alternate file. This subcommand is the same as the <b>:e #</b> subcommand.

## Editing a List of Files

Enter the following subcommands in command mode. If you need information about the format of vi subcommands, see vi General Subcommand Syntax.

Item	Description
:n	Edits the next file in the list entered on the command line. If you are using this subcommand from the ex editor, a : (colon) is not needed.
:n <i>Files</i>	Specifies a new list of files to edit. If you are using this subcommand from the ex editor, a : (colon) is not needed.

## Finding File Information

Enter the following subcommand in command mode. If you need information about the format of vi subcommands, see vi General Subcommand Syntax.

Item	Description
Ctrl-G	Shows the current file name, current line number, number of lines in the file, and percentage of the way through the file where the cursor is located.

## Other Actions

The vi editor provides the subcommands described in the following sections:

- Adjusting the screen
- Entering shell commands
- Interrupting and ending the vi editor

### Adjusting the Screen

Enter the following subcommands in command mode. An incomplete subcommand can be canceled by pressing the Esc key. If you need information about the format of vi subcommands, see vi General Subcommand Syntax.

Item	Description
Ctrl-L	Clears and redraws the screen.
Ctrl-R	Redraws the screen and eliminates blank lines marked with @ (at sign).
<i>zNumber</i>	Makes the window the specified number of lines long.

### Entering Shell Commands

The following subcommands allow you to run a command within the vi editor. Enter these subcommands in command mode. If you need information about the format of vi subcommands, see vi General Subcommand Syntax.

<b>Item</b>	<b>Description</b>
<code>:sh</code>	Enters the shell to allow you to run more than one command. You can return to the vi editor by pressing the Ctrl-D key sequence. If you are using this subcommand within the ex editor, a : (colon) is not needed.
<code>!:Command</code>	Runs the specified command and then returns to the vi editor. If you are using this subcommand within the ex editor, a : (colon) is not needed.  <b>Tip:</b> The # (alternate file), % (current file), and ! (previous command) special characters are expanded when following a !: subcommand. To prevent any of these characters from being expanded, use the \ (backslash).
<code>!!</code>	Repeats the last !:Command subcommand.
<code>Number!!Command</code>	Runs the specified command and replaces the lines specified by <i>Number</i> with the output of the command. If a number is not specified, the default value is 1. If the command expects standard input, the specified lines are used as input.
<code>!Object Command</code>	Runs the specified command and replaces the object specified by the <i>Object</i> parameter with the output of the command. If the command expects standard input, the specified object is used as input.

## Interrupting and Ending the vi Editor

Enter the following subcommands in command mode. If you need information about the format of vi subcommands, see vi General Subcommand Syntax.

<b>Item</b>	<b>Description</b>
<code>Q</code>	Enters the ex editor in command mode.
<code>ZZ</code>	Exits the vi editor, saving changes.
<code>:q</code>	Quits the vi editor. If you have changed the contents of the editing buffer, the vi editor displays a warning message and does not quit. If you are using this subcommand from the ex editor, a : (colon) is not needed.
<code>:q!</code>	Quits the vi editor, discarding the editing buffer. If you are using this subcommand from the ex editor, a : (colon) is not needed.
<code>Esc</code>	Ends text input or ends an incomplete subcommand.
<code>Ctrl-?</code>	Interrupts a subcommand.

## Exit Status

The following exit values are returned:

<b>Item</b>	<b>Description</b>
<code>0</code>	Indicates successful completion.
<code>&gt;0</code>	Indicates an error occurred.

## Input Files

Input files must be text files or files that are similar to text files except for an incomplete last line that contains no null characters.

The `.exrc` files must be text files consisting of `ex` commands.

The `$HOME/.vi_history` file is an auto-generated text file that records the last line mode command history.

By default, the vi editor reads lines from the files to be edited without interpreting any of those lines as any form of vi editor command.

### Related information:

ctags command  
ex command

tvi command  
.profile file

---

## view Command

### Purpose

Starts the vi editor in read-only mode.

### Syntax

```
view [ -cSubcommand ] [ -l ] [ -t Tag ] [ -wNumber ] [ -y ] [ -r [ File ] ] [ + [ Subcommand ] ] [ File ... ]
```

### Description

The **view** command starts the vi full-screen editor in read-only mode. The read-only mode is only advisory to prevent accidental changes to the file. To override read-only mode, use the ! (exclamation point) when executing a command. The *File* parameter specifies the name of the file you want to browse. Use vi subcommands for moving within the file. Use the **:q** subcommand to exit the **view** command. If you modify the file you can save your modifications by pressing the Esc key and wq!.

### Flags

Item	Description
-cSubcommand	Carries out the ex editor subcommand before viewing with vi begins. When a null operand is entered, as in -c '', the editor places the cursor on the last line of the file.
-l	Enters a version of the vi editor with specialized features designed for writing programs in the LISP language. In this mode, the vi editor indents appropriately for LISP programming, and the (, ), {, }, [[, and ]] subcommands are modified to act appropriately for LISP.
-r [File]	Recovers a file after an editor or system crash. If you do not specify a <i>File</i> parameter, the editor displays a list of all saved files.
-tTag	Edits the file containing the tag specified by the <i>Tag</i> parameter and positions the editor at its definition. To use this flag, you must first create a database of function names and their locations using the ctags command.
-wNumber	Sets the default window size to the value specified by the <i>Number</i> parameter. This is useful when your terminal communicates with the system running the editor over a slow communications line.
-y	Overrides the maximum line setting of 1,048,560 with any value greater than 1024.
+ [Subcommand]	Carries out the ex editor subcommand specified by the <i>Subcommand</i> parameter before viewing with vi begins. If you do not specify a subcommand, the cursor is placed on the last line of the file.

#### Related reference:

“vi or vedit Command” on page 13

#### Related information:

ctags command

---

## vmh Command

### Purpose

Starts a visual interface for use with MH commands.

### Syntax

```
vmh [ -prompt String ] [ -vmhproc CommandString | -novmhproc ]
```



## Description

The **vmh** command starts a visual interface for use with MH commands. The **vmh** command implements the server side of the MH window management protocol and maintains a split-screen interface to any program that implements the client side of the protocol.

The **vmh** command prompts for commands and sends them to the client side of the protocol. If the command produces a window with more than one screen of output, the **vmh** command prompts the user for a subcommand. The **vmh** subcommands enable you to display specific portions of the command output.

## vmh Subcommands

Item	Description
Ctrl-L	Refreshes the screen.
Space	Advances to the next screen.
[Number] Enter	Advances the specified number of lines. The default is one line.
[Number] d	Advances 10 times the specified number of lines. The default for the <i>Number</i> variable is 1, for a total of 10 lines.
[Number] g	Goes to the specified line.
[Number] G	Goes to the end of the window. If the <i>Number</i> variable is specified, this command acts like the <b>g</b> flag.
[Number] u	Goes back 10 times the specified number of lines. The default for the <i>Number</i> variable is 1, for a total of 10 lines.
[Number] y	Goes back the specified number of lines. The default is one line.
h	Displays a help message.
q	Ends output.

## Flags

Item	Description
-help	Lists the command syntax, available switches (toggles), and version information. <b>Note:</b> For MH, the name of this flag must be fully spelled out.
-novmproc	Runs the default <b>vmproc</b> without the window management protocol.
-prompt <i>String</i>	Uses the specified string as the prompt.
-vmhproc <i>CommandString</i>	Specifies the program that implements the client side of the window management protocol. The default is the <b>msh</b> program.

## Profile Entries

The following entries are entered in the *UserMhDirectory/.mh\_profile* file:

Item	Description
Path:	Specifies the user's MH directory.
mshproc:	Specifies the program used for the MH shell.

## Files

Item	Description
<code>\$HOME/mh_profile</code>	Contains the MH user profile.
<code>/usr/bin/vmh</code>	Contains the <code>vmh</code> command.

#### Related information:

msh command  
 mh\_alias command  
 mh\_profile command

---

## vmo Command

### Purpose

Manages Virtual Memory Manager tunable parameters.

### Syntax

```
vmo [ -p | -r ] [ -y ] { -o Tunable [= Newvalue] }
```

```
vmo [ -p | -r ] [ -y ] { -d Tunable }
```

```
vmo [ -p | -r ] [ -y ] -D
```

```
vmo [ -p | -r ] [ -F ] -a
```

```
vmo -h [ Tunable ]
```

```
vmo [ -F ] -L [ Tunable ]
```

```
vmo [ -F ] -x [ Tunable ]
```

**Note:** Multiple `-o`, `-d`, `-x` and `-L` are allowed.

### Description

**Note:** The `vmo` command can only be executed by root. The `vmo` command is a self-documenting command. The information about some of the flags or tunable parameters might be missing or out-of-date. You can find an up-to-date list of all the flags and tunable parameters by using the `-h`, `-L`, or `-x` flag.

Use the `vmo` command to configure Virtual Memory Manager tuning parameters. This command sets or displays current or next boot values for all Virtual Memory Manager tuning parameters. This command can also make permanent changes or defer changes until the next reboot. Whether the command sets or displays a parameter is determined by the accompanying flag. The `-o` flag specifies both actions. It can either display the value of a parameter or set a new value for a parameter.

The Virtual Memory Manager (VMM) maintains a list of free real-memory page frames. The page frames are available to hold virtual-memory pages that are needed to satisfy a page fault. When the number of pages on the free list falls below the values that are specified by the `minfree` parameter, the VMM begins to steal pages to add to the free list. The VMM continues to steal pages until the free list has at least the number of pages that are specified by the `maxfree` parameter.

If the number of file pages (permanent pages) in memory is less than the number specified by the `minperm%` parameter, the VMM steals frames from either computational or file pages, regardless of repage rates. If the number of file pages is greater than the number specified by the `maxperm%`

parameter, the VMM steals frames only from file pages. Between the two, the VMM normally only steals file pages, but if the repage rate for file pages is higher than the repage rate for computational pages, computational pages are stolen as well.

You can also modify the thresholds that are used to decide when the system is running out of paging space. The **npswarn** parameter specifies the number of paging-space pages available at which the system begins warning processes that paging space is low. The **npskill** parameter specifies the number of paging-space pages available at which the system begins stopping processes to release paging space.

**Note:** Options **-o**, **-d**, and **-D**, which attempt to change the value of a virtual memory manager tunable parameter, are not supported within a workload partition.

## Understanding the Effect of Changing Tunable Parameters

Misuse of this command can cause performance degradation or operating-system failure. Before you experiment with the **vm0** command, familiarize yourself with both Performance overview of the Virtual Memory Manager and Enhanced JFS file system cache limit with the **maxclient** parameter.

Before modifying any tunable parameter, you should first carefully read about all its characteristics in the Tunable Parameters section below, and follow any Refer To pointer, in order to fully understand its purpose.

You must then make sure that the Diagnosis and Tuning sections for this parameter truly apply to your situation and that changing the value of this parameter could help improve the performance of your system.

If the Diagnosis and Tuning sections both contain only "N/A", you should probably never change this parameter unless specifically directed by AIX development.

## Flags

Item	Description
<b>-a</b>	Displays current, reboot (when used with the <b>-r</b> option), or permanent (when used with the <b>-p</b> option) values for all tunable parameters, one per line in pairs <i>Tunable = Value</i> . For the permanent option, a value is displayed only for a parameter if its reboot and current values are equal. Otherwise, NONE is displayed as the value.
<b>-d</b> <i>Tunable</i>	Resets the <i>Tunable</i> parameter to its default value. If a <i>Tunable</i> parameter, which must be changed because it is not set to its default value, meets one or more of the following sets of criteria, a warning message is displayed and no change is made to the parameter: <ul style="list-style-type: none"><li>• The tunable parameter is of type <b>Bosboot</b> or <b>Reboot</b>.</li><li>• The tunable parameter is of type <b>Incremental</b> and was changed from its default value, and the <b>-r</b> flag is not used in combination.</li></ul>
<b>-D</b>	Resets all <i>Tunable</i> values to their default values. If <i>Tunables</i> that need to be changed because they are not set to their default values meet one or more of the following sets of criteria, a warning message is displayed and no change is made: <ul style="list-style-type: none"><li>• The tunable is of type <b>Bosboot</b> or <b>Reboot</b>.</li><li>• The tunable is of type <b>Incremental</b> and was changed from its default value, and <b>-r</b> is not used in combination.</li></ul>
<b>-F</b>	Forces display of the restricted tunable parameters when the <b>-a</b> , <b>-L</b> or <b>-x</b> options are specified alone on the command line to list all tunables. When the <b>-F</b> flag is not specified, restricted tunables are not displayed, unless these restricted tunables are specifically named with a display option.
<b>-h</b> [ <i>Tunable</i> ]	Displays help about the <i>Tunable</i> parameter if one is specified. Otherwise, displays the <b>vm0</b> command usage statement.

Item	Description
-L [ <i>Tunable</i> ]	<p>Lists the characteristics of one or all tunables, one per line, using the following format:</p> <pre> NAME                CUR    DEF    BOOT  MIN    MAX    UNIT          TYPE   DEPENDENCIES ----- memory_frames      128K      128K              4KB pages      S ----- maxfree            1088  1088   130   16    200K  4KB pages      D   minfree   memory_frames ----- minfree            960    960   122   8     200K  4KB pages      D   maxfree   memory_frames ----- ... where:   CUR = current value   DEF = default value   BOOT = reboot value   MIN = minimal value   MAX = maximum value   UNIT = tunable unit of measure   TYPE = parameter type: D (for Dynamic), S (for Static), R for Reboot),         B (for Bosboot), M (for Mount), I (for Incremental),         C (for Connect), and d (for Deprecated)   DEPENDENCIES = list of dependent tunable parameters, one per line </pre>
-o <i>Tunable</i> [= <i>Newvalue</i> ]	<p>Displays the value or sets tunable to <i>Newvalue</i>. If a tunable must be changed (the specified value is different than the current value), is of type <b>Bosboot</b> or <b>Reboot</b>, or if it is of type <b>Incremental</b> and its current value is greater than the specified value, and <b>-r</b> is not used in combination, the tunable value is not changed but a warning is displayed.</p> <p>When the <b>-r</b> flag is used in combination without a new value, the nextboot value for tunable is displayed. When <b>-p</b> is used in combination without a new value, a value is displayed only if the current and next boot values for the tunable are the same. Otherwise, NONE is displayed as the value.</p>
-p	<p>When used in combination with <b>-o</b>, <b>-d</b> or <b>-D</b>, makes changes apply to both current and reboot values, that is, turns on the updating of the <code>/etc/tunables/nextboot</code> file in addition to the updating of the current value. These combinations cannot be used on <b>Reboot</b> and <b>Bosboot</b> type parameters because their current value can't be changed.</p>
-r	<p>When used with <b>-a</b> or <b>-o</b> without specifying a new value, values are displayed only if the current and next boot values for a parameter are the same. Otherwise NONE is displayed as the value.</p> <p>When the <b>-r</b> flag is used with the <b>-a</b> or <b>-o</b> options without specifying a new value, the values are displayed only if the current and next boot values for a parameter are the same. Otherwise, NONE is displayed as the value. The <b>-r</b> flag changes the reboot values when it is used with the <b>-o</b>, <b>-d</b>, or <b>-D</b> flags. For example, you can update the <code>/etc/tunables/nextboot</code> file when you use the <b>-r</b> flag. If any parameter of type <b>Bosboot</b> is changed, the user will be prompted to run the <b>bosboot</b> command.</p>
-x [ <i>Tunable</i> ]	<p>When used with the <b>-a</b> or the <b>-o</b> options without specifying a new value, next boot values for tunables are displayed instead of current values.</p> <p>Lists characteristics of one or all tunables, one per line, using the following (spreadsheet) format:</p> <pre> tunable,current,default,reboot,min,max,unit,type,{dtunable } </pre> <p>where:</p> <pre> current = current value default = default value reboot = reboot value min = minimal value max = maximum value unit = tunable unit of measure type = parameter type: D (for Dynamic), S (for Static), R (for Reboot),         B (for Bosboot), M (for Mount), I (for Incremental),         C (for Connect), and d (for Deprecated) dtunable = list of dependent tunable parameters </pre>
-y	<p>Suppresses the confirmation prompt before running the <b>bosboot</b> command.</p>

If a restricted tunable parameter is changed, a warning message is displayed that indicates that a tunable of the restricted use type has been modified. If the **-r** or the **-p** option is specified, you are prompted to confirm the change. In addition, at system reboot, restricted tunables that are displayed in the

`/etc/tunables/nextboot` file and are changed to values that are different from their default values (using a command line specifying the `-r` or `-p` option) causes an error log entry that identifies the list of these changed tunables.

When modifying tunable, the tunable value might be specified using abbreviations such as K, M, G, T, P and E to indicate units. See the following lists for abbreviations and their correspondent values:

- K=2<sup>10</sup>
- M=2<sup>20</sup>
- G=2<sup>30</sup>
- T=2<sup>40</sup>
- P=2<sup>50</sup>
- E=2<sup>60</sup>

Thus, a tunable value of 1024 might be specified as 1K.

Any change (with `-o`, `-d` or `-D`) to a parameter of type Mount will result in a message being displayed to warn the user that the change is only effective for future mountings.

Any change (with `-o`, `-d` or `-D` flags) to a parameter of type Connect will result in `inetd` being restarted, and a message displaying a warning to the user that the change is only effective for future socket connections.

Any attempt to change (with `-o`, `-d` or `-D`) a parameter of type **Bosboot** or **Reboot** without `-r`, will result in an error message.

Any attempt to change (with `-o`, `-d` or `-D` but without `-r`) the current value of a parameter of type Incremental with a new value smaller than the current value, will result in an error message.

### Tunable Parameters Type

All the tunable parameters manipulated by the tuning commands (**no**, **nfso**, **vmo**, **ioo**, **raso**, and **schedo**) have been classified into these categories:

Item	Description
Dynamic	If the parameter can be changed at any time
Static	If the parameter can never be changed
Reboot	If the parameter can only be changed during reboot
Bosboot	If the parameter can only be changed by running <code>bosboot</code> and rebooting the machine
Mount	If changes to the parameter are only effective for future file systems or directory mounts
Incremental	If the parameter can only be incremented, except at boot time
Connect	If changes to the parameter are only effective for future socket connections
Deprecated	If changing this parameter is no longer supported by the current release of AIX.

For parameters of type Bosboot, whenever a change is performed, the tuning commands automatically prompt the user to ask if they want to execute the **bosboot** command. For parameters of type Connect, the tuning commands automatically restart the **inetd** daemon.

Note that the current set of parameters managed by the **vmo** command only includes Static, Dynamic, and Bosboot types.

### Compatibility Mode

When running in compatibility mode (controlled by the **pre520tune** attribute of **sys0**), reboot values for parameters, except those of type **Bosboot**, are not really meaningful because in this mode they are not applied at boot time. For more information, see *Performance management*.

In compatibility mode, you can set reboot values to tuning parameters by imbedding calls to tuning commands in scripts called during the boot sequence. Parameters of type **Reboot** can be set without using the **-r** flag, so that existing scripts continue to work.

## Tunable Parameters

To view the default and range of values allowed for the tunables, run the **vmo** command with the **-h** option as follows:

```
vmo -h <tunable_parameter_name>
```

Tunable	Description
<b>ame_cpus_per_pool</b>	<p><b>Purpose</b></p> <p>Determines the ratio of CPUs per compressed memory pool. For every <i>ame_cpus_per_pool</i> CPUs, at least one compressed memory pool is created.</p> <p><b>Tuning</b></p> <p>Lower ratios are used to reduce contention on compressed memory pools. This ratio is not the only factor used to determine the number of compressed memory pools (amount of memory and the layout is also considered) so certain changes to this ratio may not result in any change to the number of compressed memory pools.</p>
<b>ame_maxfree_mem</b>	<p><b>Purpose</b></p> <p>Specifies the average amount of free memory in a compressed memory pool free list at which the VMM will shrink the compressed pool.</p> <p><b>Tuning</b></p> <p>Excessive shrink and grow operations can occur if compressed memory pool size tends to change significantly. This can occur if the workload working set size frequently changes. Increase this tunable to raise the threshold at which the VMM will shrink a compressed memory pool and thus reduce the number of overall shrink and grow operations.</p>
<b>ame_min_ucpool_size</b>	<p><b>Purpose</b></p> <p>Defines the minimum size of the uncompressed pool.</p> <p><b>Tuning</b></p> <p>If compressed memory pool grows too large, there may not be enough space in memory to house uncompressed memory which can slow down application performance due to excessive use of the compressed memory pool. Increase this value to limit the size of the compressed memory pool and make more uncompressed pages available.</p>

Tunable	Description
<b>ame_minfree_mem</b>	<p><b>Purpose</b></p> <p>Specifies the amount of free memory in a compressed memory pool free list at which the VMM will grow the compressed pool.</p> <p><b>Tuning</b></p> <p>If processes are being delayed waiting for compressed memory to become available, increase <i>ame_minfree_mem</i> to improve response time. Note, that this must be at least 64 KB less than <i>ame_maxfree_mem</i>.</p>
<b>ame_mpsize_support</b>	<p><b>Purpose</b></p> <p>Enables all supported page sizes in an Active Memory Expansion (AME) environment for POWER8<sup>®</sup> processor-based servers, or later, which supports the 64 KB accelerator.</p> <p><b>Tuning</b></p> <p>A value of 0 enables legacy behavior in an AME environment. In this case, only 4 KB and 16 MB page sizes are enabled. A value of 1 enables all supported page sizes in an AME environment. You can change this tunable parameter only in POWER8 processor-based servers, or later, which supports the 64 KB accelerator.</p>
<b>ams_loan_policy</b>	<p><b>Purpose</b></p> <p>This tunable toggles the loaning behavior when shared memory mode is enabled.</p> <p><b>Tuning</b></p> <p>When the tunable is set to 0, loaning is disabled. When set to 1, loaning of file cache is enabled. When set to 2, loaning of any type of data is enabled. In response to low memory in the AMS pool, the VMM will free memory and loan it to the hypervisor.</p>
<b>force_realias_lite</b>	<p><b>Purpose</b></p> <p>If set to 0, a heuristic is used, when tearing down a <b>mmapped</b> region, to determine when to avoid locking the source <b>mmapped</b> segment</p> <p><b>Tuning</b></p> <p>This is a scalability tradeoff, controlled by <i>realias_percentage</i>, possibly costing more compute time used. If set to 1, the source segment lock is avoided whenever possible, regardless of the value of <i>realias_percentage</i>.</p>
<b>kernel_heap_psize</b>	<p><b>Purpose</b></p> <p>Specifies the default page size to use for the kernel heap.</p> <p><b>Tuning</b></p> <p>This is an advisory setting. Support for 64 KB pages is provided by POWER5+ and later machines and used when <i>vmm_mpsize_support</i> is enabled. The 16 MB pages, provided by POWER4 and later machines, should only be used for the kernel heap under high performance environments. A value of 0 indicates that the kernel will use the preferred default value of 64 KB, if that page size is supported, else 4 KB pages are used.</p>

Tunable	Description
lgpg_regions	<p><b>Purpose</b></p> <p>Specifies the number of large pages to reserve for implementing with the <code>shmget()</code> system call with the SHM_LGPAGE flag</p> <p><b>Tuning</b></p> <p>The <code>lgpg_size</code> parameter must also be used in addition to this option. The application must be modified to specify the SHM_LGPAGE flag when calling <code>shmget()</code>. This will improve performance in the case where there are many TLB misses and large amounts of memory is being accessed.</p> <p>Although this parameter is Dynamic on DLPAR-capable systems, the <code>nextboot</code> value is written into the boot image when a <code>bosboot</code> command is run so that the setting is optimally restored at reboot.</p>
lgpg_size	<p><b>Purpose</b></p> <p>Specifies the size in bytes of the hardware-supported large pages used for the implementation for the <code>shmget()</code> system call with the SHM_LGPAGE flag.</p> <p><b>Tuning</b></p> <p>Supported on systems from POWER4 onwards. Although this parameter is Dynamic on DLPAR-capable systems, the <code>nextboot</code> value is written into the boot image when a <code>bosboot</code> command is issued so that the setting is optimally restored at reboot. The <code>lgpg_regions</code> parameter must be set to a non-zero value in addition to this parameter. The application must be modified to specify the SHM_LGPAGE flag when calling the <code>shmget()</code> subroutine. This will improve the performance in the case where there are many TLB misses and large amounts of memory is being accessed.</p>
low_ps_handling	<p><b>Purpose</b></p> <p>Specifies the action to change the system behavior in relation to process termination during low paging space conditions.</p> <p><b>Tuning</b></p> <p>A value of 1 indicates current behavior of process termination on low paging space. A value of 2 indicates a new behavior where processes with SIGDANGER handler will be killed, if no other processes were found earlier to recover from low paging space condition.</p>
maxfree	<p><b>Purpose</b></p> <p>Specifies the number of frames on the free list at which page-stealing is to stop.</p> <p><b>Tuning</b></p> <p>Observe free-list-size changes with <code>vmstat -n</code> command. If the <code>vmstat -n</code> command displays the free-list size frequently driven below <code>minfree</code> by application demands, increase the <code>maxfree</code> value to reduce calls to replenish the free list. Setting the value too high causes page replacement to run for a longer period of time. The difference between <code>maxfree</code> and <code>minfree</code> should be of the order of <code>maxpagehead</code>, and no less than 8.</p>



Tunable	Description
<b>maxpin%</b>	<p><b>Purpose</b></p> <p>Specifies the maximum percentage of real memory that can be pinned.</p> <p><b>Tuning</b></p> <p>Change if cannot pin memory, although free memory is available. If this value is changed, the new value should ensure that at least 4 MB of real memory will be left unpinned for use by the kernel. The <b>vmo</b> command converts <b>maxpin%</b> to the corresponding <b>maxpin</b> absolute value, which is the value used by the kernel. Change this parameter only in extreme situations, such as maximum-load benchmarking.</p> <p>This dynamic parameter will have its <b>nextboot</b> value written into the boot image if a <b>bosboot</b> command is issued.</p>
<b>memory_frames</b>	<p><b>Purpose</b></p> <p>Number of valid memory frames.</p> <p><b>Tuning</b></p> <p>N/A</p>
<b>memplace_data</b>	<p><b>Purpose</b></p> <p>Specifies the default memory placement policy for data.</p> <p><b>Tuning</b></p> <p>Refers to the data of the main executable (initialized data, BSS), heap, shared library and object modules loaded at run-time. Data placement can be set to first-touch (value of 1), round-robin across the system (value of 2) or automatic (value of 0), where the system decides the best placement for the memory.</p>
<b>memplace_mapped_file</b>	<p><b>Purpose</b></p> <p>Specifies the default memory placement policy for files that are mapped into the address space of a process (such as through <b>shmat()</b> and <b>mmap()</b>).</p> <p><b>Tuning</b></p> <p>Default placement of memory mapped files can be set to first-touch (value of 1) or round-robin across the system (value of 2) or automatic (value of 0), where the system decides the best placement for the memory.</p>

Tunable	Description
<b>memplace_shm_anonymous</b>	<p><b>Purpose</b></p> <p>Specifies the default memory placement policy for anonymous shared memory.</p> <p><b>Tuning</b></p> <p>Anonymous shared memory refers to working storage memory, created via <b>shmget()</b> or <b>mmap()</b>, that can be accessed only by the creating process or its descendants. This memory is not associated with a name (or key). Default placement of anonymous shared memory can be set to first-touch (value of 1) or round-robin across the system (value of 2) or automatic (value of 0), where the system decides the best placement for the memory.</p>
<b>memplace_shm_named</b>	<p><b>Purpose</b></p> <p>Specifies the default memory placement policy for named shared memory.</p> <p><b>Tuning</b></p> <p>Named shared memory refers to working storage memory, created via <b>shmget()</b> or <b>shm_open()</b>, which is associated with a name (or key) that allows more than one process to access it simultaneously. Default placement of named shared memory can be set to first-touch (value of 1) or round-robin across the system (value of 2) or automatic (value of 0), where the system decides the best placement for the memory.</p>
<b>memplace_stack</b>	<p><b>Purpose</b></p> <p>Specifies the default memory placement policies for the program stack.</p> <p><b>Tuning</b></p> <p>Stack placement can be set to first-touch (value of 1) or round-robin across the system (value of 2) or automatic (value of 0), where the system decides the best placement for the memory.</p>
<b>memplace_text</b>	<p><b>Purpose</b></p> <p>Specifies the default memory placement policy for the application text.</p> <p><b>Tuning</b></p> <p>This applies only to the text of the main executable and not to its dependencies. Text placement can be set to first-touch (value of 1) or round-robin across the system (value of 2) or automatic (value of 0), where the system decides the best placement for the memory.</p>

Tunable	Description
<b>memplace_unmapped_file</b>	<p><b>Purpose</b></p> <p>Specifies the default memory placement policy for unmapped file access, such as through <b>read()/write()</b>.</p> <p><b>Tuning</b></p> <p>Default placement of unmapped file access can be set to first-touch (value of 1) or round-robin across the system (value of 2) or automatic (value of 0), where the system decides the best placement for the memory.</p>
<b>minfree</b>	<p><b>Purpose</b></p> <p>Specifies the number of frames on the free list at which the VMM starts to steal pages to replenish the free list.</p> <p><b>Tuning</b></p> <p>Page replacement occurs when the number of free frames reaches <i>minfree</i>. If the processes are being delayed by page stealing, increase <i>minfree</i> to improve response time. The difference between <i>maxfree</i> and <i>minfree</i> should be of the order of <i>maxpagehead</i>, and no less than 8.</p>
<b>minperm%</b>	<p><b>Purpose</b></p> <p>Specifies the point below which the page-stealer will steal file or computational pages regardless of repaging rates.</p> <p><b>Tuning</b></p> <p>You can decrease this parameter if large number of file pages in memory is causing working storage pages to be replaced. On the other hand, if some files are known to be read repetitively, and I/O rates do not decrease with time from startup, <i>minperm</i> may be too low.</p>
<b>nokilluid</b>	<p><b>Purpose</b></p> <p>The user IDs lower than this value will be exempt from getting killed due to low page-space conditions.</p> <p><b>Tuning</b></p> <p>A value of 0 indicates off. Useful when system is out of paging space and the system administration processes are being killed. Either set this tunable to 1 in order to protect specific user ID processes from getting killed due to low page space or ensure there is sufficient paging space available.</p>
<b>npskill</b>	<p><b>Purpose</b></p> <p>Specifies the number of free paging-space pages at which the operating system begins killing processes.</p> <p><b>Tuning</b></p> <p>The default value is the maximum of 64 and (number of paging space pages)/128. The <i>npskill</i> value must be greater than zero and less than the total number of paging space pages on the system.</p>

Tunable	Description
npswarn	<p><b>Purpose</b></p> <p>Specifies the number of free paging-space pages at which the operating system begins sending the SIGDANGER signal to processes.</p> <p><b>Tuning</b></p> <p>The default value is the maximum of 512 and <math>(4 * npskill)</math>. The value of <i>npswarn</i> must be greater than zero and less than the total number of paging space pages on the system. Increase the value if you experience processes being killed because of low paging space.</p>
numpsblks	<p><b>Purpose</b></p> <p>Total number of paging-space blocks.</p> <p><b>Tuning</b></p> <p>N/A</p>
pinnable_frames	<p><b>Purpose</b></p> <p>Number of pages available for pinning</p> <p><b>Tuning</b></p> <p>N/A</p>
relalias_percentage	<p><b>Purpose</b></p> <p>If <i>force_relalias_lite</i> is set to 0, then this specifies the factor used in the heuristic to decide whether to avoid locking the source <b>mmapped</b> segment.</p> <p><b>Tuning</b></p> <p>This is used when tearing down an <b>mmapped</b> region and is a scalability statement, where avoiding the lock may help system throughput, but, in some cases, at the cost of more compute time used. If the number of pages being unmapped is less than this value divided by 100 and multiplied by the total number of pages in memory in the source <b>mmapped</b> segment, then the source lock is avoided. A value of 0 for <i>relalias_percentage</i>, with <i>force_relalias_lite</i> also set to 0, will cause the source segment lock to always be taken. Effective values for <i>relalias_percentage</i> will vary by workload, however, a suggested value is 200.</p>
scrub	<p><b>Purpose</b></p> <p>Enables or Disables freeing of paging space disk blocks from pages in memory for Deferred Page Space Allocation Policy pages.</p> <p><b>Tuning</b></p> <p>A value of 0 disables scrubbing completely. A value of 1 enables scrubbing of in memory paging space disk blocks when the number of system free paging space blocks is below <b>npsscubmin</b>, and continues until above <b>npsscubmax</b>.</p>

Tunable	Description
v_pinshm	<p><b>Purpose</b></p> <p>If set to 1, will allow pinning of shared memory segments.</p> <p><b>Tuning</b></p> <p>A value of 0 indicates off. Change this value when the overhead is high and in pinning or unpinning of AIO buffers from shared memory segments. Useful only if the application also sets the SHM_PIN flag when doing a <code>shmget()</code> call and if doing async I/O from shared memory segments.</p>
vmm_default_pspa	<p><b>Purpose</b></p> <p>This tunable controls the default aggressiveness of page size promotion. The value is an abstract aggressiveness weighting which is treated by the operating system as the inverse of the page promotion threshold.</p> <p><b>Tuning</b></p> <p>A value of 0 for the <code>vmm_default_pspa</code> setting is equivalent to a page promotion threshold of 100%, that is, a memory range must have 100% real memory occupancy in order to be promoted. A value of 100 for the <code>vmm_default_pspa</code> setting is equivalent to a page promotion threshold of 0%, that is, a memory range should be promoted immediately on first reference to memory in the range. A value of -1 for the <code>vmm_default_pspa</code> setting is equivalent to a page promotion threshold of -1, that is, never do page promotion for a memory range. Page size promotion thresholds are only considered at segment creation time. Thus, changing <code>vmm_default_pspa</code> will only affect the page size promotion thresholds for segments created after the tunable is adjusted.</p>
wlm_memlimit_nonpg	<p><b>Purpose</b></p> <p>Selects whether non-pageable page sizes (16M, 16G) are included in the WLM <code>realmem</code> and <code>virtmem</code> counts. If 1 is selected, then non-pageable page sizes are included in the <code>realmem</code> and <code>virtmem</code> limits count. If 0 is selected, then only pageable page sizes (4K, 64K) are included in the <code>realmem</code> and <code>virtmem</code> counts. This value can only be changed when WLM Memory Accounting is off, or the change will fail.</p> <p><b>Tuning</b></p> <p>When this tunable is set to 0, WLM virtual and real memory limits will only apply to pageable pages consumed by a WLM class. Because heavy use of pageable pages is what causes paging on a system, a value of 0 provides more granular control over how much a WLM class pages when non-pageable pages are in use. This tunable should only be adjusted when WLM real or virtual memory limits are being used on a system configured with non-pageable pages.</p>

## Security

**Attention RBAC users and Trusted AIX users:** This command can perform privileged operations. Only privileged users can run privileged operations. For more information about authorizations and privileges, see Privileged Command Database in *Security*. For a list of privileges and the authorizations associated with this command, see the `lssecattr` command or the `getcmdattr` subcommand.

## Examples

1. To list the current and reboot value, range, unit, type and dependencies of all tunable parameters managed by the **vmo** command, enter:  

```
vmo -L
```
2. To turn on and reserve 16MB large pages on a POWER4 system, enter:  

```
vmo -o lpgg_regions=10 -o lpgg_size=16777216
```

This command will propose **bosboot** to the user, and warn that a reboot is necessary before the change will be effective.
3. To display help on **nokilluid**, enter:  

```
vmo -h nokilluid
```
4. To turn on **v\_pinshm** after the next reboot, enter:  

```
vmo -r -o v_pinshm=1
```
5. To permanently reset all **vmo** tunable parameters to default, enter:  

```
vmo -p -D
```
6. To list the reboot value for all virtual memory manager tuning parameters, enter:  

```
vmo -r -a
```
7. To list (spreadsheet format) the current and reboot value, range, unit, type and dependencies of all tunable parameters managed by the **vmo** command, enter:  

```
vmo -x
```

### Related information:

ioo command

schedo command

no command

nfso command

raso command

---

## vmstat Command

### Purpose

Reports virtual memory statistics.

### Syntax

```
vmstat [ -f ] [ -i ] [ -s ] [ -I[-W]] [ -t ] [ -v ] [ -h ] [ -w ] [ -l ] [-c] [ -@ wparname | ALL] [{ -p | -P }  
physicalvolume | ALL] ALL [-S power] [physicalvolume] [interval] [count]
```

**Note:** Do not use the *wparname* parameter and the **-i** flag together inside workload partitions.

### Description

The **vmstat** command reports statistics about kernel threads, virtual memory, disks, hypervisor pages, traps, and processor activity. Reports that are generated by the **vmstat** command can be used to balance system load activity. These system-wide statistics (among all processors) are calculated as averages for values that are expressed as percentages, and as sums otherwise. The **vmstat** command might return inconsistent statistics because the statistics are not read atomically.

If you run the **vmstat** command without flags, the report contains a summary of the virtual memory activity since system startup. If you specify the **-f** flag, the **vmstat** command reports the number of forks since system startup. The *physicalvolume* parameter specifies the name of the physical volume.

The *interval* parameter specifies the amount of time in seconds between each report. If you do not specify the *interval* parameter, the **vmstat** command generates a single report that contains statistics for the time since system startup and then exits. You can specify the *count* parameter only with the *interval* parameter. If you specify the *count* parameter, its value determines the number of reports that are generated and the number of seconds apart. If you specify the *interval* parameter without the *count* parameter, reports are continuously generated. Do not specify a value of zero to the *count* parameter.

The kernel maintains statistics for kernel threads, paging, and interrupt activity, which the **vmstat** command accesses by using the *perfstat* kernel extension. The disk input/output statistics are maintained by device drivers. For disks, the average transfer rate is determined by using the active time and number of times information is being transferred. The percent active time is computed from the amount of time the drive is busy during the report.

The **vmstat** command reports the number of physical processors consumed (pc), and the percentage of entitlement consumed (ec), in Micro-Partitioning<sup>®</sup> environments. These metrics display on the Micro-Partitioning environments.

The report that is generated by the **vmstat** command contains a system configuration row and column headings. If the *-@* flag is specified, the report consists of system configuration and WPAR configuration. The system configuration row has the following values:

**lcpu** Indicates the number of logical processors.

**mem** Indicates the amount of memory.

**tnem** Indicates the true memory size of the LPAR.

**Note:** This flag is available only when *-c* option is provided and Active Memory<sup>™</sup> Expansion is enabled.

**ent** Indicates the entitled capacity. Displays only when the partition is running with shared processor.

**drives** Indicates the number of disks. Displays only when physical volume name is monitored.

#### **WPARs**

Indicates the number of active workload partitions. It is displayed only when the *-@* flag is specified.

#### **memlim**

Indicates the limit of the memory resource of the workload partition. The limit is in megabytes (MB). This information is displayed only for the WPAR with enforced memory resource limit.

#### **cpulim**

Indicates the limit of processor resource of the workload partition in processor units. This information is displayed only for the WPAR with enforced processor resource limit.

**rset** Indicates the type of the **rset** registry that is associated with a WPAR. The type can be regular or exclusive. This information is displayed only for the WPARs that are associated with a **rset** registry.

#### **mmode**

Indicates memory mode. This metric is displayed automatically in a system with Active Memory Sharing enabled. This metric is also displayed when *-c* option is used.

**mpsz** Size of the memory pool in gigabytes. This metric is displayed only in shared-memory mode.

The column headings and their descriptions follow:

**WPAR:** Information about workload partitions. It displays only when the *-@* flag is specified.

**WPAR** Workload partition name.

**Notes:**

1. The *system* WPAR name indicates system-wide statistics. The *global* WPAR name indicates the statistics belong to Global only.
2. When the **vmstat** command is started with the **-@ ALL** option and the WPAR specific information is not available for a metric, then an en dash sign (-) is displayed instead of a value.
3. When the **vmstat** command is started with the **-@ wparname** or started inside a WPAR, if the WPAR information is not available for a metric, then that metric is marked with the at sign (@), and the system-wide value is displayed for that metric.
4. If a metric is not supported, then an en dash sign (-) is displayed instead of a value.

**kthr:** Information about kernel thread states.

- r** Average number of runnable kernel threads over the sampling interval. Runnable threads consist of the threads that are ready but still waiting to run, and the threads that are already running.
- b** Average number of kernel threads that are placed in the Virtual Memory Manager (VMM) wait queue (awaiting resource, awaiting input/output) over the sampling interval.

**Memory:** Information about the usage of virtual and real memory. Virtual pages are considered active if they are accessed. A page is 4096 bytes.

- avm** Active virtual pages.
- fre** Size of the free list.

**Note:** A large portion of real memory is used as a cache for file system data. It is not unusual for the size of the free list to remain small.

**Page:** Information about page faults and paging activity. This information is averaged over the interval and given in units per second.

- re** Pager input/output list.
- pi** Pages that are paged in from paging space.
- po** Pages paged out to paging space.
- fr** Pages freed (page replacement).
- sr** Pages that are scanned by page-replacement algorithm.
- cy** Clock cycles by page-replacement algorithm.

**Faults:** Trap and interrupt rate averages per second over the sampling interval.

- in** Device interrupts.
- sy** System calls.
- cs** Kernel thread context switches.

**CPU:** Breakdown of percentage usage of processor time.

- us** User time.  
If the current physical processor consumption of the uncapped partitions exceeds the entitled capacity, the percentage becomes relative to the number of physical processor consumed (pc).
- sy** System time.  
If the current physical processor consumption of the uncapped partitions exceeds the entitled capacity, the percentage becomes relative to the number of physical processor consumed (pc).



- id** Processor idle time.  
If the current physical processor consumption of the uncapped partitions exceeds the entitled capacity, the percentage becomes relative to the number of physical processor consumed (pc).
- wa** Processor idle time during which the system had outstanding disk/NFS I/O request.  
If the current physical processor consumption of the uncapped partitions exceeds the entitled capacity, the percentage becomes relative to the number of physical processor consumed (pc).
- pc** Number of physical processors used. Displayed only if the partition is running with shared processor.
- ec** The percentage of entitled capacity that is consumed. Displayed only if the partition is running with shared processor. Because the time base over which this data is computed can vary, the entitled capacity percentage can sometimes exceed 100%. This excess is noticeable only with small sampling intervals.
- rc** The percentage of processor resource that is used. This information is displayed only for the WPARs with enforced processor resource limit.

**Disk:** Provides the number of transfers per second to the specified physical volumes that occurred in the sample interval. The *physicalvolume* parameter can be used to specify one to four names. Transfer statistics are given for each specified drive in the order specified. This count represents requests to the physical device. It does not imply an amount of data that was read or written. Several logical requests can be combined into one physical request. If the *physicalvolume* parameter is used, the physical volume names are printed at the beginning of command execution.

If the **-I** flag is specified, an I/O oriented view is presented with the following column changes.

- kthr** The **p** column is displayed in addition to the **r** and **b** columns.
- p** Average number of threads waiting for I/O messages from raw devices. Raw devices are the devices that are directly attached to the system.  
If the **-W** flag is specified along with the **-I** flag, an additional **w** column is also displayed along with the **r**, **b**, and **p** flags.
- w** Number of threads per second of time that are waiting for the file system direct I/O event to occur. These events include the following types:
- Asynchronous I/O (AIO)
  - Buffer cache subsystem
  - Concurrent I/O (CIO)
  - File system direct I/O
  - NFS subsystem
  - A thread is waiting for an action from the virtual memory manager (VMM) waiting list.
- page** New **fi** and **fo** columns are displayed instead of the **re** and **cy** columns.
- fi** File page-ins per second.
- fo** File page-outs per second.

If the **-c** flag is specified, Active Memory Expansion view is presented with the following column changes.

**memory**

The columns **csz**, **cfr**, and **dxm** are displayed besides columns **avm** and **fre**.

- csz** Current compressed pool size, in 4K page units.

- cfp** Free pages available in compressed pool, in 4K page units.
- dxm** Deficit in Expanded Memory Size, in 4K page units.
- page** New columns **ci** and **co** are displayed instead of **re** and **cy** columns.
- ci** Number of page-ins per second from compressed pool.
- co** Number of page-outs per second to compressed pool.

If while the **vmstat** command is running, there is a change in system configuration that affects the output, **vmstat** prints a warning message about the configuration change. It then continues the output after printing the updated system configuration information and the header.

If the **-l** flag is specified, an additional "large-page" section is displayed with the following columns:

- alp** Indicates the number of large pages currently in use.
- flp** Indicates the number of large pages on the large page freelist.

If the **-p** option is specified, additional lines of VMM statistics are displayed for the specified page sizes. With **-I** and **-t** options, the **-p** option produces an additional line for the specified page size. This line contains the following VMM statistics relevant to the specified page size:

- **avm**
- **fre**
- **re**
- **fi**
- **fo**
- **pi**
- **po**
- **ci**
- **co**
- **fr**
- **sr**
- **cy**

#### Notes:

1. The display of the **re**, **fi**, **fo**, and **cy** options are affected by the **-I** option.
2. The display of the **re**, **ci**, **co**, and **cy** options are affected by the **-c** option.
3. If there is no resource control, then the **avm** and **fre** options are system-wide. Therefore, with the **-@** option set, both the **avm** and **fre** options will be marked with the at sign (@).

These VMM statistics are preceded by a **psz** column and followed by an **siz** column. The description of these two columns follows:

- psz** Page size (for example, 4 KB, 64 KB).
- siz** Number of frames of the specified page size that exist on the system.

With the **-s** option, the **-p** option produces a separate stanza of output that contains only the statistics relevant to the specified page size. This additional stanza is preceded by a page size header.

The **-P** option produces the following report for the specified page size:

- pgsz** Indicates the page size (for example, 4 KB, 64 KB).

#### Memory

Indicates the memory statistics for the specified page sizes.

- siz** The number of frames of the specified page size that exist on the system.
- avm** Active virtual pages applicable to the specified page size.
- fre** Size of the free list for the specified page size.
- Page** Indicates the relevant page faults and paging activity for the specified page size. The page-related columns **re**, **pi**, **po**, **fr**, **sr**, **cy**, **fi**, **fo**, **ci**, and **co** are also applicable to this report.

## Flags

**Note:** If the **-f** (or **-s**) flag is entered on the command line, then the system accepts the **-f** (or **-s**) flag and ignores other flags. If both the **-f** and **-s** flags are specified, the system accepts only the first flag and ignore the second flag.

Item	Description
<b>-@</b> <i>wparname</i>	Reports the Virtual Memory activity of a workload partition: <ul style="list-style-type: none"> <li>• The <b>-@ ALL</b> option indicates that the report pertains to the system and global environment, in addition to all of the workload partitions in the system.  <b>Note:</b> The values that are system-wide statistics are marked with dash sign (-) against the WPAR section.</li> <li>• The <b>-@ wparname</b> flag indicates that the activity is only for that workload partition. In a workload partition, if you specify the <b>-@</b> flag, system-wide statistics and workload partition statistics are displayed. The system-wide statistics are marked with the at sign (@).  <b>Note:</b> Do not use the <b>-@</b> flag with any combination of the <b>-i</b> flag.</li> </ul>
<b>-c</b>	Displays memory compression statistics with the new columns of output, <b>csz</b> , <b>cfr</b> , and <b>dxm</b> under the heading memory, and columns <b>ci</b> and <b>co</b> under the heading page instead of the columns <b>re</b> and <b>cy</b> . <b>Note:</b> This option is available only when Active Memory Expansion is enabled.
<b>-f</b>	Reports the number of forks since system startup.
<b>-i</b>	Displays the number of interrupts that are taken by each device since system startup. <b>Note:</b> The <b>-I</b> , <b>-t</b> , <b>-w</b> , and <b>-l</b> flags are ignored when they are specified with the <b>-i</b> flag.
<b>-I</b>	Displays I/O oriented view with the new columns of output, <b>p</b> under heading kthr, and columns <b>fi</b> and <b>fo</b> under heading page instead of the columns <b>re</b> and <b>cy</b> in the page heading.
<b>-l</b>	Displays an extra "large-page" section with the <b>alp</b> and <b>flp</b> columns.
<b>-p</b> <i>pagesize</i>	Appends the VMM statistics for the specified page size to the regular <b>vmstat</b> output.
<b>-P</b> <i>pagesize</i>	Displays only the VMM statistics, which are relevant for the specified page size.
<b>-s</b>	Writes to standard output the contents of the sum structure, which contains an absolute count of paging events since system initialization. The <b>-s</b> flag can only be used with the <b>-v</b> flag. These events are described as follows: <p><b>address translation faults</b></p> Incremented for each occurrence of an address translation page fault. I/O may or may not be required to resolve the page fault. Storage protection page faults (lock misses) are not included in this count.

**Item**

-s

**Description****backtracks**

Incremented for each page fault that occurs while resolving a previous page fault. (The new page fault must be resolved first and then initial page faults can be *backtracked*.)

**CPU context switches**

Incremented for each processor context switch (dispatch of a new process).

**decrementer interrupts**

Incremented on each decrementer interrupt.

**device interrupts**

Incremented on each hardware interrupt.

**executable-filled page faults**

Incremented for each instruction page fault.

**extend XPT waits**

Incremented each time that a process is waited by VMM due to a commit in progress for the segments accessed.

**free frame waits**

Incremented each time that a process requests a page frame. The free list is empty, and the process is forced to wait while the free list is replenished.

**iodones** Incremented at the completion of each VMM I/O request.

**mpc send interrupts**

Incremented on each mpc send interrupt.

**mpc receive interrupts**

Incremented on each mpc receive interrupt.

**page ins**

Incremented for each page read in by the virtual memory manager. The count is incremented for page ins from page space and file space. Along with the page-out statistic, this value represents the total amount of real I/O initiated by the virtual memory manager.

**page outs**

Incremented for each page that is written out by the virtual memory manager. The count is incremented for page outs to page space and for page outs to file space. Along with the page in statistic, this statistic represents the total amount of real I/O initiated by the virtual memory manager.

**paging space page ins**

Incremented for VMM initiated page ins from paging space only.

**paging space page outs**

Incremented for VMM initiated page outs to paging space only.

**pages examined by the clock**

VMM uses a clock-algorithm to implement a pseudo least recently used (lru) page replacement scheme. Pages are *aged* by being examined by the clock. This count is incremented for each page examined by the clock.

**pages freed by the clock**

Incremented for each page the clock algorithm selects to free from real memory.

**pending I/O waits**

Incremented each time that a process is waited by VMM for a page-in I/O to complete.

**Item****-s****Description****phantom interrupts**

Incremented on each phantom interrupt

**revolutions of the clock hand**

Incremented for each VMM clock revolution (that is after each complete scan of memory).

**start I/Os**

Incremented for each read or write I/O request that is initiated by VMM.

**syscalls** Incremented for each system call.**total reclaims**

Incremented when an address translation fault can be satisfied without initiating a new I/O request. This can occur if the page has been previously requested by VMM, but the I/O has not yet completed; or if the page was pre-fetched by VMM's read-ahead algorithm, but was hidden from the faulting segment; or if the page has been put on the free list and has not yet been reused.

**traps** Not maintained by the operating system.**zero-filled page faults**

Incremented if the page fault is to working storage and can be satisfied by assigning a frame and zero-filling it.

When the **-c** flag is specified along with the **-s** flag, the following additional metrics are displayed.**compressed pool page ins**

Number of page-ins from Compressed Pool since system boot.

**compressed pool page outs**

Number of page-outs to Compressed Pool since system boot.

**-s**When used with the **-p** *pagesize* option, the **-s** option appends the sum structure for the specified page size to the system-wide sum structure. This additional stanza is preceded by a page size header (for example, 4K pages). The following details are not be displayed in this pagesize-based stanza as these statistics are not related to page sizes:

- Processor context switches
- Device interrupts
- Software interrupts
- Decrementer interrupts
- MPC-sent interrupts
- MPC-received interrupts
- Phantom interrupts
- Traps
- Syscalls

**Notes:**

1. When the **-s** flag is used with the **-@ ALL** option, the system-wide statistics are repeated in the workload partition section.
2. When the **-s** flag is used with the *wparname* option, all metrics are reported and the system-wide statistics are marked with the at sign (@).
3. When the **-s** flag is used with the **-l** flag, the **vmstat** command displays the following metric:

**large-page hi water count**

Specifies the maximum value of the large-page inuse count.

**Item****-S** *power***Description**

Multiplies the statistics of the processor with a value of  $10^{\text{power}}$ . The default value of the power is 0.

The following statistics are scaled:

- **us**
- **sy**
- **id**
- **wa**
- **pc**
- **ec**

**Notes:**

1. Do not use the **-S** flag with the **-f**, **-s**, **-i**, **-v**, or **-p** flags.
2. When the **-S** flag is specified, the **us**, **sy**, **id**, and **wa** statistics change. By default, the **us**, **sy**, **id**, and **wa** statistics are relative to the processor consumption of WPAR. When the **-S** flag is specified with a value of power that is not equal to zero, these statistics will be relative to system-wide processor consumption.
3. The value of power for **-S** flag can be only between 0 and 3.

**-t**

Prints the time-stamp next to each line of output of **vmstat**. The time-stamp is displayed in the HH:MM:SS format.

**Note:** Time stamp is not be printed if **-f**, **-s**, or **-i** flags are specified.

**Item**  
**-v**

**Description**

Writes to standard output various statistics maintained by the Virtual Memory Manager. The **-v** flag can only be used with the **-s** and the **-h** flags.

If you specify the **-v** flag, the following statistics are displayed:

**compressed percentage**

Percentage of memory used by compressed pages.

**client filesystem I/Os blocked with no fsbuf**

Number of client filesystem I/O requests blocked because no fsbuf was available. NFS (Network File System) and VxFS (Veritas) are client filesystems. Fsbuf are pinned memory buffers used to hold I/O requests in the filesystem layer.

**client pages**

Number of client pages.

**compressed pages**

Number of compressed memory pages.

**external pager filesystem I/Os blocked with no fsbuf**

Number of external pager client filesystem I/O requests blocked because no fsbuf was available. JFS2 is an external pager client filesystem. Fsbuf are pinned memory buffers used to hold I/O requests in the filesystem layer.

**file pages**

Number of 4 KB pages that are currently used by the file cache.

**free pages**

Number of free 4 KB pages.

**filesystem I/Os blocked with no fsbuf**

Number of filesystem I/O requests blocked because no fsbuf was available. Fsbuf are pinned memory buffers used to hold I/O requests in the filesystem layer.

**lrutable pages**

Number of 4 KB pages that are considered for replacement. This number excludes the pages that are used for VMM internal pages, and the pages that are used for the pinned part of the kernel text.

**maxclient percentage**

Tuning parameter (managed using **vmo**) specifying the maximum percentage of memory, which can be used for client pages.

**maxperm percentage**

Tuning parameter (managed using **vmo**) in percentage of real memory.

**maxpin percentage**

Tuning parameter (managed using **vmo**) specifying the percentage of real memory which can be pinned.

**memory pages**

Size of real memory in number of 4 KB pages.

**memory pools**

Tuning parameter (managed using **vmo**) specifying the number of memory pools.

**minperm percentage**

Tuning parameter (managed using **vmo**) in percentage of real memory.

**numclient percentage**

Percentage of memory that is occupied by client pages.

**Item**  
**-v**

**Description**

*(Statistics that are displayed by -v, continued):*

**numperm percentage**

Percentage of memory that is currently used by the file cache.

**paging space I/Os blocked with no psbuf**

Number of paging space I/O requests that are blocked because the psbuf space is not available. The psbufs space is pinned memory buffers that are used to hold I/O requests at the virtual memory manager layer.

**pending disk I/Os blocked with no pbuf**

Number of pending disk I/O requests blocked because no pbuf was available. Pbufs are pinned memory buffers used to hold I/O requests at the logical volume manager layer

**pinned pages**

Number of pinned 4 KB pages.

**Note:** When the kernel locking feature (`vmm_klock_mode` parameter) is enabled, the pinned pages include the kernel locking (**klocked**) pages. For more information about the kernel locking feature, enter the following command: `vmo -h vmm_klock_mode`.

**remote pageouts scheduled**

Number of pageouts scheduled for client file systems.

If you specify the **-h** flag with the **-v** flag, the following additional metrics are displayed:

**Time resolving virtualized partition memory page faults**

The total time that the virtual partition is blocked to wait for the resolution of its memory page fault. The time is measured in seconds, with millisecond granularity.

**Virtualized partition memory page faults**

The total number of virtual partition memory page faults that are recorded for the virtualized partition.

**Number of 4 KB page frames loaned**

The number of the 4 KB pages of partition's memory loaned to the hypervisor.

**Percentage of partition memory loaned**

The percentage of the partition's memory loaned to the hypervisor.

**Notes:**

1. When the **-v** flag is used with the **-@ ALL** option, the system-wide statistics are not repeated in the workload partition section.
2. When the **-s** flag is used with the *wparname* option, all metrics are reported and the system-wide statistics are marked with the at sign (@).

When the **-c** flag is specified along with the **-v** flag, the following additional metrics are displayed:

**Compressed Pool Size**

Size of Compressed Pool, in 4K page unit.

**percentage of true memory that is used for compressed pool**

Percentage of unexpanded memory that is used for compressed pool.

**free pages in compressed pool (4K pages)**

Number of free pages in compressed pool, in 4K page unit.

**target memory expansion factor**

Target memory expansion factor that is configured for the LPAR.

**achieved memory expansion factor**

Current memory expansion factor achieved.



Item	Description
-h	<p>Displays the <b>hypv-page</b> section that includes the hypervisor page information. The <b>hypv-page</b> section contains the following metrics:</p> <p><b>hpi</b>      Number of hypervisor page-in per second.</p> <p><b>hpit</b>     Average time that is spent in milliseconds per hypervisor page-in.</p> <p><b>pmem</b>    Amount of physical memory that is backing the logical memory of partitions. The value is measured in gigabytes.</p> <p>If you specify the <b>-h</b> flag with the <b>-v</b> flag, the following metrics are displayed in addition to the metrics that are displayed using the <b>-v</b> flag:</p> <p><b>Time resolving virtualized partition memory page faults</b> The total time that the virtual partition is blocked to wait for the resolution of its memory page fault. The time is measured in seconds, with millisecond granularity.</p> <p><b>Virtualized partition memory page faults</b> The total number of virtual partition memory page faults that are recorded for the virtualized partition.</p> <p><b>Number of 4 KB page frames loaned</b> The number of the 4 KB pages of the memory that is loaned to the hypervisor in the partition.</p> <p><b>Percentage of partition memory loaned</b> The percentage of the memory loaned to the hypervisor in the partition.</p>
-w	Displays the report in wide mode.
-W	Displays an extra field <b>w</b> in the <b>kthr</b> section. This option is allowed only with <b>-I</b> flag.

#### Notes:

1. If Active Memory Expansion is enabled, the **vmstat** reports memory statistics in the expanded view. However, if the environment variable **AME\_MEMVIEW** is set to **TRUE**, the memory statistics represent the true view.
2. The **AME\_MEMVIEW** environment variable has no impact on memory statistics reported using the **-c** option.

#### Examples

1. To display a summary of the statistics since boot, enter the following command:  
`vmstat`
2. To display five summaries at 2-second intervals, enter the following command:  
`vmstat 2 5`
3. To display a summary of the statistics since boot including statistics for logical disks **scdisk13** and **scdisk14**, enter the following command:  
`vmstat scdisk13 scdisk14`
4. To display fork statistics, enter the following command:  
`vmstat -f`
5. To display the count of various events, enter the following command:  
`vmstat -s`
6. To display time-stamp next to each column of output of **vmstat**, enter the following command:  
`vmstat -t`
7. To display the I/O oriented view with an alternative set of columns, enter the following command:  
`vmstat -I`
8. To display all the VMM statistics available, enter the following command:  
`vmstat -vs`
9. To display the large-page section with the **alp** and **f1p** columns at 8-second intervals, enter the following command:

```
vmstat -l 8
```

10. To display the VMM statistics specific to a particular page size (in the example, 4 KB), enter the following command:

```
vmstat -p 4K
```

11. To display the VMM statistics for all page sizes that are supported on the system, enter the following command:

```
vmstat -p ALL
```

Or enter the following command:

```
vmstat -p all
```

12. To display only the VMM statistics for a particular page size (in this example, 4 KB), enter the following command:

```
vmstat -P 4K
```

13. To display only the per-page breakdown of VMM statistics for all supported page sizes, enter the following command:

```
vmstat -P ALL
```

Or enter the following command:

```
vmstat -P all
```

14. To display a summary of the statistics for all of the workload partitions after boot, enter the following command:

```
vmstat -@ ALL
```

15. To display all of the VMM statistics available for all of the workload partitions, enter the following command:

```
vmstat -vs -@ ALL
```

16. To display both WPAR and system-wide VMM statistics from a workload partition, enter the following command:

```
vmstat -@
```

17. To multiply the processor values with 10 and display the results, enter the following command:

```
vmstat -S 1
```

18. To display the statistics for the hypervisor page, enter the following command:

```
vmstat -h
```

19. To display the information about pages that are loaned to the hypervisor, enter the following command:

```
vmstat -vh
```

20. To display memory compression statistics (in an LPAR with Active Memory Expansion enabled), enter the following command:

```
vmstat -c
```

21. To display memory compression statistics specific to per-pagesize (in an LPAR with Active Memory Expansion enabled), enter the following command:

```
vmstat -c -P ALL
```

22. To append memory compression information to the statistics displayed by `-soption` (in an LPAR with Active Memory Expansion enabled), enter the following command:

```
vmstat -s -c
```

23. To append memory compression information to the statistics displayed by `-vooption` (in an LPAR with Active Memory Expansion enabled), enter the following command:

```
vmstat -v -c
```

## Files

Item	Description
/usr/bin/vmstat	Contains the <b>vmstat</b> command.

### Related reference:

“vmo Command” on page 34

### Related information:

iostat command

Memory performance

---

## vpdadd Command

### Purpose

Adds entries to the product, lpp, history, and vendor databases.

### Syntax

```
vpdadd { -c Component | -p Product | -f Feature } -v v.r.m.f [ -D Destdir ] [ -U Command ] [ -R Prereq ] [ -S Msg_Set ] [ -M Msg_Number ] [ -C Msg_Catalog ] [ -P Parent ] [ -I Description ]
```

### Description

The **vpdadd** command is for use with or by installers that wish to be listed in Vital Product Database (VPD). The VPD consists of the product, lpp, and history databases. Entries to the inventory database must be added by the **sysck** command. A new vendor database is now included to track products that use destination directories and **non-installp** uninstallers.

The **vpdadd** command uses a tree structure of *Product* at the highest level, then *Feature*, and then *Component*.

The *Component* is the lowest installable unit, but in this hierarchy, a *Component* is not selectable for install or uninstall. Therefore, if an installer is using the **vpdadd** command to update the install database, they should look at their own tree representation and add entries based on their structure. If only adding one entry per install, then adding a *Product* type rather than *Component* type would allow that entry to be listed in the uninstall SMIT interfaces. All the entries are made in the VPD, but *Components* and *Features* are filtered out in the default **lspp** listings (**-Lc**).

### Flags

Item	Description
<b>-C</b> <i>Msg_Catalog</i>	Specifies the message catalog to search for a translated description of the <i>Component</i> . The default (English) description is specified with the <b>-I</b> flag. If the message catalog is not in the standard NLSPATH, then the full path name should be given.
<b>-c</b> <i>Component</i>	Specifies the <i>Component</i> name to add to the VPD. An entry is only added if it is unique. Uniqueness is described as having a different destination directory. If the same instance of a <i>Component</i> is already in the database, then no entry is added, and an error is returned. This allows a force install (that is, reinstall).
<b>-D</b> <i>Destdir</i>	Specifies the root (prefix) path that is added to all the files in a <i>Component</i> when being installed (and when being added to the inventory database by the <b>sysck</b> command). Files in a <i>Component</i> are listed with relative path names, so the root path is allowed to change. The default destination directory is <b>/opt</b> .
<b>-f</b> <i>Feature</i>	Specifies the <i>Feature</i> name to add to the VPD. An entry is only added if it is unique. Uniqueness is described as having a different VRMF or destination directory. If the same instance of a <i>Feature</i> is already in the database, then no entry is added, and an error is not returned. This allows for a force install (that is, reinstall).

Item	Description
<b>-I</b> <i>Description</i>	Specifies the default description of the <i>Component</i> , <i>Feature</i> or <i>Product</i> . The description must be specified in double quotation marks. Single quotation marks are allowed inside the description, and double quotation marks must be prepended with a \.
<b>-M</b> <i>Msg_Number</i>	Specifies the message number for the description.
<b>-P</b> <i>Parent</i>	Specifies the parent software unit. A <i>Component</i> specifies either a <i>Feature</i> or a <i>Product</i> as its parent, depending on where it was in the tree.
<b>-p</b> <i>Product</i>	Specifies the <i>Product</i> name to add to the VPD. An entry is only added if it is unique. Uniqueness is described as having a different VRMF or destination directory. If the same instance of a <i>Product</i> is already in the database, then no entry is added, and an error is not returned. This allows a force install (that is, reinstall).
<b>-R</b> <i>Prereq</i>	Specifies a <i>Component</i> (fileset) that is a requisite of the installing <i>Component</i> . The argument must be specified in quotation marks. This flag can be used more than once to specify multiple prerequisites. Although these are treated as prerequisites at install time (by the installer), they are listed as corequisites in the <i>Product</i> database to avoid creating circular requisite chains.
<b>-S</b> <i>Msg_Set</i>	Specifies the message set (if more than one in the catalog).
<b>-U</b> <i>Command</i>	Specifies the <i>Command</i> to launch the uninstaller for this <i>Component</i> . This may be just a command path name, or it may include parameters if there is a global uninstaller. The <b>geninstall</b> command calls this uninstaller, and <b>installp</b> does not deinstall a fileset if this value is set in the VPD.
<b>-v</b> <i>v.r.m.f</i>	The VRMF of the <i>Component</i> , <i>Feature</i> or <i>Product</i> being added.

## Security

**Attention RBAC users and Trusted AIX users:** This command can perform privileged operations. Only privileged users can run privileged operations. For more information about authorizations and privileges, see Privileged Command Database in *Security*. For a list of privileges and the authorizations associated with this command, see the **Issecattr** command or the **getcmdattr** subcommand.

## Examples

1. The following example shows how the Registry service would call **vpdadd** to add a *Component* for the *Foo product*. This *Component* has two requisites, one that is specific to the operating system, and one that is listed as GUID.

```
vpdadd -c EPL2890198489F -v 1.2.3.0 -R "bos.rte.odm 4.3.3.0" -R "8KDE0KY90245686 1.1.0.0" \
-U /usr/opt/foo/uninstaller.class -p KID892KYLIE25 -I "Foo Database Component"
```

2. To add a new product devices.pci.cool.rte to the VPD, enter:

```
vpdadd -p devices.pci.cool.rte -v 5.1.0.0 -U /usr/sbin/udisetaup
```

## Files

/usr/sbin/vpdadd

### Related information:

RBAC

Trusted AIX

lspp command

geninstall command

---

## vpddel Command

### Purpose

Removes entries from the product, lpp, history, and vendor databases.

### Syntax

```
vpddel { -c Component | -p Product | -f Feature } -v V.R.M.F -D Dest_dir
```

## Description

The **vpddel** command removes entries from the product, lpp, history, and vendor databases. The *vrmf* and destination directory must be specified so that the correct entries are removed.

## Flags

Item	Description
<b>-c</b> <i>Component</i>	Removes the specified <i>Component</i> . The <i>VRMF</i> must also be included when removing a <i>Component</i> .
<b>-D</b> <i>Dest_dir</i>	Specifies the destination directory of the <i>Component</i> to remove. If a destination directory is not included, then the default <b>/opt</b> is used.
<b>-f</b> <i>Feature</i>	Specifies the <i>Feature</i> to remove from the vendor database.
<b>-p</b> <i>Product</i>	The <i>Product</i> to remove from the vendor database.
<b>-v</b> <i>V.R.M.F</i>	Specifies the version, release, modification and fix level of the <i>component</i> to delete from the VPD and vendor database.

## Security

**Attention RBAC users and Trusted AIX users:** This command can perform privileged operations. Only privileged users can run privileged operations. For more information about authorizations and privileges, see Privileged Command Database in *Security*. For a list of privileges and the authorizations associated with this command, see the **Issecattr** command or the **getcmdattr** subcommand.

## Example

To remove the *Component* EPL2890198489F from the product, history, lpp, and vendor databases, type:

```
vpddel -c EPL2890198489F -v 1.2.3.0 -D /usr/lpp/Foo
```

## Files

**/usr/sbin/vpddel**

**Related information:**

lslpp command

RBAC

Trusted AIX

---

## vsdatalst Command

### Purpose

**vsdatalst** – Displays virtual shared disk subsystem information.

### Syntax

```
vsdatalst {-g | -n | -v | -c}
```

### Description

Use this command to display one of several kinds of information to standard output.

You can use the System Management Interface Tool (SMIT) to run the **vsdatalst** command. To use SMIT, enter:

```
smit list_vsd
```

and select the option for the kind of virtual shared disk SDR information you want to see.

## Flags

Only one of the following flags can be specified with each invocation of **vsdata1st**:

- g** Displays the following global volume group data:
  - global\_group\_name*,
  - local\_group\_name*,
  - primary\_server\_node*,
  - secondary\_server\_node*. (This is only enabled with the Recoverable virtual shared disk subsystem.)
  - eio\_recovery*
  - recovery*
  - CVSD server\_list*
  
- n** Displays the following Node data:
  - node\_number*,
  - host\_name*,
  - adapter\_name*,
  - min\_buddy\_buffer\_size*,
  - max\_buddy\_buffer\_size*,
  - max\_buddy\_buffers*.
  
- v** Displays the following definition data:
  - vsd\_name*,
  - logical\_volume\_name*,
  - global\_group\_name*,
  - minor\_number*.
  
- c** Displays the following cluster information:
  - node\_number*
  - cluster\_name*

## Parameters

None.

## Security

You must have root authority to run this command.

## Exit Status

**0** Indicates the successful completion of the command.

**nonzero**

Indicates that an error occurred.

## Restrictions

You must issue this command from a node that is online in the peer domain. To bring a peer domain online, use the **startpdomain** command. To bring a particular node online in an existing peer domain, use the **startpnode** command. For more information on creating and administering an RSCT peer domain, refer to the *RSCT: Administration Guide*.

## Standard Output

Current RVSD subsystem run level.

### Examples

1. To display global volume group date, enter:

```
vsdata1st -g
```

The system displays a message similar to the following:

VSD Global Volume Group Information

Global Volume Group name	Local VG name	Server Node Numbers			recovery	server_list	vsd_type
		primary	backup	eio_recovery			
gpfs0vgg	gpfs0vg	1	2	0	0	0	VSD
gpfs1vgg	gpfs1vg	2	1	0	0	0	VSD
gpfs3vgg	gpfs3vg	1	0	0	0	1:2	CVSD

2. To display global volume group date, enter:

```
vsdata1st -n
```

The system displays a message similar to the following:

VSD Node Information

node number	host_name	VSD adapter	IP packet size	Buddy Buffer		
				minimum size	maximum size	# maxbufs
1	host1	m10	61440	4096	262144	128
2	host2	m10	61440	4096	262144	128

3. To display global volume group date, enter:

```
vsdata1st -v
```

The system displays a message similar to the following:

VSD Table

VSD name	logical volume	Global Volume Group	minor#	size_in_MB
gpfs0vsd	gpfs01v	gpfs0vgg	3	4096
gpfs1vsd	gpfs11v	gpfs1vgg	1	4096
gpfs3vsd	gpfs31v	gpfs3vgg	4	4096

## Location

/opt/rsct/vsd/bin/vsdata1st

**Related information:**

lsvsd command

updatevsdnode command

---

## vsdchgserver Command

### Purpose

**vsdchgserver** – Switches the server function for one or more virtual shared disks from the node that is currently acting as the server node to the other.

### Syntax

**vsdchgserver**

**-g** vsd\_global\_volume\_group\_name **-p** primary\_node

`[-b secondary_node] [-o EIO_recovery]`

## Description

The **vsdchgserver** command allows the serving function for a global volume group defined on a primary node to be taken over by the secondary node, or to be taken over by the primary node from the secondary node. This allows an application to continue to use virtual shared disks in situations where the cable or adapter between the physical disks and one of the attached nodes is not working.

The Recoverable virtual shared disk subsystem automatically updates the virtual shared disk devices if, and only if, the **vsdchgserver** command is used to flip the currently-defined primary node and secondary node in the global volume group specified in the **-g** flag.

## Flags

- g** Specifies the Global Volume Group name for the volume group that represents all the virtual shared disks defined on a particular node.
- p** Specifies the primary server node number for the global volume group.
- b** Specifies the secondary node number for the global volume group. If the **-b** flag is not specified, the secondary node definition will be removed.
- o** Specified as **0**, for no recovery on an EIO error, or **1**, for recovery on an EIO error.

## Parameters

None.

## Security

You must have root authority to run this command.

## Exit Status

**0** Indicates the successful completion of the command.

**nonzero**

Indicates that an error occurred.

## Restrictions

You must issue this command from a node that is online in the peer domain. To bring a peer domain online, use the **starttrpdomain** command. To bring a particular node online in an existing peer domain, use the **starttrpnode** command. For more information on creating and administering an RSCT peer domain, refer to the *RSCT: Administration Guide*.

## Standard Output

Current RVSD subsystem run level.

## Examples

To change the primary server node for the global volume group `node12vg` to node 1 and the secondary node to node 2, with EIO recovery, enter:

```
vsdchgserver -g node12vg -p 1 -b 2 -o 1
```



## Location

/opt/rsct/vsd/bin/vsdchgserver

---

## vsdelnode Command

### Purpose

Removes virtual shared disk information for a node or series of nodes.

### Syntax

```
vsdelnode node_number ...
```

### Description

This command is used to remove virtual shared disk data for a node or series of nodes.

The **vsdelnode** command makes the listed nodes no longer virtual shared disk nodes so that no virtual shared disks can be accessed from them. This command is unsuccessful for any nodes that are servers for any global volume groups.

You can use the System Management Interface Tool (SMIT) to run the **vsdelnode** command. To use SMIT, enter:

```
smit delete_vsd
```

and select the **Delete Virtual Shared Disk Node** Information option.

### Flags

- g** Specifies the Global Volume Group name for the volume group that represents all the virtual shared disks defined on a particular node.
- p** Specifies the primary server node number for the global volume group.
- b** Specifies the secondary node number for the global volume group. If the **-b** flag is not specified, the secondary node definition will be removed.
- o** Specified as **0**, for no recovery on an EIO error, or **1**, for recovery on an EIO error.

### Parameters

*node\_number*

Specifies the node number of the node whose virtual shared disk information you want to remove.

### Security

You must have **root** authority to run this command.

### Restrictions

The recoverable virtual shared disk subsystem must be stopped on the node(s) you are deleting. Otherwise, the results may be unpredictable. For more information, see *RSCT for AIX 5L™ Managing Shared Disks*.

You must issue this command from a node that is online in the peer domain. To bring a peer domain online, use the **starttrpdomain** command. To bring a particular node online in an existing peer domain,

use the **startprnode** command. For more information on creating and administering an RSCT peer domain, refer to *RSCT Administration Guide*.

## Examples

To delete virtual shared disk node information for nodes 3 and 6, enter:

```
vsdelnode 3 6
```

## Location

```
/opt/rsct/vsd/bin/vsdelnode
```

---

## vsdelvg Command

### Purpose

**vsdelvg** – Removes virtual shared disk global volume group information.

### Syntax

```
vsdelvg [-f] global_group_name ...
```

### Description

Use this command to remove virtual shared disk global volume group information. If any virtual shared disks are defined on a global volume group, the **vsdelvg** command is unsuccessful unless **-f** is specified. If **-f** is specified, any such virtual shared disks must be unconfigured and in the defined state on all the virtual shared disk nodes to be deleted.

You can use the System Management Interface Tool (SMIT) to run the **vsdelvg** command. To use SMIT, enter:

```
smit delete_vsd
```

and select the **Delete Virtual Shared Disk Global Volume Group Information** option.

### Flags

**-f** Forces the removal of any virtual shared disks defined on this global volume group.

### Parameters

*global\_group\_name*

Specifies the volume group that you no longer want to be global to the system.

### Security

You must have root authority to run this command.

### Exit Status

**0** Indicates the successful completion of the command.

**nonzero**

Indicates that an error occurred.

## Restrictions

You must issue this command from a node that is online in the peer domain. To bring a peer domain online, use the **startprdomain** command. To bring a particular node online in an existing peer domain, use the **startprnode** command. For more information on creating and administering an RSCT peer domain, refer to the *RSCT: Administration Guide*.

## Standard Output

Current RVSD subsystem run level.

## Examples

To remove the global volume group **vg1n1**, enter:

```
vsdelvg vg1n1
```

## Location

`/opt/rsct/vsd/bin/vsdelvg`

### Related reference:

“vsdvg Command” on page 71

### Related information:

undefvdsd command

---

## vsdnode Command

### Purpose

Define virtual shared disk information for a node or series of nodes.

### Syntax

#### **vsdnode**

*node\_number... adapter\_name min\_buddy\_buffer\_size*

*max\_buddy\_buffer\_size max\_buddy\_buffers*

*vsd\_max\_ip\_msg\_size [cluster\_name]*

### Description

Use this command to make the specified nodes virtual shared disk nodes and to assign their virtual shared disk operational parameters. If this information is the same for all nodes, run this command once. If the information is different for the nodes, run this command once for each block of nodes that should have the same virtual shared disk information.

You can use the System Management Interface Tool (SMIT) to run the **vsdnode** command. To use SMIT, enter:

```
smit vsd_data
```

and select the **virtual shared disk Node Information** option.

### Flags

**-f** Forces the removal of any virtual shared disks defined on this global volume group.

## Parameters

### *node\_number*

Specifies the node or nodes whose virtual shared disk information is to be set. The value you specify for *node\_number* must match a valid RSCCT remote peer domain node number.

### *adapter\_name*

Specifies the adapter name to be used for virtual shared disk communications for the nodes specified. The adapter name must already be defined to the nodes. Note that the nodes involved in virtual shared disk support must be fully connected so that proper communications can take place. Use **ml0** to specify that the virtual shared disk device driver transmits data requests over the SP Switch. The **ml0** adapter will be used the next time the virtual shared disk device driver is loaded.

### *min\_buddy\_buffer\_size*

Specifies the smallest buddy buffer a server uses to satisfy a remote request to a virtual shared disk. This value must be a power of 2 and greater than or equal to 4096. The suggested value is 4096 (4 KB). For a 512 byte request, 4 KB is excessive. However, recall that a buddy buffer is only used for the short period of time while a remote request is being processed at the server node.

### *max\_buddy\_buffer\_size*

Specifies the largest buddy buffer a server uses to satisfy a remote noncached request. This value must be a power of 2 and greater than or equal to the *min\_buddy\_buffer\_size*. The suggested value is 262144 (256 KB). This value depends on the I/O request size of applications using the virtual shared disks and the network used by the virtual shared disk software.

### *max\_buddy\_buffers*

Specifies the number of *max\_buddy\_buffer\_size* buffers to allocate. The virtual shared disk device driver will have an initial size when first loaded, and then will dynamically allocate and reclaim additional space as needed. The suggested value is 2000 256 KB buffers.

Buddy buffers are only used on the servers. On client nodes you may want to set *max\_buddy\_buffers* to 1.

**Note:** The **statvsd** command will indicate if remote requests are queueing waiting for buddy buffers.

### *vsd\_max\_ip\_msg\_size*

Specifies the maximum message size in bytes for virtual shared disks. This value must not be greater than the maximum transmission unit (MTU) size of the network. The recommended values are:

- 61440 (60KB) for a switch
- 8192 (8KB) for jumbo frame Ethernet
- 1024 (1KB) for 1500-byte MTU Ethernet

### *cluster\_name*

A cluster name must be specified for server nodes that will be serving concurrently accessed shared disks. The cluster name can be any user provided name. A node can only belong to one cluster. For example, when you have a concurrent access environment, the two servers for the CVSD must both specify the same cluster name.

**Note:** The *cluster\_name* is required only for SSA (Serial Storage Architecture) disks.

## Security

You must have **root** authority to run this command.

## Restrictions

The node specified on this command must already belong to a peer domain, and you must issue this command from a node that is online in the peer domain. To bring a peer domain online, use the **startdomain** command. To bring a particular node online in an existing peer domain, use the **startnode** command. For more information on creating and administering an RSCT peer domain, refer to *RSCT Administration Guide*.

## Examples

The following example defines information for a switch network and nodes 1 through 8.

```
vsdnode 1 2 3 4 5 6 7 8 m10 4096 262144 128 61440
```

## Location

`/opt/rsct/vsd/bin/vsdnode`

**Related information:**

updatevsdnode command

---

## vsdsklst Command

### Purpose

Produces output that shows you the disk resources used by the virtual shared disk subsystem across a peer domain.

### Syntax

```
vsdsklst [-v] [-d] [-a | -n node_number [, node_number2, ...]]
```

### Description

Use this command to check disk utilization across a peer domain.

### Flags

- v** Displays only disk utilization information about volume groups and the virtual shared disks associated with them.
- d** Displays only disk utilization information about volume groups and the physical disks associated with them.
- a** Displays specified information for all nodes in the system or system partition.
- n *node\_number***  
Lists one or more node numbers for which information is to be displayed.

### Parameters

*node\_number*

Specifies the node or nodes whose virtual shared disk information is to be set. The value you specify for *node\_number* must match a valid RSCT remote peer domain node number.

*adapter\_name*

Specifies the adapter name to be used for virtual shared disk communications for the nodes specified. The adapter name must already be defined to the nodes. Note that the nodes involved in virtual shared disk support must be fully connected so that proper communications can take

place. Use **m10** to specify that the virtual shared disk device driver transmits data requests over the SP Switch. The **m10** adapter will be used the next time the virtual shared disk device driver is loaded.

#### *min\_buddy\_buffer\_size*

Specifies the smallest buddy buffer a server uses to satisfy a remote request to a virtual shared disk. This value must be a power of 2 and greater than or equal to 4096. The suggested value is 4096 (4 KB). For a 512 byte request, 4KB is excessive. However, recall that a buddy buffer is only used for the short period of time while a remote request is being processed at the server node.

#### *max\_buddy\_buffer\_size*

Specifies the largest buddy buffer a server uses to satisfy a remote noncached request. This value must be a power of 2 and greater than or equal to the *min\_buddy\_buffer\_size*. The suggested value is 262144 (256 KB). This value depends on the I/O request size of applications using the virtual shared disks and the network used by the virtual shared disk software.

#### *max\_buddy\_buffers*

Specifies the number of *max\_buddy\_buffer\_size* buffers to allocate. The virtual shared disk device driver will have an initial size when first loaded, and then will dynamically allocate and reclaim additional space as needed. The suggested value is 2000 256KB buffers.

Buddy buffers are only used on the servers. On client nodes you may want to set *max\_buddy\_buffers* to 1.

**Note:** The **statvsd** command will indicate if remote requests are queueing waiting for buddy buffers.

#### *vsd\_max\_ip\_msg\_size*

Specifies the maximum message size in bytes for virtual shared disks. This value must not be greater than the maximum transmission unit (MTU) size of the network. The recommended values are:

- 61440 (60KB) for a switch
- 8192 (8KB) for jumbo frame Ethernet
- 1024 (1KB) for 1500-byte MTU Ethernet

#### *cluster\_name*

A cluster name must be specified for server nodes that will be serving concurrently accessed shared disks. The cluster name can be any user provided name. A node can only belong to one cluster. For example, when you have a concurrent access environment, the two servers for the CVSD must both specify the same cluster name.

**Note:** The *cluster\_name* is required only for SSA (Serial Storage Architecture) disks.

## Security

You must have **root** authority to run this command.

## Restrictions

You must issue this command from a node that is online in the peer domain. To bring a peer domain online, use the **startpdomain** command. To bring a particular node online in an existing peer domain, use the **startpnode** command. For more information on creating and administering an RSCT peer domain, refer to *RSCT Administration Guide*.

## Examples

This command:

```
vsdsklst -dv -a
```

displays the following information on a system that has volume groups and virtual shared disks defined on nodes 1 and 2.

```
c164n12.ppd.pok.ibm.com: Node Number:2; Node Name:c164n12.ppd.pok.ibm.com
c164n12.ppd.pok.ibm.com: Volume group:rootvg; Partition Size:32; Total:271; Free:168
c164n12.ppd.pok.ibm.com:     Physical Disk:hdisk0; Total:271; Free:168
c164n12.ppd.pok.ibm.com:     Volume group:testvg is not varied on.
c164n12.ppd.pok.ibm.com:     Physical Disk:hdisk5;
c164n12.ppd.pok.ibm.com: Volume group:test1vg; Partition Size:4; Total:537; Free:534
c164n12.ppd.pok.ibm.com:     Physical Disk:hdisk2; Total:537; Free:534
c164n12.ppd.pok.ibm.com:     VSD Name:vsd1n2[testnewlv21n2]; Size:1
c164n12.ppd.pok.ibm.com:     VSD Name:vsd2n2[testlv1n2]; Size:346112.25
c164n12.ppd.pok.ibm.com:     VSD Name:vsd3n2[testlv2n2]; Size:346112.25
c164n12.ppd.pok.ibm.com:     Volume group:vg1 is not varied on.
c164n12.ppd.pok.ibm.com:     Physical Disk:hdisk9;
c164n12.ppd.pok.ibm.com:     Volume group:sharkvg is not varied on.
c164n12.ppd.pok.ibm.com:     Physical Disk:hdisk7;
c164n12.ppd.pok.ibm.com:     Physical Disk:hdisk10;
c164n12.ppd.pok.ibm.com: Volume group:bdhclvg; Partition Size:32; Total:134; Free:102
c164n12.ppd.pok.ibm.com:     Physical Disk:hdisk13; Total:134; Free:102
c164n12.ppd.pok.ibm.com: Volume group:gpfs0vg; Partition Size:8; Total:536; Free:0
c164n12.ppd.pok.ibm.com:     Physical Disk:hdisk12; Total:536; Free:0
c164n12.ppd.pok.ibm.com:     VSD Name:gpfs0vsd[gpfs0lv]; Size:352256.75
c164n12.ppd.pok.ibm.com:     Not allocated physical disks:
c164n12.ppd.pok.ibm.com:     Physical disk:hdisk1
c164n12.ppd.pok.ibm.com:     Physical disk:hdisk3
c164n12.ppd.pok.ibm.com:     Physical disk:hdisk4
c164n12.ppd.pok.ibm.com:     Physical disk:hdisk6
c164n12.ppd.pok.ibm.com:     Physical disk:hdisk11
c164n12.ppd.pok.ibm.com:     Physical disk:hdisk15
c164n11.ppd.pok.ibm.com: Node Number:1; Node Name:c164n11.ppd.pok.ibm.com
c164n11.ppd.pok.ibm.com: Volume group:rootvg; Partition Size:32; Total:271; Free:172
c164n11.ppd.pok.ibm.com:     Physical Disk:hdisk0; Total:271; Free:172
c164n11.ppd.pok.ibm.com: Volume group:bdhclvg; Partition Size:32; Total:134; Free:102
c164n11.ppd.pok.ibm.com:     Physical Disk:hdisk9; Total:134; Free:102
c164n11.ppd.pok.ibm.com:     VSD Name:bdhcvsd1n1[1vbdhcvsd1n1]; Size:45056
c164n11.ppd.pok.ibm.com: Volume group:testvg; Partition Size:16; Total:134; Free:70
c164n11.ppd.pok.ibm.com:     Physical Disk:hdisk13; Total:134; Free:70
c164n11.ppd.pok.ibm.com:     Not allocated physical disks:
c164n11.ppd.pok.ibm.com:     Physical disk:hdisk1
c164n11.ppd.pok.ibm.com:     Physical disk:hdisk2
c164n11.ppd.pok.ibm.com:     Physical disk:hdisk3
```

## Location

`/opt/rsct/vsd/bin/vsdsklst`

---

## vsvdvg Command

### Purpose

Defines a virtual shared disk global volume group.

### Syntax

```
vsvdvg [-g global_volume_group] {-l server_list local_group_name | local_group_name primary_node [secondary node [eio_recovery]]}
```

### Description

Use this command to define volume groups for use by the Virtual shared disk subsystem. This is done by specifying the local volume group name, the node on which it resides, and the name by which the volume group will be known throughout the cluster.

You can use the System Management Interface Tool (SMIT) to run the **vsdvg** command. To use SMIT, enter the following command and select the **Virtual Shared Disk Global Volume Group Information** option:

```
smit vsd_data
```

## Flags

**-g** *global\_volume\_group*

Specifies a unique name for the new global volume group. This name must be unique across the system partition. It should be unique across the SP, to avoid any naming conflicts during future system partitioning operations. The suggested naming convention is **vgxxnyy**, where *yy* is the node number, and *xx* uniquely numbers the volume groups on that node. If this is not specified, the local group name is used for the global name. The length of the name must be less than or equal to 31 characters.

**-l** *server\_list*

Define the list of servers for CVSD. More than one server indicates the *global\_volume\_group* is a concurrent volume group.

## Parameters

*local\_group\_name*

Specifies the name of a volume group that you want to indicate as being used for virtual shared disks. This name is local to the host upon which it resides. The length of the name must be less than or equal to 15 characters.

*primary\_node*

Specifies the primary server node number on which the volume group resides. The length of the name must be less than or equal to 31 characters.

*secondary\_node*

Specifies the secondary server node number on which the volume group resides. The length of the name must be less than or equal to 31 characters.

*eio\_recovery*

Specifies how the Recoverable virtual shared disk subsystem will respond to EIO errors. If *eio\_recovery* is set to the value 1 (the default), an EIO error will cause the Recoverable virtual shared disk system to flip the current primary node and the secondary node and perform one more retry on the new primary node.

## Security

You must have root authority to run this command.

## Exit Status

**0** Indicates the successful completion of the command.

**nonzero**

Indicates that an error occurred.

## Restrictions

You must issue this command from a node that is online in the peer domain. To bring a peer domain online, use the **starttrpdomain** command. To bring a particular node online in an existing peer domain, use the **starttrpnode** command. For more information on creating and administering an RSCT peer domain, refer to the *RSCT: Administration Guide*.



## Standard Output

Current RVSD subsystem run level.

## Examples

1. The following example defines **gpfs1gvg** as a virtual shared disk global volume group with the local volume group **gpfs1vg** accessed from **node1** as the *primary\_node* and **node2** as the *secondary\_node*.  

```
vsdvg -g gpfs1gvg gpfs1vg 1 2
```
2. The following example defines **gpfs3gvg** as a virtual shared disk global volume group with the local volume group **gpfs3vg** concurrently accessed from **node1** and **node2**.  

```
vsdvg -g gpfs3gvg -l 1:2 gpfs3vg
```

## Location

`/opt/rsct/vsd/bin/vsdvg`

**Related reference:**

“vsdelvg Command” on page 66

**Related information:**

updatevsdvg command

---

## vsdvgts Command

### Purpose

Updates the timestamp used by the Recoverable virtual shared disk subsystem by reading the timestamp from the volume group descriptor area (VGDA) of the physical disks.

### Syntax

```
vsdvgts [-a] [ volgrp]
```

### Description

Use this command to update the timestamp that the Recoverable virtual shared disk subsystem uses to determine if a twin-tailed volume group has changed. When the subsystem detects a change, the recovery scripts export the volume group and then import the volume group.

This command can be used to avoid exporting the volume group and then importing the volume group during recovery in situations where the export and import operations are not really necessary. This command should be used very carefully.

### Flags

- a Specifies that the timestamps for this volume group for both primary and secondary nodes should be updated. If this flag is not specified, the timestamp is updated on the local node only.

### Parameters

*volgrp* Specifies a volume group. If this operand is not specified, the timestamps for all the volume groups on this node are updated.

### Security

You must have root authority to run this command.

## Exit Status

- 0 Indicates the successful completion of the command.
- 1 Indicates that the program was unable to read one or more timestamps.

## Restrictions

You must issue this command from a node that is online in the peer domain. To bring a peer domain online, use the **startdomain** command. To bring a particular node online in an existing peer domain, use the **startnode** command. For more information on creating and administering an RSCT peer domain, refer to the *RSCT: Administration Guide*.

## Standard Output

Current RVSD subsystem run level.

## Examples

To update the timestamp associated with the virtual shared disk volume group vsdvg1 for just this node, enter:

```
vsdvgtvs vsdvg1
```

## Location

**/usr/lpp/vsd/bin/vsdvgtvs**

### Related information:

updatevsdvg command

---

## W

The following AIX commands begin with the with the letter *w*.

---

### w Command

#### Purpose

Prints a summary of current system activity.

#### Syntax

```
w [ -h ] [ -u ] [ -w ] [ -l | -s [ -X ] [ -@ [ WPAR ] ] [ User ]
```

#### Description

The **w** command prints a summary of the current activity on the system. The summary includes the following:

Item	Description
WPAR	Workload partition name.
User	Who is logged on.
tty	Name of the tty the user is on.
login@	Time of day the user logged on.
idle	Number of minutes since a program last attempted to read from the terminal. <b>Note:</b> The idle time is taken from the global terminal when you log into wpar using the <b>cllogin</b> command.
JCPU	System unit time used by all processes and their children on that terminal.
PCPU	System unit time used by the currently active process.
What	Name and arguments of the current process.

The heading line of the summary shows the current time of day, how long the system has been up, the number of users logged into the system, and the load average. The load average is the number of runnable processes over the preceding 1-, 5-, 15-minute intervals.

The following examples show the different formats used for the login time field:

Item	Description
10:25am	The user logged in within the last 24 hours.
Tue10am	The user logged in between 24 hours and 7 days.
12Mar91	The user logged in more than 7 days ago.

If a user name is specified with the *User* parameter, the output is restricted to that user.

#### Flags

Item	Description
-@	<p>Prints the System activity tagged with a workload partition name:</p> <ul style="list-style-type: none"> <li>• providing the -@ option without a WPAR name indicates the global environment in addition to all WPARs active in the system, and the heading line indicates values for the global environment only</li> <li>• providing the -@ option with a WPAR name indicates the activity, and the heading line indicates values for only that WPAR</li> <li>• providing -@ Global indicates the activity, and the heading line indicates values for the global environment only.</li> </ul> <p><b>Note:</b> Not providing the -@ option indicates that the current WPAR or global environment, wherever the <b>w</b> command is running.</p>
-h	Suppresses the heading.
-l	Prints the summary in long form. This is the default.
-s	Prints the summary in short form. In the short form, the tty is abbreviated, and the login time, system unit time, and command arguments are omitted.
-u	Prints the time of day, amount of time since last system startup, number of users logged on, and number of processes running. This is the default. Specifying the -u flag without specifying the -w or -h flag is equivalent to the uptime command.
-w	The equivalent of specifying the -u and -l flags, which is the default.
-X	Prints all available characters of each user name instead of truncating to the first 8 characters. The user name is also moved to the last column of the output.

## Files

Item	Description
/etc/utmp	Contains the list of users.

### Related reference:

“who Command” on page 92

### Related information:

uptime command

## wait Command

### Purpose

Waits until the termination of a process ID.

### Syntax

```
wait [ ProcessID ... ]
```

### Description

The **wait** command waits (pauses execution) until the process ID specified by the *ProcessID* variable terminates. If the *ProcessID* variable is not specified, the **wait** command waits until all process IDs known to the invoking shell have terminated and exit with a 0 exit status. If a *ProcessID* variable represents an unknown process ID, the **wait** command treats them as known process IDs that exited with exit status 127. The **wait** command exits with the exitstatus of the last process ID specified by the *ProcessID* variable.

### Flag

Item	Description
<i>ProcessID</i>	Specifies an unsigned decimal integer process ID of a command, which the <b>wait</b> command waits on until termination.

## Exit Status

If one or more operands were specified, all of the operands terminated or were not known by the invoking shell, and the status of the last operand specified is known, then the exit status of the **wait** command is the same as the exit status information of the command indicated by the last operand specified. If the process terminated abnormally due to the receipt of a signal, then the exit status is greater than 128 and distinct from the exit status information generated by other signals, although the exact status value is unspecified (see the **kill -l** command option). Otherwise, the **wait** command exits with one of the following values:

Item	Description
0	The <b>wait</b> command was invoked with no operands and all process IDs known by the invoking shell have terminated.
1-126	The <b>wait</b> command detected an error.
127	The command identified by the last <i>ProcessID</i> operand specified is unknown.

## File

Item	Description
<i>/usr/bin/wait</i>	Contains the <b>wait</b> command.

### Related information:

kill command  
sleep command  
pause command

---

## wall Command

### Purpose

Writes a message to all users that are logged in.

### Syntax

```
wall [ -a ] [ -g Group ][ Message ]
```

### Description

The **wall** command writes a message to all users that are logged in. If the *Message* parameter is not specified, the **wall** command reads the message from standard input until it reaches an end-of-file character. The message is then sent to all logged in users. The following heading precedes the message:

```
Broadcast message from  
user@node
```

```
(tty) at hh:mm:ss ...
```

hh:mm:ss represents the hours, minutes, and seconds when the message was sent.

To override any protections set up by other users, you must operate with root user authority. Typically, the root user uses the **wall** command to warn all other users of an impending system shutdown.

### Note:

- The **wall** command only sends messages to the local node.
- Messages can contain multibyte characters.

## Flags

Item	Description
-a	Performs the default operation. This flag is provided for System V compatibility. It broadcast messages to the console and pseudo-terminals.
-g <i>Group</i>	Broadcasts to a specified group only.

## Files

Item	Description
/dev/tty	Specifies a device.

### Related reference:

“write Command” on page 122

“wallevnt Command”

### Related information:

mesg command

## wallevent Command

### Purpose

Broadcasts an event or a rearm event to all users who are logged in.

### Syntax

wallevent [-c] [-h]

### Description

The **wallevent** script broadcasts a message on an event or a rearm event to all users who are currently logged in to the host when the event or the rearm event occurs. Event or rearm event information is captured and posted by the event response resource manager in environment variables that are generated by the event response resource manager when an event or a rearm event occurs. This script can be used as an action that is run by an event response resource. It can also be used as a template to create other user-defined actions. The language in which the messages of the **wallevent** script are returned depend on the locale settings.

Messages are displayed in this format at the consoles of all users who are logged in when an event or a rearm event occurs for which this script is a response action :

Broadcast message from *user@host* (*tty*) at *hh:mm:ss...*

*severity event\_type* occurred for Condition *condition\_name*  
on the resource *resource\_name* of *resource\_class\_name* at *hh:mm:ss mm/dd/yy*  
The resource was monitored on *node\_name* and resided on {*node\_names*}.

Event information is returned about the ERRM environment variables, and also includes the following:

#### Local Time

Time when the event or rearm event is observed. The actual environment variable supplied by ERRM is ERRM\_TIME. This value is localized and converted to readable form before being displayed.

This script captures the environment variable values and uses the **wall** command to write a message to the currently logged-in user consoles.

## Flags

- c** Instructs **wallevent** to broadcast the **ERRM\_VALUE** of an ERRM event. When the **-c** flag is specified, **wallevent** broadcasts the SNMP trap message.
- h** Writes the script's usage statement to standard output.

## Parameters

*log\_file* Specifies the name of the file where event information is logged. An absolute path for the *log\_file* parameter should be specified.

The *log\_file* is treated as a circular log and has a fixed size of 64KB. When *log\_file* is full, new entries are written over the oldest existing entries.

If *log\_file* already exists, event information is appended to it. If *log\_file* does not exist, it is created so that event information can be written to it.

## Exit Status

- 0 Script has run successfully.
- 1 Error occurred when the script was run.

## Restrictions

1. This script must be run on the node where the ERRM is running.
2. The **wall** command is used to write a message to currently logged-in user consoles. Refer to the **wall** man page for more information on the **wall** command.

## Standard Output

When the **-h** flag is specified, the script's usage statement is written to standard output.

## Examples

1. Suppose the **wallevent** script is a predefined action in the critical-notification response, which is associated with the **/var space used** condition on the resource **/var**. The threshold of the event expression defined for this condition is met, and an event occurs. The critical-notification response takes place, and **wallevent** is run. The following message is displayed on the consoles of all users who are logged in:

```
Broadcast message from joe@neverland.com (pts/6) at 18:42:03...
```

```
Critical event occurred for Condition /var space used  
on the resource /var of filesys of IBM.FileSystem at 18:41:50 03/28/02  
The resource was monitored on c174n05 and resided on {c174n05}.
```

2. When a rearm event occurs for the **/var space used** condition on the resource **/var**, the following message is displayed on the consoles of all users who are logged in:

```
Broadcast message from joe@neverland.com (pts/6) at 18:42:03...
```

```
Critical rearm event occurred for Condition /var space used  
on the resource /var of filesys of IBM.FileSystem at 18:41:50 03/28/02  
The resource was monitored on c174n05 and resided on {c174n05}.
```

## Location

**/opt/rsct/bin/wallevent**

**Related reference:**

## watch Command

### Purpose

Observes a program that might be untrustworthy.

### Syntax

```
watch [-e Events] [-o File] [-X] Command [Parameter ... ]
```

### Description

The **watch** command allows the root user or a member of the audit group to observe the actions of a program that are thought to be untrustworthy. The **watch** command starts the program you specify with the *Command* parameter, with or without any *Parameter* fields, and records all audit events or the audit events you specify with the **-e** flag.

The **watch** command observes all the processes that are created while the program runs, including any child process. The **watch** command continues until all processes exit, including the process it created, to observe all the events that occur.

The **watch** command formats the audit records and writes them to standard output or to a file you specify with the **-o** flag.

For the **watch** command to work, the auditing subsystem is not configured and enabled.

### Flags

Item	Description
<b>-e</b> <i>Events</i>	Specifies the events to be audited. The <i>Events</i> parameter is a comma-separated list of audit events that are defined in the <code>/etc/security/audit/events</code> file. The default value is all events.
<b>-o</b> <i>File</i>	Specifies the path name of the output file. If the <b>-o</b> flag is not used, output is written to standard output.
<b>-X</b>	Prints long user names when used with other flags that display user names. The upper limit is determined by the <code>max_logname</code> object data manager (ODM) attribute in the predefined attribute (PdAt) and customized attributes (CuAt) object classes. If a user name is greater than the <code>max_logname</code> attribute, it is truncated to the number of characters as specified by the <code>max_logname</code> attribute, minus 1 character.

### Security

Access Control: This command grants execute (x) access to the root user and members of the audit group. The **setuid** command is set for the root user. This setting allows access to other audit subsystem commands and files, and to the **trusted computing base** attribute.

Files Accessed:



Mode	File
r	/dev/audit
x	/usr/sbin/auditstream
x	/usr/sbin/auditselect
x	/usr/sbin/auditpr

**Attention RBAC users and Trusted AIX users:** This command can perform privileged operations. Only privileged users can run privileged operations. For more information about authorizations and privileges, see Privileged Command Database in *Security*. For a list of privileges and the authorizations associated with this command, see the **lssecattr** command or the **getcmdattr** subcommand.

## Examples

1. To watch all files opened by the **bar** command, enter:

```
watch -e FILE_Open /usr/lpp/foo/bar -x
```

This command opens the audit device and issues the **/usr/lpp/foo/bar** command. It then reads all records and selects and formats the files with the event type of FILE\_Open.

2. To watch the installation of the xyzproduct program, that might be untrustworthy, enter:

```
watch /usr/sbin/installp xyzproduct
```

This command opens the audit device and issues the **/usr/sbin/installp** command. It then reads all records and formats them.

## Files

Item	Description
/usr/sbin/watch	Contains the <b>watch</b> command.
/dev/audit	Specifies the audit device from which the audit records are read.

### Related information:

audit command

RBAC

Trusted AIX

## wc Command

### Purpose

Counts the number of lines, words, bytes, or characters in a file.

### Syntax

```
wc [ -c | -m ] [ -l ] [ -w ] [ File ... ]
```

```
wc -k [ -c ] [ -l ] [ -w ] [ File ... ]
```

### Description

By default, the **wc** command counts the number of lines, words, and bytes in the files specified by the *File* parameter. The command writes the number of newline characters, words, and bytes to the standard output and keeps a total count for all named files.

When you use the *File* parameter, the **wc** command displays the file names as well as the requested counts. If you do not specify a file name for the *File* parameter, the **wc** command uses standard input.

The **wc** command is affected by the **LANG**, **LC\_ALL**, **LC\_CTYPE**, and **LC\_MESSAGES** environment variables.

The **wc** command considers a word to be a string of characters of non-zero length which are delimited by a white space (for example **SPACE** , **TAB**).

## Flags

Item	Description
<b>-c</b>	Counts bytes unless the <b>-k</b> flag is specified. If the <b>-k</b> flag is specified, the <b>wc</b> command counts characters.
<b>-k</b>	Counts characters. Specifying the <b>-k</b> flag is equivalent to specifying the <b>-klwc</b> flag. If you use the <b>-k</b> flag with other flags, then you must include the <b>-c</b> flag. Otherwise, the <b>-k</b> flag is ignored. For more information, see examples 4 and 5. <b>Note:</b> This flag is to be withdrawn in a future release.
<b>-l</b>	Counts lines.
<b>-m</b>	Counts characters. This flag cannot be used with the <b>-c</b> flag.
<b>-w</b>	Counts words. A word is defined as a string of characters delimited by spaces, tabs, or newline characters.

**Note:** If no flag is specified, **wc** by default counts the lines, words, bytes in a file or from standard input.

## Exit Status

This command returns the following exit values:

Item	Description
0	The command ran successfully.
>0	An error occurred.

## Examples

1. To display the line, word, and byte counts of a file, enter:

```
wc chap1
```

The **wc** command displays the number of lines, words, and bytes in the chap1 file.

2. To display only byte and word counts, enter:

```
wc -cw chap*
```

The **wc** command displays the number of bytes and words in each file that begins with chap. The command also displays the total number of bytes and words in these files.

3. To display the line, word, and character counts of a file, enter:

```
wc -k chap1
```

The **wc** command displays the number of lines, words, and characters in the chap1 file.

4. To display the word and character counts of a file, enter:

```
wc -kcw chap1
```

The **wc** command displays the number of characters and words in the chap1 file.

5. To use the **wc** command on standard input, enter:

```
wc -klw
```

The **wc** command displays the number of lines and words in standard input. The **-k** flag is ignored.

6. To display the character counts of a file, enter:

```
wc -m chap1
```

The **wc** command displays the number of characters in the chap1 file.

7. To use the **wc** command on standard input, enter:

```
wc -mlw
```

The **wc** command displays the number of lines, words, and characters in standard input.

## Files

Item	Description
/usr/bin/wc, /bin/wc	Contains the <b>wc</b> command.
/usr/ucb/wc	Contains the symbolic link to the <b>wc</b> command.

### Related information:

Files

Input and output redirection

---

## what Command

### Purpose

Displays identifying information in files.

### Syntax

```
what [ -s] Pathname/File.
```

### Description

The **what** command searches specified files for all occurrences of the pattern that the **get** command substitutes for the **@(#)** keyletter (see the **get** or **prs** command for a description of identification keywords). By convention, the value substituted is "**@(#)**" (double quotation marks, at sign, left parenthesis, pound sign, right parenthesis, double quotation marks). If no file is specified, the **what** command reads from standard input.

The **what** command writes to standard output whatever follows the pattern, up to but not including the first double quotation mark ("), greater than symbol (>), new-line character, backslash (\), or null character.

The **what** command should be used in conjunction with the **get** command, which automatically inserts the identifying information. You can also use the **what** command on files where the information is inserted manually.

The **what** command accommodates the compiler inserted command line options in a binary file. The command line options saved in a binary file by AIX compilers may contain the backslash (\), the greater than symbol (>), or the double quotation mark ("), within the macro definitions. The **what** command behaves in the following manner to write the command line options saved by the compiler.

At the start of a line, if the pattern **@(#)** is followed by "opt" and is with or without a blank space (" ") in between, then the **what** command writes the character till "\n" that is the end of a line.

For example,

- For C and FORTRAN AIX compilers, use **@(#)** opt (...).
- For C++ AIX compilers, use **@(#)** opt (...).

The whole line is printed after "**@(#)**".

**Note:** The **what** command may fail to find SCCS identification strings in executable files.

## Flags

Item	Description
-s	Searches for only the first occurrence of the @(#) pattern.

## Exit Status

This command returns the following exit values:

Item	Description
0	Any matches were found.
1	Otherwise.

## Examples

Suppose that the file `test.c` contains a C program that includes the line:

```
char ident[ ] = "@(#)Test Program";
```

If you compile `test.c` to produce `test.o`, then the command:

```
what test.c test.o
```

displays:

```
test.c:  
Test Program  
test.o:  
Test Program
```

**Note:** The full file path names `usr/bin/test.c` and `usr/bin/test.o` are required if the files are not in the current directory.

## Files

Item	Description
<code>/usr/bin/what</code>	Contains the <b>what</b> command.

### Related information:

List of SCCS commands

get command

prs command

---

## whatis Command

### Purpose

Describes what function a command performs.

### Syntax

```
whatis [ -M PathName ] Command ...
```

## Description

The **whatis** command looks up a given command, system call, library function, or special file name, as specified by the *Command* parameter, from a database you create using the **catman -w** command. The **whatis** command displays the header line from the manual section. You can then issue the **man** command to obtain additional information.

The **whatis** command is equivalent to using the **man -f** command.

**Note:** When the `/usr/share/man/whatis` database is built from the HTML library using the **catman -w** command, section 3 is equivalent to section 2 or 3. See the **man** command for further explanation of sections.

## Flags

Item	Description
<b>-M</b> <i>PathName</i>	Specifies an alternative search path. The search path is specified by the <i>PathName</i> parameter, and is a colon-separated list of directories in which the <b>whatis</b> command expects to find the standard manual subdirectories.

## Examples

To find out what the **ls** command does, enter:

```
whatis ls
```

This produces the following output:

```
ls(1) -Displays the contents of a directory.
```

## Files

Item	Description
<code>/usr/share/man/whatis</code>	Contains the <b>whatis</b> database.

### Related information:

catman command

ls command

man command

---

## whatnow Command

### Purpose

Starts a prompting interface for draft disposition.

### Syntax

```
whatnow [ { -draftfolder +folder | -nodraftfolder | file } { -draftmessage message | file } ] [-editor editor | -noedit ] [ -prompt string ]
```

### Description

The **whatnow** command provides an interface for the disposition of messages. By default, the interface operates on the current draft message. When you enter the **whatnow** command, the system places you in the interface and returns the following prompt:

```
What now?
```

Within the interface you can manipulate message drafts using the **whatnow** subcommands. To see a listing of the subcommands, press the Enter key at the What now? prompt. To exit the interface, press q.

If you do not specify the **-draftfolder** flag or if the Draft-Folder: entry in the \$HOME/.mh\_profile file is undefined, the **whatnow** command searches your MH directory for a **draft** file. Specifying a message after the **-draftfolder** +*folder* flag is the same as specifying the **-draftmessage** flag.

To change the default editor for the **whatnow** command, use the **-editor** flag or define the Editor: entry in the *UserMhDirectory/.mh\_profile* file.

**Note:** The **comp**, **dist**, **forw**, or **repl** commands use the same interface as the **whatnow** command.

## Flags

Item	Description
<b>-draftfolder</b> + <i>folder</i>	Specifies the folder containing the message. By default, the system uses the <i>UserMhDirectory/draft</i> file. Specifying a message after the <b>-draftfolder</b> + <i>folder</i> is the same as using the <b>-draftmessage</b> flag.
<b>-draftmessage</b> <i>message</i>	Specifies the draft message.
<b>-editor</b> <i>editor</i>	Specifies that the value of the <i>editor</i> variable is the initial editor for composing or revising the message.
<b>-help</b>	Lists the command syntax, available switches (toggles), and version information. <b>Note:</b> For MH, the name of this flag must be fully spelled out.
<i>file</i>	User selected draft file.
<i>message</i>	Specifies the message. Use the following references to specify messages:
	<i>Number</i> Number of the message.
	<b>cur</b> or <b>.</b> ( <b>period</b> ) Current message. This is the default.
	<b>first</b> First message in a folder.
	<b>last</b> Last message in a folder.
	<b>next</b> Message following the current message.
	<b>prev</b> Message preceding the current message.
	<b>-nodraftfolder</b> Places the draft in the <i>UserMhDirectory/draft</i> file.
	<b>-noedit</b> Suppresses the initial edit.
	<b>-prompt</b> <i>string</i> Uses the specified string as the prompt. The default string is What now?.

## whatnow Subcommands

The **whatnow** subcommands enable you to edit the message, direct the disposition of the message, or end the processing of the **whatnow** command.

Item	Description
<b>display</b> [ <i>flags</i> ]	Displays the message being redistributed or replied to. You can specify any <i>flags</i> parameter that is valid for the listing program. (Use the lproc: entry in the \$HOME/.mh_profile file to set a default listing program.) If you specify flags that are invalid for the listing program, the <b>whatnow</b> command does not pass the path name of the draft.
<b>edit</b> [ <i>commandstring</i> ]	Specifies with the <i>commandstring</i> parameter an editor for the message. You can specify the editor and any valid flags to that editor. If you do not specify an editor, the <b>whatnow</b> command uses the editor specified by the Editor: entry in your <i>UserMhDirectory/.mh_profile</i> file. If your Editor: entry is undefined, the <b>whatnow</b> command starts the editor used in the previous editing session.

Item	Description
<b>list</b> [ <i>flags</i> ]	Displays the draft. You can specify any <i>flags</i> parameter that is valid for the listing program. (To specify a default listing program, set a default <code>lproc:</code> entry in the <code>\$HOME/.mh_profile</code> file.) If you specify any flags that are invalid for the listing program, the <b>whatnow</b> command does not pass the path name of the draft.
<b>push</b> [ <i>flags</i> ]	Sends the message in the background. You can specify any valid flag for the <b>send</b> command.
<b>quit</b> [-delete]	Ends the <b>whatnow</b> session. If you specify the <b>-delete</b> flag, the <b>whatnow</b> command deletes the draft. Otherwise, the <b>whatnow</b> command stores the draft.
<b>refile</b> [ <i>flags</i> ] + <i>folder</i>	Files the draft in the specified folder and supplies a new draft having the previously specified form. You can specify any <i>flags</i> parameter that is valid for the command serving as the <b>fileproc</b> . (You can set a default <code>fileproc:</code> entry in the <code>\$HOME/.mh_profile</code> file.)
<b>send</b> [ <i>flags</i> ]	Sends the message. You can specify any valid flags for the <b>send</b> command.
<b>whom</b> [ <i>flags</i> ]	Displays the addresses to which the message would be sent. You can specify any valid flags for the <b>whom</b> command.

## Profile Entries

The following entries are entered in the `UserMhDirectory/.mh_profile` file:

Item	Description
Draft-Folder:	Sets the default folder for drafts.
Editor:	Sets the default editor.
fileproc:	Specifies the program used to refile messages.
LastEditor-next:	Specifies the editor used after exiting the editor specified by the <code>LastEditor</code> variable.
lproc:	Specifies the program used to list the contents of a message.
Path:	Specifies the <code>UserMhDirectory</code> .
sendproc:	Specifies the program used to send messages.
whomproc:	Specifies the program used to determine the users to whom a message would be sent.

## Security

**Attention RBAC users and Trusted AIX users:** This command can perform privileged operations. Only privileged users can run privileged operations. For more information about authorizations and privileges, see Privileged Command Database in *Security*. For a list of privileges and the authorizations associated with this command, see the **Issecattr** command or the **getcmdattr** subcommand.

## Examples

1. To display the original message when you are replying to a message, enter the following at the What now? prompt:

```
display
```

The system displays the original message. If you enter the **display** subcommand from a command other than the **dist** or **repl** command, you will receive a system message stating that there is no alternate message to display.

2. To edit the draft message with the vi editor, enter the following at the What now? prompt:

```
edit vi
```

3. To edit the draft message with the default editor specified in your `.mh_profile` file, enter the following at the What now? prompt:

```
edit
```

4. To list the contents of the draft message you have composed, enter the following at the What now? prompt:

```
list
```

The draft message you are composing is displayed.

5. To send the draft message in the background and get a shell prompt immediately, enter the following at the What now? prompt:

```
push
```

The draft message is sent and you immediately receive the shell prompt.

6. To quit composing a draft message and save it to a file so that you can later finish composing the message, enter the following at the What now? prompt:

```
quit
```

The system responds with a message similar to the following.

```
whatnow: draft left on /home/dale/Mail/draft
```

In this example, user dale's draft message is saved to the **/home/dale/Mail/draft** file.

7. To quit composing a draft message and delete the message, enter the following at the What now? prompt:

```
quit -delete
```

The shell prompt is displayed when the draft message is deleted.

8. To file the draft message you are composing before you send it, enter the following at the What now? prompt:

```
refile +tmp
```

The system responds with a message similar to the following:

```
Create folder "home/dale/Mail/tmp"?
```

In this example, if you answer yes, the draft message is filed in user dale's folder tmp.

9. To send the draft message you have composed, enter the following at the What now? prompt:

```
send
```

The shell prompt is displayed when the message is sent.

10. To verify that all addresses in the draft message are recognized by the mail delivery system, enter the following at the What now? prompt:

```
whom
```

The system responds with a message similar to the following:

```
jeanne... User unknown  
dale@venus... deliverable
```

In this example, the mail delivery system recognized dale@venus as a correct address, but did not recognize jeanne as a correct address.



## Files

Item	Description
<code>\$HOME/mh_profile</code>	Specifies the MH user profile.
<code>UserMhDirectory/draft</code>	Contains the current message draft.
<code>/usr/bin/whatnow</code>	Contains the <b>whatnow</b> command.

### Related information:

RBAC

Trusted AIX

refile command

---

## whereis Command

### Purpose

Locates source, binary, or manual for program.

### Syntax

```
whereis [ -s ] [ -b ] [ -m ] [ -u ] [ { -S | -B | -M } Directory ... ]... -f ] File ...
```

### Description

The **whereis** command locates the source, binary, and manuals sections for specified files. The supplied names are first stripped of leading path name components and any (single) trailing extension of the form *.ext* (for example, *.c*). Prefixes of **s**, resulting from use of the Source Code Control System (see **SCCS**) are also dealt with. The command then attempts to find the desired program from a list of standard locations.

A usage message is returned if a bad option is entered. In other cases, no diagnostics are provided.

### Flags

If any of the **-b**, **-s**, **-m** or **-u** flags are given, the **whereis** command searches only for binary, source, manual, or unusual sections respectively (or any two thereof).

Item	Description
<b>-b</b>	Searches for binary sections of a file.
<b>-m</b>	Searches for manual sections of a file.
<b>-s</b>	Searches for source sections of a file.
<b>-u</b>	Searches for unusual files. A file is said to be unusual if it does not have one entry of each requested type. Entering <code>whereis -m -u *</code> asks for those files in the current directory which have no documentation.

The **-B**, **-M**, and **-S** flags can be used to change or otherwise limit the places where the **whereis** command searches. Since the program uses the **chdir** subroutine to run faster, path names given with the **-M**, **-S** and **-B** flag directory list must be full; for example, they must begin with a / (slash).

Item	Description
<b>-B</b>	Like <b>-b</b> , but adds a directory to search. Change or limit the places where the <b>whereis</b> command searches for binaries.
<b>-M</b>	Like <b>-m</b> , but adds a directory to search. Change or limit the places where the <b>whereis</b> command searches for manual sections.
<b>-S</b>	Like <b>-s</b> , but adds a directory to search. Change or limit the places where the <b>whereis</b> command searches for sources
<b>-f</b>	Terminates the last <b>-M</b> , <b>-S</b> or <b>-B</b> directory list and signal the start of file names.

## Examples

To find all of the files in the `/usr/ucb` directory that either are not documented in the `/usr/man/man1` directory or do not have source in the `/usr/src/cmd` directory, enter:

```
cd /usr/ucb
whereis -u -M /usr/man/man1 -S /usr/src/cmd -f *
```

## Files

Item	Description
<code>/usr/share/man/*</code>	Directories containing manual files.
<code>/sbin, /etc, /usr/{lib,bin,ucb,lpp}</code>	Directories containing binary files.
<code>/usr/src/*</code>	Directories containing source code files.

### Related information:

chdir command

---

## which Command

### Purpose

Locates a program file, including aliases and paths.

### Syntax

```
which [ Name ... ]
```

### Description

The **which** command takes a list of program names and looks for the files that run when these names are given as commands. The **which** command expands each argument, if it is aliased, and searches for it along the user's path. The aliases and paths are taken from the `.cshrc` file in the user's home directory. If the `.cshrc` file does not exist, or if the path is not defined in the `.cshrc` file, the **which** command uses the path defined in the user's environment.

A diagnostic is given if a name is aliased to more than a single word or if an executable file with the argument name is not found in the path.

In the Korn shell, you can use the **whence** command to produce a more verbose report.

### Examples

To find the executable file associated with a command name of `lookup`:

```
which lookup
```

### Files

Item	Description
<code>\$HOME/.cshrc</code>	Contains the source of aliases and path values.

### Related information:

Korn shell or POSIX shell built-in commands

csch command

ksh command

## which\_fileset Command

### Purpose

Searches the `/usr/lpp/bos/AIX_file_list` file for a specified file name or command.

### Syntax

```
which_fileset [ File ]
```

### Description

The `which_fileset` command searches the `/usr/lpp/bos/AIX_file_list` file for a specified file name or command name, and prints out the name of the fileset that the file or command is shipped in.

The `/usr/lpp/bos/AIX_file_list` file is large and not installed automatically. You must install the `bos.content_list` fileset to receive this file.

The `File` parameter can be the command name, the full path name, or a regular expression search pattern.

### Examples

- To display which fileset the `dbx` command is shipped in, enter:

```
which_fileset dbx
```

The screen displays the following:

```
/usr/bin/dbx > /usr/ccs/bin/dbx          bos.adt.debug 4.2.1.0
/usr/ccs/bin/dbx                          bos.adt.debug 4.2.1.0
```

- To display all commands and paths containing the `sendmail` string, enter:

```
which_fileset sendmail.*
```

The screen displays the following:

```
/usr/ucb/mailq > /usr/sbin/sendmail      bos.compat.links 4.2.0.0
/usr/ucb/newaliases > /usr/sbin/sendmail  bos.compat.links 4.2.0.0
/usr/lib/nls/msg/Ca_ES/sendmail87.cat    bos.msg.Ca_Es.net.tcp.client 4.2.0.0
/usr/lib/nls/msg/ca_ES/sendmail87.cat    bos.msg.ca_Es.net.tcp.client 4.2.0.0
/usr/lib/nls/msg/cs_CZ/sendmail87.cat    bos.msg.cs_CZ.net.tcp.client 4.2.0.0
/usr/lib/nls/msg/De_DE/sendmail87.cat    bos.msg.De_DE.net.tcp.client 4.2.0.0
/usr/lib/nls/msg/de_DE/sendmail87.cat    bos.msg.de_DE.net.tcp.client 4.2.0.0
/usr/lib/nls/msg/En_US/sendmail87.cat    bos.msg.En_US.net.tcp.client 4.2.0.0
/usr/lib/nls/msg/en_US/sendmail87.cat    bos.msg.en_US.net.tcp.client 4.2.0.0
/usr/lib/nls/msg/Es_ES/sendmail87.cat    bos.msg.Es_ES.net.tcp.client 4.2.0.0
/usr/lib/nls/msg/es_ES/sendmail87.cat    bos.msg.es_ES.net.tcp.client 4.2.0.0
/usr/lib/nls/msg/Fr_FR/sendmail87.cat    bos.msg.Fr_FR.net.tcp.client 4.2.0.0
/usr/lib/nls/msg/fr_FR/sendmail87.cat    bos.msg.fr_FR.net.tcp.client 4.2.0.0
/usr/lib/nls/msg/hu_HU/sendmail87.cat    bos.msg.hu_HU.net.tcp.client 4.2.0.0
/usr/lib/nls/msg/It_IT/sendmail87.cat    bos.msg.It_IT.net.tcp.client 4.2.0.0
/usr/lib/nls/msg/it_IT/sendmail87.cat    bos.msg.it_IT.net.tcp.client 4.2.0.0
/usr/lib/nls/msg/Ja_JP/sendmail87.cat    bos.msg.Ja_JP.net.tcp.client 4.2.0.0
/usr/lib/nls/msg/ja_JP/sendmail87.cat    bos.msg.ja_JP.net.tcp.client 4.2.0.0
```

```

/usr/lib/nls/msg/ko_KR/sendmail87.cat bos.msg.ko_KR.net.tcp.client 4.2.0.0
/usr/lib/nls/msg/pl_PL/sendmail87.cat bos.msg.pl_PL.net.tcp.client 4.2.0.0
/usr/lib/nls/msg/ru_RU/sendmail87.cat bos.msg.ru_RU.net.tcp.client 4.2.0.0
/usr/lib/nls/msg/Sv_SE/sendmail87.cat bos.msg.Sv_SE.net.tcp.client 4.2.0.0
/usr/lib/nls/msg/sv_SE/sendmail87.cat bos.msg.sv_SE.net.tcp.client 4.2.0.0
/usr/lib/nls/msg/ZH_CN/sendmail87.cat bos.msg.ZH_CN.net.tcp.client 4.2.0.0
/usr/lib/nls/msg/zh_CN/sendmail87.cat bos.msg.zh_CN.net.tcp.client 4.2.0.0
/usr/lib/nls/msg/Zh_TW/sendmail87.cat bos.msg.Zh_TW.net.tcp.client 4.2.0.0
/usr/lib/nls/msg/zh_TW/sendmail87.cat bos.msg.zh_TW.net.tcp.client 4.2.0.0
/etc/sendmail.cf bos.net.tcp.client.4.2.1.0
/usr/lib/sendmail > /usr/sbin/sendmail bos.net.tcp.client.4.2.1.0
/usr/sbin/mailq > /usr/sbin/sendmail bos.net.tcp.client.4.2.1.0
/usr/sbin/newaliases > /usr/sbin/sendmail bos.net.tcp.client.4.2.1.0
/usr/sbin/sendmail bos.net.tcp.client.4.2.1.0

```

3. To find where the `/usr/sbin/which_fileset` command is shipped, enter:

```
which_fileset /usr/bin/which_fileset
```

The screen displays:

```
/usr/sbin/which_fileset bos.rte.install 4.2.1.0
```

## who Command

### Purpose

Identifies the users currently logged in.

### Syntax

```
who [ -a | -b -d -i -l -m -p -q -r -s -t -u -w -A -H -T -X ] [ File ]
```

```
who am { i | I }
```

### Description

The **who** command displays information about all users currently on the local system. The following information is displayed: login name, tty, date and time of login. Typing `who am i` or `who am I` displays your login name, tty, date and time you logged in. If the user is logged in from a remote machine, then the host name of that machine is displayed as well.

The **who** command can also display the elapsed time since line activity occurred, the process ID of the command interpreter (shell), logins, logoffs, restarts, and changes to the system clock, as well as other processes generated by the initialization process.

The general output format of the **who** command is as follows:

```
Name [State] Line Time [Activity] [Pid] [Exit] (Hostname)
```

where:

Item	Description
Name	Identifies the user's login name.
State	Indicates whether the line is writable by everyone (see the <code>-T</code> flag).
Line	Identifies the line name as found in the <code>/dev</code> directory.
Time	Represents the time when the user logged in.
Activity	Represents the hours and minutes since activity last occurred on that user's line. A . (dot) here indicates line activity within the last minute. If the line has been quiet more than 24 hours or has not been used since the last system startup, the entry is marked as old.
Pid	Identifies the process ID of the user's login shell.

Item	Description
Term	Identifies the process termination status (see the <b>-d</b> flag). For more information on the termination values, refer to the <b>wait</b> subroutine or to the <b>/usr/include/sys/signal.h</b> file.
Exit	Identifies the exit status of ended processes (see the <b>-d</b> flag).
Hostname	Indicates the name of the machine the user is logged in from.

To obtain information, the **who** command usually examines the **/etc/utmp** file. If you specify another file with the *File* parameter, the **who** command examines that file instead. This new file is usually the **/var/adm/wtmp** or **/etc/security/failedlogin** file.

If the *File* parameter specifies more than one file name, only the last file name will be used.

**Note:** This command only identifies users on the local node.

## Flags

Item	Description
<b>-a</b>	Processes the <b>/etc/utmp</b> file or the named file with all information. Equivalent to specifying the <b>-bdlprtTu</b> flags.
<b>-b</b>	Indicates the most recent system startup time and date.
<b>-d</b>	Displays all processes that have expired without being regenerated by <b>init</b> . The exit field appears for dead processes and contains the termination and exit values (as returned by <b>wait</b> ) of the dead process. (This flag is useful for determining why a process ended by looking at the error number returned by the application.)
<b>-l</b>	Lists any login process.
<b>-m</b>	Displays information about the current terminal only. The <b>who -m</b> command is equivalent to the <b>who am i</b> and <b>who am I</b> commands.
<b>-p</b>	Lists any active process that is currently active and has been previously generated by <b>init</b> .
<b>-q</b>	Prints a quick listing of users and the number of users on the local system.
<b>-r</b>	Indicates the current run-level of the process.
<b>-s</b>	Lists only the name, line, and time fields. This flag is the default; thus, the <b>who</b> and <b>who -s</b> commands are equivalent.
<b>-t</b>	Indicates the last change to the system clock by the root user using the <b>date</b> command. If the <b>date</b> command has not been run since system installation, the <b>who -t</b> command produces no output.
<b>-u</b> or <b>-i</b>	Displays the user name, tty, login time, line activity, and process ID of each current user.
<b>-A</b>	Displays all accounting entries in the <b>/etc/utmp</b> file. These entries are generated through the <b>acctwtmp</b> command.
<b>-H</b>	Displays a header (title).
<b>-T</b> or <b>-w</b>	Displays the state of the tty and indicates who can write to that tty as follows: <ul style="list-style-type: none"> <li><b>+</b> Writable by anyone.</li> <li><b>-</b> Writable only by the root user or its owner.</li> <li><b>?</b> Bad line encountered.</li> </ul>
<b>-X</b>	Prints all available characters of each user name instead of truncating to the first 8 characters. The user name is also moved to the last column of the output.

## Exit Status

This command returns the following exit values:

Item	Description
0	Successful completion.
>0	An error occurred.

## Examples

1. To display information about who is using the local system node, type:

```
who
```

Information similar to the following is displayed:

```
pts/1      Nov  9 00:20  long_username_greater_than_eight_characters  (localhost)
```

2. To display your user name, type:

```
who am i
```

Information similar to the following is displayed:

```
george lft/0 Jun  8 08:34
```

3. To display a history of logins, logouts, system startups, and system shutdowns, type:

```
who /var/adm/wtmp
```

Information similar to the following is displayed:

```
hank  lft/0   Jun   8   08:34  (ausnix5)
john  lft/0   Jun   8   08:34  (JIKey)
mary  lft/0   Jun   8   08:22  (machine.austin.ibm)
jan   pts4     Jun   8   09:19  (puff.wisc.edu)
```

4. To display the run-level of the local system node, type:

```
who -r
```

Information similar to the following is displayed:

```
. run-level 2 Jun  8 04:15 2 0 s
```

5. To display any active process that is currently active and has been previously generated by init, type:

```
who -p
```

Information similar to the following is displayed:

```
srcmstr . Jun  8 04:15 old 2896
cron . Jun  8 04:15 old 4809
uprintfd . Jun  8 04:15 old 5158
```

6. To process the `/var/adm/wtmp` file with the `-bdlprtTu` flags specified, type:

```
who -a /var/adm/wtmp
```

Information similar to the following is displayed:

```
.      system boot Jun 19 10:13
.      run-level 2 Jun 19 10:13
.      . Jun 19 10:14 old
.      . Jun 19 10:14 old
.      . Jun 19 10:14 old
.      . Jun 19 10:14 old
rc     - . Jun 19 10:13 old
.      . Jun 19 10:16 old
.      . Jun 19 10:14 old
srcmstr - . Jun 19 10:14 old
rctcpip - . Jun 19 10:14 old
rcdce  - . Jun 19 10:14 old
rccm   - . Jun 19 10:15 old
dceupdt - . Jun 19 10:15 old
rcnfs  - . Jun 19 10:15 old
cron   - . Jun 19 10:16 old
piobe  - . Jun 19 10:16 old
qdaemon - . Jun 19 10:16 old
writesrv - . Jun 19 10:16 old
uprintfd - . Jun 19 10:16 old
.      . Jun 19 10:16 old
LOGIN  - lft0 Jun 19 10:16 old
.      . Jun 19 10:16 old
.      . Jun 19 10:16 old
```

## Files

Item	Description
/etc/utmp	Contains user and accounting information.
/etc/security/failedlogin	Contains the history of all invalid logins.
/var/adm/wtmp	Contains the history of all logins since the file was last created.
/usr/include/sys/signal.h	Contains a list of termination values.

### Related information:

date command  
wait command

---

## whoami Command

### Purpose

Displays your login name.

### Syntax

whoami

### Description

The **whoami** command displays your login name. Unlike using the command **who** and specifying **am i**, the **whoami** command also works when you have root authority since it does not examine the **/etc/utmp** file.

## Files

Item	Description
/etc/passwd	Contains user IDs.

### Related reference:

“who Command” on page 92

---

## whodo Command

### Purpose

Lists the jobs being performed by users on the system.

### Syntax

whodo [ -h ] [ -l ] [ -X ] [ *User* ]

### Description

Prints information on all processes for a terminal, as well as the child processes.

By default, the output generated by the command for each active logged user will contain name of the terminal, user ID, date login time. The output is headed by the date, time and machine name. This information is followed by a record of active processes associated with that user ID. Each record shows the terminal name, process-ID, CPU minutes and seconds used, and process name.

### Flags

Item	Description
-h	Suppress the heading that is printed on the output.
-l	Produce a long form of output. A summary of the current activity on the system is printed. The summary includes the following:  <b>User</b> Who is logged on.  <b>tty</b> Name of the tty the user is on.  <b>login@</b> Time of day the user logged on.  <b>idle</b> Number of minutes since a program last attempted to read from the terminal.  <b>JCPU</b> System unit time used by all processes and their children on that terminal.  <b>PCPU</b> System unit time used by the currently active process.  <b>what</b> Name and parameters of the current process. The heading line of the summary shows the current time of day, how long the system has been up, the number of users logged into the system.
-X	Prints all available characters of each user name instead of truncating to the first 8 characters. The user name is also moved to the last column of the output.

## Parameters

Item	Description
<i>User</i>	Limits output to all sessions pertaining to the user specified with <i>User</i> . More than one user name cannot be specified at a time.

## Exit Status

- 0 The command completed successfully.
- >0 An error occurred.

## Examples

1. When the **whodo** command is invoked on host "linguist" without any flags or parameters, the output looks similar to the following:

```
Sun Jul 28 16:27:12 2002
linguist

1ft0 jeffg 8:15
? 4136 0:00 dtlogin
? 3408 4:55 dtsession
? 2072 4:37 dtwm
? 17310 0:00 dtexec
? 20904 5:53 dtterm
pts/0 22454 0:00 ksh
pts/0 4360 0:07 ksh
pts/0 25788 0:00 whodo
? 23672 0:00 dtexec
? 27536 0:00 dtterm
pts/3 21508 0:00 ksh
? 23888 0:00 dtexec
? 24384 2:49 dtterm
pts/2 24616 0:00 ksh
pts/2 25002 0:04 ksh
pts/2 26110 0:00 ksh
? 25276 0:00 dtexec
? 27090 0:31 dtterm
pts/1 24232 0:00 ksh
pts/1 23316 0:01 ksh
? 12566 4:23 dtfile
? 21458 1:35 dtfile
```



```
pts/0 jeffg 8:16
pts/0 22454 0:00 ksh
pts/0 4360 0:07 ksh
pts/0 25788 0:00 whodo
```

```
pts/1 jeffg 17:8
pts/1 24232 0:00 ksh
pts/1 23316 0:01 ksh
```

```
pts/2 jeffg 17:20
pts/2 24616 0:00 ksh
pts/2 25002 0:04 ksh
pts/2 26110 0:00 ksh
```

```
pts/3 root 16:26
pts/3 21508 0:00 ksh
```

2. The command **whodo -l** on the host linguist produces the following output:

```
04:33PM up 20 day(s), 22 hr(s), 51 mins(s) 5 user(s)
User  tty      login@      idle      JCPU      PCPU what
jeffg  lft0     08Jul02     21day(s)
jeffg  pts/0    08Jul02     14:00     7 whodo -l
jeffg  pts/1    16Jul02     10day(s)  44
jeffg  pts/2    12Jul02     11        8:39     4 /usr/bin/ksh
root   pts/3    04:26PM     7         -ksh
```

3. The command **whodo -lX** on the host kq11 produces the following output:

```
12:50AM up 3 day(s), 1 hr(s), 41 mins(s) 4 user(s)
tty      login@      idle      JCPU      PCPU what      User
tty0     Wed11PM     2day(s)
pts/0    12:12AM
pts/1    12:20AM     whodo -lX long_username_greater_than_eight_characters
pts/2    Fri05AM     2day(s)     -ksh      root
```

## Files

Item	Description
<code>/usr/sbin/whodo</code>	Contains the <b>whodo</b> command.
<code>/etc/utmp</code>	Contains the list of users.

### Related reference:

“who Command” on page 92

---

## whois Command

### Purpose

Identifies a user by user ID or alias.

### Syntax

```
whois [ -h HostName ] [ . | ! ] [ * ] Name [ . . . ]
```

**whois ?**

### Description

The `/usr/bin/whois` command searches a user name directory and displays information about the user ID or nickname specified in the *Name* parameter. The **whois** command tries to reach ARPANET host internic.net where it examines a user-name database to obtain information. The **whois** command should be used only by users on ARPANET. Refer to RFC 812 for more complete information and recent changes to the **whois** command.

**Note:** If your network is on a national network, such as ARPANET, the host name is hard-coded as internic.net.

The *Name* [ . . . ] parameter represents the user ID, host name, network address, or nickname on which to perform a directory search. The **whois** command performs a wildcard search for any name that matches the string preceding the optional ... (three periods).

## Flags

Item	Description
.	Forces a name-only search for the name specified in the <i>Name</i> parameter.
!	Displays help information for the nickname or handle ID specified in the <i>Name</i> parameter.
*	Displays the entire membership list of a group or organization. If there are many members, this can take some time.
?	Requests help from the ARPANET host.
-h <i>HostName</i>	Specifies an alternative host name. The default host name on the ARPANET is internic.net. You can contact the other major ARPANET user-name database, nic.ddn.mil, by specifying the -h <i>HostName</i> flag.

## Examples

1. To display information about ARPANET registered users by the name of Smith, enter:

```
whois Smith
```

2. To display information about ARPANET registered users that use the handle Hobo, enter:

```
whois !Hobo
```

3. To display information about ARPANET registered users with the name of John Smith, enter:

```
whois .Smith, John
```

4. To display information about ARPANET registered users whose names or handles begin with the letters HEN, enter:

```
whois HEN ...
```

5. To get help information for the **whois** command, enter:

```
whois ?
```

### Related reference:

“who Command” on page 92

### Related information:

named.conf file format

---

## whom Command

### Purpose

Manipulates Message Handler (MH) addresses.

### Syntax

```
whom [ -alias File ... ] [-nocheck | -check ] [ { -draftfolder +Folder | -nodraftholder | File } {  
-draftmessage Message | -draftFile } ]
```

## Description

The **whom** command does the following:

- Expands the headers of a message into a set of addresses.
- Lists the addresses of the proposed recipients of a message.
- Verifies that the addresses are deliverable to the transport service.

**Note:** The **whom** command does not guarantee that addresses listed as being deliverable will actually be delivered.

A message can reside in a draft folder or in a file. To specify where a message resides, use the **-draft**, **-draftfolder**, or **-draftmessage** flag.

If you do not specify the **-draftfolder** flag or if the **Draft-Folder:** entry in the **\$HOME/.mh\_profile** file is undefined, the **whom** command searches your MH directory for a **draft** file. Specifying a message after the **-draftfolder +Folder** flag is the same as specifying the **-draftmessage** flag.

## Flags

Item	Description
<b>-alias</b> <i>File</i>	Specifies a file to search for mail aliases. By default, the system searches the <b>/etc/mh/MailAliases</b> file.
<b>-draft</b>	Uses the header information in the <i>UserMhDirectory/draft</i> file if it exists.
<b>-draftfolder</b> <i>+Folder</i>	Uses the header information from the draft message in the specified folder. If you specify a draft folder that doesn't exist, the system creates one for you.
<b>-draftmessage</b> <i>Message</i>	Uses the header information from the specified draft message.
<b>-help</b>	Lists the command syntax, available switches (toggles), and version information. <b>Note:</b> For MH, the name of this flag must be fully spelled out.
<i>Message</i>	Specifies the message draft. Use the following to specify messages:  <i>Number</i> Number of the message.  <b>cur</b> or <b>.</b> ( <b>period</b> ) Current message. This is the default.  <b>first</b> First message in a folder.  <b>last</b> Last message in a folder.  <b>next</b> Message following the current message.  <b>prev</b> Message preceding the current message.
<b>-nodraftfolder</b>	Undoes the last occurrence of the <b>-draftfolder +Folder</b> flag.

**Note:** Two other flags, **-check** and **-nocheck**, are also available. These flags have no effect on how the **whom** command performs verification. The **-check** and **-nocheck** flags are provided for compatibility only.

## Profile Entries

The following entries are entered in the *UserMhDirectory/.mh\_profile* file:

Item	Description
Draft-Folder:	Sets your default folder for drafts.
postproc:	Specifies the program used to post messages.

## Security

**Attention RBAC users and Trusted AIX users:** This command can perform privileged operations. Only privileged users can run privileged operations. For more information about authorizations and privileges, see Privileged Command Database in *Security*. For a list of privileges and the authorizations associated with this command, see the **lssecattr** command or the **getcmdattr** subcommand.

## Examples

To list and verify the addresses of the proposed recipients of a message, enter the addressees and subject of the message at the respective prompt, as follows:

```
To: d77@nostromo
Subject: a test
```

When prompted again, enter the text of the message:

```
-----Enter initial text
test
-----
```

After the whatnow prompt, enter the **whom** command:

```
whatnow>>> whom
```

The address of the proposed recipients of the message is then displayed:

```
lance...
d77@nostromo... deliverable
```

## Files

Item	Description
<code>\$HOME/.mh_profile</code>	Specifies the MH user profile.
<code>/usr/bin/whom</code>	Contains the <b>whom</b> command.

### Related reference:

“whatnow Command” on page 85

### Related information:

RBAC

Trusted AIX

## wlmassign command

### Purpose

Manually assigns processes to a workload management class or cancels prior manual assignments for processes.

### Syntax

```
wlmassign [ -s | -S ] [ -u | Class_Name ] [ pid_list ] [ -g pgid_list ]
```

```
wlmassign [ -t { tag [ -i inheritance ] | -r } ] [ pid_list ] [ -g pgid_list ]
```

## Description

The **wlmassign** command:

- Assigns a set of processes specified by a list of process identifiers (*PID*) or process group identifiers (*PGIDs*) to a specified superclass or subclass or both, thus overriding the automatic class assignment or a prior manual assignment.
- Cancels a previous manual assignment for the processes specified in *pid\_list* or *pgid\_list*.
- Assigns the Workload Manager (WLM) tag process attribute to a set of processes specified by a list of *pids* or *pgids*.
- Removes the WLM tag process attribute from a set of processes specified by a list of *pids* or *pgids*.

In addition to the tag, additional inheritance suboptions can be specified and these suboptions indicates to the WLM if a child process must inherit the tag from its parent after **fork** or **exec** subroutine.

The process requires at least the **SIGPRIV** privileges or higher for tagging another process.

The WLM tag assignment remains in effect until one or more of the following conditions are true:

- The tag is removed by using the **-r** flag.
- The tagged process is ended.
- The tag is overwritten with a new tag.

When a WLM tag is assigned to a process that belongs to a class that has the inheritance property turned off, it is automatically reclassified according to the current assignment rules and the new tag is taken into account during reclassification. The WLM tag is effective if class inheritance attribute is not specified for the current process class. To override the class inheritance attribute in favor of reclassification based on tag rules, the **/usr/samples/kernel/wlmtune** command available in the **bos.adt.samples** PTF, can be used to modify the behavior of WLM. The related tunables follow:

### **tag\_override\_super**

Indicates to the WLM that superclass inheritance is bypassed in favor of the rule-based classification if a rule matches the process tag. The default value is **0**.

### **tag\_override\_sub**

Indicates to the WLM that subclass inheritance is bypassed in favor of the rule-based classification if a rule matches the process tag. The default value is **0**.

For the tag assignment rules-based classification to be effective, the tunable values must be set prior to a WLM reclassification update.

The interactions between automatic assignment (inheritance and rules), inheritance, and manual assignment are detailed in the Workload management in *Operating system and device management*.

The **wlmassign** command allows to specify processes using a list of PIDs, a list of PGDIDs, or both. The formats of these lists follow:

**pid**[,**pid**[,**pid**[...]]]

**pgid**[,**pgid**[,**pgid**[...]]]

The name of a valid superclass or subclass must be specified to manually assign the target processes to a class. If the target class is a superclass, each process is assigned to one of the subclasses of the specified superclass according to the assignment rules for the subclasses of this superclass.

A manual assignment remains in effect (and a process remains in its manually assigned class) until:

- The process terminates

- Workload Management (WLM) is stopped. When WLM is restarted, the manual assignments in effect when WLM was stopped are lost.
- The class the process has been assigned to is deleted
- A new manual assignment overrides a prior one.
- The manual assignment for the process is canceled using the **-u** flag.
- The process calls the **exec()** routine.

The name of a valid superclass or subclass must be specified to manually assign the target processes to a class. The assignment can be done or canceled at the superclass level, the subclass level or both. When a manual assignment is canceled for a process, or the process calls **exec()**, the process is then subject to automatic classification; if inheritance is enabled for the class that the process is in, it will remain in that class, otherwise the process will be reclassified according to the assignment rules.

For a manual assignment:

- If the *Class\_Name* is the name of a superclass, the processes in the list are assigned to the superclass. The subclass is then determined, for each process, using the assignment rules for the subclasses of the target superclass.
- If the class name is a subclass name (*supername.subname*), the processes by default are assigned to both the superclass and the subclass. The processes can be assigned to the superclass only by specifying the **-S** flag or the subclass only by specifying the **-s** flag.

```
wlmassign super1.sub2 -S pid1
```

is equivalent to:

```
wlmassign super1 pid1
```

To assign a process to a class or cancel a prior manual assignment, the user must have authority both on the process and on the target class. These constraints translate into the following:

- The root user can assign any process to any class.
- A user with administration privileges on the subclasses of a given superclass (that is, the user or group name matches the user or group names specified in the attributes **adminuser** and **admingroup** of the superclass) can manually reassign any process from one of the subclasses of this superclass to another subclass of the superclass.
- Users can manually assign their own processes (same real or effective user ID) to a class, for which they have manual assignment privileges (that is, the user or group name matches the user or group names specified in the attributes **authuser** and **authgroup** of the superclass or subclass).

This defines 3 levels of privilege among the persons who can manually assign processes to classes, root being the highest. For a user to modify or terminate a manual assignment, they must have at least the same level of privilege as the person who issued the last manual assignment.

**Note:** The **wlmassign** command works with currently loaded WLM configuration. If the current configuration is a set, and the assignment is made to a class which does not exist in all configurations in the set, the assignment will be lost when a configuration that does not contain the class becomes active (class is deleted).

## Flags

<b>Item</b>	<b>Description</b>
<b>-g</b> <i>pgid_list</i>	Indicates that the following list is a list of PGID.
<b>-S</b>	Specifies that the assignment is to be done or canceled at the superclass level only. This flag is used with a subclass name of the form <i>supername.subname</i> .
<b>-s</b>	Specifies that the assignment is to be done or canceled at the subclass level only. This flag is used with a subclass name of the form <i>supername.subname</i> .
<b>-u</b>	Cancel any manual assignment in effect for the processes in the <b>pid_list</b> or the <b>pgid_list</b> . If none of the <b>-s</b> or <b>-S</b> flags are used, this cancels the manual assignments for both the superclass and the subclass level.
<b>-r</b>	Removes a WLM tag from the specified process or the process group list.
<b>-t</b> <i>tag</i>	Sets a WLM tag for the specified process or the process group list.
<b>-i</b> <i>inheritance</i>	Specifies one or both tag inheritance sub-options in a comma-separated list. The following tag inheritance sub-options can be specified: <ul style="list-style-type: none"> <li><b>fork</b> Specifies that the children of this process should inherit the parent tag across fork.</li> <li><b>exec</b> Specifies that the process retains its tag after a call to exec.</li> </ul>

**Related information:**

lsclass command

rmclass command

Workload management

## wlmcheck command

### Purpose

Check automatic assignment rules and/or determines the Workload Manager class a process with a specified set of attributes would be classified in.

### Syntax

```
wlmcheck [ -d Config ] [ -a Attributes ] [ -q ]
```

### Description

The **wlmcheck** command with no arguments, gives the status of Workload Manager (WLM) and makes some coherency checks:

- Displays the current status of WLM (running/non running, active/passive, rsets bindings active, total limits enabled).
- Displays the status files that report the last loading errors, if any. If 'current' configuration is a set, this applies to all configurations in the set, and messages logged by the WLM daemon are reported.
- Checks the coherency of the attributes and assignment rules file(s) (such as, the existence of the classes, validity of user and group names, existence of application file names, etc).

If the **-d** *Config* flag is not specified, the checks are performed on the 'current' configuration.

The **wlmcheck** command can apply to a configuration set. In this case, the checks mentioned above are performed on all configurations of the set, after checking the set itself. Superclass names are reported in the form 'config/superclass' to indicate the regular configuration which they belong to.

Specifying a configuration with **-d** *Config* performs the checks on the *Config* configuration or set instead of 'current'. This does not change the reporting of status files and of the WLM daemon log, which only applies to the active configuration.

With the **-a** flag, **wlmcheck** displays the class that the process with attributes specified by *Attributes* would be assigned to, according to the rules for the current or specified configuration or configuration set. The format of the *Attributes* string is similar to an entry in the *rules* file, with the following differences:

- The class field is omitted (it is actually an output of **wlmcheck**)
- Each field can have at most one value. Exclusion (!), attribute groupings (\$), comma separated lists, and wild cards are not allowed. For the *type* field, the AND operator "+" is allowed, since a process can have several of the possible values for the type attribute at the same time. For instance a process can be a 32 bit process and call plock, or be a 64 bit fixed priority process.
- At least one field must be specified (have a value different from a hyphen (-)).

In addition, the first 2 fields are mandatory. The other fields, if not present default to a hyphen (-) which mean that any value in the corresponding field of an assignment rule is a match. When one or more of the fields in the attribute string are either not present or specified as a hyphen (-), the string is likely to match more than one rule. In this case, **wlmcheck** displays all the classes corresponding to all the possible matches.

Example of valid attribute strings:

```
$ wlmcheck -a "- root system /usr/lib/frame/framemaker - -"
$ wlmcheck -a "- - staff - 32bit+fixed"
$ wlmcheck -a "- bob"
```

## Flags

Item	Description
<b>-a</b> <i>Attributes</i>	Passes a set of values for the classification attributes of the process in order to determine which class the process would be put into. This is a way to check that the assignment rules are correct and classify processes as expected.
<b>-d</b> <i>Config</i>	Uses the WLM property files in <b>/etc/wlm/Config</b> (which may indicate a set of time-based configurations) instead of <b>/etc/wlm/current</b> .
<b>-q</b>	Suppresses the output of the status of the latest activation/update of WLM and of messages logged by the WLM daemon (quiet mode).

## Security

**Attention RBAC users and Trusted AIX users:** This command can perform privileged operations. Only privileged users can run privileged operations. For more information about authorizations and privileges, see Privileged Command Database in *Security*. For a list of privileges and the authorizations associated with this command, see the **lssecattr** command or the **getcmdattr** subcommand.

## Files

Item	Description
<b>classes</b>	Contains the names and definitions of the classes.
<b>limits</b>	Contains the resource limits enforced on the classes.
<b>rules</b>	Contains the automatic assignment rules.
<b>shares</b>	Contains the resource shares allocated to the classes.

### Related information:

RBAC

Trusted AIX

rules command



---

## wlmcntrl Command

### Purpose

Starts or stops the Workload Manager.

### Syntax

```
wlmcntrl [ [ -a | -c | -p ] [ -T [ class | proc ] [ -g ] [ -d Config_Dir ] [ -o | -q ]
```

```
wlmcntrl -u [ -S Superclass | -d Config_Dir ]
```

### Description

The **wlmcntrl** command stops, starts, updates or queries the state of Workload Manager (WLM). When starting or updating WLM, the WLM property files for the target configuration are pre-processed, and the data is loaded into the kernel. WLM can be started in two different modes:

- An active mode where WLM monitors and regulates the processor, memory and disk I/O utilization of the processes in the various classes.
- A passive mode where WLM only monitors the resource utilization without interfering with the standard operating system resource allocation mechanisms.

The active mode is the usual operating mode of WLM.

The classes, their limits and shares are described respectively in the **classes**, **limits**, and **shares** files. The automatic assignment rules are taken from the **rules** file. The class properties files for the superclasses of the WLM configuration **Config** are located in the subdirectory **/etc/wlm/Config**. The class properties files for the subclasses of the superclass **Super** of the configuration **Config** are located in **/etc/wlm/Config/Super**. The standard configuration shipped with the operating system is in **/etc/wlm/standard**. The current configuration is the one in the directory pointed to by the symbolic link **/etc/wlm/current**.

When the **-d Config\_dir** flag is not used, **wlmcntrl** uses the configuration files in the directory pointed to by the symbolic link **/etc/wlm/current**.

When the **-d Config\_dir** flag is used, **wlmcntrl** uses the configuration files in **/etc/wlm/Config\_dir** and updates the **/etc/wlm/current** symbolic link to point to **/etc/wlm/Config\_dir**, making **/etc/wlm/Config\_dir** the current configuration. This is the recommended way to make **/etc/wlm/Config\_dir** the current configuration.

When updating WLM using the **-u** flag, an empty string can be passed as **Config\_dir** with the **-d** flag:

```
wlmcntrl -u -d ""
```

will simply refresh (reload) the assignment rules of the current configuration into the kernel without reloading the class definitions. This can be useful when a prior activation of WLM detected that some application files could not be accessed. After the system administrator has fixed the problems with either the rules or the files, this command can be used to reload only the rules.

The WLM configuration **Config** may also be a set of time-based configurations, in which case the subdirectory **/etc/wlm/Config** does not contain the properties files, but a list of configurations and the times of the week when they apply. The properties files are still in the subdirectory of each regular configuration of the set. When WLM is started or updated which such a set, a daemon is responsible for switching regular configurations of the set when the applicable one changes.

**Note:** This command is not supported when executed within a workload partition.

## Flags

Item	Description																																																		
<b>-a</b>	Starts WLM in active mode or switches from passive to active mode. This is the default when no flag other than <b>-d</b> , <b>-g</b> , or <b>-T</b> is specified.																																																		
<b>-c</b>	Starts WLM in processor-only mode or switches from any mode to processor-only mode. In this mode, the WLM accounts for all resources, but only processor resource is regulated.																																																		
<b>-d</b> <i>Config_dir</i>	Uses <b>/etc/wlm/Config_dir</b> as an alternate directory for the WLM configuration (containing the classes, limits, shares and rules files) or configuration set (containing the list of configurations and the time tanges when they apply). This makes <b>/etc/wlm/Config_dir</b> the current configuration. This flag is effective when starting the WLM in active, processor-only or passive mode, or when updating the WLM. This flag cannot be used in conjunction with the <b>-o</b> and <b>-q</b> flags or when switching from a mode (among active, processor-only and passive) to another.																																																		
<b>-g</b>	Instructs WLM to ignore any potential resource set bindings. This means that all classes have access to the whole resource set of the system, regardless of whether or not they use a restricted resource set.																																																		
<b>-o</b>	Stops Workload Manager.																																																		
<b>-p</b>	Start WLM in passive mode or switches from any mode to passive mode. In this mode, the WLM accounts for all resources, but no resource is regulated.																																																		
<b>-q</b>	Queries the current state of WLM. Returns:  <table><tbody><tr><td>0</td><td>WLM is running in active mode.</td></tr><tr><td>1</td><td>WLM is not started.</td></tr><tr><td>2</td><td>WLM is running in passive mode.</td></tr><tr><td>3</td><td>WLM is running in active mode with no rset bindings.</td></tr><tr><td>4</td><td>WLM is running in passive mode with no rset bindings.</td></tr><tr><td>5</td><td>WLM is running in active mode for processor only</td></tr><tr><td>6</td><td>WLM is running in active mode for processor only with no rset bindings.</td></tr><tr><td>16</td><td>WLM is running in active mode, process total accounting is off.</td></tr><tr><td>18</td><td>WLM is running in passive mode, process total accounting is off.</td></tr><tr><td>19</td><td>WLM is running in active mode with no rset bindings, process total accounting is off.</td></tr><tr><td>20</td><td>WLM is running in passive mode with no rset bindings, process total accounting is off.</td></tr><tr><td>21</td><td>WLM is running in active mode for processor only, process total accounting is off.</td></tr><tr><td>22</td><td>WLM is running in active mode for processor only with no rset bindings, process total accounting is off.</td></tr><tr><td>32</td><td>WLM is running in active mode, class total accounting is off.</td></tr><tr><td>34</td><td>WLM is running in passive mode, class total accounting is off.</td></tr><tr><td>35</td><td>WLM is running in active mode with no rset bindings, class total accounting is off.</td></tr><tr><td>36</td><td>WLM is running in passive mode with no rset bindings, class total accounting is off.</td></tr><tr><td>37</td><td>WLM is running in active mode for processor only, class total accounting is off.</td></tr><tr><td>38</td><td>WLM is running in active mode for processor only with no rset bindings, class total accounting is off.</td></tr><tr><td>48</td><td>WLM is running in active mode, class and process total accounting are off.</td></tr><tr><td>50</td><td>WLM is running in passive mode, class and process total accounting are off.</td></tr><tr><td>51</td><td>WLM is running in active mode with no rset bindings, class and process total accounting are off.</td></tr><tr><td>52</td><td>WLM is running in passive mode with no rset bindings, class and process total accounting are off.</td></tr><tr><td>53</td><td>WLM is running in active mode for processor only, class and process total accounting are off.</td></tr><tr><td>54</td><td>WLM is running in active mode for processor only with no rset bindings, class and process total accounting are off.</td></tr></tbody></table>	0	WLM is running in active mode.	1	WLM is not started.	2	WLM is running in passive mode.	3	WLM is running in active mode with no rset bindings.	4	WLM is running in passive mode with no rset bindings.	5	WLM is running in active mode for processor only	6	WLM is running in active mode for processor only with no rset bindings.	16	WLM is running in active mode, process total accounting is off.	18	WLM is running in passive mode, process total accounting is off.	19	WLM is running in active mode with no rset bindings, process total accounting is off.	20	WLM is running in passive mode with no rset bindings, process total accounting is off.	21	WLM is running in active mode for processor only, process total accounting is off.	22	WLM is running in active mode for processor only with no rset bindings, process total accounting is off.	32	WLM is running in active mode, class total accounting is off.	34	WLM is running in passive mode, class total accounting is off.	35	WLM is running in active mode with no rset bindings, class total accounting is off.	36	WLM is running in passive mode with no rset bindings, class total accounting is off.	37	WLM is running in active mode for processor only, class total accounting is off.	38	WLM is running in active mode for processor only with no rset bindings, class total accounting is off.	48	WLM is running in active mode, class and process total accounting are off.	50	WLM is running in passive mode, class and process total accounting are off.	51	WLM is running in active mode with no rset bindings, class and process total accounting are off.	52	WLM is running in passive mode with no rset bindings, class and process total accounting are off.	53	WLM is running in active mode for processor only, class and process total accounting are off.	54	WLM is running in active mode for processor only with no rset bindings, class and process total accounting are off.
0	WLM is running in active mode.																																																		
1	WLM is not started.																																																		
2	WLM is running in passive mode.																																																		
3	WLM is running in active mode with no rset bindings.																																																		
4	WLM is running in passive mode with no rset bindings.																																																		
5	WLM is running in active mode for processor only																																																		
6	WLM is running in active mode for processor only with no rset bindings.																																																		
16	WLM is running in active mode, process total accounting is off.																																																		
18	WLM is running in passive mode, process total accounting is off.																																																		
19	WLM is running in active mode with no rset bindings, process total accounting is off.																																																		
20	WLM is running in passive mode with no rset bindings, process total accounting is off.																																																		
21	WLM is running in active mode for processor only, process total accounting is off.																																																		
22	WLM is running in active mode for processor only with no rset bindings, process total accounting is off.																																																		
32	WLM is running in active mode, class total accounting is off.																																																		
34	WLM is running in passive mode, class total accounting is off.																																																		
35	WLM is running in active mode with no rset bindings, class total accounting is off.																																																		
36	WLM is running in passive mode with no rset bindings, class total accounting is off.																																																		
37	WLM is running in active mode for processor only, class total accounting is off.																																																		
38	WLM is running in active mode for processor only with no rset bindings, class total accounting is off.																																																		
48	WLM is running in active mode, class and process total accounting are off.																																																		
50	WLM is running in passive mode, class and process total accounting are off.																																																		
51	WLM is running in active mode with no rset bindings, class and process total accounting are off.																																																		
52	WLM is running in passive mode with no rset bindings, class and process total accounting are off.																																																		
53	WLM is running in active mode for processor only, class and process total accounting are off.																																																		
54	WLM is running in active mode for processor only with no rset bindings, class and process total accounting are off.																																																		

A message indicating the current state of WLM is printed on STDOUT.

Item	Description
-S <i>Superclass</i>	Requests an update of WLM that is limited to the subclasses of the Superclass. Use this flag with the -u flag. If the running configuration is a set of time-based configurations, Superclass must be given in the form "config/Superclass" where "config" is the regular configuration of the set which the Superclass belongs to. If "config" is the currently active configuration of the set, the changes will take effect immediately, else they will take effect at the next time "config" will be made active.
-T	Disables both class and process total limits accounting and regulation.
-T <i>class</i>	Disables only class total limits accounting and regulation.
-T <i>proc</i>	Disables only process total limits accounting and regulation.
-u	Updates the WLM. A single update operation can change the attributes, limits and shares of existing classes and/or add or remove classes. If the running configuration is a set, this operation refreshes the set description along with the content of all configurations of the set. Update can be used by a user with root authority to switch to an alternate configuration or configuration set. Update can also be used by a superclass administrator to update only the subclasses of the superclass he has administrative access to (using the -S flag).

## Security

Access Control: Starting, stopping, switching from one mode to another, and updating superclasses or a configuration set requires root privileges. Updating the subclasses of a given superclass requires only admin user or admin group privileges (superclass administrator). Any user can query the state of WLM.

**Attention RBAC users and Trusted AIX users:** This command can perform privileged operations. Only privileged users can run privileged operations. For more information about authorizations and privileges, see Privileged Command Database in *Security*. For a list of privileges and the authorizations associated with this command, see the **lssecattr** command or the **getcmdattr** subcommand.

## Files

Item	Description
<b>classes</b>	Contains the names and definitions of the classes.
<b>limits</b>	Contains the resource limits enforced on the classes.
<b>rules</b>	Contains the automatic assignment rules.
<b>shares</b>	Contains the resource shares allocated to the classes.
<b>description</b>	Contains the description text for each configuration.
<b>groupings</b>	Contains attribute value groupings for the configuration

### Related information:

Workload management  
rmclass command

---

## wlmstat Command

### Purpose

Shows Workload Manager (WLM) per class resource utilization statistics.

### Syntax

```
wlmstat [-l Class | -t Tier] [-S | -s] [-@] [-c ] [-m] [-b] [-B Device] [-T] [-a] [-w] [-v] [Interval] [Count]
```

```
wlmstat [-l Class | -t Tier] [-@] [-c] [-m] [-b] [-u] [Interval] [Count]
```

```
wlmstat [-l Class | -t Tier] [-@] [-M] [-S | -s] [-w] [-v] [Interval] [Count]
```

## Description

The **wlmstat** command symbolically displays the contents of WLM data structures retrieved from the kernel. If a *Count* is specified, **wlmstat** loops *Count* times and sleeps *Interval* seconds after each block is displayed. If *Interval* and *Count* are not specified, one output report is produced. If *Interval* is specified but no *Count* is given, **wlmstat** outputs results continuously at the given interval until stopped by a signal (SIGINTR, SIGQUIT, and SIGKILL). By default, **wlmstat** displays the statistics for all the resources for every superclass and subclass. You can specify flags to narrow the focus of the statistics to a type of resource, tier, superclass, or subclass and alter the output format.

**Note:** The following should be considered when viewing the **wlmstat** output:

1. Starting with AIX 5.3, the WLM processor usage values and process priority adjustments are updated 10 times per second by default.
2. The value displayed for processor usage is not the current instantaneous usage from the last second, but is instead an average of the last *N* readings (starting with AIX 5.3, the default value for *N* is 15).
3. The Unmanaged class is used to report system interrupt time and for tracking memory usage for all of the pinned pages in the system that are not managed by the WLM. No processes are assigned to this class.

It is possible for a process with a hard limit of 50 percent to use more than 50 percent of the processor between two consecutive WLM usage updates. Each tenth of a second, every process is assigned a priority, and the scheduler then schedules all processes based on their assigned priorities. A process might receive more of the processor resources than the process hard limit between WLM updates.

By default, each instantaneous value of processor usage from each update is kept for the following 15 readings and is averaged with the other 14 readings before being displayed by **wlmstat**. This can potentially result in a value of greater than 50 percent due to a single instance of more than 50 percent usage between WLM updates.

The priority of a process will be greatly reduced and the process will be unable to run if the process consistently reaches or exceeds its hard limit. Over the long term, the resource utilization of the process must be at or under the process hard maximum. Over a short time interval, **wlmstat** may show the process using more than the process hard limit. The `/usr/samples/kernel/wlmtune` command that is available in the bos.adt.samples PTF can be used to modify the behavior of WLM in such an instance. The related tunables are:

### **schedhz**

The frequency at which the WLM scheduler recalculates class consumption and priority for processor. The default is 10. Modifying this value changes the responsiveness of WLM. Increasing this value causes WLM to update more frequently, thereby reducing the possibility of a process exceeding its hard limit during a short time interval. The trade-off for this is increased overhead, since more WLM processing occurs. This can potentially affect overall system performance.

### **cpuhist**

The number of consecutive processor consumption values used in the average calculation. The default is 15. Increasing this value further smooths the reported processor usage values by averaging over a longer period.

To make WLM more responsive so that classes do not exceed their maximums over long periods, it is recommended that you first try modifying **schedhz** until the **wlmstat** output displays the desired results. You may want to also modify **cpuhist** so that **wlmstat** averages over the same time interval. For example, if **schedhz** is 20 and **cpuhist** is 15, **wlmstat** will average over a period of 0.75 seconds (15/20), so you may want to change **cpuhist** to 30 so that **wlmstat** still averages over 1.5 seconds.

On systems with no contention for processor, an *Interval* of 5 for **wlmstat** is recommended in order to adhere to WLM limits.

## Flags

Item	Description
-@	Displays workload partition resource information.
-a	Displays subclass consumption in absolute terms. By default, the subclass consumption percentages are shown relative to the superclass consumption. With this option, subclass consumption is displayed relative to the total amount of resource available on the system (as is done for superclasses). All values are displayed with 1% precision. For instance, if a superclass has a processor target of 20% and the processor percentage shown by <b>wlmstat</b> without <b>-a</b> for a subclass is 10%, <b>wlmstat</b> with <b>-a</b> shows the processor percentage for the subclass as 2%.
-b	Displays only disk I/O statistics.
-B <i>Device</i>	Displays disk I/O device statistics. Passing an empty string (-B "") displays the statistics for all the disks accessed by the class.
-c	Shows only processor statistics.
-l <i>Class</i>	Displays statistics for <i>Class</i> name. If not specified, all classes display along with a summary for appropriate fields.
-m	Shows only physical memory statistics.
-M	Displays the Real/Virtual Memory statistics. Use of the <b>-M</b> option adds the following columns in the output:  <b>RMSIZ</b> Utilized real memory size for the class <b>VMSIZ</b> Utilized virtual memory size for the class <b>RMLIM</b> Real memory limit for the class <b>VMLIM</b> Virtual memory limit for the class <b>LGPGSIZ</b> Utilized large pages in the class <b>LGPGLIM</b> Large page limit for the class <b>Note:</b> A - will be displayed for the <b>RMLIM</b> , <b>VMLIM</b> , and <b>LGPGLIM</b> fields if the limit is undefined. When the <b>-M</b> and <b>-w</b> options are used together, <b>RMSIZ</b> and <b>VMSIZ</b> fields contain the high watermarks for these attributes instead of the actual utilized values. In addition, the <b>LGPGSIZ</b> and <b>LGPGLIM</b> fields is turned off.
-s	Displays only subclass statistics.
-S	Displays only superclasses statistics.
-t <i>Tier</i>	Displays statistics only for the specified <i>Tier</i> .
-T	Displays the total numbers for resource utilization since WLM was started or the class was created, whichever is the latter. The units are:  <b>CPU</b> The total processor time, in milliseconds, consumed by a class <b>MEM</b> Unused <b>DKIO</b> The total number of 512 byte blocks sent/received by a class for all the disk devices accessed.

**Item****Description****-v**

Specifies verbose mode. This flag, intended for trouble shooting, also displays some class attributes, resource shares and limits and other WLM parameters, including internal parameter values intended for AIX support personnel. The following information can be of interest for users:

**Column Header**

Description

**CLASS** Class name.**tr** tier number (0 to 9)**i** Value of the inheritance attribute: 0 = no, 1 = yes.**#pr** Number of processes in the class. If a class has no (0) process assigned to it, the values shown in the other columns might not be significant.**CPU** Processor utilization of the class (%).**MEM** Physical memory utilization of the class (%).**DKIO** Disk IO bandwidth utilization for the class (%).**sha** Number of shares ('-' is represented as -1)**min** Resource minimum limit (%)**smx** Resource soft maximum limit (%)**hmx** Resource hard maximum limit (%)**des** (desired): percentage goal (target) calculated by WLM using the shares numbers (%)**npg** Number of memory pages owned by the class.

The other columns are for internal use only and bear no meaning for administrators and end users. This format is better used with a resource selector (**-c**, **-m**, or **-b**), otherwise the lines might be too long to fit into a line of a display terminal.

**-w**

Displays the memory *high water mark*, that is the maximum number of pages that a class had in memory at any given time since WLM was started or the class was created (whichever happened last).

**-u**

Displays per-tier and total unused resources.

**Display**

Results are tabulated, with the following fields:

<b>Name</b>	<b>Class name</b>
CPU	Percentage of total processor time consumed by the class.
MEM	Percentage of physical memory consumed by the class.
DKIO	Percentage of the disk IO bandwidth consumed by the class. This number is the average of the disk bandwidth on all the disk devices accessed by the class, and is usually not significant. For instance if a class consumes 80% of the bandwidth of one disk and 5% of the bandwidth of two other disks, the DKIO column shows 30%. For details on the per device utilization, use the -B device option.

**Examples**

- To get a printout of WLM activity right now, enter:

```
wlmstat
```

This produces the following output:

```

      CLASS CPU MEM DKIO
Unclassified  0  0  0
  Unmanaged  0  0  0
    Default  0  0  0
    Shared   0  0  0
    System   0  0  0
    class1   12  0  0
class1.Default  4  0  0

```

```

class1.Shared      0  0  0
class1.subclass1   4  0  0
class1.subclass2   4  0  0
class2              12  0  0
class2.Default     4  0  0
class2.Shared      0  0  0
class2.subclass1   4  0  0
class2.subclass2   4  0  0

```

2. To get a report for superclass **class1**, enter:

```
wlmstat -l class1
```

This produces the following output:

```

CLASS CPU MEM DKIO
class1 12  0  0
class1.Default  4  0  0
class1.Shared   0  0  0
class1.subclass1 4  0  0
class1.subclass2 4  0  0

```

3. To get a report for subclass **class1.subclass2** updated every 10 seconds, for one minute, enter:

```
wlmstat -l class1.subclass2 10 6
```

This produces the following output:

```

CLASS CPU MEM DKIO
class1.subclass2  4  0  0
class1.subclass2  4  0  0
class1.subclass2  4  0  0
class1.subclass2  4  0  0
class1.subclass2  4  0  0
class1.subclass2  4  0  0

```

4. To display virtual/real memory statistics, enter:

```
wlmstat -M
```

This produces the following output:

CLASS	RMSIZ	RMLIM	VMSIZ	VMLIM	LGPGSIZ	LGPGGLIM
Unmanaged	1024	4096	4096	8192	0	-
Default	0	-	0	-	0	-
Shared	0	-	0	-	0	-
System	23567	50000	819234	1000000	0	-

5. To display the memory high water mark, enter:

```
wlmstat -M -w
```

This produces the following output:

CLASS	RMSIZ	RMLIM	VMSIZ	VMLIM
Unmanaged	1024	4096	4096	8192
Default	0	-	0	-
Shared	0	-	0	-
System	23567	50000	819234	1000000

## Errors

A warning message is issued by **wlmstat** if WLM is not started.

### Related reference:

“wlmcntrl Command” on page 105

---

## wol command

### Purpose

Wakes up one or more hosts that are connected to a network in suspend mode by sending a Magic Packet.

### Syntax

To send a Magic Packet to a subnet-directed broadcast address:

```
wol { [-m MACAddress [ [-h Host -s SubnetMask ] | -i Interface ] | -f File } [ -v ]
```

To send a Magic Packet to a multicast address:

```
wol { -m MACAddress -M MulticastAddress [ -p Port ] [ -i Interface ] | -f File } [ -v ]
```

### Description

The **wol** command wakes up one or more hosts that are connected to a network in suspend mode by sending a Magic Packet to the specified address or addresses on the specified subnet.

If the user doesn't specify either the **-h**, nor **-s** flag, the **wol** manager will broadcast the Magic Packet as follows:

- If the user specifies the interface name (**-i Interface**), the Magic Packet will be broadcast from the specified interface.
- If the user doesn't specify the interface name, then the **wol** manager will loop through each network interface installed on the machine. If an interface is up, it will broadcast the Magic Packet from that interface, and then continue to the next interface until it goes through the entire interface list on the machine.

The file specified with **-f File** contains the list of hosts which need to be awakened. This file consists of one or more lines, each line containing the following information in this format:

MacAddress; Hostname/IPaddress; SubnetMask; Multicast; Port; Interface

For example, the file might look like this:

```
00:20:35:7a:7:89a;      9.41.86.19;      255.255.255.0 ; ; ;
00:04:ac:17:c0:9f ;      obiwan.aoot.austin.ibm.com;      255.255.255. 224; ; ;
00:07:be:4a:2:394; ; ; ; en0
00:06:38:6b:7e:8f ;      ; ; 234.5.6.7; 12345 ;
```

A line starting with a "#" character is a comment and is ignored. Each line contains 6 tokens separated by ";" character. The MAC address is mandatory. The other tokens are optional, but the ";" character must be used to separate unused tokens.

### Flags



Item	Description
<b>-i</b> <i>Interface</i>	Specifies the interface to use on the host where the <b>wol</b> command is being run
<b>-f</b> <i>File</i>	Specifies the name of a file containing a group list. This allows the user to wake a specified group of hosts.
<b>-h</b> <i>Host</i>	Specifies a host to wake, either as a hostname or as an IPv4 address in dot string representation (for example, 10.0.0.3).
<b>-m</b> <i>MACAddress</i>	Specifies the a 48 bits MAC address of the host in hex representation (for example, 00:20:35:7a:78:9a).
<b>-M</b> <i>MulticastAddress</i>	Specifies an IPv4 multicast address.
<b>-p</b> <i>Port</i>	Specifies the port to use on the multicast machine.
<b>-s</b> <i>SubnetMask</i>	Specifies an IPv4 subnet mask in dot string representation (for example, 255.255.255.0).
<b>-v</b>	Specifies verbose mode.

## Exit Status

Item	Description
0	The command completed successfully.
>0	An error occurred.

## Security

**Attention RBAC users and Trusted AIX users:** This command can perform privileged operations. Only privileged users can run privileged operations. For more information about authorizations and privileges, see Privileged Command Database in *Security*. For a list of privileges and the authorizations associated with this command, see the **lssecattr** command or the **getcmdattr** subcommand.

## Location

/usr/sbin/wol

Related information:

RBAC

Trusted AIX

## wparerr Command

### Purpose

Logs an error message for a specific Workload partition (WPAR).

### Syntax

From global

```
wparerr -w WparName [{"-c <cat file> -s <set no> -n <msg no>}] -S default_string args
```

From WPAR

```
wparerr [ -c <cat file> -s <set no> -n <msg no>}] -S default_string args
```

### Description

The **wparerr** command provides a mechanism to log error messages for a given WPAR. Each WPAR can hold up to 1 KB of error messages. If there is enough space to log a new message, the command logs the message; otherwise, it fails. The **-w** option should not be used inside a WPAR. Everything after **-S** flag is treated as arguments for the message.

## Flags

Item	Description
-w	Specifies the name of the workload partition for which the message should be logged.
-c	Specifies the catalog file name to be used for translation.
-s	Specifies the message set number of the error message in the catalog file.
-n	Specifies the message number of the error message.
-S	Specifies the default message string. Follows the same syntax as the <b>printf</b> subroutine <i>Format</i> parameter. Floating point is not supported.
args	Specifies the arguments to the message if any.

## Security

### Attention RBAC users and Trusted AIX users

This command can perform privileged operations. Only privileged users can run privileged operations. For more information about authorizations and privileges, see **Privileged Command Database** in the security. For a list of privileges and the authorizations associated with this command, see the **lssecattr** command or the **getcmdattr** sub commands.

### Example

- To log a message for the WPAR *mywpar* from global:  
`wparerr -w mywpar -c wparerrs.msg -s 1 -n 12 -S " %s failed with return value %d\n" lswpar -1`
- To log a command failure from inside a WPAR:  
`wparerr -c wparerrs.msg -s 1 -n 5 -S " %s application does not allow checkpoint\n" myapplication`

### Related information:

kwwpar\_err command

lssecattr command

wpar\_log\_err command

wpar\_print\_err command

---

## wparexec Command

### Purpose

Creates an application workload partition or specification file.

### Syntax

```
wparexec [ -a ] [ -c ] [ -1 ] [ -F ] [ -h hostname ] [ -H architecture ] [ -i ] [ -I attribute=value ... ] ... [ -M attribute=value... ] ... [ -N attribute=value... ] ... [ -R attribute=value... ] [ -u userscript ] [ -v ] [ -x ] { -n wparname [ -e existingwpar | -f infile ] [ -o outfile [ -w ] ] | -f infile [ -n wparname ] [ -o outfile [ -w ] ] | -w -o outfile [ -n wparname ] [ -e existingwpar | -f infile ] } [ -- ] [ var=value ... ] /path/to/command [ arg ... ]
```

### Note:

Regardless of locale, only ASCII characters are allowed as arguments to the **wparexec** command.

In addition to this, there are more restrictions for a WPARs name:

- May not be more than 25 bytes.
- May not contain white space or any of the following symbols:  
`= : / ! ; ` ' " < > ~ & ( ) * + [ ] , . ^ 0 { } | \`

- May not start with '-' or '0'.

## Description

The **wparexec** command builds and starts an application workload partition, or creates a specification file to simplify the creation of future application workload partitions.

An application workload partition is an isolated execution environment that might have its own network configuration and resource control profile. Although the partition shares the global file system space, the processes that are running therein are only visible to other processes in the same partition. This isolated environment allows process monitoring, gathering of resource, accounting, and auditing data for a predetermined cluster of applications.

The **wparexec** command starts and monitors a single application within this isolated environment. The **wparexec** command returns synchronously with the return code of this tracked process only when all of the processes in the workload partition terminate. For example, if the tracked process creates a daemon and exits with the 0 return code, the **wparexec** command will block until the daemon and all of its children terminate, and then exit with the 0 return code, regardless of the return code of the daemon or its children.

## Flags

Item	Description
<b>-1</b>	Creates the configuration only. Causes the <b>wparexec</b> command to stop after creating the configuration of the application WPAR. The <b>startwpar</b> command must then be used to start the WPAR. Only advanced users can use the <b>-1</b> option.
<b>-a</b>	Automatically resolves conflicting static settings if required. Resolvable settings are name, hostname, and network configuration.
<b>-c</b>	Enables this workload partition to be checkpointed. This option is only valid when additional checkpoint-restart software is installed and configured. When this option is used, any file systems associated with this workload partition (for example, with the <b>-M</b> option) must be remote (for example, <b>vfs=nfs</b> ).
<b>-e existingwpar</b>	Uses an existing application workload partition as the source for specification data. Do not use the <b>-e</b> flag with the <b>-f</b> flag. Any values specified by other <b>wparexec</b> flags override those values from the existing workload partition.
<b>-f infile</b>	Indicates the specification file to read default values from. Do not use with the <b>-e</b> flag. Any values specified by other <b>wparexec</b> flags override those values from the loaded specification file.
<b>-F</b>	Suppresses or overrides most error conditions. With the <b>-F</b> flag, the <b>wparexec</b> command continues with a warning.
<b>-h hostname</b>	Specifies a host name for this workload partition. If not specified, the <b>wparexec</b> command uses the workload partition name as host name.
<b>-H architecture</b>	Creates an architecture-compatible workload partition. Valid architecture values are <b>pwr4</b> , <b>ppc970</b> , <b>pwr5</b> , <b>pwr6</b> , and <b>pwr7</b> . The architecture value must be earlier, or equal to, the system hardware version. The applications in the workload partition are presented with the lowest common denominator of the specified architecture. If the workload partition can create a checkpoint, the workload partition is able to migrate between systems with hardware levels greater than, or equal to, the workload partition architecture. <b>Note:</b> The POWER5 processor-based systems and BladeCenter JS21 Express systems are not compatible with each other. You cannot create a JS21-compatible (ppc970) WPAR on a POWER5 processor-based system even though the JS21 using an earlier processor than a POWER5 processor-based systems.
<b>-i</b>	Enables WPAR-specific routing for the workload partition. By default, outgoing network traffic from a workload partition is routed like it is being sent from the global environment, notably in the following ways: <ul style="list-style-type: none"> <li>• Traffic between addresses that were hosted on the same global system is sent through the loopback interface.</li> <li>• Routing table entries that are configured in the global system, including the default route, are used to transmit workload partition traffic.</li> </ul> <p>If you enable WPAR-specific routing by specifying the <b>-i</b> flag, the workload partition creates and uses its own routing table for outgoing traffic. Routing entries are created automatically for each of the network addresses of the workload partition to accommodate broadcast, loopback, and subnet routes. For more information about the network attributes, see the <b>-N</b> flag. You can create explicit additions to the routing table of the workload partition using the <b>-I</b> flag. In particular, you can use the <b>-I</b> flag to configure the default route, as no default route is created automatically.</p>

Item	Description
<b>-I</b> <i>attribute=value ...</i>	<p>Adds routing table entries to the entries that are automatically created when WPAR-specific routing is in effect. You can specify more than one <b>-I</b> flag to configure multiple routes. Using the <b>-I</b> flag automatically enables WPAR-specific routing as described under the <b>-i</b> flag.</p> <p>You can specify the following attributes with the <b>-I</b> flag. The <b>rtdest</b> attribute and the <b>rtgateway</b> attribute are required to be specified.</p> <p><b>rtdest=destination</b> Identifies the host or network to which you are directing the route. You can specify the value using either a symbolic name or a numeric address. You can use the keyword <b>default</b> to specify a default route. For more information about the route <b>rtdest</b> attribute, see the <i>Destination</i> parameter of the <b>route</b> command.</p> <p><b>rtgateway=gateway</b> Identifies the gateway to which packets are addressed. You can specify the value using either a symbolic name or a numeric address.</p> <p><b>rtnetmask=A.B.C.D</b> Specifies the network mask to the destination address.</p> <p><b>rtprefixlen=n</b> Specifies the length of a destination prefix, which is the number of bits in the netmask. The value must be a positive integer.</p> <p><b>rttype={net host}</b> Forces the <b>rtdest</b> attribute to be interpreted as the specified type.</p> <p><b>rtinterface=if</b> Specifies the interface, for example, en0, to associate with the route so that packets are sent using the interface when the route is chosen.</p> <p><b>61Vrtfamily={inet inet6}</b> Specifies the address family. For information about the parameters of the <b>rtfamily</b> flag, see the parameter section of the route command.</p>
<b>-M</b> <i>directory=dir [ vfs=type ] [ dev=devicepath ] [ host=remotehost ] [ mountopts=mountopts ]</i>	<p>Specifies file system dependencies only. Attributes must be space-separated. By default, an application workload partition has the same level of access to all of the global file systems and mounts as the user who created the workload partition. Use the <b>-M</b> flag with the <b>directory</b> attribute that is set to the file system name to specify additional file systems. More than one <b>-M</b> flag can be specified.</p> <p><b>Note:</b> All of the mounts and all of the directories are created and available at global level. File systems that are based on disk, such as the <b>vfs=jfs</b> and the <b>vfs=jfs2</b>, will not be created for application workload partitions.</p> <p>A local file system dependency can be added by defining only the <b>directory</b> attribute. However, the <b>directory</b> specified must exist in the <b>/etc/filesystems</b>.</p> <p>If an error occurs during the process of creating the workload partition, any file systems mounted by the <b>wparexec</b> command are unmounted. After the creation succeeds, the file systems are not unmounted, regardless of the return status of the user application.</p> <p>The following are the valid values for the <i>vfs</i> attribute for application workload partitions:</p> <p><b>nos</b> The <b>directory</b> specified by the <b>dev</b> attribute on the system specified by the <b>host</b> attribute is mounted at the location that is specified by the <b>directory</b> attribute. If the mount point does not exist, it will be created. The only other attributes that are applicable to an <b>nfs</b> mount are the <b>mountopt</b> attributes, corresponding to the <b>-o</b> option of the <b>mount</b> command or the <b>options</b> attribute in an <b>/etc/filesystems</b> stanza. If not specified, no mount options are used by default. Acceptable option values correspond to the <b>-o</b> options to the <b>mount</b> command. For more information, see the <b>mount</b> Command in <i>Commands Reference, Volume 3</i>.</p> <p><b>namefs</b> The global <b>directory</b> specified by the <b>dev</b> attribute is mounted over the <b>directory</b> specified by the <b>directory</b> attribute.</p> <p>The only other attributes that are applicable to a <b>namefs</b> mount are the <b>mountopt</b> attributes.</p> <p><b>directory</b> The global <b>directory</b> specified by the <b>directory</b> attribute is created if it does not exist. No mounting is performed.</p>
<b>-n</b> <i>wparname</i>	<p>Specifies the name for the workload partition to be created. If no name is supplied, a name is generated.</p>

**Item****-N** *attribute=value***Description**

Allows specification of the following network configuration attributes:

- **interface**= *if* or **interface**=*namemappedif*
- **address**=*A.B.C.D*
- **netmask**=*A.B.C.D*
- **broadcast**=*A.B.C.D*
- **address6**=*S:T:U:V:W:X:Y:Z*
- **prefixlen**=*n*

The name-mapped interface is defined in the **/etc/wpars/devmap** file. You can specify the mapping between the name-mapped interface and the system interface as follows:

```
# The comments start with '#'
# Each line contains a pair of name-mapped interface
# and real interface separated by tab or blank spaces.
foo en0
goo en1
soo en2
```

The *attribute=value* pairs must be separated by spaces. More than one **-N** flag can be used to configure multiple IP addresses. At minimum, the **address** or the **address6** attribute must be specified. The **wparexec** command collects any other values that are not specified from the global system settings. If no **-N** flag is specified, the **wparexec** command will attempt to discover an appropriate IP address for the workload partition by running the **gethostbyname** subroutine on the workload partition name specified with the **-n** flag. If an address is found on the same subnet as any global interface, the settings of that interface will be used with the resolved IP address to create the default network entry.

To define an IPv6 network configuration, specify the **-N** flag with the **address6** attribute, the **prefixlen** attribute, and the **interface** attribute:

- The **address6** attribute is a 128-bit address. The address is represented by eight 16-bit integers that are separated by colons. Each integer is represented by four hex digits. Leading zeros can be skipped, and consecutive null 16-bit integers can be replaced by two colons (one time per address).
- The **prefixlen** attribute is the number of high-order bits that are used to mask the IPv6 address and to comprise the prefix. The value of the **prefixlen** attribute ranges from 0 through 128. Each **-N** flag can accept either IPv4 attributes, or IPv6 attributes, but not both.

**-o** *outfile*

Indicates an output path and file name to write specification data to. This specification file can be used to create an application workload partition later with the **-f** flag.

**Item**

**-R** *attribute=value*

**Description**

Allows specification of resource control attributes. Only one **-R** flag can be specified. Most resource controls are similar to those supported by the Workload Manager (WLM). See the listed WLM pages for descriptions of these attributes. Valid attributes are as follows:

**active** Allows resource control definitions to be retained, but rendered inactive. This attribute can take the *yes* or *no* values.

**rset** Configures this workload partition to use a resource set created by the **mkrset** command. For more information about the **mkrset** command, see **mkrset** Command in *Commands Reference, Volume 3*.

**shares\_CPU**

The number of processor shares available to this workload partition. For more information about processor shares, see Workload Manager shares File in *Files Reference*.

**CPU** The percentage of processor limits for this workload partition's processes. This attribute uses the following format to define the limits values:

CPU=<m>%-<SM>%,<HM>%

The *m* value represents the minimum limit. The *SM* value represents the soft maximum limit. The *HM* value represents the hard maximum limit. For more information about limited values, see the Workload Manager limits File in *Files Reference*.

**shares\_memory**

The number of memory shares available to this workload partition. For more information about memory shares, see Workload Manager shares File in *Files Reference*.

**memory** The percentage of memory limits for this workload partition's processes. For more information about memory limit, see the Workload Manager limits File in *Files Reference*.

**procVirtMem**

The maximum amount of virtual memory that a single process can use. Processes that exceed the specified limit are terminated. The valid units are megabytes (M or MB), gigabytes (G or GB), and terabytes (T or TB). The minimum limit allowed is 1M. The maximum limit that can be specified is 8796093022207M, 8589934591G, or 8388607T. If the value is set to -1 (no units), the limit is disabled. For more information about limit values, see Workload Manager limits File in *Files Reference*.

**totalVirtMem**

The maximum amount of virtual memory that can be used by the WPAR as a whole. Processes that cause the specified limit to be exceeded will be terminated. The valid range and units are the same as for **procVirtMem**. If the value is set to '-1' (no units), the limit is disabled. See Workload Manager limits File in *Files Reference*.

**totalProcesses**

The total number of processes that are allowed in this workload partition. For more information about allowed processor number, see Workload Manager limits File in *Files Reference*.

Item

Description

(Attributes for the -R flag, continued):

totalPTYs=*n*

The total number of pseudo terminals that are allowed in the workload partition. For more information about the allowed pseudo terminals, see *pty Special File*.

totalLargePages=*n*

The number of large pages that can be allowed for the workload partition. For more information about the allowed large pages, see *Large Pages*.

pct\_msgIDs=*n*%

The percentage of the maximum number of message queue IDs of the system that are allowed in the workload partition. For more information about the allowed number of message queue IDs, see *Message Queue Kernel Services*.

pct\_semIDs=*n*%

The percentage of the maximum number of semaphore IDs of the system that are allowed in the workload partition.

pct\_shmIDs=*n*%

The percentage of the maximum number of shared memory IDs of the system that are allowed in the workload partition. For more information about the allowed number of shared memory IDs, see *Shared Memory*.

pct\_pinMem=*n*%

The percentage of the maximum pinned memory of the system that can be allocated to the workload partition. For more information about pinned memory, see *Support for pinned memory*.

totalThreads

The total number of threads that are allowed in this workload partition. For more information about allowed number of threads, see *Workload Manager limits File* in *Files Reference*.

-u *userscript*

Specifies the path to a user script to be run by the workload partition commands at various administration points. The parameter of the -u flag can be a quoted string including additional arguments to be passed to the script. In all cases, the first component of the parameter of the -u flag must be an absolute path to an existing executable file. The script is invoked as follows:

`/path/to/userScript<action><WAPR>`

The *action* argument indicates the administrative action being performed, as follows:

- **WPAR\_LOAD:** A script runs in the global environment after kernel configuration, before the tracked process is created. If the script returns a value of non-zero, the workload partition will not be started.
- **WPAR\_START:** A script runs in the global environment after the workload partition becomes active. For application workload partitions, the script runs once the tracked process is started.
 

**Note:** This code path can be run asynchronously by a dissociated process with its standard I/O streams closed or redirected. Internal messaging must be handled accordingly, and the script must account for the fact that short-lived workload partitions might be stopped or stopping at any point during the execution of the script.

If the script returns a value of nonzero, a warning is logged, but no other behavior changes.
- **WPAR\_STOP:** A script runs in the global environment after all of the workload partition processes finish, and before the kernel is unconfigured.

**Note:** This code path can be executed by a dissociated process with its standard I/O streams closed. Internal messaging should be handled accordingly.

If the script returns a value of non-zero, a warning will be logged, but no other behavior will change.

-v

Specifies the command to run in the verbose mode.

-x

Allows access to cross-WPAR semaphores and shared memory segments.

-w

Writes the specification file. When it is used with the -o flag, the -w flag causes the **wparexec** command to quit after writing the new specification file without actually creating the workload partition.

Item	Description
<pre>[--] [var=value] /path/to/command [arg ...]</pre>	<p>Specifies the application (tracked process) to be run within the workload partition, along with any necessary environment variable settings and arguments.</p> <p>The command is required, either by this command-line syntax or the <b>general.application</b> attribute in the specification file, unless the command is only creating a specification file (with <b>-w</b> flag). When it is started, the command line provided is always shell-expanded within the workload partition. When using the command line, shell metacharacters should be escaped appropriately to prevent premature expansion.</p> <p>The special double-minus separator (<b>--</b>) is used to signify that all subsequent command-line arguments comprise the tracked process. For example, use this separator to remove ambiguity between attributes to the <b>-N</b> flag and assignment of environment variables to the tracked process.</p> <p>Only one tracked process per workload partition is supported, but this application might create other processes. The workload partition is automatically stopped and removed when all of the processes therein terminate. A workload partition might be stopped and removed prematurely by the <b>stopwpar</b> command.</p>

## Security

Access Control: Only the root user can run this command.

**Attention RBAC users and Trusted AIX users:** This command can perform privileged operations. Only privileged users can run privileged operations. For more information about authorizations and privileges, see Privileged Command Database in *Security*. For a list of privileges and the authorizations associated with this command, see the **lssecattr** command or the **getcmdattr** subcommand.

## Examples

1. To create an application workload partition that is running a benchmark program, enter:

```
wparexec -n tpcc -N address=192.168.0.51 /u/tpcc/benchmark -f /tmp/logfile
```

**Note:** The **-f** flag is passed to the **/u/tpcc/benchmark** file and is not processed by the **wparexec** command as a flag.

2. To create a workload partition based on an existing specification file, enter:

```
wparexec -f /tmp/wparexec1.spec
```

3. To override the default minimal *PATH* variable provided by the **wparexec** command, enter:

```
wparexec PATH=/usr/opt/bin:/usr/bin:/usr/sbin /home/joe/runapp
```

## Files

Item	Description
<b>/usr/samples/wpars/sample.spec</b>	An annotated workload partition specification file.

### Related information:

RBAC

Trusted AIX

startwpar command

route command

stopwpar command

---

## wpar\_reg\_script Command

### Purpose

Allows a user to register scripts to be run during different phases of a system WPAR live application mobility. Currently, only the WPAR restart phase is supported.



## Syntax

```
/usr/lib/wpars/wpar_reg_script [ -r | -u ] [ -p phase ] [ -s /path/to/script ]
```

## Description

The `wpar_reg_script` command allows a user to register scripts to be run during specific phases of WPAR live application mobility.

Registering a script will place an entry in a new ODM class called `CuWscr` that contains the full path to the script and its arguments.

This script can be used with kernel extensions and mobility to load kernel extensions on the arrival node before other processes start.

## Restrictions

The `wpar_reg_script` command can only be run from inside a WPAR.

Only the file systems that are part of the WPARs config are available at the time of execution (not privately mounted file systems). These file systems are only available read-only.

No processes can be left running when the script exits.

Loading kernel extensions with this script for use with mobility is restricted to `SYS_SINGLELOAD`.

## Flags

Item	Description
-r	Register an entry including its script and phase to the <code>CuWscr</code> class.
-u	Unregister an entry from the <code>CuWscr</code> class.
-p	The phase that the command will be executed during. Value of 1 corresponds to WPAR restart. More may be added in the future. If no value is specified, it will use a default value of 1.
-s	The path to the script and its arguments that will be executed. Surround the full command with "" to include command line parameters.

## Examples

1. To register the script `/usr/sbin/foo/` to execute:  

```
/usr/lib/wpars/wpar_reg_script -r -s /usr/sbin/foo
```
2. To register the script `/usr/sbin/foo/` with arguments:  

```
/usr/lib/wpars/wpar_reg_script -r -s "/usr/sbin/foo args"
```
3. To Unregister all instances of the script `/usr/sbin/foo/` foo:  

```
/usr/lib/wpars/wpar_reg_script -u -s /usr/sbin/foo
```
4. To unregister all scripts that run during phase 1:  

```
/usr/lib/wpars/wpar_reg_script -u -p 1
```

---

## wparprnterr Command

### Purpose

Displays error messages specific to a Workload partition (WPAR).

## Syntax

From global

```
wparprnterr wparName
```

## Description

The **wparprnterr** command prints all the error messages that are logged for a WPAR by using the **wparerr** command, **wpar\_log\_err**, and **kwpar\_err** subroutines on the standard output.

## Security

### Attention RBAC users and Trusted AIX users

This command can perform privileged operations. Only privileged users can run privileged operations. For more information about authorizations and privileges, see **Privileged Command Database** in Security. For a list of privileges and the authorizations associated with this command, see the **lssecattr** command or the **getcmdattr** subcommand.

## Examples

To display messages of the WPAR *mywpar*:

```
wparprnterr mywpar
```

### Related information:

**kwpar\_err** command

**lssecattr** command

**wpar\_log\_err** command

**wpar\_print\_err** command

---

## write Command

### Purpose

Opens a line of communication to send messages to other users on the system in real time.

### Syntax

To query all messages awaiting replies from users on a host and display them with their handles, type the following:

```
write -q [ -n Host ]
```

To Reply to a Message Sent by a Utility or a Shell Script, or Redisplay the Message Associated with a Given handle, type the following:

```
write -hHandle, { ok | cancel | query } [ -n Host ]
```

To send messages to a user, optionally on another host or a particular device, type the following:

```
write [ -r ] { [ -n Host ] User | User@Host } [ Line ]
```

## Description

The **write** command enables message sending over the system in real time. It provides conversation-like communication with another logged-in user. Each user alternately sends and receives short messages from the other workstation. Long messages can be sent by putting the complete message in a file and then redirecting that file as input to the **write** command.

For another user (as specified by the *User* parameter) to receive a message, that user must be logged in and must not have refused message permission. When a message is sent to a user who is not logged in, the message user not logged in appears. If the message is sent to a user who has refused message permission by setting the **mesg** command to no, the message write: permission denied appears.

When the **write** command is issued, it immediately sends the following message, along with an attention-getting sound (the ASCII BEL character) to the message recipient or target:

```
Message from SenderID on SenderHostname (tty) [Date] ...
```

With a successful connection, the **write** command sends two ASCII BEL characters to both workstations. The beep alerts the sender that the message can begin and it alerts the receiving user that a message is coming.

Sending occurs one line at a time as the Enter key is pressed. The communication link from the sender to the receiver remains open and sending continues until the Ctrl-D key sequence ends the sending link. Then an end-of-text character (<EOT>) is sent to the target workstation and the **write** command mode is terminated.

The receiving or target user can respond by sending a **write** command to the originating user. This opens a line of communication from the receiver back to the sender, enabling message responses in return. For this type of exchange, the following convention is useful: When you first write to others, wait for a response before sending any text. End a message with a signal such as o (over) to alert the other person to reply. Use oo (over and out) when the conversation is finished.

If the character ! (exclamation point) is found at the beginning of a line, the **write** command calls the shell to execute the rest of the line as a command.

When you write to a user who is logged in at more than one workstation or multi-using more than one process, the **write** command uses the first login instance found in the */etc/utmp* file as the message delivery point (usually the login or console shell), and you get the message:

```
UserID is logged on more than one place.  
You are connected to "Workstation".  
Other locations are:  
Workstation
```

When this message is received, if you wish to send the message to a location other than the initial login location, the target user can be contacted at a different location by specifying the *Line* of the location (tty00, for example).

Permission to write to another user is granted or denied by the individual user with the **mesg** command. Some commands deny message permission while they are running to prevent interference with their output. A user with root user authority can write to any workstation regardless of the workstation's message permission.

You can use the **write** command to converse with users on other hosts. You can identify a user on a remote host by using the **-nHostName** flag or the *User@Host* parameter. In order to write to a user on a remote host, the **writesrv** daemon must be running on both the current host and the remote host.

The **write** command is also used by the **qdaemon** daemon to send messages to users on other hosts and to wait for replies. There are only three valid replies:

Item	Description
ok	The original write exits with a status of 0.
cancel	The original write exits with a status of 1.
query	The message associated with the given handle is displayed.

## Parameters

Item	Description
<i>User</i>	Specifies the user ID of the person to receive the message text.
<i>User@Host</i>	Specifies the user ID and remote host of the person to receive the message text.
<i>Line</i>	Contacts the target user at another location (tty00, for example).

## Flags

Item	Description
<b>-h</b> <i>Handle,Reply</i>	Replies to a message sent by a utility or shell script using write with the reply option. The value to be used for the <i>Handle</i> variable is generated internally and supplied to the user in the text of the original message. The reply can be ok, cancel, or query.
<b>-nHost</b>	Specifies a remote host. The <i>Host</i> variable may be a nickname or an internet address.
<b>-q</b>	Queries all messages awaiting replies from users on a host and displays them with their handles.
<b>-r</b>	Generates a message handle, places it in the message header, sends the message, and waits for a reply. This flag is used by the <b>qdaemon</b> daemon for operator messages and can be put in shell scripts. It is not used for interactive conversations. An exit status of 0 indicates that the reply was ok, a status of 1 indicates that the reply was cancel, and an exit status of 2 indicates that the user could not be contacted.

## Requirements:

- The **writesrv** daemon must be running on the target host in order for any of the flags to work. If you are not using either the **-n** flag or *@Host*, but using **-h**, **-q**, or **-r**, the **writesrv** daemon must be running on your host.
- If TCP/IP is not installed on your machine but the *HostName* is set, in order to converse with users on the local host using the **write** command with the **-h**, **-q**, or **-r** flag, you must append your host name to the end of the loopback entry in the */etc/hosts* file. The original entry should read:

```
127.0.0.1 loopback LocalHostName
```

The new entry should read:

```
127.0.0.1 loopback LocalHostName HostName
```

## Exit Status

This command returns the following exit values:

Item	Description
0	Successful completion.
>0	The addressed user either is not logged on or denies permission.

## Examples

1. To write a message to a user who is logged in, enter:

```
write june
```

Press the Enter key and type,

```
I need to see you! Meet me in the computer room at 12:30.
```

Then press the Ctrl-D key sequence to terminate the **write** command mode.

If your user ID is karen and you are using workstation tty3, june's workstation displays:

```
Message from karen on trek tty3 Aug 17 11:55:24 ...
I need to see you! Meet me in the computer room at 12:30.
<EOT>
```

2. To hold a conversation, enter:

```
write june
```

Press the Enter key and type,

```
Meet me in the computer room at 12:30.
o
```

This starts the conversation. The o at the beginning of the next line means the message is over. It tells June that you are waiting for a response. Do not press Ctrl-D if you wish to continue.

Now June replies by typing:

```
write karen
```

Presses the Enter key and types,

```
I'm running tests at 12:30. Can we meet at 3?
o
```

And you might respond:

```
OK--the computer room at 3.
oo
```

The oo means *over and out*, telling June that you have nothing more to say. If June is also finished oo, then you both press Ctrl-D to end the conversation.

3. To write someone a prepared message, enter:

```
write june < message.text
```

This writes the contents of the **message.text** file to june's workstation.

4. To write to the person using a certain workstation, enter:

```
write -n console
```

Press the Enter key and type,

```
The printer in building 998 has jammed.
Please send help.
```

Then press the Ctrl-D key sequence.

This writes the message to the person logged in at the workstation /dev/console.

5. To send a message to user spuds at host partya, enter:

```
write -n partya spuds
```

Press the Enter key and type,

```
Your new tape has just arrived,
come see me to pick it up.
Thanks!
```

Then press the Ctrl-D key sequence.

OR

```
write spuds@party
```

Press the Enter key and type,

Your new tape has just arrived,  
come see me to pick it up.  
Thanks!

Then press the Ctrl-D key sequence.

6. Here is an example of a message sent by the **qdaemon** daemon:

```
Message from mary on trek (tty10) Aug 17 10:03:34 ...  
Use "write -h 6398492,reply" to reply  
Please insert tape number 5 into rmt0.  
<EOT>
```

To reply in the affirmative, enter:

```
write -h 6398492,ok
```

Then press the Ctrl-D key sequence.

To reply in the negative, enter:

```
write -h 6398492,cancel
```

Then press the Ctrl-D key sequence.

With the **-h** flag, there is no need to supply the host name or user ID. This information is tracked with the handle.

## Files

Item	Description
/etc/hosts	Contains TCP/IP host information.
/etc/utmp	Contains user and accounting information for the <b>who</b> , <b>write</b> , and <b>login</b> commands.

### Related information:

mesg command

“writesrv Daemon”

---

## writesrv Daemon

### Purpose

Allows users to send messages to and receive messages from a remote system.

### Syntax

```
writesrv
```

### Description

The **writesrv** daemon allows users to send messages to users on a remote system and receive responses from users on a remote system with the **write** command.

The **writesrv** utility receives incoming requests from a **write** command and creates a server process to handle the request. This server process communicates with the client process (**write**) and provides whatever services are requested.

To perform these services, the **writesrv** daemon creates a socket that is attached to the port defined in the **/etc/services** file. All requests for service are sent as messages to this socket.

**Note:** If the **writesrv** daemon terminates abnormally (such as a system crash, power failure, or the **kill -9** command), the **/var/spool/writesrv** directory must be manually cleaned out to remove any files left behind by the **writesrv** daemon.

## Examples

1. To start the **writesrv** daemon from the **/etc/rc** script, enter:

```
/usr/sbin/writesrv
```

The **writesrv** daemon is started from the **/etc/rc** script. This is the usual way the daemon is started.

2. To start the **writesrv** daemon using the System Resource Controller (SRC), enter:

```
startsrc -s writesrv &
```

The **writesrv** daemon is started using SRC.

## Files

Item	Description
<b>/etc/services</b>	Contains the Network Services directory.

### Related reference:

“write Command” on page 122

### Related information:

kill command

---

## wtmpfix Command

### Purpose

Manipulates connect-time accounting records by correcting date and time stamp inconsistencies.

### Syntax

```
/usr/sbin/acct/wtmpfix [ File ... ]
```

### Description

The **wtmpfix** command is called by the **runacct** procedure to examine standard input or *Files* that contain records in **wtmp** format, and correct problems that could make the **acctcon1** or **acctcon2** commands fail. The **wtmpfix** command corrects date and time stamp inconsistencies, and writes the corrected records to standard output. If the date and time stamps are not consistent when the **acctcon1** command runs, the **acctcon1** command generates an error and stops.

The **wtmpfix** command also checks the validity of the name field to ensure that it consists only of alphanumeric characters, a \$ (dollar sign), or spaces. If the name is invalid, the **wtmpfix** command changes the login name to **INVALID** and writes a diagnostic message to standard error. In this way, the **wtmpfix** command reduces the chance that the **acctcon2** command will fail.

Each time the date is set (on system startup or with the **date** command), a pair of date change records is written to the **/var/adm/wtmp** file. The first record is the old date, denoted by the *old time* string. The *old time* string is placed in the line field and the **OLD\_TIME** flag is placed in the type field. The second record is the new date, denoted by the string *new time*. The *new time* string is placed in the line field and the **NEW\_TIME** flag is placed in the type field. The **wtmpfix** command uses these records to synchronize all date and time stamps in the file.

## Flags

None.

## Parameters

Item	Description
<i>File</i>	Specifies the file to examine that contains records in <b>wtmp</b> format.

## Security

Access Control: These commands should grant execute (x) access only to members of the **adm** group.

## Examples

1. To correct date and time stamp inconsistencies in **/var/adm/wtmp** and write the corrected records to **dummy.file**, enter:

```
/usr/sbin/acct/wtmpfix /var/adm/wtmp > /tmp/dummy.file
```

## Files

Item	Description
<b>/usr/sbin/acct/wtmpfix</b>	Contains the <b>wtmpfix</b> command.
<b>/var/adm/wtmp</b>	Contains records of date changes that include an old date and a new date.
<b>/usr/include/utmp.h</b>	Contains history records that include a reason, date, and time.

### Related information:

acctcon1 or acctcon2 command

System accounting

---

## wump Command

### Purpose

Starts the hunt the wumpus game.

### Syntax

**wump**

### Description

A wumpus is a creature living in a cave with many rooms interconnected by tunnels. You move among the rooms trying to shoot the wumpus with an arrow and trying to avoid being eaten by the wumpus or falling into bottomless pits. There are also super bats that may pick you up and drop you in some randomly selected room. For moving among the rooms and shooting arrows, the game prompts you with appropriate questions and follows your instructions. For example:

```
You are in room 14.  
I feel a draft.  
There are tunnels to 1 13 18.  
Move or shoot? (m-s) m  
Which room? 1  
You are in room 1.  
I feel a draft.  
There are tunnels to 14 17 18.  
Move or shoot? (m-s) m
```



Which room? 17  
You are in room 17.  
You fell into a pit!  
Another game? (y-n)

In the above example, you start out in room 14. The computer displays I feel a draft. This is the hint that a pit is nearby. You choose to move to room 1. Again you are warned of the pit. You then choose to move to room 17 where you fall into a pit and die.

At the beginning of the game, you are prompted Instructions? (y-n). Choosing y provides an explanation of the warnings, how to move, and how to shoot.

The game ends and you are prompted Another game? (y-n) if:

- You kill the wumpus.
- The wumpus eats you.
- You fall into a bottomless pit.
- You run out of arrows.

To quit the game at any time, press the interrupt (Ctrl-C) key sequence.

## Files

Item	Description
<code>/usr/games</code>	Contains the location of the system's games.



---

## X

The following AIX commands begin with the with the letter *x*.

---

### X Command

#### Purpose

Starts the X Server.

#### Syntax

```
X [ -a Number ][ -auth FileName ][ -bc | +bc ][ -bp Color ][ -broadcast ][ -bs | -nobs ][  
-c Volume ][ -cc VisualType [ :Display ] ][ -class DisplayClass ][ -co File ][  
-cookie XDMAAuthenticationBit ][ -D File ][ -d Depth [ :Display ] ] -displayID DisplayID ][  
-damage | +damage ][ -evie | +evie ][ -f Number ][ -fc Font ][ -fixes | +fixes ][ -fn Font ][  
-fp Font ][ -help ][ -I ][ -indirect HostName ][ -layer # [ :Display ] ][ -logo | nologo ][  
-n :Number ][ -once ][ -P RowColumn Display } ][ -pbuffer level [:display name | :display number]  
][ -p Number ][ -port PortNumber ][ -query HostName ][ -r | r ][ +render | -render ][  
-s Number ][ -secIP [PermissionCode] ][ -secLocal [PermissionCode] ][ -secSMT [PermissionCode] ][ -stereo  
[:Display] ][ -su ][ -T ][ -t Number ][ -to Number ][ -v ][ -vfb ][ -wm ][ -wp Color ][  
-wrap | [ -wrapx ][ -wrapy ] ][ -x ExtensionName ][ -xkbdir Directory ][ -xkbmap FileName ][  
+|-]accessx ][ -ar1 Milliseconds ][ -ar2 Milliseconds ][ -sp FileName ][ +/- xinerama FileName ]
```

#### Description

The **X** command starts the X Server, a display server that runs on bitmapped terminals. The X Server distributes input and output requests to or from programs located on either the host system or systems connected to it through a network.

End an Enhanced X-Windows session by using the Ctrl+Alt+Backspace key sequence.

You can specify one or more display devices. If none are specified, the default is all. The default configuration order is determined by the adapter slot order. The adapter in the first slot is initialized as the left most screen, the adapter in the second slot is the next screen to the right. To rearrange the layout of the screens, use the **-P** flag. The **-P** flag associates the row and column of the device with the device name. You can determine the device name by using the **lsdisp** command.

The two displays are arranged either vertically or horizontally. The following example shows **-P** flags specifying a horizontal arrangement:

```
-P11 ppr0 -P12 ppr1
```

The 2 in the right position of the second **-P** flag indicates that the second monitor view is along the x-axis. This produces the horizontal arrangement:

```
Display          Display  
1                2
```

To see two monitors in a vertical arrangement, the **-P** flags should read:

```
-P11 ppr0 -P21 ppr1
```

The 2 in the first position indicates that the monitors are in a vertical configuration along the y-axis:

```
Display 1
Display 2
```

In the horizontal configuration, when a mouse is traveling from left to right in Display 1 and reaches the border of Display 1 and 2, the cursor continues into Display 2 at the same y-axis position. When it reaches the edge of Display 2 and the **-wrapx** flag is set, it appears at the leftmost edge of Display 1 in the same y-axis position. If the **-wrapx** flag is not set, the mouse stops at the far edge of Display 2.

In a vertical configuration, when the mouse is traveling from top to bottom in Display 1 and reaches the border of Display 1 and Display 2, the cursor continues into Display 2 at the same x-axis position. When the cursor reaches the bottom of the display 2 and the **-wrapy** flag is set, the cursor appears at the top edge of Display 1 in the same x-axis position. If the **-wrapy** flag is not set, the mouse stops at the bottom of Display 2.

In addition, information and error messages (for example, a message indicating that an extension not able to load) are listed in the **/tmp/xlogfile** file. This file can provide useful information in cases when the X Server encounters a problem. This file is re-written every time the X Server is instantiated. This file provides additional error and non-error information but is not a complete error log for the X Server.

When X-Server is started, it comes up using the default color class. Depending on the driver, the X-Server may default to using the PseudoColor or TrueColor class.

The PseudoColor class uses a colormap to display the colors on the screen. Many graphic adapters only support one hardware colormap. In this case, if the default color class is PseudoColor and an application is using a colormap that is different from the default colormap, incorrect colors may be displayed on the screen. Only the window that has focus will display the correct colors. It is advisable to run the X-Server in the TrueColor mode to prevent wrong colors from being displayed on the screen.

The **-cc X-Server** command flag can be used to bring the X-Server up using the TrueColor class. The **/usr/lpp/X11/defaults/xserverrc** file can be modified to allow this as shown in the following example.

As a root user, edit the **/usr/lpp/X11/defaults/xserverrc** file. Update the **EXTENSIONS=""** variable as shown in the following example:

```
#-----
# Start the X server in True Color mode
#-----
EXTENSIONS="$EXTENSIONS -cc 4"
```

Restart X Server by logging out of CDE and clicking reset.

**Note:** The **xdpyinfo** command can be used to verify the default color class.

## Flags

Item	Description										
<b>-a</b> <i>Number</i>	Specifies the acceleration multiplier for mouse movement. For example, a value of 5 causes the cursor to move five times as fast as the mouse. The default is 4 pixels; any value specified must be a positive value greater than 0.										
<b>-auth</b> <i>FileName</i>	Specifies to X the file from which to read the MIT (Massachusetts Institute of Technology) magic cookie.										
<b>-bc</b>	Turns off backward compatibility with Enhanced X-Windows version 1.1.										
<b>+bc</b>	Turns on backward compatibility with Enhanced X-Windows version 1.1. This is the default.										
<b>-bp</b> <i>Color</i>	Specifies a black pixel color for the display. The default is display dependent.										
<b>-bs</b>	Enables backing store support on all screens. Backing store support is disabled by default.										
<b>-c</b> <i>Volume</i>	Specifies key click volume.										
<b>-cc</b> <i>VisualType</i> [ <i>:Display</i> ]	Specifies the type of visual to use for the root window of the screen specified by the display name. Not all visual types are available on all adapters at all depths. The <i>:Display</i> parameter is optional, but useful when using the multihead option. The <i>:Display</i> parameter is the name of the display as shown in the <b>lsdisp</b> command. If no display number or name is supplied, the specified visual is selected for all screens.  To specify the visual type and depth for the default visual, use the <b>-cc</b> and <b>-d</b> flags, respectively.  Values for the <i>VisualType</i> parameter are specified as a string or a number as follows:  <b>String    Numeric equivalent</b> <b>StaticGray</b> 0 <b>GrayScale</b> 1 <b>StaticColor</b> 2 <b>PseudoColor</b> 3 <b>TrueColor</b> 4 <b>DirectColor</b> 5										
<b>-co</b> <i>File</i>	Sets the name of the red, green, and blue (RGB) color database. This is the default flag for the color database.										
<b>-D</b> <i>File</i>	Specifies the full path name of the color definition database file. The default is <b>/usr/lib/X11/rgb</b> .										
<b>-d</b> <i>Depth</i> [ <i>:Display</i> ]	Specifies the root depth for the screen specified by the display name. Not all visual types will be available on all adapters at all depths.  The <i>:Display</i> parameter is optional, but useful when using the multihead option and must correspond to the values passed with the <b>-P</b> flag. The <i>:Display</i> parameter is the name of the display as shown in the <b>lsdisp</b> command. In the absence of the <i>:Display</i> parameter, the specified depth is selected for all the selected displays in the multihead option, as specified in the <b>-P</b> flag.										
<b>-damage</b>	Disables the X Damage extension.										
<b>+damage</b>	Enables the X Damage extension.										
<b>-evie</b>	Disables the X Event Interception extension.										
<b>+evie</b>	Enables the X Event Interception extension.										
<b>-f</b> <i>Number</i>	Specifies the beep volume. The default is -1 or medium. The supported values are as follows:  <table border="0" style="margin-left: 20px;"> <tr> <td><b>Value</b></td> <td><b>Setting</b></td> </tr> <tr> <td><b>0</b></td> <td>Off</td> </tr> <tr> <td><b>1-33</b></td> <td>Low</td> </tr> <tr> <td><b>-1 or 34-66</b></td> <td>Medium</td> </tr> <tr> <td><b>67-100</b></td> <td>High</td> </tr> </table>	<b>Value</b>	<b>Setting</b>	<b>0</b>	Off	<b>1-33</b>	Low	<b>-1 or 34-66</b>	Medium	<b>67-100</b>	High
<b>Value</b>	<b>Setting</b>										
<b>0</b>	Off										
<b>1-33</b>	Low										
<b>-1 or 34-66</b>	Medium										
<b>67-100</b>	High										
<b>-fc</b> <i>Font</i>	Specifies the cursor font for cursor glyphs and cursor masks. The default depends on the operating system and the display.										
<b>-fixes</b>	Disables the X Fixes extension.										
<b>+fixes</b>	Enables the X Fixes extension.										

Item	Description
<b>-fn</b> <i>Font</i>	Specifies the default text font. The default depends on the operating system and the display.
<b>-fp</b> <i>Font</i>	Specifies the font path.
<b>-I</b>	Causes all remaining command line arguments to be ignored. (Uppercase i)
<b>-help</b>	Prints a usage message.
<b>-layer</b> #[: <i>Display</i> ]	Specifies that the default visual should be in the # layer. The <i>:Display</i> parameter is the name of the display as shown in the <b>lsdisp</b> command. Specifying this flag for an adapter that does not have overlays, or has less than 8 bits of overlay, has no effect. Specifying this flag with a # higher than the number of supported layers results in the default visual residing in the default layer of the screen (as if no <b>-layer</b> flag had been used).
<b>-logo</b>	Turns on the X Window System logo display in the screen saver. There is currently no way to change this from a client.
<b>-n</b> <i>:Number</i>	Specifies the connection number. Valid values for the <i>Number</i> parameter are 0 to 255. The default is the next available number. The <i>Number</i> parameter is used by programs to communicate with a specific X Server. For example, the command:  X -n :18  specifies that communication to the activated X Server takes place by unix:18 or by <i>Hostname</i> :18.
<b>-nobs</b>	Disables backing store support on all screens. This is the default.
<b>nologo</b>	Turns off the X Window System logo display in the screen saver. There is currently no way to change this from a client.
<b>-once</b>	Instructs the server to exit after the first session ends. Normally, the server starts sessions automatically.
<b>-P</b> <i>RowColumn Display</i>	Specifies the physical positioning of the displays in a multihead configuration. The <i>Row</i> parameter indicates the row in which the display is located. The <i>Column</i> parameter indicates the column in which the display is located.  The <i>Display</i> parameter is the device name of the display as shown in the first column of output from the <b>lsdisp</b> command. The first <b>-P</b> <i>RowColumn Display</i> occurrence on the command line describes screen 0 to the X Server, the second describes screen 1, and so on.
<b>-pbuffer</b> <i>level</i> [ <i>:display name</i>   <i>:display number</i> ]	The <b>-P</b> flag is for use with multiple head support. Specifies the <b>pbuffer</b> memory allocation level for the screen specified by <i>:display</i> . This flag is only useful when used in conjunction with the GLX extension.  The <i>level</i> parameter indicates the relative amount of frame buffer memory to be reserved for pbuffers. Specified values must be in the range of [0..2]. A value of 0 indicates that no memory should be reserved for pbuffers. A value of 1 indicates that a low amount of memory should be reserved. A value of 2 indicates that a high amount of memory should be reserved. Not all adapters support pbuffers. For those that do, not all screen configurations support pbuffers. The actual amount of frame buffer memory reserved for pbuffers is device-dependent, and may be influenced by other factors, such as screen resolution or default pixel depth.  The <i>:display</i> parameter is optional, but useful when using the multihead option. The <i>:display</i> parameter is the name of the display as shown in the <b>lsdisp</b> command. If no <i>display number</i> or <i>name</i> is supplied, the specified <b>pbuffer</b> width is selected for all screens.
<b>-p</b> <i>Number</i>	Specifies the time interval, in minutes, between changes of the X Window System logo position. This flag is used with the <b>-s</b> (screen saver timeout) flag to control the blanking of the screen.
<b>-r</b>	Disables autorepeat. The default is autorepeat enabled.

Item	Description
<b>r</b>	Turns on autorepeat.
<b>+render</b>	Enables the X Render extension. By default, the X Render extension is disabled. <b>Note:</b> X Render Extension can be activated on the X Server only with GXT135P, GXT145, GXT4500P, and GXT6500P graphics adapters. To check the available adapter on the system, run the <b>lsdisp</b> command. Use the <b>-vfb</b> flag with the virtual frame buffer along with the <b>+render</b> flag.
<b>-render</b>	Disables the X Render extension.
<b>-s</b> <i>Number</i>	Specifies the number of minutes to wait before blanking the screen. The default is 10 minutes. If this value is set to 0, the screen-saver is disabled.
<b>-secIP</b> [ <i>PermissionCode</i> ]	Sets local access control on the internet socket. The <i>PermissionCode</i> is 3 octal digits which can set read, write, and execute bits. If no <i>PermissionCode</i> is specified after a security flag, then permission is defaulted to 0 for that socket.
<b>-secLocal</b> [ <i>PermissionCode</i> ]	Sets access control on the unix socket. The <i>PermissionCode</i> is 3 octal digits which can set read, write, and execute bits. If no <i>PermissionCode</i> is specified after a security flag, then permission is defaulted to 0 for that socket.
<b>-secSMT</b> [ <i>PermissionCode</i> ]	Sets access control on the shared memory transport socket. The <i>PermissionCode</i> is 3 octal digits which can set read, write, and execute bits. If no <i>PermissionCode</i> is specified after a security flag, then permission is defaulted to 0 for that socket.
<b>-stereo</b> [ <i>:Display</i> ]	Configures the graphics adapter for optimum stereo support for the screen specified by <i>Display</i> .  Supported screens will configure the adapter to provide the best available support for stereo. This may decrease other resources such as texture memory. The actual amount of memory affected is device-dependent, and may be influenced by other factors, such as screen resolution or default pixel depth.  The <i>Display</i> parameter is optional, but useful when using the multihead option. The <i>Display</i> parameter is the name of the display as shown in the <b>lsdisp</b> command. If no display number or name is supplied, the <b>-stereo</b> flag pertains to all supported screens.  Unsupported screens will ignore the <b>-stereo</b> flag.
<b>-su</b>	Disables save under support on all screens.
<b>-T</b>	Disables the Ctrl+Alt+Backspace key sequence that, by default, ends the AIXwindows session and all windows opened from it.
<b>-t</b> <i>Number</i>	Specifies the mouse threshold. The default is 2 pixels. Acceleration takes effect only if the mouse is moved beyond the mouse threshold in one time interval and only applies to the amount beyond the threshold.
<b>-to</b> <i>Number</i>	Specifies the number of minutes to elapse between connection checks. The default is 60 minutes. A specified value must be greater than 0.
<b>-v</b>	Specifies that the display be replaced with the current background color after the time specified by the <b>-s</b> flag expires. By default, if the <b>-v</b> flag is not used, the entire display is painted with the background tile after the time specified by the <b>-s</b> flag expires.
<b>-vfb</b>	Starts the X Server with Virtual Frame Buffer (VFB), without initializing any graphics adaptor.
<b>-wm</b>	Forces the default backing store of all windows to have the <b>WhenMapped</b> value. This is a convenient way of applying backing store to all windows.
<b>-wp</b> <i>Color</i>	Specifies a white pixel display color. The default depends on the display.

Item	Description
<b>-wrap</b>	<p>Specifies the behavior of the mouse when its hotspot reaches the left or right border or the top or bottom of any root window. If this flag is set and the hotspot of the mouse reaches the left border of the leftmost root window, the mouse is automatically positioned at the right border of the rightmost root window at the same y position.</p> <p>Conversely, if this flag is set and the hotspot of the mouse reaches the right border of the rightmost root window, the mouse is automatically positioned at the left border of the leftmost root window at the same y position. If this flag is not set, the mouse stops at the left or right border of any root window.</p> <p>If this flag is set and the hotspot of the mouse reaches the top border of the topmost root window, the mouse is positioned at the bottom border of the bottommost root window at the same x position.</p> <p>Conversely, if this flag is set and the hotspot of the mouse reaches the bottom border of the bottommost root window, the mouse is positioned at the top border of the topmost root window at the same x position.</p>
<b>-wrapx</b>	<p>The <b>-wrap</b> flag is for use with multiple head support.</p> <p>Specifies the behavior of the mouse when its hotspot reaches the left or right border of any root window. If this flag is set and the hotspot of the mouse reaches the left border of the leftmost root window, the mouse is positioned at the right border of the rightmost root window at the same y position. Conversely, if this flag is set and the hotspot of the mouse reaches the right border of the rightmost root window, the mouse is positioned at the left border of the leftmost root window at the same y position. If this flag is not set, the mouse stops at the left or right border of any root window.</p>
<b>-wrapy</b>	<p>The <b>-wrapx</b> flag is for use with multiple head support.</p> <p>Specifies the behavior of the mouse when its hotspot reaches the top or bottom border of any root window. If this flag is set and the hotspot of the mouse reaches the top border of the topmost root window, the mouse is positioned at the bottom border of the bottommost root window at the same x position.</p> <p>Conversely, if this flag is set and the hotspot of the mouse reaches the bottom border of the bottommost root window, the mouse is positioned at the top border of the topmost root window at the same x position. If this flag is not set, the mouse stops at the top or bottom border of any root window.</p>
<b>-x <i>ExtensionName</i></b>	<p>The <b>-wrapy</b> flag is for use with multiple head support.</p> <p>Specifies that the extension name should be loaded when the server is initialized. This is particularly useful for large extensions, such as the Display PostScript Level 2 (<b>dps</b>). This flag can be specified more than once with multiple extension names.</p>
<b>-query <i>HostName</i></b>	<p>Enables Enhanced X-Windows Display Manager Control Protocol (<b>XDMCP</b>) and sends a <b>Query</b> packet to the specified host.</p>
<b>-broadcast</b>	<p>The <b>-query</b> flag is for use with <b>XDMCP</b>.</p> <p>Enables <b>XDMCP</b> and broadcasts <b>BroadcastQuery</b> packets to the network. The first responding display manager is chosen for the session.</p>
<b>-indirect <i>HostName</i></b>	<p>The <b>-broadcast</b> flag is for use with <b>XDMCP</b>.</p> <p>Enables <b>XDMCP</b> and sends <b>IndirectQuery</b> packets to the specified host.</p>
	<p>The <b>-indirect</b> flag is for use with <b>XDMCP</b>.</p>



**Item****-port** *PortNumber***Description**

Specifies an alternative port number for XDMCP. This flag must be specified before any **-query**, **-broadcast**, or **-indirect** flags. Normally, the server starts sessions one after another. This flag causes the server to exit after the first session ends.

**-class** *DisplayClass*

The **-port** flag is for use with XDMCP.

Sets the value for an additional display qualifier used by XDMCP in resource lookup for display-specific options.

**-cookie** *XDMAAuthenticationBits*

The **-class** flag is for use with XDMCP.

Specifies a private key to be shared between the server and the manager when testing XDM-AUTHENTICATION-1.

**-displayID** *DisplayID*

The **-cookie** flag is for use with XDMCP.

Allows the display manager to identify each display so that it can locate the shared key specified by the **-cookie** flag.

**+/- xinerama**

The **-displayID** flag is for use with XDMCP.

Enable/Disable panoramic screen or Virtual Large Screen (VLS). Allows users to treat all heads in a multihead environment as a large screen.

**Xkeyboard Flags****Item****-xkbdir** *Directory***Description**

Specifies the base directory for the keyboard layout files.

**-xkbmap** *FileName*

Specifies the keyboard description to load on startup.

**[+|-]accessx**

Enables (+) or disables (-) AccessX key sequences.

**-ar1** *Milliseconds*

Sets the length of time in milliseconds that a key must be pressed before autorepeat starts.

**-ar2** *Milliseconds*

Sets the length of time in milliseconds that should elapse between autorepeat generated keystrokes.

**Security Extension Flags**

**Item****-sp** *FileName***Description**

Causes the server to attempt to read and interpret *FileName* as a security policy file with the format described below. The file is read at server startup and reread at each server reset.

The syntax of the security policy file is as follows. Notation: "\*" means zero or more occurrences of the preceding element, and "+" means one or more occurrences. To interpret *foo|bar*, ignore the text after the *|*; it is used to distinguish between instances of *foo* in the next section.

```

policy file ::= version line other line*
version line ::= string/v '\n'
other line ::= comment | access rule | site policy | blank line
comment ::= # not newline* '\n'
blank line ::= space '\n'
site policy ::= sitepolicy string/sp '\n'
access rule ::= property property/ar window perms '\n'
property ::= string
window ::= any | root | required property
required property ::= property/rp | property with value
property with value ::= property/rpv = string/rv
perms ::= [ operation | action | space ]*
operation ::= r | w | d
action ::= a | i | e
string ::= dbl quoted string | single quoted string | unquoted string
dbl quoted string ::= space " not dqoute* " space
single quoted string ::= space ' not squote* ' space
unquoted string ::= space not space+ space
space ::= [ ' | '\t' ]*

```

Character sets:

```

not newline ::= any character except '\n'
not dqoute ::= any character except "
not squote ::= any character except '
not space ::= any character except those in space

```

The semantics associated with the previously described syntax are as follows.

*version line*

The first line in the file, specifies the file format version. If the server does not recognize the version *string/v*, it ignores the rest of the file. The version string for the file format described here is *version-1*.

Once past the *version line*, lines that do not match the above syntax are ignored.

*comment* Lines are ignored.

*sitepolicy* Lines are currently ignored. They are intended to specify the site policies used by the XC-QUERY-SECURITY-1 authorization method.

*access rule*

Lines specify how the server should react to untrusted client requests that affect the X Window property named *property/ar*. The rest of this section describes the interpretation of an *access rule*.

For an *access rule* to apply to a given instance of *property/ar*, *property/ar* must be on a window that is in the set of windows specified by *window*. If *window* is **any**, the rule applies to *property/ar* on any window. If *window* is **root**, the rule applies to *property/ar* only on root windows.

If *window* is *required property*, the following apply. If *required property* is a *property/rp*, the rule applies when the window also has that *property/rp*, regardless of its value. If *required property* is a *property with value*, *property/rpv* must also have the value specified by *string/rv*. In this case, the property must have type STRING and format 8, and should contain one or more null-terminated strings. If any of the strings match *string/rv*, the rule applies.

The definition of string matching is simple case-sensitive string comparison with one elaboration: the occurrence of the character '\*' in *string/rv* is a wildcard meaning "any string." A *string/rv* can contain multiple wildcards anywhere in the string. For example, *x\** matches strings that begin with *x*, *\*x* matches strings that end with *x*, *\*x\** matches strings containing *x*, and *x\*y\** matches strings that start with *x* and subsequently contain *y*.

There may be multiple *access rule* lines for a given *property/ar*. The rules are tested in the order that they appear in the file. The first rule that applies is used.

Item	Description
<i>perms</i>	Specify operations that untrusted clients may attempt, and the actions that the server should take in response to those operations.
<i>operation</i>	Can be <b>r</b> (read), <b>w</b> (write), or <b>d</b> (delete). The following information shows how X Protocol property requests map to these operations in the X Consortium server implementation. <p><b>GetProperty</b>  <b>r</b>, or <b>r</b> and <b>d</b> if delete = True</p> <p><b>ChangeProperty</b>  <b>w</b></p> <p><b>RotateProperties</b>  <b>r</b> and <b>w</b></p> <p><b>DeleteProperty</b>  <b>d</b></p> <p><b>ListProperties</b>  <b>none</b>, untrusted clients can always list all properties</p>
<i>action</i>	Can be <b>a</b> (allow), <b>i</b> (ignore), or <b>e</b> (error). <p><b>Allow</b> Executes the request as if it had been issued by a trusted client.</p> <p><b>Ignore</b> Treats the request as a no-op. In the case of GetProperty, ignore means return an empty property value if the property exists, regardless of its actual value.</p> <p><b>Error</b> Specifies not to execute the request and return a BadAtom error with the atom set to the property name. Error is the default action for all properties, including those not listed in the security policy file.</p> <p>An <i>action</i> applies to all <i>operations</i> that follow it, until the next <i>action</i>&gt; is encountered. Thus, irwad means ignore read and write, allow delete.</p> <p>GetProperty and RotateProperties might do multiple operations (<b>r</b> and <b>d</b>, or <b>r</b> and <b>w</b>). If different actions apply to the operations, the most severe action is applied to the whole request; there is no partial request execution. The severity ordering is: allow &lt; ignore &lt; error. Thus, if the <i>perms</i> for a property are <i>ired</i> (ignore read, error delete), and an untrusted client attempts GetProperty on that property with delete = True, an error is returned, but the property value is not. Similarly, if any of the properties in a RotateProperties do not allow both read and write, an error is returned without changing any property values.</p>

Item	Description
------	-------------

An example a security policy file follows:

```
version-1

# Allow reading of application resources, but not writing.
property RESOURCE_MANAGER    root    ar iw
property SCREEN_RESOURCES    root    ar iw

# Ignore attempts to use cut buffers. Giving errors causes apps to crash,
# and allowing access may give away too much information.
property CUT_BUFFER0         root    irw
property CUT_BUFFER1         root    irw
property CUT_BUFFER2         root    irw
property CUT_BUFFER3         root    irw
property CUT_BUFFER4         root    irw
property CUT_BUFFER5         root    irw
property CUT_BUFFER6         root    irw
property CUT_BUFFER7         root    irw

# If you are using Motif, you probably want these.

property _MOTIF_DEFAULT_BINDINGS    rootar iw
property _MOTIF_DRAG_WINDOW    root    ar iw
property _MOTIF_DRAG_TARGETS    any    ar iw
property _MOTIF_DRAG_ATOMS    any    ar iw
property _MOTIF_DRAG_ATOM_PAIRS    any ar iw

# The next two rules let xwininfo -tree work when untrusted.
property WM_NAME                any    ar

# Allow read of WM_CLASS, but only for windows with WM_NAME.
# This might be more restrictive than necessary, but demonstrates
# the required property facility, and is also an attempt to
# say "top level windows only."
property WM_CLASS                WM_NAME ar

# These next three let xlsclients work untrusted. Think carefully
# before including these; giving away the client machine name and command
# may be exposing too much.
property WM_STATE                WM_NAME ar
property WM_CLIENT_MACHINE        WM_NAME ar
property WM_COMMAND                WM_NAME ar

# To let untrusted clients use the standard colormaps created by
# xstdcmap, include these lines.
property RGB_DEFAULT_MAP        root    ar
property RGB_BEST_MAP            root    ar
property RGB_RED_MAP            root    ar
property RGB_GREEN_MAP          root    ar
property RGB_BLUE_MAP           root    ar
property RGB_GRAY_MAP           root    ar
```

Item	Description
------	-------------

```
# To let untrusted clients use the color management database created
# by xcmsdb, include these lines.
property XDCCC_LINEAR_RGB_CORRECTION    rootar
property XDCCC_LINEAR_RGB_MATRICES     rootar
property XDCCC_GRAY_SCREENWHITEPOINT   rootar
property XDCCC_GRAY_CORRECTION         rootar

# oddball property names and explicit specification of error conditions
property "property with spaces"        'property with "aw er ed

# Allow deletion of Woo-Hoo if window also has property OhBoy with value
# ending in "son". Reads and writes will cause an error.
property Woo-Hoo                        OhBoy = "*son"ad
```

## Example

To start the X Server with X Render extension, enter the following command:

```
$X -T -force :0 -vfb -d 32 +render
```

In this example, the X Server will use the Virtual Frame Buffer (VFB) for rendering instead of using the physical graphics adaptor.

### Related information:

lsdisp command  
aixterm command

---

## x\_add\_fs\_fpe Command

### Purpose

Adds a network font server to a font path.

### Syntax

```
x_add_fs_fpe Host Port Position TypeName
```

### Description

The `x_add_fs_fpe` command adds a font path element to the font path of the selected network type name for a font server to access fonts.

Item	Description
<i>Host</i>	Specifies the name of the system where the font server resides.
<i>Port</i>	Specifies the number of the font server port. This number must be in the <code>/etc/services</code> file and specified in decimal.
<i>Position</i>	Specifies where to insert this element in the font path.
<i>TypeName</i>	Specifies the name of the network type. Each network type has a font path consisting of one or more font path elements. Specify the name of the network type to which the font path element will be added, or choose to have it added to all network type names by specifying <code>All</code> . If a font path element is added to <code>All</code> network types, will be placed at the end of each font path.

### Security

Access Control: Only the root user should have execute (x) access to this command.

## Example

To add the font server to the start of the font path for network type `x_st_mgr.ether`, enter:

```
x_add_fs_fpe winter 7500 1 x_st_mgr.ether
```

In this example, the font server on host `winter` has been added to the start of the font path for network type `x_st_mgr.ether`. The font server port is `7500`.

## Files

Item	Description
<code>/usr/lpp/x_st_mgr/bin/x_add_fs_fpe</code>	Contains the <code>x_add_fs_fpe</code> command.
<code>/etc/x_st_mgr/ether.cf</code>	Contains the network type <code>x_st_mgr.ether</code> configuration file (sample).

### Related reference:

“`x_add_nfs_fpe` Command”

“`x_rm_fpe` Command” on page 143

---

## x\_add\_nfs\_fpe Command

### Purpose

Adds a NFS/TFTP accessed font directory to a font path.

### Syntax

```
x_add_nfs_fpe Host Directory Method Position TypeName
```

### Description

The `x_add_nfs_fpe` command adds a font path element to the font path of the selected network type name. This font directory will be accessed using Network File System (NFS) or Trivial File Transfer Protocol (TFTP).

Item	Description
<i>Host</i>	Specifies the system name to access for the font directory.
<i>Directory</i>	Specifies the complete path to the directory that contains the fonts.
<i>Method</i>	Specifies either <code>nfs</code> or <code>tftp</code> to be used to access the fonts.
<i>Position</i>	Specifies where to insert this element in the font path.
<i>TypeName</i>	Specifies the name of the network type. Each network type has a font path consisting of one or more font path elements. Specify the name of the network type to which the font path element will be added, or choose to have it added to all network type names by specifying <code>All</code> . If a font path element is added to <code>All</code> network types, it will be placed at the end of each font path.

### Security

Access Control: Only the root user should have execute (x) access to this command.

### Example

To add the fonts in `/usr/lib/X11/fonts/100dpi` to the network type `x_st_mgr.ether`, enter:

```
x_add_nfs_fpe cedar /usr/lib/X11/fonts/100dpi nfs Last \ x_st_mgr.ether
```

In this the font path element `/usr/lib/X11/fonts/100dpi` is added to the end of the font path for network type `x_st_mgr.ether`. The font directory is on the host `cedar`, which is accessed using NFS.

## Files

Item	Description
<code>/usr/lpp/x_st_mgr/bin/x_add_nfs_fpe</code>	Contains the <code>x_add_nfs_fpe</code> command.
<code>/etc/x_st_mgr/ether.cf</code>	Contains the network type <code>x_st_mgr.ether</code> configuration file (sample).

### Related reference:

“`x_add_fs_fpe` Command” on page 141

“`x_rm_fpe` Command”

---

## x\_rm\_fpe Command

### Purpose

Removes a font path element from a font path.

### Syntax

`x_rm_fpe` *TypeName* *Position* *Method* *Host* *Port* *Directory*

### Description

The `x_rm_fpe` command removes a font path element from the font path of the selected network type name.

Item	Description
<i>TypeName</i>	Specifies from which network type name the element is to be removed.
<i>Position</i>	Specifies where the element is in the font path.
<i>Method</i>	Specifies the method used to access the font path element. The valid options are: <b>tcp</b> for Network Font Server; <b>default</b> for initial default font path element; <b>nfs</b> for NFS; and <b>tftp</b> for TFTP.
<i>Host</i>	Specifies the name of the system specified in the font path element. For elements using the default method, specify <b>None</b> .
<i>Port</i>	Specifies the number of the server port specified in the font path element. For elements using the <b>nfs</b> or <b>tftp</b> method, specify <b>None</b> .
<i>Directory</i>	Specifies the complete path to the directory that contains the fonts. For a Network Font Server element, specify <b>None</b> .

### Security

Access Control: Only the root user should have execute (x) access to this command.

### Examples

To remove the font element `/usr/lib/X11/fonts/100dpi` from the font path for network type `x_st_mgr.ether`, enter:

```
x_rm_fpe x_st_mgr.ether 3 nfs waco None /usr/lib/X11/fonts/100dpi
```

In this example, the font path element `/usr/lib/X11/fonts/100dpi` that is accessed on host `waco` using NFS has been removed from the third position of the font path for network type `x_st_mgr.ether`. Because a port number is not used for NFS, this parameter was set to `None`.

## Files

Item	Description
/usr/lpp/x_st_mgr/bin/x_rm_fpe	Contains the <b>x_rm_fpe</b> command.
/etc/x_st_mgr/ether.cf	Contains the network type <b>x_st_mgr.ether</b> configuration file (sample).

### Related reference:

“x\_add\_nfs\_fpe Command” on page 142

“x\_rm\_fpe Command” on page 143

---

## xargs Command

### Purpose

Constructs parameter lists and runs commands.

### Syntax

```
xargs [-p t x ] [-e [EOFString]] [-E EOFString] [-i [ReplaceString]] [-I ReplaceString | -L Number | -n Number] [-I [Number]] [-s Size] [Command [Argument ... ]]
```

**Note:** Do not put a blank space between the lowercase flags and the parameter.

### Description

The generated command line length is the sum of the size, in bytes, of the *Command* and each *Argument* treated as strings, including a null byte terminator for each of these strings. The **xargs** command limits the command line length. When the constructed command line runs, the combined *Argument* and environment lists can not exceed **ARG\_MAX** bytes. Within this constraint, if you do not specify the **-n** or the **-s** flags, the default command line length is at least the value specified by **LINE\_MAX**.

### Flags

Item	Description
<b>-e</b> [ <i>EOFString</i> ]	Obsolete flag. Use the <b>-E</b> flag.  Uses the <i>EOFString</i> parameter as the logical EOF string. If you do not specify the <b>-e</b> or the <b>-E</b> flags, underscore ( <code>_</code> ) is assumed for the logical EOF string. If you do not specify the <i>EOFString</i> parameter, the logical EOF string capability is disabled, and underscores are taken literally. The <b>xargs</b> command reads from standard input until either EOF or the specified string is reached.
<b>-E</b> <i>EOFString</i>	Specifies a logical EOF string to replace the default underscore ( <code>_</code> ). The <b>xargs</b> command reads standard input until either EOF or the specified string is reached.
<b>-i</b> [ <i>ReplaceString</i> ]	Obsolete flag. Use the <b>-I</b> (uppercase i) flag.  If you do not specify the <i>ReplaceString</i> parameter, the string <code>{}</code> is used. <b>Note:</b> The <b>-I</b> (uppercase i), <b>i</b> , <b>-L</b> (uppercase l), <b>l</b> , and <b>-n</b> flags are mutually exclusive. The last flag specified takes effect.
<b>-I</b> <i>ReplaceString</i>	(Uppercase i). Inserts each line of standard input as an argument for the <i>Command</i> parameter, inserting it in <i>Argument</i> for each occurrence of <i>ReplaceString</i> . <i>ReplaceStrings</i> can not be used in more than 5 arguments. Blank characters at the beginning of each standard input line are ignored. Each <i>Argument</i> can contain one or more <i>ReplaceStrings</i> , but may not be larger than 255 bytes. The <b>-I</b> flag also turns on the <b>-x</b> flag.  The <b>-I</b> (uppercase i) flag means <b>-L1</b> . Therefore, only one standard input line can be substituted as an argument at a time. If the replaced string appears more than once in the command parameter, the same standard input line is substituted for each occurrence of the replaced string. <b>Note:</b> The <b>-I</b> (uppercase i), <b>i</b> , <b>-L</b> (uppercase l), <b>l</b> , and <b>-n</b> flags are mutually exclusive. The last flag specified takes effect.



Item	Description
-l [ <i>Number</i> ]	<p>(Lowercase L). Obsolete flag. Use the -L flag.</p> <p>If you do not specify the <i>Number</i> parameter, a value of 1 is used. The -l flag also turns on the -x flag.</p> <p><b>Note:</b> The -I (uppercase i), i, -L (uppercase l), -l, and -n flags are mutually exclusive. The last flag specified takes effect.</p>
-L <i>Number</i>	<p>Runs the <i>Command</i> parameter with the specified number of nonempty parameter lines read from standard input. The last invocation of the <i>Command</i> parameter can have fewer parameter lines if fewer than the specified <i>Number</i> remain. A line ends with the first new-line character unless the last character of the line is a space or a tab. A trailing space indicates a continuation through the next nonempty line.</p> <p>Reads the <i>Number</i> lines from the standard input and places them at the end of the command line.</p> <p><b>Note:</b> The -I (uppercase i), i, -L (uppercase l), -l, and -n flags are mutually exclusive. The last flag specified takes effect.</p>
-n <i>Number</i>	<p>Runs the <i>Command</i> parameter using as many standard input arguments as possible, up to the maximum specified by the <i>Number</i> parameter.</p> <p>Reads the maximum of <i>Number</i> arguments from the standard input and places them at the end of the command line. The <b>xargs</b> command uses fewer arguments if:</p> <ul style="list-style-type: none"> <li>• The accumulated command line length exceeds the bytes specified by the -s <i>Size</i> flag.</li> <li>• The last iteration has fewer value than the value specified by the <i>Number</i> argument, but not zero, arguments remaining.</li> </ul> <p><b>Note:</b> The -I (uppercase i), i, -L (uppercase l), -l, and -n flags are mutually exclusive. The last flag specified takes effect.</p>
-p	<p>Asks whether to run the <i>Command</i> parameter. It displays the constructed command line, followed by a ? . . . (question mark, ellipsis) prompt. Enter an affirmative response specific to the locale to run the <i>Command</i> parameter. Any other response causes the <b>xargs</b> command to skip that particular invocation of the parameter. You are asked about each invocation. The -p flag also turns on the -t flag.</p>
-s <i>Size</i>	<p>Sets the maximum total size of the constructed <i>Command</i> line. The <i>Size</i> parameter must be a positive integer. Fewer arguments are used if:</p> <ol style="list-style-type: none"> <li>1. The total number of arguments exceeds those specified by the -n flag.</li> <li>2. The total number of lines exceeds those specified by the -L or -l (Lowercase L) flags.</li> <li>3. EOF is reached before the number of bytes specified by the <i>Size</i> parameter are accumulated.</li> </ol>
-t	<p>Enables the trace mode and echoes the constructed <i>Command</i> line to standard error before running.</p>
-x	<p>Stops running the <b>xargs</b> command if any <i>Command</i> line is greater than the number of bytes specified by the -s <i>Size</i> flag. This -x flag is turned on if you specify either the -I (Uppercase i) or -l (Lowercase L) flag. If you do not specify the -i, -I (Uppercase i), -l (Lowercase L), -L, or -n flag, the total length of the <i>Command</i> line must be within the limit specified by the -s <i>Size</i> flag.</p>

## Exit Status

This command returns the following exit values:

Item	Description
0	All invocations of the <i>Command</i> parameter returned exit status 0.
1-125	A command line meeting the specified requirements could not be assembled, one or more of the invocations of the <i>Command</i> parameter returned a non-zero exit status, or some other error occurred.
126	<i>Command</i> was found but could not be invoked.
127	<i>Command</i> could not be found.

If a command line meeting the specified requirements cannot be assembled, the command cannot be invoked, an invocation of the command is terminated by a signal, or an invocation of the command exits

with exit status 255. The **xargs** command will write a diagnostic message and exit without processing any remaining input.

## Examples

1. To use a command on files whose names are listed in a file, type:

```
xargs lint -a <cfiles
```

If the `cfiles` file contains the following text:

```
main.c readit.c
gettoken.c
putobj.c
```

the **xargs** command constructs and runs the following command:

```
lint -a main.c readit.c gettoken.c putobj.c
```

If the `cfiles` file contains more file names than fit on a single shell command line (up to **LINE\_MAX**), the **xargs** command runs the **lint** command with the file names that fit. It then constructs and runs another **lint** command using the remaining file names. Depending on the names listed in the `cfiles` file, the commands might look like the following:

```
lint -a main.c readit.c gettoken.c . . .
lint -a getisx.c getprp.c getpid.c . . .
lint -a fltadd.c fltmult.c fltdiv.c . . .
```

This command sequence is not quite the same as running the **lint** command once with all the file names. The **lint** command checks cross-references between files. However, in this example, it cannot check between the `main.c` and the `fltadd.c` files, or between any two files listed on separate command lines.

For this reason you may want to run the command only if all the file names fit on one line. To specify this to the **xargs** command use the **-x** flag by typing:

```
xargs -x lint -a <cfiles
```

If all the file names in the `cfiles` file do not fit on one command line, the **xargs** command displays an error message.

2. To construct commands that contain a certain number of file names, type:

```
xargs -t -n 2 diff <<EOF
starting chap1 concepts chap2 writing
chap3
EOF
```

This command sequence constructs and runs **diff** commands that contain two file names each (**-n 2**):

```
diff starting chap1
diff concepts chap2
diff writing chap3
```

The **-t** flag causes the **xargs** command to display each command before running it, so you can see what is happening. The **<<EOF** and **EOF** pattern-matching characters define a *here document*, which uses the text entered before the end line as standard input for the **xargs** command.

3. To insert file names into the middle of command lines, type:

```
ls | xargs -t -I {} mv {} {}.old
```

This command sequence renames all files in the current directory by adding `.old` to the end of each name. The `-I` flag tells the `xargs` command to insert each line of the `ls` directory listing where `{}` (braces) appear. If the current directory contains the files `chap1`, `chap2`, and `chap3`, this constructs the following commands:

```
mv chap1 chap1.old
mv chap2 chap2.old
mv chap3 chap3.old
```

4. To run a command on files that you select individually, type:

```
ls | xargs -p -n 1 ar r lib.a
```

This command sequence allows you to select files to add to the `lib.a` library. The `-p` flag tells the `xargs` command to display each `ar` command it constructs and to ask if you want to run it. Type `y` to run the command. Press the any other key if you do not want to run the command.

Something similar to the following displays:

```
ar r lib.a chap1 ?...
ar r lib.a chap2 ?...
ar r lib.a chap3 ?...
```

5. To construct a command that contains a specific number of arguments and to insert those arguments into the middle of a command line, type:

```
ls | xargs -n6 | xargs -I{} echo {} - some files in the directory
```

If the current directory contains files `chap1` through `chap10`, the output constructed will be the following:

```
chap1 chap2 chap3 chap4 chap5 chap6 - some files in the directory
chap7 chap8 chap9 chap10 - some files in the directory
```

## File

Item	Description
<code>/usr/bin/xargs</code>	Contains the <code>xargs</code> command.

### Related information:

`ar` command  
`diff` command  
`lint` command

---

## xauth Command

### Purpose

Edits and displays the authorization information used in connecting to the X server.

### Syntax

```
xauth [ -f AuthFile ] [ -v | -q ] [ -i ] [ -b ] [ CommandArgument ... ]
```

### Description

The `xauth` command is usually used to edit and display the authorization information used in connecting to the X server. This program extracts authorization records from one machine and merge them into another (for example, when using remote logins or granting access to other users).

The following commands can be entered interactively, on the `xauth` command line, or in scripts. Note that this program does not contact the X server.

## Item

**add** *DisplayName ProtocolName Hexkey*

**extract** *FileName DisplayName...*

**generate** *DisplayName ProtocolName [trusted | untrusted] [timeout seconds] [group group-id] [data hexdata]*

**list** [*DisplayName...*]

**merge** [*FileName...*]

## Description

An authorization entry is added to the authorization file for the indicated display using the given protocol and key data. The data is specified as an even-length string of hexadecimal digits, each pair representing one octet. The first digit of each pair gives the most significant 4 bits of the octet, and the second digit of the pair gives the least significant 4 bits. For example, a 32-character hexkey would represent a 128-bit value. A protocol name consisting of just a single period is treated as an abbreviation for **MIT-MAGIC-COOKIE-1**.

Authorization entries for each of the specified displays are written to the indicated file. The extracted entries can be read back in using the **merge** and **nmerge** commands. If the file name consists of just a single dash, the entries are written to the binary output.

This command is similar to **add**. The main difference is that instead of requiring the user to supply the key data, it connects to the server specified in *displayname* and uses the SECURITY extension in order to get the key data to store in the authorization file. If the server cannot be contacted or if it does not support the SECURITY extension, the command fails. Otherwise, an authorization entry for the indicated display using the given protocol is added to the authorization file. A protocol name consisting of just a single period is treated as an abbreviation for **MIT-MAGIC-COOKIE-1**.

If the *trusted* option is used, clients that connect using this authorization will have full run of the display, as usual. If *untrusted* is used, clients that connect using this authorization will be considered untrusted and prevented from stealing or tampering with data belonging to trusted clients. See the SECURITY extension specification for full details on the restrictions imposed on untrusted clients. The default is *untrusted*.

The *timeout* option specifies how long in seconds this authorization will be valid. If the authorization remains unused (no clients are connected with it) for longer than this time period, the server purges the authorization, and future attempts to connect using it will fail. Note that the purging done by the server does not delete the authorization entry from the authorization file. The default *timeout* is 60 seconds.

The *group* option specifies the application group that clients connecting with this authorization should belong to. See the application group extension specification for more details. The default is to not belong to an application group.

The *data* option specifies data that the server should use to generate the authorization. Note that this is not the same data that gets written to the authorization file. The interpretation of this data depends on the authorization protocol. The *hexdata* is in the same format as the *hexkey* described in the **add** command. The default is to send no data.

Authorization entries for each of the specified displays (or all displays if none are named) are printed on the standard output in a textual format. Key data is always displayed in the hexadecimal format given in the description of the **add** command.

Authorization entries are read from the specified files and are merged into the authorization database, superseding any matching existing entries. If a file name consists of just a single dash, the binary input is read if it has not been read before.

<b>Item</b>	<b>Description</b>
<b>[n]extract</b> <i>Filename DisplayName...</i>	Authorization entries for each of the specified displays are written to the indicated file. The entries are written in a numeric format suitable for non-binary transmission (such as secure electronic mail). The extracted entries can be read back in using the <b>merge</b> and <b>nmerge</b> commands. If the file name consists of just a single dash, the entries are written to the standard output.
<b>[n]list</b> <i>[DisplayName...]</i>	Authorization entries for each of the specified displays (or all displays if none are named) are printed on the standard output in the numeric format used by the <b>nextract</b> command. Key data is always displayed in the hexadecimal format given in the description of the <b>add</b> command.
<b>[n]merge</b> <i>[FileName...]</i>	Authorization entries are read from the specified files and are merged into the authorization database, superseding any matching existing entries. The numeric format given in the description of the <b>extract</b> command is used. If a file name consists of just a single dash, the standard input is read if it has not been read before.
<b>remove</b> <i>DisplayName...</i>	Authorization entries matching the specified displays are removed from the authority file.
<b>source</b> <i>FileName</i>	The specified file is treated as a script containing <b>xaauth</b> commands to execute. Blank lines and lines beginning with a # (pound sign) are ignored. A single dash can be used to indicate the standard input, if it has not already been read.
<b>info</b>	Information describing the authorization file, whether or not any changes have been made, and from where <b>xaauth</b> commands are being read is printed on the standard output.
<b>exit</b>	If any modifications have been made, the authority file is written out (if allowed), and the program exits. An end of file is treated as an implicit exit command.
<b>quit</b>	The program exits, ignoring any modifications. This may also be accomplished by pressing the interrupt character.
<b>help</b> <i>[String]</i>	A description of all commands that begin with the given string (or all commands if no string is given) is printed on the standard output.
<b>?</b>	A short list of the valid commands is printed on the standard output.

Display names for the **add**, **[n]extract**, **[n]list**, **[n]merge**, and **remove** commands use the same format as the **DISPLAY** environment variable and the common *display* command-line argument. Display-specific information (such as the screen number) is unnecessary and is ignored. Same-machine connections (such as local-host sockets, shared memory, and the Internet Protocol *HostName LocalHost*) are referred to as *HostName/unix:DisplayNumber* so that local entries for different machines can be stored in one authority file.

**Note:** Users that have unsecure networks should take care to use encrypted file transfer mechanisms to copy authorization entries between machines. Similarly, the MIT-MAGIC-COOKIE-1 protocol is not very useful in unsecure environments. Sites that are interested in additional security may need to use encrypted authorization mechanisms such as Kerberos. Spaces are currently not allowed in the protocol name. Quoting could be added.

## Flags

The following options are used with the **xaauth** command. They can be given individually (for example, **-q -i**) or combined (for example, **-qi**).

Item	Description
<b>-f</b> <i>AuthFile</i>	Specifies the name of the authority file to use. By default, <b>xauth</b> uses the file specified by the <b>XAUTHORITY</b> environment variable or <b>.xauthority</b> in the user's home directory.
<b>-v</b>	Indicates that <b>xauth</b> should operate verbosely and print status messages indicating the results of various operations (for example, how many records have been read in or written out). This is the default if <b>xauth</b> is reading commands from its standard input and its standard output is directed to a terminal.
<b>-q</b>	Indicates that <b>xauth</b> should operate quietly and not print unsolicited status messages. This is the default if an <b>xauth</b> command is given on the command line or if the standard output is not directed to a terminal.
<b>-i</b>	Indicates that <b>xauth</b> should ignore any authority file locks. Normally, <b>xauth</b> refuses to read or edit any authority files that have been locked by other programs (usually <b>xdm</b> or another <b>xauth</b> ).
<b>-b</b>	Indicates that <b>xauth</b> should attempt to break any authority file locks before proceeding. Use this option only to clean up stale locks.

## Example

The most common use for the **xauth** command is to extract the entry for the current display, copy it to another machine, and merge it into the user's authority file on the remote machine:

```
% xauth extract \- $DISPLAY | rsh otherhost xauth merge \-
```

## Files

Item	Description
<b>\$HOME/.Xauthority</b>	Contains the default authority file if the <b>XAUTHORITY</b> environment variable is not defined.

---

## xclock Command

### Purpose

Continuously displays the current time of day.

### Syntax

```
xclock [ -Xtoolkitoption ... ] [ -analog | -digital ] [ -chime ] [ -hd Color ] [ -help ] [ -hl Color ] [ -padding Number ] [ -update Seconds ]
```

### Description

The **xclock** command gets the time from the system clock, then displays and updates it in the form of a digital or analog clock. Select the **-analog** or **-digital** flag to display the clock in analog or digital formats. You can also select flags to specify the presentation of the clock, including chime and update frequency, colors, and border width.

This command uses the Athena clock widget, which understands core resource names and classes. To specify these resources, you need to know the hierarchy of the widgets that comprise the **xclock** command. In the following example, the indented items indicate the hierarchical structure. The widget class name is given first, followed by the widget instance name:

```
XClock xclock
Clock cclock
```

The following examples demonstrate the possible ways to specify resources for this client:

```
xclock.cclock.background
XClock*background
xclock*background
```

**Note:** Specifying resources as `xclock.background` which worked with the previous version of `xclock` will not work with this version.

## Flags

Item	Description
<code>-Xtoolkitoption</code>	The <b>xclock</b> command accepts all of the standard X Toolkit command-line option flags in addition to the specific flags listed.
<code>-analog</code>	You can view standard X Toolkit command-line option flag the in the <b>custom</b> command. Sets the analog display mode, which is the default mode. Draws a conventional 12-hour clock face with ticks for each minute and stroke marks on each hour.
<code>-chime</code>	Specifies the sounding of a chime once on the half hour and twice on the hour.
<code>-digital</code>	Sets the 24-hour digital display mode. Displays the date and time in digital form.
<code>-hd Color</code>	Specifies the color of the hands in analog mode on color displays. The default is black.
<code>-help</code>	Prints a brief summary of the allowed options.
<code>-hl Color</code>	(lowercase HL) Specifies the highlight color of the edges of the hands of the analog clock. The default is black.
<code>-padding Number</code>	Specifies the width in pixels of the padding between the window border and the clock text or picture. The default is 8.
<code>-update Seconds</code>	Specifies the frequency in seconds that the <b>xclock</b> command updates its display. If the <b>xclock</b> window is obscured and then exposed, the <b>xclock</b> command redisplay immediately. The specification of an update frequency less than 30 seconds enables the second hand in the analog mode. The default update frequency is 60 seconds.

## .Xdefaults Keywords

Use the following keywords to set the defaults for the **xclock** command.

Item	Description
<b>analog</b> (class Boolean)	Specifies an analog clock instead of a digital clock. The default is true.
<b>chime</b> (class Boolean)	Specifies whether a bell sounds on the hour and half hour.
<b>fontSet</b> (class FontSet)	Specifies the fontset for the digital clock. Variable-width fonts do not always display correctly.
<b>foreground</b> (class Foreground)	Specifies the color of tick marks on color displays. If <b>reverseVideo</b> is specified, the default is white, otherwise the default is black.
<b>hands</b> (class Foreground)	Specifies the color on the inside of the hands in the analog clock on color displays. If <b>reverseVideo</b> is specified, the default is white, otherwise the default is black.
<b>highlight</b> (class Foreground)	Specifies the color used to highlight the clock's hands. If <b>reverseVideo</b> is specified, the default is white, otherwise the default is black.
<b>height</b> (class Height)	Specifies the height of the clock. The default for the analog clock is 164 pixels. The default for the digital clock is whatever is required to hold the clock when displayed in the chosen font.
<b>padding</b> (class Margin)	Specifies the amount of internal padding in pixels. The default is 8.
<b>update</b> (class Interval)	Specifies the frequency in seconds in which the <b>xclock</b> command updates its display.
<b>width</b> (class Width)	Specifies the width of the clock. The default for the analog clock is 164 pixels. The default for the digital clock is whatever is needed to hold the clock when displayed in the chosen font.

## Environment Variables

Item	Description
DISPLAY	Gets the default host and display number.
XENVIRONMENT	Gets the name of a resource file that overrides the global resources stored in the RESOURCE_MANAGER property.

## Examples

- To specify a digital clock display, enter:  

```
xclock -digital
```
- To specify red hands on an analog clock, enter:  

```
xclock -hd red
```

## File

Item	Description
<code>/usr/lib/X11/app-defaults/XClock</code>	Specifies the required resources.

### Related information:

custom command

## xcmsdb Command

### Purpose

Loads, queries, or removes Screen Color Characterization Data stored in properties on the root window of the screen.

**Note:** The `xcmsdb` command is only supported in X11R5 (AIXwindows Version 1.2.3).

### Syntax

```
xcmsdb [ -display Display ] [ [ -query ] [ -remove ] [ -color ] ] | [ -format 32 | 16 | 8 ] [ FileName ]
```

### Description

The `xcmsdb` command is used to load, query, or remove Screen Color Characterization Data stored in properties on the root window of the screen. Screen Color Characterization Data is an integral part of `Xlib`, which is necessary for proper conversion between device-independent and device-dependent color specifications. `Xlib` uses the `XDCCC_LINEAR_RGB_MATRICES` and `XDCCC_LINEAR_RGB_CORRECTION` properties to store color characterization data for color monitors. It uses `XDCCC_GRAY_SCREENWHITEPOINT` and `XDCCC_GRAY_CORRECTION` properties for gray scale monitors. Because `Xlib` allows the addition of Screen Color Characterization Function Sets, added function sets may place their Screen Color Characterization Data on other properties. This utility is unaware of these other properties; therefore, you will need to use a similar utility provided with the function set, or use the example `xprop` utility.

The ASCII readable contents of the `FileName` parameter (or the standard input if no input file is given) are appropriately transformed for storage in properties, provided the `-query` or `-remove` flag options are not specified.

**Note:** The Xcms API in `libX11.a` is supported; however, the client side color name data base, `/usr/lib/X11/Xcms.txt`, and a device color characterization file, `/usr/lib/X11/XcmsIBM5081.dcc`, are provided as unsupported samples.



## Flags

Item	Description
<code>-display</code> <i>Display</i>	Specifies the server to which you are converting.
<code>-query</code>	Reads or attempts to read the XDCCC properties off the screen's root window. If successful, it transforms the data into a more readable format, and then sends the data to standard output.
<code>-remove</code>	Removes or attempts to remove the XDCCC properties on the screen's root window.
<code>-color</code>	Sets the <code>-query</code> and <code>-remove</code> options to only check for the <code>XDCCC_LINEAR_RGB_MATRICES</code> and <code>XDCCC_LINEAR_RGB_CORRECTION</code> properties. If the <code>-color</code> option is not set, the <code>-query</code> and <code>-remove</code> options check for all the properties.
<code>-format 32   16   8</code>	Specifies the property format (32, 16, or 8 bits per entry) for the <code>XDCCC_LINEAR_RGB_CORRECTION</code> property. Precision of encoded floating-point values increases with the increase in bits per entry. The default is 32 bits per entry.

## Parameter

Item	Description
<code>FileName</code>	Specifies the ASCII readable contents of a Screen Color Characterization Data file.

## Examples

1. Use the following example to put Screen Color Characterization Data on the root window by telling the `xcmsdb` command to read it from a file:  

```
xcmsdb /usr/lib/X11/XcmsIBM5081.dcc
```
2. Use the following example after you have already put Screen Color Characterization Data on the root window to tell the `xcmsdb` command to read the data back if it exists:  

```
xcmsdb -query
```

---

## xdm Command

### Purpose

Manages a collection of X Displays with support for XDMCP.

### Syntax

```
xdm [ -config ConfigurationFile] [ -debug DebugLevel ] [ -nodaemon ] [ -error ErrorLogFile ] [ -resources ResourceFile ] [ -server ServerEntry ] [ -udpPort PortNumber ] [ -session SessionProgram ] [ -xrm ResourceSpecification ]
```

### Description

The `xdm` (X Display Manager) command manages a collection of X displays, which may be on the local host or remote servers. The design of the `xdm` command was guided by the needs of X terminals as well as the X Consortium standard XDMCP, the *X Display Manager Control Protocol*. The `xdm` command provides services similar to those provided by the `init`, `getty`, and `login` commands on character terminals: prompting for login name and password, authenticating the user, and running a session.

A *session* is defined by the lifetime of a particular process; in the traditional character-based terminal world, it is the user's login shell. In the `xdm` context, it is an arbitrary session manager. This is because in a windowing environment, a user's login shell process does not necessarily have any terminal-like interface with which to connect. When a real session manager is not available, a window manager or terminal emulator is typically used as the *session manager*, meaning that ending this process ends the user's session.

When the session is ended, **xdm** resets the X server and (optionally) restarts the whole process.

When the **xdm** command receives an **Indirect** query by way of XDMCP, it can run a **chooser** process to perform an XDMCP **BroadcastQuery** (or an XDMCP Query to specified hosts) on behalf of the display and offer a menu of possible hosts that offer XDMCP display management. This feature is useful with X terminals that do not offer a host menu themselves.

Because the **xdm** command provides the first interface that users see, it is designed to be simple to use and easy to customize to the needs of a particular site.

## Typical Usage

The **xdm** command is designed to operate in a wide variety of environments.

First, set up the **xdm** configuration file. Make a directory (usually **/usr/lib/X11/xdm**) to contain all of the relevant files. The following is a reasonable configuration file, which could be named **xdm-config**:

```
DisplayManager.servers:      /usr/lib/X11/xdm/Xservers
DisplayManager.errorLogFile: /usr/lib/X11/xdm/xdm-errors
DisplayManager*resources:   /usr/lib/X11/xdm/Xresources
DisplayManager*startup:     /usr/lib/X11/xdm/Xstartup
DisplayManager*session:     /usr/lib/X11/xdm/Xsession
DisplayManager.pidFile:     /usr/lib/X11/xdm/xdm-pid
DisplayManager._0.authorize: true
DisplayManager*authorize:   false
```

This file contains references to other files. Some of the resources are specified with an \* (asterisk) separating the components. These resources can be made unique for each display by replacing the \* (asterisk) with the display name, but typically this is not useful. See the Resources section on the next page for a complete discussion.

The first file, **/usr/lib/X11/xdm/Xservers**, contains the list of displays to manage that are not using XDMCP. Most workstations have only one display, numbered 0 (zero), so the file looks something like this:

```
:0 Local local /usr/bin/X11/X -force
```

This keeps **/usr/bin/X11/X** running on this display and manages a continuous cycle of sessions.

The **/usr/lib/X11/xdm/xdm-errors** file contains error messages from **xdm** and anything output to standard error by **Xsetup**, **Xstartup**, **Xsession** or **Xreset** scripts. If you have trouble starting the **xdm** command, check the **/usr/lib/X11/xdm/xdm-errors** file to see if the **xdm** command has any clues to the trouble.

The next configuration entry, **/usr/lib/X11/xdm/Xresources**, is loaded onto the display as a resource database using the **xrdb** command. As the authentication widget reads this database before starting, it usually contains parameters for that widget.

## Flags

All of these options (except **-config**) specify values that can also be specified in the configuration file as resources.

Item	Description
<b>-config</b> <i>ConfigurationFile</i>	Names the configuration file, which specifies resources to control the behavior of the <b>xdm</b> command. The <code>/usr/lib/X11/xdm/xdm-config</code> file is the default.
<b>-debug</b> <i>DebugLevel</i>	Specifies the numeric value for the <b>DisplayManager.debugLevel</b> resource. A nonzero value causes <b>xdm</b> to print debugging statements to the terminal and disables the <b>DisplayManager.daemonMode</b> resource, forcing <b>xdm</b> to run synchronously. These error messages may be unclear. To interpret them, check the X11R4 source code for the <b>xdm</b> command.
<b>-nodaemon</b>	Specifies False as the value for the <b>DisplayManager.daemonMode</b> resource. This suppresses the usual daemon behavior, in which the <b>xdm</b> command closes all file descriptors, disassociates itself from the controlling terminal, and puts itself in the background when it first starts up.
<b>-error</b> <i>ErrorLogFile</i>	Specifies the value for the <b>DisplayManager.errorLogFile</b> resource. This file contains errors from <b>xdm</b> as well as anything written to standard error by the various scripts and programs run during the progress of the session.
<b>-resources</b> <i>ResourceFile</i>	Specifies the value for the <b>DisplayManager*resources</b> resource. This file is loaded using the <b>xrdb</b> command to specify configuration parameters for the authentication widget.
<b>-server</b> <i>ServerEntry</i>	Specifies the value for the <b>DisplayManager.servers</b> resource. See the section Server Specification for a description of this resource.
<b>-udpPort</b> <i>PortNumber</i>	Specifies the value for the <b>DisplayManager.requestPort</b> resource. This sets the port number that the <b>xdm</b> command monitors for <b>XDMCP</b> requests. <b>XDMCP</b> uses the registered well-known UDP port 177. Do not change this resource except when debugging.
<b>-session</b> <i>SessionProgram</i>	Specifies the value for the <b>DisplayManager*session</b> resource. This indicates the program to run as the session after the user has logged in.
<b>-xrm</b> <i>ResourceSpecification</i>	Allows an arbitrary resource to be specified, as in most X Toolkit applications.

## Resources

At many stages, the actions of **xdm** can be controlled through the use of its configuration file, which is in the X resource format. Some resources modify the behavior of **xdm** on all displays, while others modify its behavior on a single display. When actions relate to a specific display, the display name is inserted into the resource name between "DisplayManager" and the final resource name segment. For example, **DisplayManager.expo\_0.startup** is the name of the resource that defines the startup shell file on the "expo:0" display. Because the resource manager uses colons to separate the name of the resource from its value and dots to separate resource name parts, **xdm** substitutes underscores for both dots and colons when generating the resource name.

Item	Description
<b>DisplayManager.servers</b>	Specifies either a file name full of server entries, one per line (if the value starts with a slash), or a single server entry. See the section Server Specification for details.
<b>DisplayManager.requestPort</b>	Indicates the UDP port number that the <b>xdm</b> command uses to listen for incoming <b>XDMCP</b> requests. Unless you need to debug the system, leave this with its default value of 177.
<b>DisplayManager.errorLogFile</b>	Redirects error messages to go to the named file rather than to the console. This file also contains any output directed to standard error by the <b>Xsetup</b> , <b>Xstartup</b> , <b>Xsession</b> , and <b>Xreset</b> files, so it will contain descriptions of problems in those scripts as well.
<b>DisplayManager.debugLevel</b>	If the integer value of this resource is greater than 0 (zero), the <b>xdm</b> command outputs a large amount of debugging information. It also disables daemon mode, which would discard the information and allow nonroot users to run the <b>xdm</b> command that would typically not be useful.
<b>DisplayManager.daemonMode</b>	The <b>xdm</b> command attempts to make itself into a daemon process unassociated with any terminal. This is accomplished by forking and leaving the parent process to exit, and then closing file descriptors and releasing the controlling terminal. In some environments this is not desired (in particular, when debugging). Setting this resource to False disables this feature.

Item	Description
<code>DisplayManager.pidFile</code>	The file name specified is created to contain an ASCII representation of the process ID of the main <code>xdm</code> process. The <code>xdm</code> command also uses file locking on this file to attempt to eliminate multiple daemons running on the same machine, which would have unpredictable results.
<code>DisplayManager.lockPidFile</code>	Controls whether the <code>xdm</code> command uses file locking to keep multiple display managers from running simultaneously.
<code>DisplayManager.authDir</code>	Names a directory in which the <code>xdm</code> command stores authorization files while initializing the session. The default value is <code>/usr/lib/X11/xdm</code> .
<code>DisplayManager.autoRescan</code>	A Boolean value that controls whether the <code>xdm</code> command rescans the configuration, servers, access control, and authentication keys files after a session ends and the files have changed. By default the value is True. You can force the <code>xdm</code> daemon to reread these files by sending a <code>SIGHUP</code> signal to the main process.
<code>DisplayManager.removeDomainname</code>	When computing the display name for XDMCP clients, the name resolver typically creates a fully qualified host name for the terminal. As this is sometimes confusing, the <code>xdm</code> command removes the domain name portion of the host name if it is the same as the domain name of the local host when this variable is set. The default value is True.
<code>DisplayManager.keyFile</code>	XDM-AUTHENTICATION-1 style XDMCP authentication requires that a private key be shared between the <code>xdm</code> daemon and the terminal. This resource specifies the file containing those values. Each entry in the file consists of a display name and the shared key. By default, the <code>xdm</code> command does not include support for XDM-AUTHENTICATION-1 because it requires the data encryption method (DES), which is not generally distributable because of United States export restrictions.
<code>DisplayManager.accessFile</code>	To prevent unauthorized XDMCP service and to allow forwarding of XDMCP <code>IndirectQuery</code> requests, this file contains a database of host names that are allowed direct access to this machine or have a list of hosts to which queries should be forwarded. The format of this file is described in the XDMCP Access Control section.
<code>DisplayManager.exportList</code>	A whitespace-separated list of additional environment variables to pass on to the <code>Xsetup</code> , <code>Xstartup</code> , <code>Xsession</code> , and <code>Xreset</code> programs.
<code>DisplayManager.randomFile</code>	A file to checksum to generate the seed of authorization keys. This should be a file that changes frequently. The default is <code>/dev/mem</code> .
<code>DisplayManager.choiceTimeout</code>	Number of seconds to wait for the display to respond after a user has selected a host from the chooser. If the display sends an XDMCP <code>IndirectQuery</code> within this time, the request is forwarded to the chosen host. Otherwise, it is assumed to be from a new session and the chooser is offered again. The default is 15.
<code>DisplayManager.DISPLAY.resources</code>	Specifies the name of the file to be loaded by the <code>xrdb</code> command as the resource database onto the root window of screen 0 of the display. The <code>Login</code> widget, <code>Xsetup</code> , and <code>chooser</code> programs use the resources set in this file. This resource data base is loaded just before the authentication procedure is started, so it can control the appearance of the login window. See the section Authentication Client , which describes the various resources that are appropriate to place in this file. There is no default value for this resource, but <code>/usr/lib/X11/xdm/Xresources</code> is the conventional name.
<code>DisplayManager.DISPLAY.chooser</code>	Specifies the program run to offer a host menu for indirect queries redirected to the special host name <code>CHOOSER</code> . <code>/usr/lib/X11/xdm/chooser</code> is the default. See the sections XDMCP Access Control and <code>Chooser</code> .
<code>DisplayManager.DISPLAY.xrdb</code>	Specifies the program used to load the resources. By default, the <code>xdm</code> command uses <code>/usr/bin/X11/xrdb</code> .
<code>DisplayManager.DISPLAY.cpp</code>	Specifies the name of the C preprocessor that is used by the <code>xrdb</code> command.
<code>DisplayManager.DISPLAY.setup</code>	Specifies a program that is run (as root) before offering the login window. This resource may be used to change the appearance of the screen around the login window or to put up other windows (for example, you may want to run <code>xconsole</code> here). By default, no program is run. The conventional name for a file used here is <code>Xsetup</code> . See the section Setup Program.

Item	Description
<code>DisplayManager.DISPLAY.startup</code>	Specifies a program that is run (as root) after the authentication process succeeds. By default, no program is run. The conventional name for a file used here is <b>Xstartup</b> . See the section Startup Program.
<code>DisplayManager.DISPLAY.session</code>	Specifies the session to be run (when not running as root). By default, <code>/usr/bin/X11/xterm</code> is run. The conventional name is the <b>Xsession</b> script. See the section Session Program.
<code>DisplayManager.DISPLAY.reset</code>	Specifies a program that is run (as root) after the session ends. By default, no program is run. The conventional name is the <b>Xreset</b> script. See the section Reset Program.
<code>DisplayManager.DISPLAY.openDelay</code>	Controls the behavior of the <b>xdm</b> command when attempting to open intransigent servers by specifying the length of the pause (in seconds) between successive attempts.
<code>DisplayManager.DISPLAY.openRepeat</code>	Controls the behavior of the <b>xdm</b> command when attempting to open intransigent servers by specifying the number of attempts to make.
<code>DisplayManager.DISPLAY.openTimeout</code>	Controls the behavior of the <b>xdm</b> command when attempting to open intransigent servers by specifying the number of seconds to wait while actually attempting the open (that is, the maximum time spent in the <code>connect(2)</code> system call).
Item	Description
<code>DisplayManager.DISPLAY.startAttempts</code>	Controls the behavior of the <b>xdm</b> command when attempting to open intransigent servers by specifying the number of times that the entire process is completed before giving up on the server. After the number of attempts specified by the Display Manager <b>openRepeat</b> resource have been made, or if the number of seconds specified by the Display Manager <b>openTimeout</b> resource elapse in any particular attempt, the <b>xdm</b> command ends and restarts the server, attempting to connect again. This process is repeated <i>startAttempts</i> times, at which point the display is declared inactive and disabled. Although this behavior may seem arbitrary, it has been empirically developed and works well on most systems. The default is a value of <b>5</b> for <i>openDelay</i> , a value of <b>5</b> for <i>openRepeat</i> , a value of <b>30</b> for <i>openTimeout</i> , and a value of <b>4</b> for <i>startAttempts</i> .
<code>DisplayManager.DISPLAY.pingInterval</code>	To discover when remote displays disappear, the <b>xdm</b> command occasionally pings them, using an X connection and <b>XSsync</b> calls. This resource specifies the time (in minutes) between ping attempts. By default, it is set to 5 minutes. If you frequently use X terminals, which can become isolated from the managing host, you may want to increase this value. <b>Note:</b> AIXwindows sessions may continue to exist after the terminal has been accidentally disabled. The <b>xdm</b> command does not ping local displays. A workstation session can be ended if the server hangs for NFS service and does not respond to the ping.
<code>DisplayManager.DISPLAY.pingTimeout</code>	To discover when remote displays disappear, the <b>xdm</b> command occasionally pings them, using an X connection and <b>XSsync</b> calls. This resource specifies the maximum amount of time (in minutes) to wait for the terminal to respond to the request. If the terminal does not respond, the session is declared inactive and ended. By default, it is set to 5 minutes. If you frequently use X terminals, which can become isolated from the managing host, you may want to increase this value. <b>Note:</b> AIXwindows sessions may continue to exist after the terminal has been accidentally disabled. The <b>xdm</b> command does not ping local displays. A workstation session could be ended if the server hangs for NFS service and does not respond to the ping.
<code>DisplayManager.DISPLAY.terminateServer</code>	Specifies whether the X server should be canceled when a session ends (instead of resetting it). This option can be used when the server tends to grow without bound over time, to limit the amount of time the server is run. The default value is <b>False</b> .
<code>DisplayManager.DISPLAY.userPath</code>	The <b>xdm</b> command sets the <b>PATH</b> environment variable for the session to this value. It should be a list of directories separated by colons; see the <b>sh</b> command in <i>Commands Reference</i> for a full description. <code>:/bin:/usr/bin:/usr/bin/X11:/usr/ucb</code> is a common setting. The default value can be specified at build time in the AIXwindows system configuration file with the <b>DefaultUserPath</b> resource.

Item	Description
DisplayManager.DISPLAY.systemPath	The <b>xdm</b> command sets the <b>PATH</b> environment variable for the startup and reset scripts to the value of this resource. The default for this resource is specified at build time by the <b>DefaultSystemPath</b> resource entry in the system configuration file; <b>/etc/bin/usr/bin/usr/bin/X11/usr/ucb</b> is a common choice. Note the absence of <b>.</b> (period) (the current directory) from this entry. This is a good practice to follow for root; it avoids many common "Trojan Horse" system penetration schemes.
DisplayManager.DISPLAY.systemShell	The <b>xdm</b> command sets the <b>SHELL</b> environment variable for the startup and reset scripts to the value of this resource. It is <b>/bin/sh</b> by default.
DisplayManager.DISPLAY.failSafeClient	If the default session fails to run, the <b>xdm</b> command returns to this program. This program is run with no arguments, using the same environment variables as the session would have had (see the section Session Program). By default, <b>/usr/bin/X11/xterm</b> is used.
DisplayManager.DISPLAY.grabServer	To improve security, the <b>xdm</b> command grabs the server and keyboard while reading the login name and password. The <b>grabServer</b> resource specifies if the server should be held for the duration of the name/password reading. When set to <b>False</b> , the server is ungrabbed after the keyboard grab succeeds, otherwise the server is grabbed until just before the session begins. The default value is <b>False</b> . The <b>grabTimeout</b> resource specifies the maximum time that the <b>xdm</b> command waits for the grab to succeed. The grab may fail if some other client has the server grabbed, or possibly if the network latencies are high. This resource has a default value of 3 seconds; be cautious when raising it, as a user may be confused by a look-alike window on the display. If the grab fails, the <b>xdm</b> command becomes inactive and restarts the server (if possible) and the session.
DisplayManager.DISPLAY.grabTimeout	
DisplayManager.DISPLAY.authorize	The <b>authorize</b> is a Boolean resource that controls whether the <b>xdm</b> command generates and uses authorization for the local server connections. If authorization is used, the <b>xdm</b> command uses the authorization mechanisms indicated as a whitespace-separated list as the value of the <b>authName</b> resource. <b>XDMCP</b> connections dynamically specify which authorization mechanisms are supported, so the <b>authName</b> resource is ignored in this case. When the <b>authorize</b> resource is set for a display and authorization is not available, the user is informed by a different message displayed in the Login widget. By default, the <b>authorize</b> resource is <b>True</b> ; <b>authName</b> is <b>MIT-MAGIC-COOKIE-1</b> .
DisplayManager.DISPLAY.authName	
DisplayManager.DISPLAY.authFile	Indicates the file is used to communicate the authorization data from the <b>xdm</b> command to the server, using the <b>-auth</b> server command-line option. It should be kept in a directory with restricted write permissions as it could easily be removed, disabling the authorization mechanism in the server.
DisplayManager.DISPLAY.authComplain	If set to a value of <b>False</b> , this disables the use of the <b>unsecureGreeting</b> in the login window. See the section Authentication Client . The default is a value of <b>True</b> .
DisplayManager.DISPLAY.resetSignal	The number of the signal that the <b>xdm</b> command sends to reset the server. See the section Controlling the Server . The default is <b>1(SIGHUP)</b> .
DisplayManager.DISPLAY.termSignal	The number of the signal that the <b>xdm</b> command sends to end the server. See the section Controlling the Server . The default is <b>15(SIGTERM)</b> .
DisplayManager.DISPLAY.resetForAuth	Causes the <b>xdm</b> command to send <b>SIGHUP</b> to the server after setting up the authorization file, causing an additional server reset to occur, during which time the new authorization information is read. The default is a value of <b>False</b> , which works for all AIXwindows servers.
DisplayManager.DISPLAY.userAuthDir	When the <b>xdm</b> command is unable to write to the usual user authorization file ( <b>\$HOME/.Xauthority</b> ), it creates a unique file name in this directory and sets the <b>XAUTHORITY</b> environment variable to the name of the created file. It uses <b>/tmp</b> by default.

## XDMCP Access Control

The database file specified by the **DisplayManager.accessFile** resource provides information that the **xdm** command uses to control access from displays requesting XDMCP service. This file contains three types of entries:

- Entries that control the response to **Direct** and **Broadcast** queries.
- Entries that control the response to **Indirect** queries.
- Macro definitions.

**Direct** query entries contain either a host name or a pattern, which is distinguished from a host name by the inclusion of one or more pattern-matching characters. An \* (asterisk) matches any sequence of 0 (zero) or more characters, and a ? (question mark) matches any single character. These are compared against the host name of the display device. If the entry is a host name, all comparisons are done using network addresses, so that any name that converts to the correct network address may be used. For patterns, only actual host names are used in the comparison, so ensure that you do not attempt to match aliases. Preceding either a host name or a pattern with an ! (exclamation point) causes hosts that match that entry to be excluded.

An Indirect entry also contains a host name or pattern, but follows it with a list of host names or macros to which **indirect** queries should be sent.

A macro definition contains a macro name and a list of host names and other macros that the macro expands to. To distinguish macros from host names, macro names start with a % (percent) character. Macros may be nested.

Indirect entries may also specify to have the **xdm** command run the **chooser** command to offer a menu of hosts to which to connect. For more information, see Chooser.

When checking access for a particular display host, each entry is scanned in turn and the first matching entry determines the response. For example, a **Direct** query entry is ignored when scanning for an **Indirect** entry. A **Broadcast** query entry is ignored when scanning for a **Direct** entry.

Blank lines are ignored. The # character is treated as a comment delimiter causing the rest of that line to be ignored, and a \ (backslash) at the end of the line causes the new line to be ignored, allowing indirect host lists to span multiple lines.

The following is an example **Xaccess** file:

```
#
# Xaccess - XDMCP access control file
#
#
# Direct/Broadcast query entries
#
!xtra.lcs.mit.edu      # disallow direct/broadcast service for xtra
bambi.ogi.edu        # allow access from this particular display
*.lcs.mit.edu        # allow access from any display in LCS
#
# Indirect query entries
#
%HOSTS                expo.lcs.mit.edu xenon.lcs.mit.edu \
                    excess.lcs.mit.edu kanga.lcs.mit.edu
extract.lcs.mit.edu   xenon.lcs.mit.edu #force extract to contact xenon
!xtra.lcs.mit.edu     dummy          #disallow indirect access
*.lcs.mit.edu        %HOSTS          #all others get to choose
```

## Chooser

For X terminals that do not offer a host menu for use with **Broadcast** or **Indirect** queries, the **chooser** program can do this for them. In the **Xaccess** file, specify **CHOOSER** as the first entry in the Indirect host list. The **chooser** program sends a **Query** request to each of the remaining host names in the list and offers a menu of all the hosts that respond.

The list may consist of the word **BROADCAST**, in which case **chooser** sends a **Broadcast** query instead, again offering a menu of all hosts that respond.

The following is an example **Xaccess** file using **chooser**:

```
extract.lcs.mit.edu      CHOOSER  %HOSTS      #offer a menu of these hosts
xtra.lcs.mit.edu        CHOOSER  BROADCAST   #offer a menu of all hosts
```

The program to use for **chooser** is specified by the **DisplayManager.DISPLAY.chooser** resource. Resources for this program can be put into the file named by the **DisplayManager.DISPLAY.resources** resource.

The **chooser** has been implemented using a Motif **SelectionBoxWidget**. Refer to the **XmSelectionBoxWidget Class** documentation for a description of resources and widget or gadget names.

## Server Specification

The resource **DisplayManager.servers** gives a server specification or, if the values starts with a / (slash), the name of a file containing server specifications, one per line.

Each specification indicates a display that should constantly be managed and that is not using **XDMCP**. Each consists of at least three parts:

- Display name
- Display class
- Display type
- For local servers, a command line to start the server.

A typical entry for local display number 0 would be:

```
:0 IBM-GT local /usr/bin/X11/X :0
```

The display types are:

Item	Description
local	local display: <code>\fIXdm\fP</code> must run the server
foreign	remote display: <code>\fIXdm\fP</code> opens an X connection to a running server

The display name must be something that can be passed in the **-display** option to an X program. This string is used to generate the display-specific resource names, so be careful to match the names (for example, use `":0 local /usr/bin/X11/X :0"` instead of `"localhost:0 local /usr/bin/X11/X :0"` if your other resources are specified as `"DisplayManager._0.session"`). The display class portion is also used in the display-specific resources as the class of the resource. This is useful if you have a large collection of similar displays (like a corral of X terminals) and would like to set resources for groups of them. When using **XDMCP**, the display is required to specify the display class, so the manual for your particular X terminal should document the display class string for your device. If it does not, you can run the **xdm** command in debug mode and look at the resource strings that it generates for that device, which will include the class string.



## Setup Program

The **Xsetup** file is run after the server is reset, but before the login window is offered. The file is typically a shell script. It is run as root, so be careful about security. This is the place to change the root background or bring up other windows that should be displayed on the screen along with the Login widget. Because **xdm** grabs the keyboard, other windows will not be able to receive keyboard input. They will be able to interact with the mouse, however; beware of potential security holes here. If **DisplayManager.DISPLAY.grabServer** is set, **Xsetup** will not be able to connect to the display at all. Resources for this program can be put into the file named by **DisplayManager.DISPLAY.resources**.

In addition to any specified by **DisplayManager.exportList**, the following environment variables are passed:

Item	Description
<b>DISPLAY</b>	Specifies the associated display name.
<b>PATH</b>	Specifies the value of <b>DisplayManager.DISPLAY.systemPath</b> .
<b>SHELL</b>	Specifies the value of <b>DisplayManager.DISPLAY.systemShell</b> .
<b>XAUTHORITY</b>	Specifies that it may be set to an authority file.

## Authentication Client

The MIT authentication widget has been replaced by an authentication client composed of standard Motif widgets. The following is a list of the widget names (and their widget class):

```
outframe(xmFrameWidget)
  inframe(xmFrameWidget)
    main(XmFormWidget)
      tframe(xmFrameWidget)
        greeting(xmLabelGadget)
        logoLine(xmFormWidget)
        dpyname(xmLabelWidget)
        userLine(xmRowColumnWidget)
        userLabel(xmLabelWidget)
        username(xmTextWidget)
        passLabel(xmLabelWidget)
        password(xmTextWidget)
        failsafeLine(xmFormWidget)
        failsafe(xmToggleButtonWidget)
        cancelLine(xmFormWidget)
        cancel(xmPushButtonWidget)
        message(xmLabelWidget)
```

The authentication client reads a name/password pair from the keyboard. Put resources for this client into the file named by **DisplayManager.DISPLAY.resources**. All of these have reasonable default values, so it is unnecessary to specify any of them. See **/usr/lib/X11/xdm/Xresources** for more information on default values for authentication client resources as well as the appropriate widget class documentation. The following resources are also supported by the authentication client:

Item	Description
<b>Xlogin*foreground</b>	Specifies the color used for the foreground.
<b>Xlogin*background</b>	Specifies the color used for the background.
<b>Xlogin*greeting</b>	Specifies a string that identifies this window. The default is AIXwindows environment.
<b>Xlogin*greetFont</b>	Specifies the font used to display the greeting.
<b>Xlogin*frameColor</b>	Specifies the background color used to display the greeting.
<b>Xlogin*titleFont</b>	Specifies the font used to display the title.
<b>Xlogin*namePrompt</b>	Specifies the string displayed to prompt for a user name. The Xrdb program strips trailing white space from resource values. Add spaces escaped with backslashes at the end of the prompt. The default is "login:".
<b>Xlogin*promptFont</b>	Specifies the font used to display both prompts.
<b>Xlogin*failFont</b>	Specifies the font used for the failsafe button.
<b>Xlogin*cancelFont</b>	Specifies the font used for the cancel button.

Item	Description
<code>Xlogin*messageFontlist</code>	Specifies the font used to display the failure message.
<code>Xlogin*failColor</code>	Specifies the color used to display the failure message.
<code>Xlogin*failTimeout</code>	Specifies the number of seconds that the failure message is displayed. The default is thirty seconds.
<code>Xlogin*sessionArgument</code>	Specifies the argument to be passed to the session program.
<code>Xlogin*XmText.translations</code>	This specifies the translations use for the authentication client. Refer to the X Toolkit documentation for a complete discussion on translations. The default translation table is: <pre> Ctrl&lt;Key&gt;b: backward-character()\n\ Ctrl&lt;Key&gt;a: beginning-of-line()\n\ Ctrl&lt;Key&gt;e: end-of-line()\n\ Ctrl&lt;Key&gt;f: forward-character()\n\ Ctrl&lt;Key&gt;d: kill-next-character()\n\ Ctrl&lt;Key&gt;k: kill-to-end-of-line()\n\ Ctrl&lt;Key&gt;u: kill-to-start-of-line()\n\ </pre> <p>You may setup XDM to use the standard XDM translations by replacing the <code>XmText</code> translations as defined in <code>Xresources</code>:</p> <p><b>Note:</b> Use <code>&lt;Key&gt;osfHelp</code> instead of <code>&lt;Key&gt;F1</code> due to the Motif default virtual bindings.)</p> <pre> Xlogin*XmText.translations: #override\n\ &lt;Key&gt;osfHelp: set-session-argument(failsafe) finish-field()\n\ Ctrl&lt;Key&gt;Return: set-session-argument(failsafe) finish-field()\n\ Ctrl&lt;Key&gt;H: delete-previous-character() \n\ Ctrl&lt;Key&gt;D: delete-character() \n\ Ctrl&lt;Key&gt;B: move-backward-character() \n\ Ctrl&lt;Key&gt;F: move-forward-character() \n\ Ctrl&lt;Key&gt;A: move-to-beginning() \n\ Ctrl&lt;Key&gt;E: move-to-end() \n\ Ctrl&lt;Key&gt;K: erase-to-end-of-line() \n\ Ctrl&lt;Key&gt;U: erase-line() \n\ Ctrl&lt;Key&gt;X: erase-line() \n\ &lt;Key&gt;Return: finish-field() \n\ &lt;Key&gt;BackSpace: delete-previous-character() \n\ &lt;Key&gt;Delete: delete-previous-character() \n\ </pre>
	In addition to the typical <code>XmText</code> actions, the following actions are also supported by the client to be compatible with the standard XDM translations:
	<b>delete-previous-character</b> Erases the character before the cursor.
	<b>delete-character</b> Erases the character after the cursor.
	<b>move-backward-character</b> Moves the cursor backward.
	<b>move-forward-character</b> Moves the cursor forward.
	<b>move-to-beginning</b> Moves the cursor to the beginning of the editable text.
	<b>move-to-end</b> Moves the cursor to the end of the editable text.
	<b>erase-to-end-of-line</b> Erases all text after the cursor.
	<b>erase-line</b> Erases the entire text.
	<b>finish-field</b> If the cursor is in the name field, proceeds to the password field; if the cursor is in the password field, checks the current name/password pair. If the name/password pair is valid, <code>xdm</code> starts the session. Otherwise the failure message is displayed and the user is prompted again.
	<b>insert-char</b> Inserts the character typed.
	<b>set-session-argument</b> Specifies a single word argument that is passed to the session at startup. See the sections <code>Session Program</code> and <code>Typical Usage</code> .

## Startup Program

The **Xstartup** file is typically a shell script. Because it is run as the root user, be careful about security when it runs. It usually contains commands that add entries to **/etc/utmp**, mount users' home directories from file servers, display the message of the day, or cancel the session if logins are not allowed.

In addition to the environment variables specified by **DisplayManager.exportList**, the following variables are passed:

Item	Description
<b>DISPLAY</b>	Specifies the associated display name.
<b>HOME</b>	Specifies the initial working directory of the user.
<b>USER</b>	Specifies the user name.
<b>PATH</b>	Specifies the value of <b>DisplayManager.DISPLAY.systemPath</b> .
<b>SHELL</b>	Specifies the value of <b>DisplayManager.DISPLAY.systemShell</b> .
<b>XAUTHORITY</b>	May be set to an authority file.

No arguments are passed to the script. The **xdm** command waits until this script exits before starting the user session. If the exit value of this script is nonzero, the **xdm** command discontinues the session and starts another authentication cycle.

## Session Program

The **Xsession** program establishes the style of the user's session. It is run with the permissions of the authorized user.

In addition to any specified by **DisplayManager.exportList**, the following environment variables are passed:

Item	Description
<b>DISPLAY</b>	Specifies the associated display name.
<b>HOME</b>	Specifies the initial working directory of the user.
<b>USER</b>	Specifies the user name.
<b>PATH</b>	Specifies the value of <b>DisplayManager.DISPLAY.userPath</b> .
<b>SHELL</b>	Specifies the user's default shell (from <b>getpwnam</b> ).
<b>XAUTHORITY</b>	May be set to a nonstandard authority file.

At most installations, the **Xsession** program should look in the user's home directory (**\$HOME**) for a file **.xsession**, which contains the commands that the user would like to use as a session. The **Xsession** program should also implement a system default session if no user-specified session exists. See the section Typical Usage.

An argument may be passed to this program from the authentication widget using the ``set-session-argument'` action. This can be used to select different styles of session. Usually, this feature is used to allow the user to escape from the ordinary session when it fails. This allows users to repair their own **.xsession** if it fails, without requiring administrative intervention. The section Typical Usage demonstrates this feature.

## Reset Program

The **Xreset** script is run after the user session has ended. Run as root, it should contain commands that undo the effects of commands in **Xstartup** by removing entries from **/etc/utmp** or unmounting directories from file servers. The environment variables that are passed to **Xstartup** are also passed to **Xreset**. This program is symmetrical with the **Xstartup** program.

## Controlling the Server

The **x<sub>dm</sub>** command controls local servers using POSIX signals. The **SIGHUP** signal is expected to reset the server, closing all client connections and performing other cleanup duties. The **SIGTERM** signal is expected to cancel the server. If these signals do not perform the expected actions, the resources **DisplayManager.DISPLAY.resetSignal** and **DisplayManager.DISPLAY.termSignal** can specify alternate signals.

To control remote terminals that are not using **XDMCP**, the **x<sub>dm</sub>** command searches the window hierarchy on the display and uses the protocol request **KillClient** in an attempt to clean up the terminal for the next session. This may not actually cause all of the clients to become inactive, because only those that have created windows will be noticed. **XDMCP** provides a more sure mechanism; when the **x<sub>dm</sub>** command closes its initial connection, the session is over and the terminal is required to close all other connections.

## Controlling XDM

The **x<sub>dm</sub>** command responds to two signals: **SIGHUP** and **SIGTERM**. When sent a **SIGHUP**, **x<sub>dm</sub>** rereads the configuration file, the access control file, and the servers file. For the servers file, it notices if entries have been added or removed. If a new entry has been added, the **x<sub>dm</sub>** command starts a session on the associated display. Entries that have been removed are disabled immediately, meaning that any session in progress is ended without notice and no new session is started.

When sent a **SIGTERM**, the **x<sub>dm</sub>** command stops all sessions in progress and exits. This can be used when shutting down the system.

The **x<sub>dm</sub>** command attempts to mark its various subprocesses for use by the **ps** command by editing the command-line argument list in place. Because the **x<sub>dm</sub>** command cannot allocate additional space for this task, it is useful to start the **x<sub>dm</sub>** command with a reasonably long command line (using the full path name should be enough). Each process that is servicing a display is marked **-display**.

## Other Possibilities

You can use the **x<sub>dm</sub>** command to run a single session at a time, using the **xinit** command options or other suitable daemons by specifying the server on the command line:

```
xdm -server ":0 local /usr/bin/X11/X :0 -force"
```

It might also run a file server and a collection of X terminals. The configuration for this is identical to the previous sample, except the **Xservers** file would look like the following:

```
extol:0 VISUAL-19 foreign
exalt:0 NCD-19 foreign
explode:0 NCR-TOWERVIEW3000 foreign
```

This directs the **x<sub>dm</sub>** command to manage sessions on all three of these terminals. See the section Controlling XDM for a description of using signals to enable and disable these terminals.

**Note:** The **x<sub>dm</sub>** command does not coexist well with other window systems. To use multiple window systems on the same hardware, use the **xinit** command.

## Examples

1. The sample **xstartup** script that follows prevents login while the file **/etc/nologin** exists. As there is no provision for displaying any messages here (there is no core X client that displays files), the setup in this example is not recommended because the login would fail without explanation. Thus, this is not a complete example, but a demonstration of the available functionality.

```
#!/bin/sh
#
# Xstartup
#
# This program is run as root after the user is verified
#
if [ \-f /etc/nologin ]; then
    exit 1
fi
exit 0
```

2. This **Xsession** script recognizes the special **failsafe** mode, specified in the translations in the preceding **Xresources** file, to provide an escape from the ordinary session:

```
#!/bin/sh
exec > $HOME/.xsession-errors 2>&1
case $# in
1)
    case $1 in failsafe)
        exec aixterm -geometry 80x24-0-0
        ;;
    esac
esac
startup=$HOME/.xsession
resources=$HOME/.Xresources
if [ -f /usr/bin/X11/startx ]; then
    exec /usr/bin/X11/startx -t -wait
elif [ -f $startup ]; then
    exec $startup
else
    if [ -f $resources ]; then
        xrdp -load $resources
    fi
    mwm &
    exec aixterm -geometry 80x24+10+10 -ls
fi
```

3. To have **xdm** come up from system startup, as root type the following:

```
/usr/lib/X11/xdm/xdmconf
```

4. To disable **xdm** on reboot, as root type the following:

```
/usr/lib/X11/xdm/xdmconf -d
```

5. When using **xdm** to manage your display, an authentication procedure ensures that only clients that are allowed can connect to your display. Clients that are built using X11 R4 and X11 R5 libraries understand this protocol. Clients that are built with X11 R3 or earlier libraries do not support this authentication protocol and are not allowed to connect to the Xserver unless **xhost** permission is granted. You can connect local clients by typing the following:

```
xhost =localhost
```

or

```
xhost =machine
```

where *machine* is the hostname of the local client.

## Files

Item	Description
<code>/usr/lib/X11/xdm/xdm-config</code>	The default configuration file.
<code>/usr/lib/X11/xdm/Xaccess</code>	The default access file, listing authorized displays.
<code>/usr/lib/X11/xdm/Xservers</code>	The default server file, listing non-XDMCP servers to manage.
<code>\$(HOME)/.Xauthority</code>	User authorization file where <code>xdm</code> stores keys for clients to read.
<code>/usr/lib/X11/xdm/chooser</code>	The default chooser.
<code>/usr/bin/X11/xrdb</code>	The default resource database loader.
<code>/usr/bin/X11/X</code>	The default server.
<code>/usr/bin/X11/xterm</code>	The default session program and failsafe client.
<code>/usr/lib/X11/xdm/A&lt;host&gt;\-&lt;suffix&gt;</code>	The default place for authorization files.

#### Related reference:

“X Command” on page 131

“xinit Command” on page 172

#### Related information:

startx command

## xfindproxy Command

### Purpose

Locates proxy services.

### Syntax

```
xfindproxy -manager ManagerAddress -name ServiceName -server ServerAddress [ -auth ] [ -host HostAddress ] [ -options Options ]
```

### Description

`xfindproxy` is a program used to locate available proxy services. It utilizes the Proxy Management Protocol to communicate with a proxy manager. The proxy manager keeps track of all available proxy services, starts new proxies when necessary, and makes sure that proxies are shared whenever possible.

If `xfindproxy` is successful in obtaining a proxy address, it will print it to stdout. The format of the proxy address is specific to the proxy service being used. For example, for a proxy service of LBX, the proxy address would be the X display address of the proxy (e.g, `blah.x.org:63`).

If `xfindproxy` is unsuccessful in obtaining a proxy address, it will print an error to `stderr`.

### Flags

Item	Description
<code>-manager</code>	This argument is required, and it specifies the network address of the proxy manager. The format of the address is a standard ICE network id (for example, <code>tcp/blah.x.org:6500</code> ).
<code>-name</code>	This argument is required, and it specifies the name of the desired proxy service (for example, LBX). The name is case insensitive.
<code>-server</code>	This argument is also required, and it specifies the address of the target server. The format of the address is specific to the proxy service specified with the <code>-name</code> argument. For example, for a proxy service of LBX, the address would be an X display address (e.g, <code>blah.x.org:0</code> ).

<b>Item</b>	<b>Description</b>
<b>-auth</b>	This argument is optional. If specified, <b>xfindproxy</b> will read 2 lines from standard input. The first line is an authorization/authentication name. The second line is the authorization/authentication data in hex format (the same format used by <b>xauth</b> ). <b>xfindproxy</b> will pass this auth data to the proxy, and in most cases, will be used by the proxy to authorize/authenticate itself to the target server.
<b>-host</b>	This argument is optional. If <b>xfindproxy</b> starts a new proxy service, it will pass the host specified. The proxy may choose to restrict all connections to this host. In the event that <b>xfindproxy</b> locates an already existing proxy, the host will be passed, but the semantics of how the proxy uses this host are undefined.
<b>-options</b>	This argument is optional. If <b>xfindproxy</b> starts a new proxy service, it will pass any options specified. The semantics of the options are specific to each proxy server and are not defined here. In the event that <b>xfindproxy</b> locates an already existing proxy, the options will be passed, but the semantics of how the proxy uses these options are undefined.

**Related information:**

proxymngr command

## **xfs Command**

### **Purpose**

Supplies fonts to X Window System display servers.

### **Syntax**

```
xfs [ -config ConfigurationFile ] [ -ls ListenSocket ] [ -port Number ]
```

### **Description**

**xfs** is the AIXwindows font server. It supplies fonts to AIXwindows display servers.

The **xfs** server responds to the following signals:

<b>Item</b>	<b>Description</b>
<b>SIGTERM</b>	Causes the font server to exit cleanly.
<b>SIGUSR1</b>	Causes the server to re-read its configuration file.
<b>SIGUSR2</b>	Causes the server to flush any cached data it may have.
<b>SIGHUP</b>	Causes the server to reset, closing all active connections and re-reading the configuration file.

The server is usually run by a system administrator, and started by way of boot files such as **/etc/rc.tcpip**. Users may also wish to start private font servers for specific sets of fonts.

The configuration language is a list of keyword and value pairs. Each keyword is followed by an = (equal sign) and the desired value.

The following list shows recognized keywords and the types and descriptions of valid values:

Item	Description
#	A comment character when located in the first column.
catalogue (List of string)	Ordered list of font path element names. The current implementation only supports a single catalogue ("all"), containing all of the specified fonts.
alternate-servers (List of string)	List of alternate servers for this font server.
client-limit (Cardinal)	Number of clients that this font server will support before refusing service. This is useful for tuning the load on each individual font server.
clone-self (Boolean)	Whether this font server should attempt to clone itself when it reaches the client-limit.
default-point-size (Cardinal)	The default point size (in decipoints) for fonts that do not specify.
default-resolutions (List of resolutions)	Resolutions the server supports by default. This information may be used as a hint for pre-rendering and substituted for scaled fonts which do not specify a resolution.  A resolution is a comma-separated pair of x and y resolutions in pixels per inch. Multiple resolutions are separated by commas.
error-file (String)	Filename of the error file. All warnings and errors are logged here.
port (Cardinal)	TCP port on which the server will listen for connections. The default is 7100.
use-syslog (Boolean)	Whether the <b>syslog</b> function (on supported systems) is to be used for errors.
deferglyphs (String)	Set the mode for delayed fetching and caching of glyphs. Value is none, meaning deferred glyphs is disabled. all, meaning deferred glyphs is enabled for all fonts, and 16, meaning deferred glyphs is enabled only for 16-bit fonts.

One of the following forms can be used to name a font server that accepts TCP connections:

```
tcp/hostname:port
tcp/hostname:port/cataloguelist
```

The hostname specifies the name (or decimal numeric address) of the machine on which the font server is running. The port is the decimal TCP port on which the font server is listening for connections. The cataloguelist specifies a list of catalogue names, with '+' as a separator. The following are some examples:

```
tcp/expo.lcs.mit.edu:7100, tcp/18.30.0.212:7101/all
```

One of the following forms can be used to name a font server that accepts DECnet connections:

```
decnet/nodename::font$objname
decnet/nodename::font$objname/cataloguelist
```

The nodename specifies the name (or decimal numeric address) of the machine on which the font server is running. The objname is a normal, case-insensitive DECnet object name. The cataloguelist specifies a list of catalogue names, with '+' as a separator.

## Flags



<b>Item</b>	<b>Description</b>
<b>-config</b> <i>ConfigurationFile</i>	Specifies the configuration file the font server will use.
<b>-ls</b> <i>ListenSocket</i>	Specifies a file descriptor that is already set up to be used as the listen socket. This option is only intended to be used by the font server itself when automatically spawning another copy of itself to handle additional connections.
<b>-port</b> <i>Number</i>	Specifies the TCP port number on which the server will listen for connections.

## Examples

```
#
# sample font server configuration file
#
# allow a max of 10 clients to connect to this font server
client-limit = 10
# when a font server reaches its limit, start up a new one
clone-self = on
# alternate font servers for clients to use
alternate-servers = hansen:7101,hansen:7102
# where to look for fonts
# the first is a set of Speedo outlines, the second is a set of
# misc bitmaps and the last is a set of 100dpi bitmaps
#
catalogue = /usr/lib/fonts/type1,
           /usr/lib/X11/ncd/fonts/misc,
           /usr/lib/X11/ncd/fonts/100dpi/
# in 12 points, decipoints
default-point-size = 120
# 100 x 100 and 75 x 75
default-resolutions = 100,100,75,75
```

## Files

<b>Item</b>	<b>Description</b>
<i>/usr/lib/X11/fs/config</i>	The default configuration file.

---

## xget Command

### Purpose

Receives secret mail in a secure communication channel.

### Syntax

**xget**

### Description

The **xget** command is used to receive secret mail in a secure communication channel. The messages can be read only by the intended recipient. The **xget** command asks for your password and enables you to read your secret mail.

The **xget** command is used with the **enroll** command and the **xsend** command to send and receive secret mail. The **enroll** command sets up the password used to receive secret mail. The **xsend** command sends mail that can be read only by the intended recipient.

When you issue the **xget** command, you are prompted for your encryption key. Enter the password you previously set up using the **enroll** command.

The prompt for the **xget** command is a ? (question mark). The following subcommands control message disposition:

Item	Description
<b>q</b> (quit)	Writes any mail not yet deleted to the user's mailbox and exits. Pressing End Of File (Ctrl-D) has the same effect.
<b>n</b> (delete) or <b>d</b> (delete) or <b>Enter</b>	Deletes the current message and displays the next message.
<b>!</b> <i>Command</i>	Runs the specified workstation command.
<b>s</b> [ <i>Filename</i> ]	Saves the message in the named <i>File</i> parameter instead of in the default mail file, <b>mbx</b> .
<b>w</b> [ <i>Filename</i> ]	Saves the message, without its header, in the specified <i>File</i> parameter instead of in the default mail file <b>mbx</b> .
<b>?</b> (help)	Displays a subcommand summary.

## Examples

1. To receive secret mail, enter:

```
xget
```

You are prompted for the password, established with the **enroll** command. After entering your password, the **xget** command prompt (?) and a listing of any secret mail is displayed.

2. To display your secret mail, at the **xget** prompt (?), press the Enter key.

After the most recent message is displayed, a ? (question mark) indicates the **xget** command is waiting for one of the **xget** subcommands. Enter **help** or a ? (question mark) to list the subcommands available.

3. To save a message or a file to the default mail file, enter:

```
xget
```

Press the Enter key after the ? (question mark) prompt until the desired file is displayed. When the appropriate file is displayed, enter:

```
s
```

In this example, the file is saved in the default mail file, **mbx**.

4. To save a message or a file to a specific file, enter:

```
xget
```

Press the Enter key after the ? (question mark) prompt until the desired file is displayed. When the appropriate file is displayed, enter:

```
s mycopy
```

In this example, the file is saved in a file named **mycopy**, instead of the default mail file.

5. To delete a message, enter:

```
xget
```

Press the Enter key after the ? (question mark) prompt until the desired file is displayed. When the appropriate file is displayed, enter:

```
d
```

In this example, the current file is deleted.

## Files

Item	Description
<code>/var/spool/secretmail/User.key</code>	Contains the encrypted key for <i>User</i> .
<code>/var/spool/secretmail/User.[0-9]</code>	Contains the encrypted mail messages for <i>User</i> .
<code>/usr/bin/xget</code>	Contains executable files.

### Related reference:

“xsend Command” on page 211

### Related information:

enroll command  
mail command  
Mail management

---

## xhost Command

### Purpose

Controls who accesses Enhanced X-Windows on the current host machine.

### Syntax

```
xhost [ + | - ] [ Name ]
```

### Description

The **xhost** command adds or deletes host names on the list of machines from which the X Server accepts connections.

This command must be run from the machine with the display connection. You can remove a name from the access list by using the *-Host* parameter. Do not remove the current name from the access list. If you do, log off the system before making any corrections.

Entering the **xhost** command with no variables shows the current host names with access your X Server and a message indicating whether or not access is enabled.

For security, options that affect access control may only be run from the *controlling host*. For workstations, this is the same machine as the server. For X terminals, it is the login host.

To enable a remote name by default, the name can be defined in the `/etc/X?.hosts` file, where ? is the display number to which you enable access.

For example, the display `jeanne:0` can be accessed by systems defined in the `/etc/X0.hosts` file on a system that uses the default host name of `jeanne`. In both the display name and the file name, 0 indicates the display number that the defined remote systems are allowed to access through Enhanced X-Windows.

### Flags

Item	Description
+ <i>Name</i>	Defines the host name (the plus sign is optional) to be added to the X Server access list.
- <i>Name</i>	Defines the host name to be removed from the X Server access list. Existing connections are not broken, but new connection attempts will be denied. Note that you can remove the current machine; however, further connections (including attempts to add it back) are not permitted. The only way to allow local connections again is to reset the server (thereby breaking all connections).
+	Specifies that access is unlimited. Access control is turned off.
-	Turns access control on.
	The complete <i>Name</i> has a the following <i>family:name</i> syntax:
<b>inet</b>	Internet host
<b>local</b>	Contains only one name, the empty string
	<b>Note:</b> The family is case sensitive. The format of the name varies with the family.

## xinit Command

### Purpose

Initializes the X Window System.

### Syntax

```
xinit [ [ Client ] Options ] [ - - [ Server ] [ Display ] Options ]
```

### Description

The **xinit** command starts the AIXwindows server and a first client program on systems that cannot start X directly from `/etc/init` or in environments that use multiple window systems. When this first client exits, the **xinit** command stops the X server and then ends.

If no specific client program is given on the command line, the **xinit** command looks for a file to run to start up client programs. The **xinit** command looks for the `$XINITRC` environment variable. If the file is not there, it then looks for the `$HOME/.xinitrc` file. If it still does not find the file, it follows these steps:

1. The **xinit** command looks next to `/usr/lib/X11/$LANG/xinitrc`.
2. Next, it looks to `/usr/lpp/X11/defaults/$LANG/xinitrc`.
3. And finally, it looks to `/usr/lpp/X11/defaults/xinitrc`.

If no such file exists, **xinit** uses the following as a default:

```
aixterm \-geometry +1+1 \-n login \-display :0
```

If no specific server program is given on the command line, the **xinit** command follows these steps:

1. The **xinit** command looks for a file to run as a shell script to start up the server. The **xinit** command looks for files first in the `$XSERVERRC` environment variable.
2. If the file is not there, it looks for the `$HOME/.xserverrc` file.
3. If it still does not find the `$HOME/.xserverrc` file, it looks next to `/usr/lpp/X11/defaults/xserverrc` file.
4. And finally, if it does not find any of the previous files, the **xinit** command runs the **X** command to start the X server and uses the following as a default:

```
X :0
```

Note that this assumes that there is a program named **X** in the current search path. However, servers are usually named `Xdisplaytype` where *displaytype* is the type of graphics display which is driven by this server. The site administrator should, therefore, make a link to the appropriate type of server on the machine, or create a shell script that runs the **xinit** command with the appropriate server.

**Note:** If you attempt to start AIXwindows without an available pointer device, such as a mouse or a tablet, AIXwindows will not open. Some devices can be plugged in but not defined and thus not available to the system, as well as the reverse.

An important point is that programs which are run by `.xinitrc` should be run in the background if they do not exit right away, so that they do not prevent other programs from starting up. However, the last long-lived program started (usually a window manager or terminal emulator) should be left in the foreground so that the script does not exit (which indicates that the user is done and that xinit should exit).

An alternate client and/or server may be specified on the command line. The desired client program and its arguments should be given as the first command line arguments to `xinit`. To specify a particular server command line, add a `—` (double dash) to the `xinit` command line (after any client and arguments) followed by the desired server command.

Both the client program name and the server program name must begin with a `/` (slash) or a `.` (period). Otherwise, they are treated as an arguments to be added to their respective startup lines. This makes it possible to add arguments (for example, foreground and background colors) without having to retype the whole command line.

If a clear server name is not given and the first argument following the `—` (double dash) is a `:` (colon) followed by a number, `xinit` uses that number as the display number instead of zero. All remaining arguments are added to the server command line.

The following environment variables are used with the `xinit` command:

Item	Description
<code>DISPLAY</code>	This variable gets set to the name of the display to which clients should connect.
<code>XINITRC</code>	This variable specifies an init file containing shell commands to start up the initial windows. By default, <code>.xinitrc</code> in the home directory is used.
<i>Options</i>	List any option you wish that is available to the client you specified.
<i>Client</i>	Specify the client with which you are working. For example, <code>xterm</code> or <code>aixterm</code> . The client you specify must begin with a <code>.</code> (dot) or a <code>/</code> (slash).
<i>Server</i>	Use any valid xserver. The server you specify must begin with a <code>.</code> (dot) or a <code>/</code> (slash).

## Examples

1. To start up a server named `X` and run the user's `xinitrc` program, if it exists, or else start an `aixterm` command enter:

```
xinit
```

2. To start a specific type of server on an alternate display, enter:

```
xinit -- /usr/bin/X11/X qdss:1
```

3. To start up a server named `X`, and add the given arguments to the default `xinitrc` or `aixterm` command, enter:

```
xinit -geometry =80x65+10+10 -fn 8x13 -j -fg white -bg navy
```

4. To use the command `/Xsun -l -c` to start the server and add the arguments `-e widgets` to the default `xinitrc` or `aixterm` command, enter:

```
xinit -e widgets -- ./Xsun -l -c
```

5. To start a server named `X` on display 1 with the arguments `-a 2 -t 5`, then start a remote shell on the machine `fasthost` in which it runs the command `cpupig`, telling it to display back on the local workstation, enter:

```
xinit /usr/ucb/rsh fasthost cpupig -display ws:1 -- :1 -a 2 -t 5
```

6. The following sample of the `.xinitrc` script starts a clock, several terminals, and leaves the window manager running as the last application. Assuming that the window manager has been configured properly, the user then chooses the `Exit` menu item to end the AIXwindows session.

```
xrdb -load $HOME/.Xresources
xsetroot -solid gray &
xclock -g 50x50-0+0 -bw 0 &
xload -g 50x50-50+0 -bw 0 &
xterm -g 80x24+0+0 &
xterm -g 80x24+0-0 &
mwm
```

7. Sites that want to create a common startup environment could simply create a default `.xinitrc` script that references a site-wide startup file:

```
#!/bin/sh . /usr/local/lib/site.xinitrc
```

8. Another approach is to write a script that starts the `xinit` command with a specific shell script. Such scripts are usually named `x11`, `xstart`, or `startx` and are a convenient way to provide a simple interface for novice users:

```
#!/bin/sh xinit /usr/local/lib/site.xinitrc -- /usr/bin/X11/X bc
```

## Files

Item	Description
<code>.xinitrc</code>	Contains the default client script files.
<code> aixterm</code>	Contains the command the client runs if <code>.xinitrc</code> does not exist.
<code>.xserverrc</code>	Contains the default server script.
<code>X</code>	Contains the command the server runs if <code>.xserverrc</code> does not exist.

### Related reference:

“X Command” on page 131

### Related information:

`startx` command

## xkbcomp Command

### Purpose

Compiles XKB keyboard description.

### Syntax

```
xkbcomp [ -a ] [ -C ] [ -dfmts ] [ -I Directory ] [ -l ] [ -m Name ] [ -merge ] [ -o OutputFile ] [ -opt Parts ] [ -R Directory ] [ -synch ] [ -w Level ] [ -xkb ] [ -xkm ] Source [ Destination ]
```

### Description

The `xkbcomp` command is a keymap compiler that converts a description of an XKB keymap into one of several output formats. The most common use for `xkbcomp` is to create a compiled keymap file (`.xkm` extension) which can be read directly by XKB-capable X servers or utilities. The keymap compiler can also produce C header files or XKB source files. The C header files produced by `xkbcomp` can be included by X servers or utilities that need a built-in default keymap. The XKB source files produced by `xkbcomp` are fully resolved and can be used to verify that the files which typically make up an XKB keymap are merged correctly or to create a single file which contains a complete description of the keymap.

The *Source* may specify an X display, or an `.xkb` or `.xkm` file; unless explicitly specified, the format of *destination* depends on the format of the source. Compiling a `.xkb` (keymap source) file generates a `.xkm` (compiled keymap file) by default. If the source is a `.xkm` file or an X display, `xkbcomp` generates a keymap source file by default.

If the *Destination* is an X display, the keymap for the display is updated with the compiled keymap.

The name of the *destination* is usually computed from the name of the source, with the extension replaced as appropriate. When compiling a single map from a file which contains several maps, **xkbcomp** constructs the destination file name by appending an appropriate extension to the name of the map to be used.

## Flags

Item	Description
<b>-a</b>	Shows all keyboard information, reporting implicit or derived information as a comment. Only affects <b>.xkb</b> format output.
<b>-C</b>	Produces a C header file as output ( <b>.h</b> extension).
<b>-dflts</b>	Computes the defaults for any missing components, such as key names.
<b>-I Directory</b>	Specifies the top-level directories to be searched for files included by the keymap description.
<b>-l</b>	List maps that specify the <i>map</i> pattern in any files listed on the command line.
<b>-m Name</b>	Specifies a map to be compiled from an file with multiple entries.
<b>-merge</b>	Merges the compiled information with the map from the server.
<b>-o Name</b>	Specifies a name for the generated output file. The default is the name of the source file with an appropriate extension for the output format.
<b>-opt Parts</b>	Specifies a list of optional parts. Compilation errors in any optional parts are not fatal. <i>Parts</i> may consist of any combination of the letters <b>c, g, k, s, t</b> which specify the compatibility map, geometry, keycodes, symbols, and types, respectively.
<b>-R Directory</b>	Specifies the root directory for relative path names.
<b>-synch</b>	Forces synchronization for X requests.
<b>-w Level</b>	Controls the reporting of warnings during compilation. A warning level of 0 disables all warnings; a warning level of 10 enables them all.
<b>-xkb</b>	Generates a source description of the keyboard as output ( <b>.xkb</b> extension).
<b>-xkm</b>	Generates a compiled keymap file as output ( <b>.xkm</b> extension).

---

## xkbevd Daemon

### Purpose

XKB event daemon.

### Syntax

```
xkbevd [ -help ] [ -cfg File ] [ -sc Command ] [ -sd Directory ] [ -display Display ] [ -bg ] [ -synch ] [ -v ]
```

### Description

The **xkbevd** event daemon listens for specified XKB events and executes requested commands if they occur. The configuration file consists of a list of event specification/action pairs and/or variable definitions.

An event specification consists of a short XKB event name followed by a string or identifier which serves as a qualifier in parentheses; empty parenthesis indicate no qualification and serve to specify the default command which is applied to events which do not match any of the other specifications. The interpretation of the qualifier depends on the type of the event:

- Bell events match using the name of the bell.
- Message events match on the contents of the message string.

- Slow key events accept any of **press**, **release**, **accept**, or **reject**.

No other events are recognized.

An action consists of an optional keyword followed by an optional string argument. **xkbevd** recognizes the following actions:

- **none**
- **ignore**
- **echo**
- **printEvent**
- **sound**
- **shell**

If the action is not specified, the string is taken as the name of a sound file to be played unless it begins with an exclamation point, in which case it is taken as a shell command.

Variable definitions in the argument string are expanded with fields from the event in question before the argument string is passed to the action processor. The general syntax for a variable is either:

**\$c**

or

**\$(str)**

where *c* is a single character and *str* is a string of arbitrary length. All parameters have both single-character and long names. The list of recognized parameters varies from event to event.

The **ignore**, **echo**, **printEvent**, **sound**, and **shell** actions do what you would expect commands named **ignore**, **echo**, **printEvent**, **sound**, and **shell** to do, except that the **sound** command has only been implemented and tested for SGI machines.

The only currently recognized variables are *soundDirectory* and *soundCommand*.

## Flags

Item	Description
<b>-bg</b>	Tells <b>xkbevd</b> to fork itself and run in the background.
<b>-cfg</b> <i>File</i>	Specifies the configuration file to read. If no configuration file is specified, <b>xkbevd</b> looks for <i>~/xkb/xkbevd.cf</i> and <b>\$(LIBDIR)/xkb/xkbevd.cf</b> in that order.
<b>-display</b> <i>Display</i>	Specifies the display to use. If not present, <b>xkbevd</b> uses <b>\$DISPLAY</b> .
<b>-help</b>	Prints a usage message.
<b>-sc</b> <i>Command</i>	Specifies the command used to play sounds.
<b>-sd</b> <i>Directory</i>	Specifies a top-level directory for sound files.
<b>-synch</b>	Forces synchronization of all X requests. Slow.
<b>-v</b>	Prints more information, including debugging messages. Multiple specifications of <b>-v</b> causes more output.

---

## xkbprint Command

### Purpose

Prints an XKB keyboard description.



## Syntax

```
xkbprint [ -? | -help ] [ -color ] [ -dflts ] [ -diffs ] [ -eps ] [ -fit ] [ -full ] [ -grid Resolution ] [ -if  
FontName ] [ -label Type ] [ -lc Locale ] [ -level1 ] [ -level2 ] [ -lg Group ] [ -ll Level ] [ -mono ] [ -n Number  
] [ -nkg Number ] [ -npk Number ] [ -o File ] [ -R Directory ] [ -pict Which ] Source [ OutputFile ]
```

## Description

The **xkbprint** command generates a printable or encapsulated PostScript description of the XKB keyboard description specified by *Source*. The *Source* can be any compiled keymap, **.xkm** file, that includes a geometry description or an X display specification. If an *OutputFile* is specified, **xkbprint** writes to it. Otherwise, **xkbprint** creates the output file, replacing the extension of the source file with **.ps** or **.eps** depending on the requested format. If the source is a non-local X display, for example **:0**, **xkbprint** appends the appropriate prefix to the display specification, replacing the colon with a - (dash). For a local display, **xkbprint** uses **server-n** where *n* is the number of the display.

## Flags

Item	Description
<b>-?   -help</b>	Prints a usage message.
<b>-color</b>	Prints using the colors specified in the geometry file; by default, <b>xkbprint</b> prints a black-and-white image of the keyboard.
<b>-dflts</b>	Attempts to compute default names for any missing components, such as keys.
<b>-diffs</b>	Shows symbols only where they are explicitly bound.
<b>-eps</b>	Generates an encapsulated PostScript file.
<b>-fit</b>	Fits the keyboard image on the page, this is the default.
<b>-full</b>	Prints the keyboard at full size.
<b>-grid Resolution</b>	Prints a grid with <i>Resolution</i> mm resolution over the keyboard.
<b>-if FontName</b>	Specifies an internal PostScript type 1 font to dump to the specified output file or to <i>fontName.pfa</i> , if no output file is specified. No keyboard description is printed if an internal font is dumped.
<b>-label Type</b>	Specifies the labels to be printed on keys. Valid types are: <ul style="list-style-type: none"><li>• <b>none</b></li><li>• <b>name</b></li><li>• <b>code</b></li><li>• <b>symbols</b></li></ul>
<b>-lc Locale</b>	Specifies a locale in which KeySyms should be resolved.
<b>-level1</b>	Generates a level 1 PostScript.
<b>-level2</b>	Generates a level 2 PostScript.
<b>-lg Group</b>	Prints symbols in keyboard groups starting from <i>Group</i> .
<b>-ll Level</b>	Prints symbols starting from shift level <i>Level</i> .
<b>-mono</b>	Generates a black-and-white image of keyboard, this is the default.
<b>-n Number</b>	Prints <i>Number</i> of copies.
<b>-nkg Number</b>	Prints the symbols in <i>Number</i> keyboard groups.
<b>-npk Number</b>	Specifies the <i>Number</i> of keyboard images to print on each page. For EPS files, this specifies the total number of keyboard images to print.
<b>-o File</b>	Writes the output to <i>File</i> .
<b>-R Directory</b>	Use <i>Directory</i> as the root directory; all path names are interpreted relative to <i>Directory</i> .
<b>-pict Which</b>	Controls the use of pictographs instead of keysym names where available. Valid values for <i>Which</i> are: <ul style="list-style-type: none"><li>• <b>all</b></li><li>• <b>none</b></li><li>• <b>common</b> (default).</li></ul>
<b>-synch</b>	Forces synchronization for X requests.

Item	Description
<code>-w</code> <i>Level</i>	Sets warning level. <ul style="list-style-type: none"> <li>• 0 for no warning</li> <li>• 10 for all warnings</li> </ul>

**Related reference:**

“xkbcomp Command” on page 174

## xlock Command

### Purpose

Locks the local X display until a password is entered.

### Syntax

```
xlock [ -batchcount Number ] [ -bg Color ] [ -delay Users ] [ -display Display ] [ -fg Color ] [
-font FontName ] [ -info TextString ] [ -invalid TextString ] [ -mode ModeName ] [ +mono |
-mono ] [ -username TextString ] [ -nice Level ] [ +nolock | -nolock ] [ -password TextString ] [
+remote | -remote ] [ +allowaccess | -allowaccess ] [ +allowroot | -allowroot ] [ +echokeys |
-echokeys ] [ +enablesaver | -enablesaver ] [ -help ] [ -saturation Value ] [ -timeout Seconds ] [
+usefirst | -usefirst ] [ +v | -v ] [ -validate TextString ]
```

### Description

The **xlock** command locks the X server until the user enters a password at the keyboard. While the **xlock** command is running, all new server connections are refused. The screen saver is disabled, the mouse cursor is turned off, the screen is blanked, and a changing pattern is displayed. If a key or a mouse button is pressed, a prompt asks for the password of the user who started the **xlock** command.

If the correct password is typed, the screen is unlocked and the X server is restored. When typing the password, Ctrl-U and Ctrl-H are active as kill and erase, respectively. To return to the locked screen, click in the small icon version of the changing pattern.

To function properly, **xlock** needs to run with root permission since the operating system restricts access to the password and access control files. To give **xlock** root permission, perform the following steps:

1. Log in as root.
2. Go to the directory that contains the **xlock** program file.
3. Run these two commands:
  - a. **chown root xlock**
  - b. **chmod u+s xlock**

### Flags

Item	Description
<b>-batchcount</b> <i>Number</i>	Sets the number of things to do per batch. <i>Number</i> refers to different things depending on the mode: <p><b>qix</b> Refers to the number of lines rendered in the same color.</p> <p><b>hop</b> Refers to the number of pixels rendered in the same color.</p> <p><b>image</b> Refers to the number of sunlogos on screen at once.</p> <p><b>swarm</b> Refers to the number of bees</p>
<b>-bg</b> <i>Color</i>	<b>life and blank</b> Does not apply.
<b>-delay</b> <i>Number</i>	Sets the color of the background on the password screen. Sets the speed at which a mode operates to the number of microseconds to delay between batches of <b>hop</b> along pixels, <b>qix</b> lines, <b>life</b> generations, <b>image</b> bits, and <b>swarm</b> motions.
<b>-display</b> <i>Display</i>	In the <b>blank</b> mode, it is important to set this to a small number because the keyboard and mouse are only checked after each delay. A delay of zero would needlessly consume the processing unit while checking for mouse and keyboard input in a tight loop since the <b>blank</b> mode has no work to do. Sets the X11 display to lock. The <b>xlock</b> command locks all available screens on the server and restricts you to locking only a local server, such as <b>unix:0</b> , <b>localhost:0</b> , or <b>:0</b> (unless you set the <b>-remote</b> flag).
<b>-fg</b> <i>Color</i>	Sets the color of the text on the password screen.
<b>-font</b> <i>FontName</i>	Sets the font to be used on the prompt screen.
<b>-help</b>	Prints a brief description of available options.
<b>-info</b> <i>TextString</i>	Defines an informational message. The default is Enter password to unlock; select icon to lock.
<b>-invalid</b> <i>TextString</i>	Specifies an password message. The default is Invalid login.
<b>-mode</b> <i>ModeName</i>	Specifies one the following six display modes: <p><b>blank</b> Displays a black screen.</p> <p><b>hop</b> Displays the real plane fractals from the September, 1986 issue of <i>Scientific American</i>.</p> <p><b>image</b> Displays several randomly appearing sun logos.</p> <p><b>life</b> Displays Conway's game of life.</p> <p><b>qix</b> Displays spinning lines.</p> <p><b>swarm</b> Displays a swarm of bees following a wasp.</p>
<b>-nice</b> <i>NiceLevel</i>	Sets system nicelevel of the <b>xlock</b> process.
<b>-password</b> <i>TextString</i>	Specifies the password prompt string. The default is Password:.
<b>-saturation</b> <i>Value</i>	Sets saturation of the color ramp. A value of 0 (zero) is grayscale and a value of 1 is very rich color. A value of 0.4 is a medium pastel.
<b>-timeout</b> <i>Seconds</i>	Sets the number of seconds before the password screen times out.
<b>-username</b> <i>TextString</i>	Specifies the message shown in front of the user name. The default is Name:.
<b>-validate</b> <i>TextString</i>	Specifies the message that is shown while validating the password. The default is Validating login....
<b>-/+allowaccess</b>	Allows the disabling of the access control list, but still causes the local server to prompt for a password. If <b>xlock</b> is killed using the <b>-KILL</b> command, the access control list is not lost.  This flag is also needed when running the <b>xlock</b> command remotely since access to the control list is restricted.
<b>-/+allowroot</b>	Allows the root password to unlock the server as well as the user who started the <b>xlock</b> command.
<b>-/+echokeys</b>	Causes the <b>xlock</b> command to echo to screen a '?' (question mark) character for each key typed into the password prompt. The default is no echo.
<b>+/-enablesaver</b>	Enables the default screensaver. It is possible to set delay parameters long enough to cause phosphor burn on some displays. This flag can be used as an added precaution.
<b>+/-mono</b>	Causes the <b>xlock</b> command to display monochrome (black and white) pixels rather than the default colored ones on color displays.

Item	Description
<code>+/-nolock</code>	Causes the <b>xlock</b> command to only draw the patterns and not to lock the display. A keypress or a mouse click terminates the screen saver.
<code>+/-remote</code>	Allows remote locking of X11 servers. This flag should be used with care. It is intended mainly to lock X11 terminals that cannot run the <b>xlock</b> command locally. If you lock a workstation other than your own, that person will need your password to unlock it. The <b>-remote</b> option does not disable your ability to toggle to another shell.
<code>+/-usefirst</code>	Allows using the keystroke which obtained the password screen as the first input character in the password. The default ignores the first keystroke.
<code>+/-v</code>	Minus prefix enables the verbose mode to tell which options the <b>xlock</b> command is going to use. The plus prefix is the default.

---

## xlsfonts Command

### Purpose

Displays the font list for X-Windows.

### Syntax

```
xlsfonts [ -display Host:Display ] [ -l [ l [ l ] ] ] [ -m ] [ -C ] [ -l ] [ -w Width ] [ -n Columns ] [ -u ] [ -o ] [ -fn Pattern ]
```

### Description

The **xlsfonts** command lists the fonts that match a specified *Pattern* parameter. Use the wildcard character "\*" (asterisk) to match any sequence of characters (including none), and the "?" (question mark) to match any single character. If no pattern is given, "\*" is assumed.

**Note:** The "\*" and "?" characters must be placed within quotation marks to prevent them from being expanded by the shell.

You can use flags to specify servers, number and width of columns to print, size of font listings, whether the output should be sorted, and whether to use **OpenFont** instead of **ListFonts**.

### Flags

**Note:** Using the **-l** (lowercase L) flag of the **xlsfonts** command can tie up your server for a long time. This is typical of single-threaded non-preemptable servers, and not a program error.

Item	Description
<b>-l</b>	Indicates that listings should use a single column. This flag is the same as the <b>-n 1</b> flag.
<b>-C</b>	Indicates that listings should use multiple columns. This flag is the same as the <b>-n 0</b> flag.
<b>-display</b> <i>Host:Display</i>	Identifies the X Server to contact by specifying the host name and display number.
<b>-fn</b> <i>Pattern</i>	Specifies the fontname <i>Pattern</i> that <b>xlsfonts</b> will list.
<b>-l</b> [ <b>l</b> [ <b>l</b> ] ]	(lowercase L) Indicates that medium, long, and very long listings, respectively, should be generated for each font.
<b>-m</b>	Indicates that long listings should also print the minimum and maximum bounds of each font.
<b>-n</b> <i>Columns</i>	Specifies the number of columns to use to display the output. By default, the <b>xlsfonts</b> command tries to fit as many columns of font names into the number of characters specified by the <b>-w</b> <i>Width</i> flag.
<b>-o</b>	Instructs the <b>xlsfonts</b> command to perform <b>OpenFont</b> (and <b>QueryFont</b> , if appropriate) instead of <b>ListFonts</b> . The <b>-o</b> flag is useful if the <b>ListFonts</b> or <b>ListFontsWithInfo</b> fails to list a known font, as is the case with some scaled font systems.

Item	Description
<code>-u</code>	Indicates that the output should remain unsorted.
<code>-w Width</code>	Specifies the width in characters that should be used to determine how many columns to print. The default is 79.

## Environment Variable

Item	Description
<code>DISPLAY</code>	Gets the default host and display to use.

## Examples

- To specify a medium-sized list of each font, use a lowercase L and enter:  

```
xlsfonts -l
```
- To specify a three-column list of each font, enter:  

```
xlsfonts -n 3
```
- To display all fonts with the string `iso8859` within their names, enter:  

```
xlsfonts -ll "*"iso8859"*
```
- To list all fonts with `rom1` plus one following character in their names, enter:  

```
xlsfonts rom1"?"
```

This obtains a listing similar to:

```
rom10 rom11 rom14 rom16 rom17
```

### Related reference:

“xset Command” on page 212

## xmbind Command

### Purpose

Configures virtual key bindings.

### Syntax

```
xmbind [ -display Host:Display:ScreenID ] [ FileName ]
```

### Description

The **xmbind** command is an X Windows System client that configures the virtual key bindings for AIXwindows applications. This action is performed by the **mwm** command at its startup, so the **xmbind** client is only needed when **mwm** is not in use or when you want to change bindings without restarting **mwm**. If a file is specified, its contents are used as the virtual key bindings. If a file is not specified, the **.motifbind** file in the user's home directory is used. If this file is not found, the **xmbind** command loads the default virtual key bindings.

### Flags

Item	Description
<b>-display</b> <i>Host:Display:ScreenID</i>	Specifies the display to use. The <b>-display</b> option has the following parameters:
<i>Host</i>	Specifies the host name of a valid system on the network. Depending on the situation, this could be the host name of the user or the host name of a remote system.
<i>Display</i>	Specifies the number (usually 0) of the display on the system on which the output is to be displayed.
<i>ScreenID</i>	Specifies the number of the screen where the output is to be displayed. This number is 0 for single-screen systems.

## Parameters

Item	Description
<i>FileName</i>	Specifies the file containing bindings for virtual mouse and key events.

## Exit Status

This command returns the following exit values:

Item	Description
0	Indicates successful completion.
>0	Indicates an error occurred.

---

## xmkmf Command

### Purpose

Creates a **Makefile** from an **Imakefile**.

### Syntax

```
xmkmf [ -a ] [ TopDir [ CurDir ] ]
```

### Description

The **xmkmf** command creates a **Makefile** from an **Imakefile** shipped with third-party software. When invoked with no arguments or variables in a directory containing an **Imakefile** file, the **imake** command runs with arguments appropriate for your system (configured into **xmkmf** when X was built) and generates a **Makefile**.

### Flag

Item	Description
<b>-a</b>	First builds the <b>Makefile</b> in the current directory, then automatically executes <b>make Makefiles</b> , <b>make includes</b> , and <b>make depend</b> . This is how to configure software that is outside of the MIT X build tree.

### Variables

Specify *TopDir* and *CurDir* if you are working inside the MIT X build tree (highly unlikely unless you are an X developer).

Item	Description
------	-------------

<i>TopDir</i>	Specify as the relative path name from the current directory to the top of the build tree.
---------------	--

<i>CurDir</i>	Specify as a relative path name from the top of the build tree to the current directory.
---------------	--

The *CurDir* variable is required if the current directory has subdirectories; otherwise, the **Makefile** will not be able to build the subdirectories. If a *TopDir* variable is given in its place, **xmkmf** assumes nothing is installed on your system and searches for files in the build tree instead of using the installed versions.

### Related information:

imake command

make command

---

## xmwlm Command

### Purpose

Provides recording of system performance or WLM metrics.

### Syntax

```
xmwlm [ -d recording_dir ] [ -n recording_name ] [ -t trace_level ] [ -L ]
```

### Description

The **xmwlm** agent provides recording capability for a limited set of local system performance metrics. These include common CPU, memory, network, disk, and partition metrics typically displayed by the **topas** command. Daily recordings are stored in the **/etc/perf/daily** directory. The **topasout** command is used to output these recordings in raw ASCII or spreadsheet format. The **xmwlm** agent can also be used to provide recording data from Workload Management (WLM). This is the default format used when **xmwlm** is run without any flags. Daily recordings are stored in the **/etc/perf/wlm** directory. The **wlmon** command can be used to process WLM-related recordings. The **xmwlm** agent can be started from the command line, from a user script, or can be placed near the end of the **/etc/inittab** file. All recordings cover 24-hour periods and are only retained for seven days by default. You can configure the **retain** value in the **/usr/lpp/perfagent/daily.cf** file to change the default recording duration.

### Flags

Item	Description
------	-------------

<b>-d</b> <i>recording_dir</i>	Specifies the output directory for the recording files. The default location is <b>/etc/perf/wlm</b> when <b>xmwlm</b> is run without any flags and <b>/etc/perf/daily</b> when <b>xmwlm</b> is run with the <b>-L</b> flag.
--------------------------------	--

<b>-L</b>	Specifies the collection of topas-like metrics. The metric set is not user configurable.
-----------	--

<b>-n</b> <i>recording_name</i>	Specifies the name for the recording file. By default, <b>xmwlm</b> creates recording files named in an <b>xmwlm.YYMMDD</b> format. For example, if <b>-n myrecording</b> is specified, the recording files will be named <b>myrecording.YYMMDD</b> .
---------------------------------	---

<b>-t</b> <i>trace_level</i>	Specifies a trace level. <b>xmwlm</b> prints various information to a log file in the appropriate <b>/etc/perf</b> subdirectory. The trace level can be set from 1 to 9. More trace data is generated at higher trace levels. This trace data is useful to determine <b>xmwlm</b> recording status and for debugging purposes. The log file name is <b>xmwlm.log1</b> or <b>xmwlm.log2</b> . <b>xmwlm</b> cycles between these two files after a file reaches the maximum allowable size.
------------------------------	---

## Session Recovery by xmwlm

If the **xmwlm** agent is terminated and restarted, **xmwlm** examines the recording files in the appropriate `/etc/perf` subdirectory or in the directory specified by the `-d` flag. If a recording file exists with the current date, **xmwlm** appends data to this file and continues to write to the recording file. Otherwise, a new recording file is created.

### Location

`/usr/bin/xmwlm`

### Files

Item	Description
<code>/usr/bin/xmwlm</code>	Contains the <b>xmwlm</b> agent. The agent is part of the <code>perfagent.tools</code> fileset.

### Related information:

`topas` command

`topasout` command

---

## xmodem Command

### Purpose

Transfers files with the **xmodem** protocol, detecting data transmission errors during asynchronous transmission.

### Syntax

```
xmodem { -s | -r } FileName
```

### Description

The **xmodem** shell command is used with the Asynchronous Terminal Emulation (ATE) program to transfer a file, designated by the *FileName* parameter, using the **xmodem** protocol.

The **xmodem** protocol is an 8-bit transfer protocol to detect data transmission errors and retransmit the data. The workstation sending data waits until the remote system sends a signal indicating it is ready to receive data.

After the receiving system get data, it returns an acknowledgment to the sending system. In the ATE program the receiving system times out if data is not received within 90 seconds after the file transfer is initiated.

Sending and receiving with the **xmodem** command are complementary operations. One system must be set to send while the other is set to receive. Use the **xmodem** command on the remote system in combination with the **send** subcommand or the **receive** subcommand from the ATE Connected Main Menu on the local system.

To interrupt an **xmodem** file transfer, press the Ctrl-X key sequence.

### Note:

1. The DOS operating system terminates each line in an ASCII file with a newline character and a carriage return (Ctrl-M) character. UNIX terminates each line in an ASCII file only with a newline



character. The carriage return characters are preserved when a DOS file is transferred to AIX. The **vi** text editor can be used to remove spurious Ctrl-M characters using the subcommand

```
:%s/<Ctrl-V><Ctrl-M>//
```

where <Ctrl-V> and <Ctrl-M> each represent a single control character that is typed. However, since Ctrl-V is the default ATE MAINMENU\_KEY, the ATE defaults must be altered in order to issue the **vi** subcommand while logged in via ATE.

2. The **xmodem** file transfer process adds Ctrl-Z characters to the last packet transferred to make the packet 128 bytes long. Most files transferred will, therefore, have Ctrl-Z characters appended to the end. The DOS operating system terminates an ASCII file with a Ctrl-Z character. Every file transferred from DOS to AIX will, therefore, end with at least one Ctrl-Z character. These extra Ctrl-Z characters can be removed with the **vi** text editor.

## Flags

Item	Description
-r	Receives data from the local workstation.
-s	Sends data to the local workstation.

## Examples

### Sending a File with the xmodem Protocol

To send the file `myfile` with the **xmodem** protocol, use the **ate** command and the **connect** or **directory** subcommand to establish a connection to the remote system.

1. After logging in to the remote system and before pressing the MAINMENU\_KEY (usually the Ctrl-V key sequence) to return to ATE on the local system, enter:

```
xmodem -r myfile
```

at the shell command line. The **xmodem** protocol starts receive mode on the remote system.

2. Press the MAINMENU\_KEY to return to ATE on the local system.  
The ATE Connected Main Menu displays.
3. Enter the **send** subcommand at the prompt on the ATE Connected Main Menu:

```
s myfile
```

The **send** subcommand instructs the local system to send `myfile` to the remote system. After transferring the file, the ATE Connected Main Menu displays.

### Receiving a File with the xmodem Protocol

Receive the file `infile` from a remote system using **xmodem** protocol with the **ate** command and the **connect** or **directory** subcommand establishing a connection to the remote system.

1. After logging in to the remote system and before pressing the MAINMENU\_KEY (usually the Ctrl-V key sequence) to return to ATE on the local system, enter:

```
xmodem -s infile
```

at the shell command line. The **xmodem** protocol starts, in send mode, on the remote system.

2. Press the MAINMENU\_KEY to return to ATE on the local system.  
The ATE Connected Main Menu displays.
3. Enter the **receive** subcommand at the prompt on the ATE Connected Main Menu:

```
r infile
```

The **receive** subcommand instructs the local system to receive infile from the remote system. After transferring the file, the ATE Connected Main Menu displays.

## File

Item	Description
<b>ate.def</b>	Contains ATE default values.

### Related information:

ate command

Editing the ATE default file

Asynchronous Terminal Emulation

---

## xmodmap Command

### Purpose

Modifies keymaps in the X Server.

### Syntax

```
xmodmap [ -display Display ] [ -e Expression ] [ -grammar | -help ] [ -n ] [ -pk ] [ -pke ] [ -pm ] [ -pp ] [ -quiet | -verbose ] [ FileName ]
```

### Description

The **xmodmap** command edits and displays the keyboard modifier map and keymap table that client applications use to convert event keycodes into key symbols. It is usually run from the session startup script to configure the keyboard according to the personal tastes of the user.

Every time a keycode expression is evaluated, the server generates a **MappingNotify** event on every client. All of the changes should be batched together and done at one time. Clients that receive keyboard input and ignore **MappingNotify** events will not notice any changes made to keyboard mappings.

The *FileName* parameter specifies a file containing the **xmodmap** command expressions to be run. This file is usually kept in the home directory of the user with a name like **.xmodmaprc**. If no file is specified, input is taken from **stdin**.

The **xmodmap** command program reads a list of expressions and parses them all before attempting to run any of them. This makes it possible to refer to key symbols that are being naturally redefined without having to worry as much about name conflicts.

Item	Description
<b>add</b>	The key symbol names are evaluated as the line is read. This permits you to remove keys from a modifier without worrying about whether they were reassigned.
<b>add</b> <i>ModifierName</i> = <i>KeySymbolName</i> ...	Adds the given key symbols to the indicated modifier map. The key symbol names are evaluated after all input expressions are read to make it easy to write expressions to swap keys.
<b>clear</b> <i>ModifierName</i>	Removes all entries in the modifier map for the given modifier, where the valid names are <b>Shift</b> , <b>Lock</b> , <b>Control</b> , <b>Mod1</b> , <b>Mod2</b> , <b>Mod3</b> , <b>Mod4</b> , and <b>Mod5</b> (case does not matter in modifier names, although it does matter for all other names). For example, <b>clear Lock</b> removes all keys bound to the <b>shift lock</b> modifier.
<b>keycode</b> <i>Number</i> = <i>KeySymbolName</i> ...	Assigns the list of key symbols to the indicated keycode (which can be specified in decimal, hex, or octal and be determined by running the <b>xev</b> program in the <b>/usr/lpp/X11/Xamples/demos</b> directory). Usually only one key symbol is assigned to a given code.

Item	Description
<b>keysym</b> <i>KeySymbolName = KeySymbolName...</i>	The <i>KeySymbolName</i> on the left hand side is translated into matching keycodes used to perform the corresponding set of <b>keycode</b> expressions. The list of keysym names can be found in the keysym database <i>/usr/lib/X11/XKeysymDB</i> or the header file <i>X11/keysymdef.h</i> (without the <i>XK_</i> prefix). Note that if the same keysym is bound to multiple keys, the expression is run for each matching keycode.
<b>pointer = default</b>	Sets the pointer map back to its default settings (such as, button 1 generates a code of 1, button 2 generates a 2, and so forth).
<b>pointer =</b> <i>Button1 Button2 Button3...</i>	Sets the pointer map to contain the indicated button codes. The list always starts with the first physical button.
<b>remove</b> <i>ModifierName = KeySymbolName...</i>	Removes all keys containing the given keysyms from the indicated modifier map. Unlike <b>add</b> , the keysym names are evaluated as the line is read in. This allows for the removal of keys from a modifier without having to worry about whether or not they have been reassigned.

Lines that begin with an ! (exclamation point) are taken as comments.

If you want to change the binding of a modifier key, you must also remove it from the appropriate modifier map.

## Flags

Item	Description
<b>-display</b> <i>Display</i>	Specifies the host and display to use.
<b>-e</b> <i>Expression</i>	Specifies an expression to be run. You can specify any number of expressions from the command line.
<b>-grammar</b>	Prints a help message describing the expression grammar used in files and with the <b>-e Expressions</b> flag prints to standard error.
<b>-help</b>	Prints a brief description of the command line arguments to standard error. This is done whenever an unhandled argument is given to the <b>xmodmap</b> command.
<b>-n</b>	Indicates that the <b>xmodmap</b> command should not change the mappings, but should display what it would do when given this flag.
<b>-pk</b>	Indicates that the current keymap table should print on the standard output.
<b>-pke</b>	Indicates that the current keymap table should be printed on the standard output in the form of expressions that can be fed back to <b>xmodmap</b> . This flag is specific to X11R5.
<b>-pm</b>	Indicates that the current modifier map should print on the standard output.
<b>-pp</b>	Indicates that the current pointer map should print on the standard output.
<b>-quiet</b>	Turns off the verbose logging. This is the default.
<b>-verbose</b>	Indicates that the <b>xmodmap</b> command should print logging information as it parses its input.

## Examples

- The following command reverses the button codes that get generated so that the primary button is pressed using the index finger of the left hand on a 3 button pointer:

```
xmodmap -e "pointer = 1 2 3 4 5"
```
- The following command attaches meta to the multi-language key (sometimes labeled Compose Character). It also takes advantage of the fact that applications that need a Meta key simply need to get the keycode and do not require the key symbol to be in the first column of the keymap table. This means that applications that are looking for a Multi\_key (including the default modifier map) will not notice any change.

```
keysym Multi_key = Multi_key Meta_L
```
- To automatically generate less than and greater than characters when the comma and period keys are shifted, reset the bindings for the comma and period with the following scripts:

```
!
! make shift-, be < and shift-. be >
!
keysym comma = comma less
keysym period = period greater
```

4. To swap the location of the Control and Shift Lock keys, use the following script:

```
!  
! Swap Caps_Lock and Control_L  
!  
remove Lock = Caps_Lock  
remove Control = Control_L  
keySYM Control_L = Caps_Lock  
keySYM Caps_Lock = Control_L  
add Lock = Caps_Lock  
add Control = Control_L
```

---

## xmpeek Command

### Purpose

The **xmpeek** command allows you to query any host about the status of its **xmtopas** daemon.

### Syntax

```
xmpeek [ -a | -l ] [ hostname ]
```

### Description

The **xmpeek** command is used to list down the data consumers that currently have instruments (stat sets) defined with the **xmtopas** daemon, and list down all known data consumers by the **xmtopas** daemon . The **xmpeek** command is also used to print down all the available SPMI statistics for any given host.

### Flags

Item	Description
-a	If this flag is specified, one line is listed for each data consumer known by the daemon. If this flag is not used, only data consumers that currently have instruments (stat sets) defined with the daemon are listed. This flag is optional.
-l	Lists all the available statistics of the remote host. This flag is optional.
hostname	If the hostname is specified, the daemon on the named host is asked. If no host name is specified, the daemon on the local host is asked.

### Examples

The following is an example of the output from the **xmpeek** program:

```
Statistics for xmtopas daemon on *** birte ***  
Instruments currently defined: 1  
Instruments currently active: 1  
Remote monitors currently known: 2  
--Instruments--- Values Packets  
  
Defined Active Active Sent Internet Protocol  
Address Port Hostname  
-----  
1 1 16 3,344 129.49.115.208 3885 xtra
```

Output from **xmpeek** can take two forms.

The first form is a line that informs you that the **xmtopas** daemon is not feeding any data-consumer programs. This form is used if no statsets are defined with the daemon and no command flags are supplied.

The second form includes at least as much as is shown in the preceding example, except that the single detail line for the data consumer on host **xtra** is shown only if either the **-a** flag is used or if the data

consumer has at least one instrument (statset) defined with the daemon. Note that **xmpeek** itself appears as a data consumer because it uses the Remote Statistics Interface (RSI) API to contact the daemon. Therefore, the output always shows at least one known monitor.

In the fixed output, first the name of the host where the daemon is running is shown. Then follows three lines giving the totals for current status of the daemon. In the above example, you can see that only one instrument is defined and that it's active. You can also see that two data consumers are known by the daemon, but that only one of them has an instrument defined with the daemon in **birte**. Obviously, this output was produced without the **-a** flag.

An example of more activity is shown in the following sample output from **xmpeek**. The output is produced with the command:

```
xmpeek -a birte
```

Notice that some detail lines show zero instruments defined. Such lines indicate that an **are\_you\_there** message was received from the data consumer but that no states were ever defined or that any previously defined states were erased.

```
Statistics for smeared daemon on *** birte ***
  Instruments currently defined: 16
  Instruments currently active: 14
  Remote monitors currently known: 6
--Instruments--- Values Packets  Internet Protocol
Defined Active  Active  Sent      Address      Port  Hostname
  8      8      35    10,232    129.49.115.203 4184  birte
  6      4      28     8,322    129.49.246.14 3211  umbra
  0      0      0       0       129.49.115.208 3861  xtra
  1      1      16    3,332    129.49.246.14 3219  umbra
  0      0      0       0       129.49.115.203 4209  birte
  1      1      16     422     129.49.115.208 3874  xtra
-----
 16      14      95    22,308
```

Notice that the same host name may appear more than once. This is because every running copy of **xmperf** and every other active data-consumer program is counted and treated as a separate data consumer, each identified by the port number used for UDP packets as shown in the **xmpeek** output.

The second detail shows that one particular monitor on host **umbra** has six instruments defined but only four active. This would happen if a remote **xmperf** console has been opened but is now closed. When you close an **xmperf** console, it stays in the Monitor menu of the **xmperf** main window and the definition of the instruments of that console remains in the tables of the data-supplier daemon but the instruments are not active.

---

## xmscheck Command

### Purpose

The **xmscheck** command is available to pre-parse a recording configuration file and to determine how the running **xmtopas** daemon is configured for recording.

### Syntax

```
xmscheck [ file_name ]
```



---

## xmtopas Command

### Purpose

The **xmtopas** daemon acts as a recording facility and is controlled through the **xmtopas.cf** configuration file. This daemon simultaneously provides near real-time network-based data monitoring and local recording on a given node.

### Syntax

```
xmtopas [ -v ] [ -b UDP_buffer_size ] [ -i min_remote_interval ] [ -l remove_consumer_timeout ] [ -m supplier_timeout ] [ -p trace_level ] [ -s max_logfile_size ] [ -t keep_alive_limit ] [ -x xmtopas_execution_priority ]
```

### Description

The **xmtopas** command is always started from **inetd** daemon. Therefore, command line options must be specified on the line defining the **xmtopas** command to the **inetd** daemon in the **/etc/inetd.conf** file.

#### Note:

1. If you modify the **xmtopas** entry in the **inetd.conf** file on the Virtual IO Server (VIOS), the new entry is added after migration. You must replace the **xmtopas** entry to a default type, such as **xmtopas -p3**, before migration.
2. For any dynamic configuration changes to the logical partition, the **xmtopas** must be restarted to reflect the changes.

### Flags

Item	Description
<b>-v</b>	Causes parsing information for the <b>xmtopas</b> recording configuration file to be written to the <b>xmtopas</b> log file.
<b>-b</b> <i>UDP_buffer_size</i>	Defines the size of the buffer used by the daemon to send and receive UDP packets. The buffer size must be specified in bytes and can be from 4,096 to 16,384 bytes. The buffer size determines the maximum number of data values that can be sent in one <b>data_feed</b> packet. The default buffer size is 4096 bytes, which allows for up to 124 data values in one packet.
<b>-i</b> <i>min_remote_interval</i>	Defines the minimum interval in milliseconds with which data feeds can be sent. Default is 500 milliseconds. A value between 100 and 5,000 milliseconds can be specified. Any value specified is rounded to a multiple of 100 milliseconds. Whichever minimum remote interval is specified causes all requests for data feeds to be rounded to a multiple of this value.
<b>-l</b> <i>remove_consumer_timeout</i>	Sets the <b>time_to_live</b> after feeding of statistics data has ceased as described in section Life and Death of <b>xmtopas</b> . Must be followed by a number of minutes. A value of 0 (zero) minutes causes the daemon to stay alive forever. The default <b>time_to_live</b> is 15 minutes.
<b>-m</b> <i>supplier_timeout</i>	When a dynamic data-supplier is active, this value sets the number of seconds of inactivity from the DDS before the SPMI assumes the DDS is dead. When the timeout value is exceeded, the <b>SiShGoAway</b> flag is set in the shared memory area and the SPMI disconnects from the area. If this flag is not given, the timeout period is set to 90 seconds.  The size of the timeout period is kept in the SPMI common shared memory area. The value stored is the maximum value requested by any data consumer program, including the <b>xmtopas</b> command.
<b>-p</b> <i>trace_level</i>	Sets the trace level, which determines the types of events written to the <b>/var/perf/xmtopas.log1</b> log file or the <b>/var/perf/xmtopas.log2</b> log file. This flag must be followed by a digit from 0 to 9, with 9 being the most detailed trace level. Default trace level is 0 (zero), which disables tracing and logging of events but logs error messages.

Item	Description
<code>-s max_logfile_size</code>	Specifies the approximate maximum size of the log files. At least every <b>time_to_live</b> minutes, it is checked if the currently active log file is bigger than the <i>max_logfile_size</i> value. If so, the current log file is closed and logging continues to the alternate log file, which is first reset to zero length. The two log files are <code>/var/perf/xmtopas.log1</code> and <code>/var/perf/xmtopas.log2</code> . Default maximum file size is 100,000 bytes. You cannot make the <i>max_logfile_size</i> value smaller than 5,000 or larger than 10,000,000 bytes.
<code>-t keep_alive_limit</code>	Sets the <i>keep_alive_limit</i> value must be followed by a number of seconds from 60 to 900 (1 to 15 minutes). The default is 300 seconds (5 minutes).
<code>-x xmtopas_execution_priority</code>	Sets the execution priority of the <b>xmtopas</b> command. Use this option if the default execution priority of the <b>xmtopas</b> command is unsuitable in your environment. Generally, the daemon should be given as high execution priority as possible (a smaller number gives a higher execution priority).  On systems other than IBM RS/6000® systems, the <code>-x</code> flag is used to set the nice priority of the <b>xmtopas</b> command. The nice priority is a value from -20 to 19. The default is -20.

## Files used by the xmtopas command

You can specify the following entries in the **xmtopas.res** file:

```
docec:<arguments>
docluster:cluster=<cluster configuration file>
```

### Example

```
docec: availmem=5 unavailmem=2
docluster: cluster=/etc/perf/xmtopasagg.cf
```

The following new fields are added to the **docec** entry in the **xmtopas.res** file to get the Hardware Management Console (HMC) details:

```
managedsys=[Managed system name under which this partition is configured]
hmc=[HMC name under which this partition is configured]
```

If the HMC platform cannot be configured for automatic queries, the global data fields that are not available to the local partition can be set by using the following options:

```
availmem = [Total amount of memory allocated to all partitions, in GB]
unavailmem = [Total amount of memory unallocated from the HMC, in GB]
availprocessor = [Total number of physical processors allocated for all partitions]
unavailprocessor = [Total number of physical processors unallocated from the HMC]
poolsize = [Defined Pool Size required if HMC Processor Utilization Authority restricts access]
partitions = [Number of partitions defined on the HMC]
reconfig = [Number of seconds between checking for HMC configuration changes.
Allowed values are 30, 60, 90, 120, 180, 240, 300 seconds.
The default is 60 seconds.]
```

### Example

```
docec: hmc=hmc.mac.in.ibm.com managedsys=cec1
```

---

## xntpd Daemon

### Purpose

Starts the Network Time Protocol (NTP) daemon.



## Syntax

```
xntpd [ -a ] [ -b ] [ -d ] [ -D Level ] [ -m ] [-x] [ -c ConfigFile ] [ -e AuthenticationDelay ] [ -f DriftFile ] [ -k KeyFile ] [ -l LogFile ] [ -o TraceFile ] [ -p pidFile ] [ -r BroadcastDelay ] [ -s StatsDirectory ] [ -t TrustedKey ] [ -v SystemVariable ] [ -V SystemVariable ]
```

## Description

The **xntpd** daemon sets and maintains a UNIX system time-of-day in compliance with Internet standard time servers. The **xntpd** daemon is a complete implementation of the Network Time Protocol (NTP) version 3 standard, as defined by RFC 1305, and also retains compatibility with version 1 and 2 servers as defined by RFC 1059 and RFC 1119, respectively. The **xntpd** daemon does all computations in fixed point arithmetic and does not require floating point code.

The **xntpd** daemon reads from a configuration file (*/etc/ntp.conf* is the default) at startup time. You can override the configuration file name from the command line. You can also specify a working, although limited, configuration entirely on the command line, eliminating the need for a configuration file. Use this method when configuring the **xntpd** daemon as a broadcast or multicast client, that determines all peers by listening to broadcasts at runtime. You can display the **xntpd** daemon internal variables with the **ntpq** command (Network Time Protocol (NTP) query program). You can alter configuration options with the **xntpd** command.

The **xntpd** daemon operates in several modes, including symmetric active/passive, client/server and broadcast/multicast. A broadcast/multicast client can automatically discover remote servers, compute one-way delay correction factors and configure itself automatically. This mode makes it possible to deploy a group of workstations without specifying a configuration file or configuration details specific to its environment.

**Note:** When operating in a client mode, the **xntpd** daemon will exit with an error if no configured servers are within 1000 seconds of local system time. Use the **date** or **ntpdate** command to set the time of a bad skewed system before starting **xntpd**.

## Flags

Item	Description
<b>-a</b>	Runs in authenticate mode
<b>-b</b>	Listens for broadcast NTP and synchronizes to them if available.
<b>-c ConfigFile</b>	Specifies the name of an alternate configuration file.
<b>-d</b>	Specifies debugging mode. This flag may occur multiple times (maximum of 10), with each occurrence indicating greater detail of display.
<b>-D Level</b>	Specifies debugging level directly (value from 1 to 10).
<b>-e AuthenticationDelay</b>	Specifies the time, in seconds, it takes to compute the NTP encryption field on this computer.
<b>-f DriftFile</b>	Specifies the location of the drift file.
<b>-k KeyFile</b>	Specifies the location of the file which contains the NTP authentication keys.
<b>-l LogFile</b>	(lowercase L) Specifies the use of a log file instead of logging to syslog.
<b>-m</b>	Listens for multicast messages and synchronizes to them if available. Assumes multicast address 224.0.1.1.
<b>-o TraceFile</b>	Specifies trace file name (default is <b>stderr</b> ).
<b>-p pidFile</b>	Specifies the name of the file to record the daemon's process id. There is no default.
<b>-r BroadcastDelay</b>	Specifies the default delay (in seconds) if the calibration procedure fails. Normally, the <b>xntpd</b> daemon automatically compensates for the network delay between the broadcast/multicast server and the client.
<b>-s StatsDirectory</b>	Specifies the directory to use for creating statistics files.
<b>-t TrustedKey</b>	Adds the specified key number to the trusted key list.
<b>-v SystemVariable</b>	Adds the specified system variable
<b>-V SystemVariable</b>	Adds the specified system variable listed by default.
<b>-x</b>	Makes small time adjustments. (SLEWING)

## Reference Clock Support

For the purposes of configuration, the **xntpd** daemon treats reference clocks in a manner analogous to normal NTP peers as much as possible. It refers to reference clocks by address, same as a normal peer is, though it uses an invalid IP address to distinguish them from normal peers. AIX supports one type of reference clock, based on the system clock (type 1).

Reference clock addresses are of the form `127.127.Type.Unit` where *Type* is an integer denoting the clock type and *Unit* indicates the type-specific unit number. You configure reference clocks by using a server statement in the configuration file where the *HostAddress* is the clock address. The *key*, *version* and *ttl* options are not used for reference clock support.

Reference clock support provides the **fudge** command, which configures reference clocks in special ways. This command has the following format:

```
fudge 127.127.Type.Unit [ time1 Seconds ] [ time2 Seconds ] [ stratum Integer ] [ refid Integer ]  
[ flag1 0 | 1 ] [ flag2 0 | 1 ] [ flag3 0 | 1 ] [ flag4 0 | 1 ]
```

The **time1** and **time2** options are in fixed point seconds and used in some clock drivers as calibration constants.

The **stratum** option is a number in the range zero to 15 and used to assign a nonstandard operating stratum to the clock. Since the **xntpd** daemon adds one to the stratum of each peer, a primary server ordinarily displays stratum one. In order to provide engineered backups, use the **stratum** option to specify the reference clock stratum as greater than zero. Except where noted, this option applies to all clock drivers.

The **refid** option is an ASCII string in the range one to four characters and used to assign a nonstandard reference identifier to the clock.

The binary flags: **flag1**, **flag2**, **flag3** and **flag4** are for customizing the clock driver. The interpretation of these values, and whether they are used at all, is a function of the needs of the particular clock driver.

## Exit Status

This command returns the following exit values:

Item	Description
0	Successful completion.
>0	An error occurred.

## Security

Access Control: You must have root authority to run this command.

Auditing Events: N/A

## Examples

1. To start the **xntpd** daemon, enter:  
`startsrc -s xntpd`
2. To stop the **xntpd** daemon, enter:  
`stopsrc -s xntpd`
3. To use the authentication key file `/etc/ntp.new.keys` when running the **xntpd** daemon, enter:  
`/usr/sbin/xntpd -k /etc/ntp.new.keys`

## Files

Item	Description
<code>/usr/sbin/xntpd</code>	Contains the <code>xntpd</code> daemon.
<code>/etc/ntp.conf</code>	Contains the default configuration file.
<code>/etc/ntp.drift</code>	Contains the default drift file.
<code>/etc/ntp.keys</code>	Contains the default key file.

### Related reference:

“xntpd Command”

### Related information:

ntpq command

---

## xntpd Command

### Purpose

Starts the query/control program for the Network Time Protocol daemon, `xntpd`.

### Syntax

```
xntpd [ -i ] [ -l ] [ -n ] [ -p ] [ -s ] [ -c SubCommand ] [ Host ... ]
```

### Description

The `xntpd` command queries the `xntpd` daemon about its current state and requests changes to that state. It runs either in interactive mode or by using command-line arguments. The `xntpd` command interface displays extensive state and statistics information. Nearly all the configuration options that can be specified at start-up using the `xntpd` daemon's configuration file, can also be specified at run-time using the `xntpd` command.

If you enter the `xntpd` command with one or more request flags, the NTP servers running on each of the hosts specified (or defaults to local host) receive each request. If you do not enter any request flags, the `xntpd` command tries to read commands from standard input and run them on the NTP server running on the first host specified or on the local host by default. It prompts for subcommands if standard input is the terminal.

The `xntpd` command uses NTP mode 7 packets to communicate with the NTP server and can query any compatible server on the network that permits it.

The `xntpd` command makes no attempt to retransmit requests, and will time-out requests if the remote host does not respond within a suitable time.

Specifying a flag other than `-i` or `-n` sends the queries to the specified hosts immediately. Otherwise, the `xntpd` command attempts to read interactive format commands from standard input.

### Flags

Item	Description
<b>-c</b> <i>SubCommand</i>	Specifies an interactive format command. This flag adds <i>SubCommand</i> to the list of commands to run on the specified hosts. You can enter multiple <b>-c</b> flags.
<b>-i</b>	Specifies interactive mode. Standard output displays prompt and standard input reads commands.
<b>-l</b>	(lowercase L) Displays a list of the peers known to the servers. This is the same as the <b>listpeers</b> subcommand.
<b>-n</b>	Displays all host addresses in dotted decimal format (0.0.0.0) rather than the canonical host names.
<b>-p</b>	Displays a list of the peers known to the server and a summary of their state. This is the same as the <b>peers</b> subcommand.
<b>-s</b>	Displays a list of the peers known to the server and a summary of their state but in a format different from the <b>-p</b> flag. This is the same as the <b>dmpeers</b> subcommand.

## Parameters

Item	Description
<i>Host ...</i>	Specifies the hosts.

## xntpdc Internal Subcommands

You can run a number of interactive format subcommands entirely within the **xntpdc** command that do not send NTP mode 7 requests to a server. The following subcommands can only be used while running the **xntpdc** query program.

### Interactive Format Subcommands

Interactive format subcommands consist of a keyword followed by zero to four arguments. You only need to type enough characters of the full keyword to uniquely identify the subcommand. The output of a subcommand goes to standard output, but you can redirect the output of individual subcommands to a file by appending a greater-than sign (>), followed by a file name, to the command line.

Item	Description
<b>?</b> [ <i>SubCommand</i> ]	Displays command usage information. When used without <i>SubCommand</i> , displays a list of all the <b>xntpdc</b> command keywords. When used with <i>SubCommand</i> , displays function and usage information about the command.
<b>help</b> [ <i>SubCommand</i> ]	Same as the <b>?</b> [ <i>SubCommand</i> ] subcommand.
<b>delay</b> <i>Milliseconds</i>	Specifies the time interval to add to timestamps included in requests that require authentication. This subcommand enables unreliable server reconfiguration over long delay network paths or between machines whose clocks are unsynchronized. If you enter this subcommand without an argument, it prints the current setting for this subcommand.
<b>host</b> <i>HostName</i>	Specifies the host to send queries to. <i>HostName</i> may be either a host name or a numeric address. If you enter this subcommand without an argument, it prints the current setting for this subcommand.
<b>hostnames</b> <b>yes</b>   <b>no</b>	Specifies whether to display the host name ( <b>yes</b> ) or the numeric address ( <b>no</b> ). The default is <b>yes</b> unless the <b>-n</b> flag is used. If you enter this subcommand without an argument, it prints the current setting for this subcommand.
<b>keyid</b> <i>Number</i>	Specifies the server key number to use to authenticate configuration requests. If you enter this subcommand without an argument, it prints the current setting for this subcommand.
<b>passwd</b>	Prompts you to type in the NTP server authentication password to use to authenticate configuration requests.
<b>quit</b>	Exits the <b>xntpdc</b> query program.
<b>timeout</b> <i>Milliseconds</i>	Specifies the time-out period for responses to server queries. The default is 8000 milliseconds. If you enter this subcommand without an argument, it prints the current setting for this subcommand.

## Query Subcommands

The **xntpdc** query subcommands result in sending NTP mode 7 packets containing requests to the server. These subcommands are read-only (they do not modify the server configuration state).

<b>Item</b>	<b>Description</b>
<b>clkbug</b> <i>ClockPeerAddress</i> [ <i>Addr2</i> ] [ <i>Addr3</i> ] [ <i>Addr4</i> ]	Displays debugging information for a reference clock driver. Some clock drivers provide this information that is mostly undecodable without a copy of the driver source in hand.
<b>clockbug</b> <i>ClockPeerAddress</i> [ <i>Addr2</i> ] [ <i>Addr3</i> ] [ <i>Addr4</i> ]	Displays information concerning a peer clock. The values obtained provide information on the setting of fudge factors and other clock performance information.
<b>dmpeers</b>	Displays a list of peers for which the server is maintaining state, along with a summary of that state. Identical to the output of the <b>peers</b> subcommand except for the character in the leftmost column. Characters only are displayed beside peers that were included in the final stage of the clock selection algorithm.  The possible character in the leftmost column are: <ul style="list-style-type: none"> <li>. Indicates that this peer was cast off in the falseticker detection.</li> <li>+ Indicates that the peer made it through.</li> <li>* Denotes the peer the server is currently synchronizing with.</li> </ul>
<b>iostats</b>	Displays statistics counters maintained in the input-output module.
<b>kerninfo</b>	Displays kernel phase-lock loop operating parameters. This information is available only if the kernel of the hosts being generated has been specially modified for a precision timekeeping function.
<b>listpeers</b>	Displays a brief list of the peers for which the server is maintaining state. These include all configured peer associations as well as those peers whose stratum is such that the server considers them to be possible future synchronization candidates.
<b>loopinfo</b> [ <b>oneline</b>   <b>multiline</b> ]	Displays the values of selected loop filter variables. The loop filter is the part of NTP that adjusts the local system clock. The <b>offset</b> is the last offset given to the loop filter by the packet processing code. The <b>frequency</b> is the frequency error of the local clock in parts-per-million (ppm). The <b>poll adjust</b> controls the stiffness (resistance to change) of the phase-lock loop and the speed at which it can adapt to oscillator drift. The <b>watchdog timer</b> is the number of elapsed seconds since the last sample offset given to the loop filter. The <b>oneline</b> and <b>multiline</b> options specify the format to display this information. The <b>multiline</b> option is the default.
<b>memstats</b>	Displays statistics counters related to memory allocation code.
<b>monlist</b>	Displays traffic counts collected and maintained by the monitor facility.

Item	Description
<b>peers</b>	<p>Displays a list of peers for which the server is maintaining state, along with a summary of that state. Summary information includes:</p> <ul style="list-style-type: none"> <li>• address of the remote peer,</li> <li>• reference ID (0.0.0.0 for an unknown reference ID),</li> <li>• the stratum of the remote peer (a stratum of 16 indicates the remote peer is unsynchronized),</li> <li>• the polling interval (seconds),</li> <li>• the reachability register (octal), and</li> <li>• the current estimated delay, offset and dispersion of the peer (seconds).</li> </ul> <p>The character in the left margin indicates the mode this peer entry is in:</p> <ul style="list-style-type: none"> <li>+ symmetric active.</li> <li>- symmetric passive.</li> <li>= remote server polled in client mode.</li> <li>^ server is broadcasting to this address.</li> <li>~ remote peer is sending broadcasts.</li> <li>* marks the peer the server is currently synchronizing to.</li> </ul> <p>The contents of the host field may be a host name, an IP address, a reference clock implementation name with its parameter or REFCLK (<i>ImplementationNumber, Parameter</i>). Only IP addresses display when using <b>hostnames no</b>.</p>
<b>pstats</b> <i>PeerAddress</i> [ <i>Addr2</i> ] [ <i>Addr3</i> ] [ <i>Addr4</i> ]	<p>Displays per-peer statistic counters associated with the specified peers.</p>
<b>reslist</b>	<p>Displays the server's restriction list which may help to understand how the restrictions are applied.</p>
<b>sysinfo</b>	<p>Displays a variety of system state variables related to the local server. All except the last four lines are described in the NTP Version 3 specification, RFC 1305. The system flags show various system flags, some of which can be set and cleared by the <b>enable</b> and <b>disable</b> configuration statements. The <b>stability</b> is the residual frequency error remaining after applying the system frequency correction. You use it for maintenance and debugging. In most architectures, this value will initially decrease from as high as 500 ppm to a nominal value in the range .01 to 0.1 ppm. If it remains high for some time after starting the daemon, something may be wrong with the local clock, or the value of the kernel variable <i>Tick</i> may be incorrect. The <b>broadcastdelay</b> shows the default broadcast delay, as set by the <b>broadcastdelay</b> configuration statement, while the <b>authdelay</b> shows the default authentication delay, as set by the <b>authdelay</b> configuration statement.</p>
<b>sysstats</b>	<p>Displays statistics counters maintained in the protocol module.</p>
<b>timerstats</b>	<p>Displays statistics counters maintained in the timer/event queue support code.</p>

## Runtime Configuration Requests Subcommands

The server authenticates all requests that cause state changes in the server by using a configured NTP key. The server can also disable this facility by not configuring a key. You must make the key number and the corresponding key known to the **xtnpdc** command. You can do this by using the **keyid** and **passwd** subcommands, which prompts at the terminal for a password to use as the encryption key. The **xtnpdc** command will also prompt you automatically for both the key number and password the first

time you give a subcommand that would result in an authenticated request to the server. Authentication not only verifies that the requester has permission to make such changes, but also protects against transmission errors.

Authenticated requests always include a timestamp in the packet data, as does the computation of the authentication code. The server compares this timestamp to the time at which it receives the packet.

The server rejects the request if they differ by more than 10 seconds. This makes simple replay attacks on the server, by someone able to overhear traffic on your LAN, much more difficult. It also makes it more difficult to request configuration changes to your server from topologically remote hosts. While the reconfiguration facility works well with a server on the local host, and may work adequately between time-synchronized hosts on the same LAN, it works very poorly for more distant hosts. So, if you choose reasonable passwords, take care in the distribution and protection of keys and apply appropriate source address restrictions, the run-time reconfiguration facility should provide an adequate level of security.

The following subcommands all make authenticated requests.

Item	Description
<b>addpeer</b> <i>PeerAddress</i> [ <i>Keyid</i> ] [ <i>Version</i> ] [ <b>prefer</b> ]	Adds a configured peer association operating in symmetric active mode at the specified address. You may delete an existing association with the same peer or simply convert an existing association to conform to the new configuration when using this subcommand. If the <i>Keyid</i> is a nonzero integer, all outgoing packets to the remote server will have an authentication field attached encrypted with this key. To specify no authentication, enter <i>Keyid</i> as 0 or leave blank. The values for <i>Version</i> can be 1, 2 or 3, with 3 as the default. The <b>prefer</b> option indicates a preferred peer used primarily for clock synchronization if possible. The preferred peer also determines the validity of the PPS signal. If the preferred peer is suitable for synchronization, so is the PPS signal.
<b>addserver</b> <i>PeerAddress</i> [ <i>Keyid</i> ] [ <i>Version</i> ] [ <b>prefer</b> ]	Same as the <b>addpeer</b> subcommand, except that the operating mode is client.
<b>addtrap</b> <i>Address</i> [ <i>Port</i> ] [ <i>Interface</i> ]	Sets a trap for asynchronous messages at the specified address and port number for sending messages with the specified local interface address. If you do not specify the port number, the value defaults to 18447. If you do not specify the interface address, the value defaults to the source address of the local interface.
<b>authinfo</b>	Displays information concerning the authentication module, including known keys and counts of encryptions and decryptions performed.
<b>broadcast</b> <i>PeerAddress</i> [ <i>Keyid</i> ] [ <i>Version</i> ]	Same as the <b>addpeer</b> subcommand, except that the operating mode is broadcast. The <i>PeerAddress</i> can be the broadcast address of the local network or a multicast group address assigned to NTP (224.0.1.1).
<b>clrtrap</b> <i>Address</i> [ <i>Port</i> ] [ <i>Interface</i> ]	Clears a trap for asynchronous messages at the specified address and port number for sending messages with the specified local interface address. If you do not specify the port number, the value defaults to 18447. If you do not specify the interface address, the value defaults to the source address of the local interface.
<b>delrestrict</b> <i>Address Mask</i> [ <b>ntpport</b> ]	Deletes the matching entry from the restrict list.
<b>disable</b> <i>Option ...</i>	Disables various server options. Does not affect options not mentioned. The <b>enable</b> subcommand describes the options.
<b>enable</b> <i>Option ...</i>	Enables various server options. Does not affect options not mentioned. You can specify one or more of the following values for <i>Option</i> :
<b>auth</b>	Causes the server to synchronize with unconfigured peers only if the peer has been correctly authenticated using a trusted key and key identifier. The default for this option is disable (off).
<b>bclient</b>	Causes the server to listen for a message from a broadcast or multicast server, following which an association is automatically instantiated for that server. The default for this argument is disable (off).
<b>monitor</b>	Enables the monitoring facility, with default enable (on).
<b>pll</b>	Enables the server to adjust its local clock, with default enable (on). If not set, the local clock free-runs at its intrinsic time and frequency offset. This option is useful when the local clock is controlled by some other device or protocol and NTP is used only to provide synchronization to other clients.
<b>stats</b>	Enables statistics facility filegen, with default enable (on).

<b>Item</b>	<b>Description</b>
<b>fudge</b> <i>PeerAddress</i> [ <i>Time1</i> ] [ <i>Time2</i> ] [ <i>Stratum</i> ] [ <i>Refid</i> ]	Provides a way to set certain data for a reference clock.  <i>Time1</i> and <i>Time2</i> are in fixed point seconds and used in some clock drivers as calibration constants.  <i>Stratum</i> is a number in the range zero to 15 and used to assign a nonstandard operating stratum to the clock.  <i>Refid</i> is an ASCII string in the range one to four characters and used to assign a nonstandard reference identifier to the clock.
<b>monitor</b> <b>yes</b>   <b>no</b>	Enables or disables the monitoring facility. A <b>monitor no</b> subcommand followed by a <b>monitor yes</b> subcommand is a good way of resetting the packet counts.
<b>readkeys</b>	Purges the current set of authentication keys and obtains a new set by rereading the keys file specified in the <b>xntpd</b> configuration file. This allows you to change encryption keys without restarting the server.
<b>reset</b> <i>Module</i>	Clears the statistics counters in various modules of the server. You can specify one or more of the following values for <i>Module</i> : <b>io</b> , <b>sys</b> , <b>mem</b> , <b>timer</b> , <b>auth</b> , <b>allpeers</b> .
<b>restrict</b> <i>Address Mask Option ...</i>	Adds the values of <i>Option</i> to an existing restrict list entry, or adds a new entry to the list with the specified <i>Option</i> . The <b>mask</b> option defaults to 255.255.255.255, meaning that <i>Address</i> is treated as the address of an individual host. You can specify one or more of the following values for <i>Option</i> :  <b>ignore</b> Ignore all packets from hosts that match this entry. Does not respond to queries nor time server polls.  <b>limited</b> Specifies that these hosts are subject to client limitation from the same net. Net in this context refers to the IP notion of net (class A, class B, class C, etc.). Only accepts the first <b>client_limit</b> hosts that have shown up at the server and that have been active during the last <b>client_limit_period</b> seconds. Rejects requests from other clients from the same net. Only takes into account time request packets. Private, control, and broadcast packets are not subject to client limitation and therefore do not contribute to client count. The monitoring capability of the <b>xntpd</b> daemon keeps a history of clients. When you use this option, monitoring remains active. The default value for <b>client_limit</b> is 3. The default value for <b>client_limit_period</b> is 3600 seconds.  <b>lowpriotrap</b> Declare traps set by matching hosts to low priority status. The server can maintain a limited number of traps (the current limit is 3), assigned on a first come, first served basis, and denies service to later trap requestors. This parameter modifies the assignment algorithm by allowing later requests for normal priority traps to override low priority traps.  <b>nomodify</b> Ignore all NTP mode 6 and 7 packets that attempt to modify the state of the server (run time reconfiguration). Permits queries that return information.  <b>nopeer</b> Provide stateless time service to polling hosts, but not to allocate peer memory resources to these hosts.  <b>noquery</b> Ignore all NTP mode 6 and 7 packets (information queries and configuration requests) from the source. Does not affect time service.  <b>noserve</b> Ignore NTP packets whose mode is not 6 or 7. This denies time service, but permits queries.  <b>notrap</b> Decline to provide mode 6 control message trap service to matching hosts. The trap service is a subsystem of the mode 6 control message protocol intended for use by remote event logging programs.  <b>notrust</b> Treat these hosts normally in other respects, but never use them as synchronization sources.  <b>ntpport</b> Match the restriction entry only if the source port in the packet is the standard NTP UDP port (123).
<b>setprecision</b> <i>Precision</i>	Sets the precision that the server advertises. <i>Precision</i> should be a negative integer in the range -4 through -20.
<b>traps</b>	Displays the traps set in the server.
<b>trustkey</b> <i>Keyid ...</i>	Adds one or more keys to the trusted key list. When you enable authentication, authenticates peers with trusted time using a trusted key.



Item	Description
<b>unconfig</b> <i>PeerAddress</i> [ <i>Addr2</i> ] [ <i>Addr3</i> ] [ <i>Addr4</i> ]	Removes the configured bit from the specified peers. In many cases deletes the peer association. When appropriate, however, the association may persist in an unconfigured mode if the remote peer is willing to continue on in this fashion.
<b>unrestrict</b> <i>Address Mask Option</i> ...	Removes the specified options from the restrict list entry indicated by <i>Address</i> and <i>Mask</i> . The <b>restrict</b> subcommand describes the values for <i>Option</i> .
<b>untrustkey</b> <i>Keyid</i> ...	Removes one or more keys from the trusted key list.

## Exit Status

This command returns the following exit values:

Item	Description
0	Successful completion.
>0	An error occurred.

## Security

Access Control: You must be part of the system group to run this command.

**Attention RBAC users and Trusted AIX users:** This command can perform privileged operations. Only privileged users can run privileged operations. For more information about authorizations and privileges, see Privileged Command Database in *Security*. For a list of privileges and the authorizations associated with this command, see the **lssecattr** command or the **getcmdattr** subcommand.

Auditing Events: N/A

Displays per-peer statistic counters associated with the specified peers.

## Examples

- To start the query/control program for the Network Time Protocol daemon, enter:  
xntpd
- To display the statistic counters of the peer at address 127.127.1.0 on host 9.3.149.107, enter:  
xntpd -c "pstats 127.127.1.0" 9.3.149.107

Output similar to the following is displayed:

```
remote host: LOCAL(0)
local interface: 127.0.0.1
time last received: 49s
time until next send: 15s
reachability change: 818s
packets sent: 13
packets received: 13
bad authentication: 0
bogus origin: 0
duplicate: 0
bad dispersion: 4
bad reference time: 0
candidate order: 1
```

## Files

Item	Description
/usr/sbin/xntpd	Contains the <b>xntpd</b> command.

#### Related information:

ntp command

“xntpd Daemon” on page 192

RBAC

Trusted AIX

## xpr Command

### Purpose

Formats a window dump file for output to a printer.

### Syntax

```
xpr [ -append FileName [ -noff ] | -output FileName ] [ -landscape | -portrait ] [ -compact ] [
-cutoff Level ] [ -density Dpi ] [ -gray { 2 | 3 | 4 } ] [ -header String ] [ -height Inches ] [
-left Inches ] [ -nosition ] [ -plane PlaneNumber ] [ -psfig ] [ -report ] [ -rv ] [ -scale Scale ] [
-split Number ] [ -top Inches ] [ -trailer String ] [ -width Inches ] [ -device Device ] [ ImageFile ]
```

### Description

The **xpr** command uses a window dump file produced by the **xwd** utility as input and formats the dump file for output on all printers supported by the hardware. If you do not specify a file argument, the **xpr** command uses standard input. By default, the **xpr** command prints the largest possible representation of the window on the output page.

The **xpr** command options allow you to add headers and trailers, specify margins, adjust the scale and orientation, and append multiple window dumps to a single output file. Output is to standard output unless the **-output** flag is specified.

### Flags

Item	Description
<b>-append</b> <i>FileName</i>	Specifies a file name previously produced by the <b>xpr</b> command to which the window is to append. (This flag is not supported on PostScript printers.)
<b>-compact</b>	Uses simple run-length encoding for compact representation of windows with many white pixels. This flag compresses white space but not black space, so it is not useful for reverse-video windows.  (This flag supports PostScript, LIPS II+, and LIPSIII output only.)
<b>-cutoff</b> <i>Level</i>	Changes the intensity level where colors are mapped to black or white for monochrome output on a LaserJet printer. The <i>Level</i> variable is expressed as a percentage of full brightness. Fractions are acceptable.

<b>Item</b>	<b>Description</b>
<b>-device</b> <i>Device</i>	Specifies the device on which the file prints. The <b>xpr</b> command supports the following printers:  <b>3812 or pp</b> IBM PP3812  <b>4207</b> Proprinter  <b>5201</b> IBM Quietwriter 1 model 2  <b>5202</b> IBM Quietwriter 2  <b>jprinter</b> IBM Japanese Printer (Japanese data stream)  <b>ljet</b> HP LaserJet and IBM Laser Printer  <b>ps</b> PostScript printers (this is the default)  <b>lips2</b> Canon LaserShot LIPS II+ mode  <b>lips3</b> Canon LaserShot LIPS III mode
<b>-density</b> <i>Dpi</i>	Indicates the dots-per-inch (dpi) density that the HP printer uses. 300 dpi is the default. Allowable densities are 300, 150, 100, and 75 dpi.
<b>-gray</b> <i>Number</i>	Specifies gray-scale conversion to a color image, rather than mapping to a black-and-white image. The <i>Number</i> variable must be one of the following:  <b>2</b> 2 x 2 conversion  <b>3</b> 3 x 3 conversion  <b>4</b> 4 x 4 conversion  This conversion doubles, triples, or quadruples, respectively, the effective width and height of the image. <b>Note:</b> This option is valid only for PostScript printers.
<b>-header</b> <i>String</i>	Specifies a header string to print above the window.
<b>-height</b> <i>Inches</i>	Specifies the maximum height of the page.
<i>ImageFile</i>	Contains the captured bitmap of the image. If you do not specify the <i>ImageFile</i> parameter, the <b>xpr</b> command uses standard input.
<b>-landscape</b>	Forces the window to print in landscape mode. (The display is laid out with the windows being wider than they are high.) By default, a window prints so that its longest side follows the long side of the paper.
<b>-left</b> <i>Inches</i>	Specifies the left margin in inches. Fractions are acceptable. By default, this flag prints the window on the center of the page.
<b>-noff</b>	When specified in conjunction with the <b>-append</b> flag, the window is displayed on the same page as the previous window. (This flag is not supported on PostScript printers.)
<b>-noposition</b>	Causes the header, trailer, and image positioning command generation to be bypassed for the LaserJet printer.
<b>-output</b> <i>FileName</i>	Specifies an output file name. If you do not specify this option, the <b>xpr</b> command uses standard output.
<b>-plane</b> <i>PlaneNumber</i>	Specifies which bit plane to use in an image. The default uses the entire image and maps values into black and white based on color intensities. This option is not supported for the LaserJet printer.
<b>-portrait</b>	Forces the window to print in portrait mode. (The display is laid out with the windows being higher than they are wide.) By default, a window prints so that its longest side follows the long side of the paper.
<b>-psfig</b>	Suppresses translation of the PostScript picture to the center of the page.
<b>-report</b>	Prints out statistics to standard error about the window <i>ImageFile</i> parameter.
<b>-rv</b>	Forces the window to print in reverse video.
<b>-scale</b> <i>Scale</i>	Affects the size of the window on the page. PostScript printers are able to translate each bit in a window pixel map into a grid of a specified size. For example, each bit might translate into a 3 x 3 grid. To specify a 3 x 3 grid, enter <b>-scale 3</b> . By default, a window prints with the largest scale that fits on the page for the specified orientation. If you do not specify a device, the aspect ratio can vary.
<b>-split</b> <i>Number</i>	Splits a window into several pages. This might be necessary for very large windows that would otherwise cause the printer to overload and print the page in an obscure manner. (This flag is not supported on PostScript or HP Laserjet printers.)

<b>Item</b>	<b>Description</b>
<b>-top</b> <i>Inches</i>	Specifies the top margin for the window in inches. Fractions are acceptable. By default, this flag prints the window on the center of the page.
<b>-trailer</b> <i>String</i>	Specifies a trailer string to print below the window.
<b>-width</b> <i>Inches</i>	Specifies the maximum width of the page. <b>Note:</b> The 4207, 5201, and 5202 printers' images must be recorded by the <b>xwd</b> utility in XYPixmap or XYBitmap format. XYPixmap images are converted into bitmaps using a thresholding algorithm. For the HP Laserjet printer, multiplane images must be recorded in ZPixmap format. Single plane images may be either XYPixmap, XYBitmap, or ZPixmap formats.

**Related reference:**

“xwd Command” on page 241

“xwud Command” on page 242

## xpreview Command

### Purpose

Displays troff files on an X display.

### Syntax

```
xpreview [ -BackingStore BackingStoreType ] [ -page Number ] [ ToolKitFlag ... ] { File | - }
```

### Description

The **xpreview** command is an AIXwindows 2.1- and Motif2.1-based application that displays output from the **troff** command on an AIXwindows display. The **troff** command output file must be prepared for the devX100 device.

The user interface contains the standard AIXwindows interface controls for calling the root menu, iconifying the window, and setting the window to full screen size. The interface also includes a main window with a scrollable display area for text. Use the pushbuttons for Next, Previous, Goto Page, Print Page, Print File, and Newfile to manipulate the viewing document.

Mouse button three actuates a popup menu for configuring print capabilities. The menu includes an option to set the command line and another to select a printer queue. The command line dialog box expects command line input through the **troff** command. For example,

```
pic -Tibm3816 troff-input-file |tbl|troff -mm -Tibm3816
```

is an acceptable command line. The printer queue option displays a list of configured printer queues. If this option is not selected, the **xpreview** command uses the system-defined default queue.

When you are previewing an input file, the Print Page and Print File buttons require command line input. Note that once a printer queue is selected, it remains selected for the duration of the viewing session, or until an alternate printer queue is selected.

Fonts supported for the devX100 device in European locales are:

- Times New Roman in normal, italic, and bold
- Courier in normal and bold
- Helvetica in normal and bold
- Symbol

The **xpreview** command supports the following font sizes: 8, 10, 14, 18, 24, 30, and 36.

The **xpreview** command does not display files resulting from the **troff** command constructed for a device other than those described in this document.

To preview a file on a certain device, the **xpreview** command requires the fonts found in the following directories:

- **/usr/lib/X11/fonts** directory for files formatted for font files other than Japanese
- **/usr/lib/X11/fonts/JP** for Japanese font files

## Multibyte Support

The **xpreview** command supports multibyte locales. Also, to display Japanese characters, Japanese 16-dot fonts (part of the Japanese BSL package) and 24- and 32-dot fonts (part of the AIXwindows font package) must be installed. To display Korean characters, Korean fonts (part of the Korean BSL package) must be installed.

Japanese support currently includes the following font sets:

- In 16-dot: RomanKn12, Kanji12, and IBM\_JPN12
- In 24-dot: RomanKn17, Kanji17, and IBM\_JPN17
- In 32-dot: RomanKn23, Kanji23, and IBM\_JPN23, or RomanKn23G, Kanji23G, and IBM\_JPN23G

Korean support currently includes the following font sets:

- In 16-dot, EnglHg16 and Hangul16
- In 24-dot, EnglHg24 and Hangul24

## Flags

The **xpreview** command accepts the standard X Toolkit command line flags, as well as the following flags:

Item	Description
-	Requires input to be read from standard input.
<b>-help</b>	Indicates that a brief summary of the allowed command line flags should be printed.
<b>-BackingStore</b> <i>BackingStoreType</i>	The <b>-BackingStore</b> flag causes the server to save the window contents so that when it is scrolled around the viewport, the window is painted from contents saved in server backing store. Redisplays of the drawing window can take up to a second or so. The <i>BackingStoreType</i> parameter can have one of the following values: <b>Always</b> , <b>WhenMapped</b> or <b>NotUseful</b> . <b>Tip:</b> Enter a space between the <b>-BackingStore</b> flag and its <i>BackingStoreType</i> parameter. <b>Requirement:</b> Use of this flag requires that the server be started with backing store enabled.
<b>-page</b> <i>Number</i>	Specifies the page number of the document to be first displayed.

**Item***ToolkitFlag***Description**

The following standard X Toolkit flags are commonly used with the **xpreview** command:

**-bg** *Color*

Specifies the color to use for the background of the window. The default is white.

**-bg** *Color*

Specifies the color to use for the background of the window. The default is white.

**-fg** *Color*

Specifies the color to use for displaying text. The default is black.

**-geometry** *Geometry*

Specifies the preferred size and position of the window.

**-display** *Host:Display*

Specifies the X server to contact.

**-xrm** *ResourceString*

Specifies a resource string to be used.

*File*

Specifies the file to be printed.

**Examples**

1. To build files output by the **troff** command into files that are suitable for use with the **xpreview** command, enter the following commands:

```
troff-TX100 troff-input | xpreview
pic -TX100 pic-troff-input | tbl | troff -man -TX100 | xpreview
```

2. To build files output by the **troff** command into files that are suitable for use with the Japanese language version of the **xpreview** command, enter the following commands:

```
LANG=ja_JP
troff -TX100 troff-input | xpreview -
pic -TX100 pic-troff-input | tbl | troff -man -TX100 \
| xpreview -
```

**Files****Item***/usr/lib/X11/app-defaults/XPreview***Description**

Contains user-configurable applications defaults file.

*/usr/lib/X11/Ja\_JP/app-defaults/XPreview*

Contains user-configurable applications default file for the Japanese (IBM-943) locale.

*/usr/lib/X11/ja\_JP/app-defaults/XPreview*

Contains user-configurable applications default file for the Japanese (IBM-eucJP) locale.

*/usr/lib/X11/ko\_KR/app-defaults/XPreview*

Contains user-configurable applications default file for the Korean locale.

*/usr/lib/X11/zh\_TW/app-defaults/XPreview*

Contains user-configurable applications default file for the Traditional Chinese locale.

*/usr/lib/font/devX100*

Contains troff fonts for devX100 devices.

*/usr/lib/X11/fonts*

Contains X fonts for 100 dpi devices.

*/usr/lib/X11/fonts/JP*

Contains X fonts for multi-byte characters.

*/usr/lib/X11/fonts/JP*

Contains X fonts for Japanese characters.

**Related information:**

troff command

---

# xprofiler Command

## Purpose

Starts Xprofiler, a GUI-based AIX performance profiling tool.

## Syntax

```
xprofiler [ program ] [ -b ] [ -s ] [ -z ] [ -a path ] [ -c file ] [ -L pathname ] [ [ -e function ]... ] [ [ -E function ]... ] [ [ -f function ]... ] [ [ -F function ]... ] [ -disp_max number_of_functions ] [ [ gmon.out ]... ]
```

xprofiler -h | -help

## Description

The **xprofiler** command invokes Xprofiler, a GUI-based AIX performance profiling tool. Xprofiler is used to analyze the performance of both serial and parallel applications. Xprofiler uses data collected by the **-pg** compiling option and presents a graphical representation of the functions in the application in addition to providing textual data in several report windows. These presentation formats are intended to identify the functions which are most processor-intensive.

## Flags

Item	Description
<b>-a</b>	To specify an alternate search path or paths for library files and source code files. If more than one path is specified, the paths must be embraced by "," and each path should be separated by either ":" or space.
<b>-b</b>	Suppresses the printing of the field descriptions for the Flat Profile, Call Graph Profile, and Function Index reports when they are written to a file with the Save As option of the File menu.
<b>-c</b>	Loads a configuration file that contains information to be used to determine which functions will be displayed when Xprofiler is brought up.
<b>-disp_max</b>	Sets the number of function boxes that Xprofiler initially displays in the function call tree. The value supplied with this flag can be any integer between 0 and 5,000. Xprofiler displays the function boxes for the most processor-intensive functions through the number you specify. For instance, if you specify 50, Xprofiler displays the function boxes for the 50 functions in your program that consume the most processor. After this, you can change the number of function boxes that are displayed via the Filter menu options. This flag has no effect on the content of any of the Xprofiler reports.
<b>-e</b>	De-emphasizes the general appearance of the function box or boxes for the specified functions in the function call tree, and limits the number of entries for these function in the Call Graph Profile report. This also applies to the specified function's descendants, as long as they have not been called by non-specified functions. In the function call tree, the function boxes for the specified functions appear greyed-out. Its size and the content of the label remain the same. This also applies to descendant functions, as long as they have not been called by non-specified functions. In the Call Graph Profile report, an entry for the specified function only appears where it is a child of another function, or as a parent of a function that also has at least one non-specified function as its parent. The information for this entry remains unchanged. Entries for descendants of the specified function do not appear unless they have been called by at least one non-specified function in the program.
<b>-E</b>	Changes the general appearance and label information of the function box or boxes for the specified functions in the function call tree. Also limits the number of entries for these functions in the Call Graph Profile report, and changes the processor data associated with them. These results also apply to the specified function's descendants, as long as they have not been called by non-specified functions in the program. In the function call tree, the function box for the specified function appears greyed-out, and its size and shape also changes so that it appears as a square of the smallest allowable size. In addition, the processor time shown in the function box label, appears as 0 (zero). The same applies to function boxes for descendant functions, as long as they have not been called by non-specified functions. This option also causes the processor time spent by the specified function to be deducted from the left side processor total in the label of the function box for each of the specified function's ancestors. In the Call Graph Profile report, an entry for the specified function only appears where it is a child of another function, or as a parent of a function that also has at least one non-specified function as its parent. When this is the case, the time in the self and descendants columns for this entry is set to 0 (zero). In addition, the amount of time that was in the descendants column for the specified function is subtracted from the time listed under the descendants column for the profiled function. As a result, be aware that the value listed in the % time column for most profiled functions in this report will change.

Item	Description
-f	De-emphasizes the general appearance of all function boxes in the function call tree, except for that of the specified function(s) and its descendant(s). In addition, the number of entries in the Call Graph Profile report for the non-specified functions and non-descendant functions is limited. The <b>-f</b> flag overrides the <b>-e</b> flag. In the function call tree, all function boxes except for that of the specified function(s) and its descendant(s) appear greyed-out. The size of these boxes and the content of their labels remain the same. For the specified function(s), and its descendants, the appearance of the function boxes and labels remain the same. In the Call Graph Profile report, an entry for a non-specified or non-descendant function only appears where it is a parent or child of a specified function or one of its descendants. All information for this entry remains the same.
-F	Changes the general appearance and label information of all function boxes in the function call tree except for that of the specified function(s) and its descendants. In addition, the number of entries in the Call Graph Profile report for the non-specified and non-descendant functions is limited, and the processor data associated with them is changed. The <b>-F</b> flag overrides the <b>-E</b> flag. In the function call tree, the function box for the specified function appears greyed-out, and its size and shape also changes so that it appears as a square of the smallest allowable size. In addition, the processor time shown in the function box label, appears as 0 (zero). In the Call Graph Profile report, an entry for a non-specified or non-descendant function only appears where it is a parent or child of a specified function or one of its descendants. The time in the self and descendants columns for this entry is set to 0 (zero). When this is the case, the time in the self and descendants columns for this entry is set to 0 (zero). As a result, be aware that the value listed in the % time column for most profiled functions in this report will change.
-h -help	Writes the Xprofiler usage to STDERR and then exits. The information includes <b>xprofiler</b> command line syntax and a description of Xprofiler runtime options.
-L	Uses an alternate path name for locating shared libraries. If you plan to specify multiple paths, use the Set File Search Paths option of the File menu on the Xprofiler GUI.
-s	If multiple <b>gmon.out</b> files are specified when Xprofiler is started, produces the <b>gmon.sum</b> profile data file. The <b>gmon.sum</b> file represents the sum of the profile information in all the specified profile files. Note that if you specify a single <b>gmon.out</b> file, the <b>gmon.sum</b> file contains the same data as the <b>gmon.out</b> file.
-z	Includes functions that have both zero processor usage and no call counts in the Flat Profile, Call Graph Profile, and Function Index reports. A function will not have a call count if the file that contains its definition was not compiled with the <b>-pg</b> option, which is common with system library files.

## Example

To use **xprofiler**, you must first compile your program (for example, **foo.c**) with **-pg**:

```
xlc -pg -o foo foo.c
```

1. When the program **foo** is executed, one **gmon.out** file will be generated for each processor involved in the execution. To invoke **xprofiler**, enter:

```
xprofiler foo [[gmon.out]...]
```

## Files

Item	Description
<code>/usr/lib/X11/app-defaults/Xprofiler</code>	Location of the <b>xprofiler</b> command.

---

## xrdb Command

### Purpose

X Server resource database utilities.

### Syntax

```
xrdb [ -display Display ] [ -help ] [ -quiet ] [ -retain ] [ -cpp FileName | -nocpp ] [
-D Name=Value ] [ -I Directory ] [ -U Name ] [ -all | -global | -screen | -screens ] [ -n ] [
-edit FileName | -backup String ] | -merge [ FileName ] | -load [ FileName ] | -query |
-remove | symbols ] -override ]
```



## Description

The **xrdb** command gets or sets the contents of the **RESOURCE\_MANAGER** property on the root window of screen 0 or the **SCREEN\_RESOURCES** property on the root window of any or all screens, or everything combined. You normally run this program from your X startup file.

Most X clients use the **RESOURCE\_MANAGER** and **SCREEN\_RESOURCES** properties to get user preferences about color, fonts, and so on for applications. Having this information in the server (where it is available to all clients) instead of on disk solves the problem in previous versions of X that required you to maintain *defaults* files on every machine that you might use. It also allows for dynamic changing of defaults without editing files.

The **RESOURCE\_MANAGER** property specifies resources that apply to all screens of the display. The **SCREEN\_RESOURCES** property on each screen specifies additional (or overriding) resources to be used for that screen. (When there is only one screen, **SCREEN\_RESOURCES** is normally not used; all resources are placed in the **RESOURCE\_MANAGER** property.)

For compatibility, if there is no **RESOURCE\_MANAGER** property defined (either because the **xrdb** command was not run or if the property was removed), the resource manager looks for a file called **.Xdefaults** in your home directory.

The file name (or the standard input if - or no file name is given) is optionally passed through the C preprocessor with the following symbols defined, based on the capabilities of the server being used:

Item	Description
<b>SERVERHOST</b> = <i>Hostname</i>	Specifies the hostname portion of the display to which you are connected.
<b>SRVR</b> <sub><i>name</i></sub>	Turns the <b>SERVERHOST</b> hostname string into a legal identifier. For example <i>my-dpy.lcs.mit.edu</i> becomes <i>SRVR_my_dpy_lcs_mit_edu</i> .
<b>HOST</b> = <i>Hostname</i>	Specifies the hostname portion of the display to which you are connected.
<b>DISPLAY_NUM</b> = <i>num</i>	Specifies the number of the display on the server host.
<b>CLIENTHOST</b> = <i>Hostname</i>	Specifies the name of the host on which xrdb is running.
<b>CLNT</b> <sub><i>name</i></sub>	Turns the <b>CLIENTHOST</b> hostname string into a legal identifier. For example <i>expo.lcs.mit.edu</i> becomes <i>CLNT_expo_lcs_mit_edu</i> .
<b>WIDTH</b> = <i>Number</i>	Specifies the width of the default screen in pixels.
<b>HEIGHT</b> = <i>Number</i>	Specifies the height of the default screen in pixels.
<b>X_RESOLUTION</b> = <i>Number</i>	Specifies the x resolution of the default screen in pixels per meter.
<b>Y_RESOLUTION</b> = <i>Number</i>	Specifies the y resolution of the default screen in pixels per meter.
<b>PLANES</b> = <i>Number</i>	Specifies the number of bit planes (the depth) of the root window of the default screen.
<b>RELEASE</b> = <i>Number</i>	Specifies the vendor release number for the server. The interpretation of this number varies depending on <b>VENDOR</b> .
<b>REVISION</b> = <i>Number</i>	Specifies the X protocol minor version supported by this server (currently 0).
<b>VERSION</b> = <i>Number</i>	Specifies the X protocol major version supported by this server (should always be 11).
<b>VENDOR</b> = <i>Vendor</i>	A string specifying the vendor of the server.
<b>VNDR</b> <sub><i>name</i></sub>	Turns the <b>VENDOR</b> name string into a legal identifier. For example <i>MIT X Consortium</i> becomes <i>VNDR_MIT_X_Consortium</i> .
<b>EXT</b> <sub><i>name</i></sub>	Turns each extension string into a legal identifier. A symbol is defined for each protocol extension supported by the server. For example <i>X3D-PEX</i> becomes <i>EXT_X3D_PEX</i> .
<b>NUM_SCREEN</b> S= <i>num</i>	Specifies the total number of screens.
<b>SCREEN_NUM</b> = <i>num</i>	Specifies the number of current screen. from 0 (zero).
<b>BITS_PER_RGB</b> = <i>Number</i>	Specifies the number of significant bits in an RGB color specification. This is the log base 2 of the number of distinct shades of each primary that the hardware can generate. Note that it is not related to <b>PLANES</b> .
<b>CLASS</b> = <i>VisualClass</i>	Specifies the visual class of the root window of the default screen which is one of the following:

Item	Description
<code>CLASS_visualclass=visualid</code>	Specifies the visual class of the root window in a form can <code>#ifdef</code> on. The value is the numeric id of the visual.
<code>CLASS_visualclass_depth=num</code>	<b>DirectColor, GrayScale, PseudoColor, StaticColor, StaticGray, TrueColor</b> A symbol is defined for each visual supported for the screen. The symbol includes the class of the visual and its depth; the value is the numeric id of the visual. (If more than one visual has the same class and depth, the numeric id of the first one reported by the server is used.)S
<code>COLOR</code>	Defined only if <code>CLASS</code> is one of <b>StaticColor, PseudoColor, TrueColor,</b> or <b>DirectColor.</b>

Comment lines begin with an ! (exclamation mark) and are ignored.

Since `xrdb` can be read from standard input, use it to change the contents of properties directly from a terminal or from a shell script.

## Flags

Item	Description
<code>-all</code>	Indicates that operation is performed on the screen-independent resource property (RESOURCE_MANAGER), as well as the screen-specific property (SCREEN_RESOURCES) on every screen of the display. For example, when used in conjunction with <code>-query</code> , the contents of all properties are output. For <code>-load</code> and <code>-merge</code> , the input file is processed once for each screen. The resources that occur in common in the output for every screen are collected and applied as the screen-independent resources. The remaining resources are applied for each individual per-screen property. This is the default mode of operation. This option is specific to X11R5.
<code>-backup String</code>	Specifies a suffix to append to the file name. Use it with <code>-edit</code> to generate a backup file. <code>-edit</code> is a prerequisite for <code>-backup String</code> .
<code>-cpp FileName</code>	Specifies the pathname of the C preprocessor program to use. Although the <code>xrdb</code> command was designed to use CPP, any program that acts as a filter and accepts the <code>-D</code> , <code>-I</code> , and <code>-U</code> flags can be used.
<code>-DName=Value</code>	Passes through to the preprocessor and defines symbols to use with conditionals such as <code>#ifdef</code> .
<code>-display Display</code>	Specifies the X Server to use. It also specifies the screen to use for the <code>-screen</code> option, and it specifies the screen from which preprocessor symbols are derived for the <code>-global</code> option.
<code>-edit FileName</code>	Indicates that the contents of the specified properties should be edited into the given file, replacing any values listed there. This allows you to put changes you made to your defaults back into your resource file, preserving any comments or preprocessor lines.
<code>-global</code>	Indicates that the operation should only be performed on the screen-independent RESOURCE_MANAGER property. This option is specific to X11R5.
<code>-help</code>	Prints a brief description of the allowed flags.
<code>-IDirectory</code>	( uppercase i ) Passes through to the preprocessor and specifies a directory to search for files referenced with <code>#include</code> .
<code>-load</code>	Indicates that the input is loaded as the new value of the specified properties, replacing the old contents. This is the default action.
<code>-merge</code>	Indicates that the input merges with, instead of replaces, the current contents of the specified properties. This option performs a lexicographic sorted merge of the two inputs, which is probably not what you want, but remains for backward compatibility.
<code>-n</code>	Indicates that changes to the specified properties (when used with <code>-load</code> or <code>-merge</code> ) or to the resource file (when used with <code>-edit</code> ) should be shown on the standard output, but should not be performed. This option is specific to X11R5.
<code>-nocpp</code>	Indicates that the <code>xrdb</code> command should not run the input file through a preprocessor before loading it into properties.
<code>-override</code>	Indicates that the input should be added to, instead of replacing, the current contents of the specified properties. New entries override previous entries.
<code>-query</code>	Indicates that the current contents of the specified properties should print onto the standard output. Note that since preprocessor commands in the input resource file are part of the input file, not part of the property, they do not appear in the output from this flag.
<code>-quiet</code>	Indicates that a warning about duplicate entries should not display. This option is specific to X11R5.

Item	Description
<b>-remove</b>	Indicates that the specified properties should be removed from the server.
<b>-retain</b>	Indicates that the server should be instructed not to reset if the <b>xrdb</b> command is the first client. This should never be necessary under normal conditions, since the <b>xdm</b> and <b>xinit</b> commands always act as the first client. This option is specific to X11R5.
<b>-screen</b>	Indicates that the operation should only be performed on the <b>SCREEN_RESOURCES</b> property of the default screen of the display. This option is specific to X11R5.
<b>-screens</b>	Indicates that the operation should be performed on the <b>SCREEN_RESOURCES</b> property of each screen of the display. For <b>-load</b> and <b>-merge</b> , the input file is processed once for each screen. This option is specific to X11R5.
<b>-symbols</b>	Indicates that the symbols defined for the preprocessor should be printed onto the standard output.
<b>-UName</b>	Passes through to the preprocessor and removes any definitions of this symbol.

## Examples

1. To load a file into the database:  

```
xrdb -load myfile
```
2. To take the contents of the database just loaded and edit or put it into newfile:  

```
xrdb -edit newfile
```

## Files

The **xrdb** command generalizes the `~/.Xdefaults` files.

## xsend Command

### Purpose

Sends secret mail in a secure communication channel.

### Syntax

```
xsend User
```

### Description

The **xsend** command sends messages that can be read only by the intended recipient. This command is similar to the **mail** command, but the mail sent with this command is intended to be secret.

The **xsend** command is used with the **enroll** command and the **xget** command to send secret mail. The **enroll** command sets up the password used to receive secret mail. The **xget** command uses that password to receive the mail.

The **xsend** command reads standard input until an EOF (Ctrl-D) or a . (period) is entered. It then encrypts this text along with some header information and sends it. After sending the encrypted message, the **xsend** command mails a standard mail message to the recipient informing them they have received secret mail.

**Note:** Secret mail can only be sent to local users.

### Examples

1. To send secret mail, enter:  

```
xsend ron
```

When you have issued the **xsend** command with the recipient's name, the mail system is used to enter the text of the message. When you finish entering the message to user ron, press the Enter key, then Ctrl-D or a . (period) to exit the mail editor and send the message. The **xsend** command encrypts the message before it is sent.

- To send a file to another user, enter:

```
xsend lance <proposal
```

In this example, the file proposal is sent to user lance.

## Files

Item	Description
<code>/var/spool/secretmail/*.keys</code>	Contains the encrypted key for <i>User</i> .
<code>/var/spool/secretmail/*.{0-9}</code>	Contains the encrypted mail messages for <i>User</i> .
<code>/usr/bin/xsend</code>	Contains the command executable files.

### Related reference:

“xget Command” on page 169

### Related information:

bellmail command

enroll command

mail command

Mail applications

Sending and receiving secret mail

## xset Command

### Purpose

Sets options for your X-Windows environment.

### Syntax

```
xset [ -display Display ] [ b [ Volume [ Pitch [ Duration ] ] ] | -b | b on | b off ] [
bc | -bc ] c [ Volume ] | -c | c on | c off ] [ [ - | + ] fp [ - | + | = ] Path [ ,Path,
[ ... ] ] ] [ fp default ] [ fp rehash ] [ [ - ] led [ Integer ] ] [ led on | led off ] [ m
[ Accelerator ] [ Threshold ] ] ] [ m [ ouse ] default ] [ p Pixel Color ] [ [ - ] r ] [ r on |
r off ] [ s [ Length [ Period ] ] ] [ s blank | s noblank ] [ s expose | s noexpose ] [ s on |
s off ] [ s activate ] [ s reset ] [ s default ] [ q ]
```

### Description

The **xset** command customizes your X-Windows environment.

### Flags

Item	Description
<b>-display</b> <i>Host:Display</i>	Specifies the X server to use. For more information about servers, see the X command.
<b>b</b> or <b>b on</b>	Turns the bell on. This is the default setting. <b>Note:</b> Not all hardware is able to vary the bell characteristics, but for that which can, all of the <b>b</b> flag permutations and its variables are available.
<b>b</b> [ <i>Volume</i> [ <i>Pitch</i> [ <i>Duration</i> ] ] ]	Specifies the bell volume, pitch, and duration. This flag accepts up to three numeric values.  <i>Volume</i> If only one numeric is given then it is assumed to be <i>Volume</i> . The bell volume is set to that numeric as a percentage of the bell's maximum possible volume dependent on current hardware capabilities.  <i>Pitch</i> The second numeric in hertz values, is the tonal sound of the bell.  <i>Duration</i> The third numeric in milliseconds, is the length of time that the bell rings.
<b>-b</b> or <b>b off</b>	Turns the bell off.
<b>bc</b> or <b>-bc</b>	Controls bug compatibility mode in the server, if possible. A preceding - (dash) disables this mode; otherwise, bug compatibility mode is enabled. The server must support the MIT-SUNDRY-NONSTANDARD protocol extension for the <b>bc</b> flag to work.  New application development should be performed with bug compatibility mode disabled.  The <b>bc</b> flag is provided for pre-X11 Release 4 (X11R4) clients. Some pre-X11R4 clients pass illegal values in various protocol requests. Such clients, when run with an X11R4 server, end abnormally or otherwise fail to operate correctly.  This flag explicitly reintroduces certain bugs into the X server so that such clients still can be run.
<b>c</b> or <b>c on</b> <b>c</b> <i>Volume</i>	Turns on the click. System default. A numeric from 0 to 100 that specifies a percentage of the click's maximum possible volume dependent on current hardware capabilities.
<b>-c</b> or <b>c off</b> <b>fp=Path,...</b>	Turns off the click. Sets the font path to the directories given in the <i>Path</i> parameter. The directories are interpreted by the server, not by the client, and are server-dependent. The server ignores directories that do not contain font databases created by the <b>mkfontdir</b> command. All of the options and variables supported by the <b>fp</b> flag are available.
<b>fp-</b> or <b>-fp</b>	Deletes the font path specified by the <i>Path</i> parameter from the end of the current font path if the - (dash) precedes <b>fp</b> and from the front of the font path if the - (dash) follows <b>fp</b> .
<b>fp+</b> or <b>+fp</b>	Adds the font path specified by the <i>Path</i> parameter to the bottom of font list if the - (dash) precedes <b>fp</b> and from the end of the font path if the - (dash) follows <b>fp</b> .
<b>fp default</b> <b>fp rehash</b>	Resets the font path to the server's default. Causes the server to reread the font databases in the current font path. Usually used only when adding new fonts to a font directory after running <b>mkfontdir</b> to recreate the font database.
<b>led orled on</b> <b>-led</b> <i>Integer</i>	Turns all LEDs on. Turns the LED specified by <i>Integer</i> off. Valid values are between 1 and 32.
<b>led</b> <i>Integer</i>	Turns the LED specified by <i>Integer</i> on. Valid values are between 1 and 32.

Item	Description
-led or led off	Turns all LEDs off. <b>Note:</b> Not all hardware assigns the same <i>Integer</i> variables to the same LED functions.
m	Allows you to control the precision of the mouse or other pointing device. If no variable or the <b>default</b> argument is specified, the system defaults are used. This flag accepts the following optional arguments and parameters:  <i>Acceleration</i> Sets the multiplier for the mouse movement. The value can be specified as an integer or a fraction.  <i>Threshold</i> Sets the minimum number of pixels needed to invoke a movement of the mouse. The value is specified in pixels.  If only one parameter is given, it will be interpreted as the <i>Acceleration</i> parameter.
p	<b>default</b> Uses the system defaults. Controls pixel color values. The root background colors may be changed on some servers by altering the entries for BlackPixel and WhitePixel. Although these values are often 0 and 1, they need not be.  Also, a server may choose to allocate those colors privately, in which case the <b>xset</b> command generates an error. The <b>xset</b> command also generates an error if the map entry is a read-only color.  Valid parameters are:  <i>Pixel</i> Specifies the color map entry number in decimal.  <i>Color</i> Specifies a color. Enables autorepeat. Disables autorepeat. Sets screen saver parameters to the default screen-saver characteristics. Specifies the length of time the server must be inactive for the screen saver to activate. <i>Period</i> specifies the period in which the background pattern must be changed to avoid burn in. The values of <i>Length</i> and <i>Period</i> are specified in seconds. If only one numerical parameter is given, it is read as a <i>Length</i> parameter. Turns the screen saver functions on and off, respectively. Causes the screen saver to activate, even if it has been turned off. Causes the screen saver to deactivate if it was activated. Sets the preference to blank the video (if the hardware can do so) rather than display a background pattern. Sets the preference to display a pattern rather than blank the video. Sets the preference to allow window exposures (the server can freely discard window contents). Sets the preference to disable screen saver unless the server can regenerate the screens without causing exposure events.
r or r on -r or r off s or s default	
s [Length[Period]]	
s on or s off s activate	
s reset s blank	
s noblank	
s expose	
s noexpose	
q	Reports information on the current settings.  These settings will be reset to default values when you log out. <b>Note:</b> Not all X implementations are guaranteed to honor all of these options.

## Examples

1. To set the bell volume to medium, the tone to 50 hertz, and length of time the bell rings to 50 milliseconds:  
`xset b 50,50,50`
2. To set the font path to the `/usr/lib/X11/fonts` directory:  
`xset fp= /usr/lib/x11/fonts`
3. To cause the server to reread the font databases in the current font path:  
`xset fp rehash`
4. To see information on the current settings:  
`xset q`

which produces output similar to the following:

```
Keyboard Control:
  auto repeat: on    key click percent: 0    LED mask: 00000000
  auto repeating keys: 0000000000000000
                      0000000000000000
                      0000000000000000
                      0000000000000000
  bell percent: 50   bell pitch: 400    bell duration: 100

Pointer Control:
  acceleration: 2 = 2 / 1    threshold: 4

Screen Saver:
  prefer blanking: no    allow exposures: no
  timeout: 0    cycle: 0

Colors:
  default colormap: 0x8006e    BlackPixel: 0    WhitePixel: 1

Font Path:
  /usr/lib/X11/fonts/,/usr/lib/X11/fonts/75dpi/,/usr/lib/X11/fonts/100dpi/,/usr/
lib/X11/fonts/oldx10/,/usr/lib/X11/fonts/oldx11/,/usr/lib/X11/fonts/bmug/,/usr/l
ib/X11/fonts/info-mac/,/usr/lib/X11/fonts/JP/,/usr/lib/X11/fonts/misc/
```

---

## xsetroot Command

### Purpose

Sets the root window parameters for the X command.

### Syntax

```
xsetroot [ -bg Color ] [ -cursor CursorFile MaskFile ] [ -cursor_name CursorName ] [ -def ] [
-display Display ] [ -fg Color ] [ -help ] [ -name String ] [ -rv ] [ -bitmap FileName | -gray |
-grey | -mod X Y | -solid Color ]
```

### Description

The `xsetroot` command allows you to tailor the appearance of the background (root) window on a workstation display running X. Normally, you experiment with the `xsetroot` command until you find a personalized look that you like, then put the `xsetroot` command that produces it into your X startup file. If no options are specified or if the `-def` flag is specified, the window is reset to its default state. The `-def` flag can be specified with other flags and only the unspecified characteristics are reset to the default state.

Only one of the background color (tiling) changing flags (`-bitmap`, `-solid`, `-gray`, `-grey`, or `-mod`) can be specified at a time.

## Flags

Item	Description
<b>-bg</b> <i>Color</i>	Uses the <i>Color</i> parameter as the background color.
<b>-bitmap</b> <i>FileName</i>	Uses the bitmap specified in the file to set the window pattern. You can make your own bitmap files (little pictures) using the <b>bitmap</b> program. The entire background is made of repeated tiles of the bitmap.
<b>-cursor</b> <i>CursorFile MaskFile</i>	Changes the pointer cursor to what you want when it is outside of any window. Cursor and mask files are bitmaps (little pictures) that can be made with the <b>bitmap</b> program. You probably want the mask file to be all black until you get used to the way masks work.
<b>-cursor_name</b> <i>CursorName</i>	Changes the pointer cursor to one of the standard cursors from the cursor font.
<b>-def</b>	Resets unspecified attributes to the default values. (Restores the background to the familiar gray mesh and the cursor to the hollow x shape.)
<b>-display</b> <i>Display</i>	Specifies the server connection. See the <b>X</b> command.
<b>-fg</b> <i>Color</i>	Uses the <i>Color</i> parameter as the foreground color. Foreground and background colors are meaningful only with the <b>-cursor</b> , <b>-bitmap</b> , or <b>-mod</b> flags.
<b>-gray</b>	Makes the entire background gray.
<b>-grey</b>	Makes the entire background grey.
<b>-help</b>	Prints a usage message and exits.
<b>-mod</b> <i>X Y</i>	Makes a plaid-like grid pattern on your screen. The <i>X</i> and <i>Y</i> parameters are integers ranging from 1 to 16. Zero and negative numbers are taken as 1.
<b>-name</b> <i>String</i>	Sets the name of the root window to the <i>String</i> parameter. There is no default value. Usually a name is assigned to a window so that the window manager can use a text representation when the window is iconified. This flag is not used because you cannot iconify the background.
<b>-rv</b>	Exchanges the foreground and background colors. Normally the foreground color is black and the background color is white.
<b>-solid</b> <i>Color</i>	Sets the background of the root window to the specified color. This flag is only used on color servers.

---

## xss Command

### Purpose

Improves the security of unattended workstations.

### Syntax

```
xss [ -e CommandString ][ -timeout Seconds ][ -display DisplayPtr ][ -v ][ -fg Color ][ -bg Color ][ -geometry wxh+x+y ]
```

### Description

The **xss** command works with the newly added Massachusetts Institute of Technology (MIT) Screen Saver Extensions in order to implement a user controllable screen saver/lock. This command is designed to improve the security of unattended workstations.

The **xss** command executes a user-specified command string when it receives a screen saver timeout message, or when the user activates the pushbutton. When no user-specified command is given, the **xss** command defaults to the **xlock** command.

**Note:** The **xss** command only uses the newly added MIT Screen Saver Extensions. The **xss** command does not work on an older X server, or when using an older X extension library.



## Flags

Item	Description
<code>-e CommandString</code>	Sets the <code>xss</code> command to execute when either the screen saver times out, or the user activates the pushbutton. Note that if the <code>CommandString</code> parameter value is longer than one word, it must be surrounded by " " (double quotations).
<code>-timeout Seconds</code>	Sets the number of seconds of user inactivity before the screen saver times out, and causes the <code>xss</code> command to run the <code>CommandString</code> parameter.
<code>-display DisplayPtr</code>	Sets the connection to the X11 display.
<code>-v</code>	Turns on verbose mode.
<code>-fg Color</code>	Sets the foreground color of the pushbutton.
<code>-bg Color</code>	Sets the background color of the pushbutton.
<code>-geometry wxh+x+y</code>	Specifies the size and location of the client window.

## Examples

When running remotely and using the `-display` flag for the `xss` command, remember that you may also have to use the `-display` flag option for the command that will be executed by the `xss` command. See the following running remote example:

1. Running remote:  

```
xss -display myhost:0 -e "xlock -remote -display myhost:0"
```
2. Screen saver only:  

```
xss -e "xlock -nolock"
```
3. Simple example:  

```
xss -e xlock
```

---

## xstr Command

### Purpose

Extracts strings from C programs to implement shared strings.

### Syntax

```
xstr [ -v ] [ -c ] [ - ] [ File ]
```

### Description

The `xstr` command maintains a file `strings` into which strings in component parts of a large program are hashed. These strings are replaced with references to this array. This serves to implement shared constant strings, most useful if they are also read-only.

The command:

```
xstr -c File
```

extracts the strings from the C source in the `File` parameter, replacing string references by expressions of the form (`&xstr[number]`) for some number. An appropriate declaration of the `xstr` array is prepended to the file. The resulting C text is placed in the file `x.c`, to then be compiled. The strings from this file are appended into the `strings` file if they are not there already. Repeated strings and strings which are suffixes of existing strings do not cause changes to the file `strings`.

If a string is a suffix of another string in the file but the shorter string is seen first by the `xstr` command, both strings are placed in the file `strings`.

After all components of a large program have been compiled, a file **xs.c** declaring the common **xstr** array space can be created by a command of the form:

```
xstr
```

This **xs.c** file should then be compiled and loaded with the rest of the program. If possible, the array can be made read-only (shared), saving space and swap overhead.

The **xstr** command can also be used on a single file. The command:

```
xstr File
```

creates files **x.c** and **xs.c** as before, without using or affecting any **strings** file in the same directory.

It may be useful to run the **xstr** command after the C preprocessor if any macro definitions yield strings or if there is conditional code which contains strings which may not, in fact, be needed.

The **xstr** command reads from its standard input when the **-** (minus sign) flag is given and does not alter the **strings** file unless the **-c** flag is specified also.

An appropriate command sequence for running the **xstr** command after the C preprocessor is:

```
cc -E name.c | xstr -c -
cc -c x.c
mv x.o name.o
```

The **xstr** command does not touch the file **strings** unless new items are added, thus the **make** command can avoid remaking the **xs.o** file unless truly necessary.

## Flags

Item	Description
<b>-c</b>	Extracts strings from the specified file, and places them in the <b>strings</b> file.
<b>-v</b>	Verbose mode. Tells when strings are found, or new in the <b>strings</b> file.
<b>-</b>	Reads from standard input.

## Examples

1. To extract the strings from the C source in the *File.c* parameter, replacing string references by expressions of the form (**&xstr[number]**):

```
xstr -c File.c
```

An appropriate declaration of the **xstr** array is prepended to the file. The resulting C text is placed in the file **x.c**, to then be compiled.

2. To declare the common **xstr** array space in the **xs.c** file:

```
xstr
```

## Files

Item	Description
<b>strings</b>	File which contains the extracted strings.
<b>x.c</b>	Massaged C source.
<b>xs.c</b>	C source for definition of array xstr.
<b>/tmp/xs*</b>	Temporary file when <b>xstr</b> command does not touch the <b>strings</b> file.
<b>/usr/ccs/bin/mkstr</b>	Contains an executable file.
<b>/usr/ccs/bin/mkstr</b>	Contains an executable file for Berkeley environment.

#### Related information:

mkstr command

## xterm Command

### Purpose

Provides a terminal emulator for the X Window System.

**Note:** The **xterm** command is ported from the Massachusetts Institute of Technology (MIT) X Window System, Version 11, Release 6 with no functional enhancements. The **xterm** command does not have support for globalization. For the localized and internationalized terminal emulator, the user can use the **aixterm** or **dterm** commands.

### Syntax

```
xterm [ --Xtoolkitoption... ] [ -Option ... ]
```

### Description

The **xterm** program is a terminal emulator for the X Window System. It provides DEC VT102 and Tektronix 4014 compatible terminals for programs that cannot use the window system directly. If the underlying operating system supports terminal resizing capabilities, the **xterm** program uses the facilities to notify programs running in the window whenever it is resized.

The VT102 and Tektronix 4014 terminals each have their own window so that you can edit text in one and look at graphics in the other at the same time. To maintain the correct aspect ratio (height/width), Tektronix graphics are restricted to the largest box with a 4014 aspect ratio that will fit in the window. This box is located in the upper left area of the window.

Although both windows might be displayed at the same time, one of them is considered the *active window* for receiving keyboard input and terminal output. This is the window that contains the text cursor. The active window can be chosen through escape sequences, the VT Options menu in the VT102 window, and the Tek Options menu in the 4014 window.

### Emulations

The VT102 emulation is fairly complete, but does not support smooth scrolling, VT52 mode, the flashing character attribute, or the double-wide and double-size character sets. The **termcap** file entries that work with the **xterm** command include **xterm**, **vt102**, **vt100** and **``ansi,"** and the **xterm** command automatically searches the **termcap** file in this order for these entries and then sets the **TERM** and the **TERMCAP** environment variables.

Many of the special **xterm** features might be modified under program control through a set of escape sequences different from the standard VT102 escape sequences.

The Tektronix 4014 emulation is also fairly good. It supports 12-bit graphics addressing, scaled to the window size. Four different font sizes and five different lines types are supported. There is no write-thru or defocused mode support.

The Tektronix text and graphics commands are recorded internally by the **xterm** command and may be written to a file by sending the COPY escape sequence (or through the Tektronix menu, as described in the following sections). The name of the file will be **COPYyy-MM-dd.hh:mm:ss**, where *yy*, *MM*, *dd*, *hh*, *mm*, and *ss* are the year, month, day, hour, minute, and second when the copy is performed (the file is created in the directory that the **xterm** command is started in, or the home directory for a login **xterm**).

## Other Features

The **xterm** command automatically highlights the text cursor when the pointer enters the window (selected) and unhighlights it when the pointer leaves the window (unselected). If the window is the focus window, the text cursor is highlighted no matter where the pointer is located.

In VT102 mode, there are escape sequences to activate and deactivate an alternate screen buffer, which is the same size as the display area of the window. When activated, the current screen is saved and replaced with the alternate screen. Saving of lines scrolled off the top of the window is disabled until the usual screen is restored.

The **termcap** file entry for the **xterm** command allows the **vi** command editor to switch to the alternate screen for editing and to restore the screen on exit.

In either VT102 or Tektronix mode, there are escape sequences to change the name of the windows.

## Options

The **xterm** terminal emulator accepts all of the standard X Toolkit command-line options as well as the following (if the option begins with a + instead of a -, the option is restored to its default value):

Item	Description
<b>-help</b>	Causes the <b>xterm</b> command to print out a message describing its options.
<b>-132</b>	Usually, the VT102 DECCOLM escape sequence that switches between 80- and 132-column mode is ignored. This option causes the DECCOLM escape sequence to be recognized, and the xterm window will resize appropriately.
<b>-ah</b>	Indicates that the <b>xterm</b> command should always highlight the text cursor. By default, the <b>xterm</b> command will display a hollow text cursor whenever the focus is lost or the pointer leaves the window.
<b>+ah</b>	Indicates that the <b>xterm</b> command should do text cursor highlighting based on focus.
<b>-b</b> <i>Number</i>	Specifies the size of the inner border (the distance between the outer edge of the characters and the window border) in pixels. The default is 2.
<b>-cc</b> <i>CharacterClassRange:Value[,...]</i>	Sets classes indicated by the given ranges for use in selecting by words.
<b>-cn</b>	Indicates that newlines should not be cut in line-mode selections.
<b>+cn</b>	Indicates that newlines should be cut in line-mode selections.
<b>-cr</b> <i>Color</i>	Specifies the color to use for the text cursor. The default is to use the same foreground color that is used for text.
<b>-cu</b>	Indicates that the <b>xterm</b> command should work around a bug in the <b>more</b> program that causes it to incorrectly display lines that are exactly the width of the window and are followed by a line beginning with a tab (the leading tabs are not displayed). This option is so named because it was originally thought to be a bug in the <b>urses</b> function cursor motion package.
<b>+cu</b>	Indicates that <b>xterm</b> should not work around the <b>more</b> function bug previously mentioned.

Item	Description
<b>-e</b> <i>Program [Arguments]</i>	Specifies the program (and its command-line arguments) to be run in the xterm window. It also sets the window title and icon name to be the base name of the program being run if neither the <b>-T</b> nor the <b>-n</b> option is given on the command line. <b>Note:</b> This must be the last option on the command line.
<b>-fb</b> <i>Font</i>	Specifies a font to be used when displaying bold text. This font must be the same height and width as the normal font. If only one of the normal or bold fonts is specified, it will be used as the normal font and the bold font will be produced by overstriking this font. The default is to do overstriking of the normal font.
<b>-i</b>	Turns on the <b>useInsertMode</b> resource.
<b>+i</b>	Turns off the <b>useInsertMode</b> resource.
<b>-j</b>	Indicates that the <b>xterm</b> command should do jump scrolling. Usually, text is scrolled one line at a time; this option allows the <b>xterm</b> command to move multiple lines at a time so that it does not fall as far behind. Its use is strongly recommended because it makes the <b>xterm</b> command much faster when scanning through large amounts of text. The VT100 escape sequences for enabling and disabling smooth scrolling as well as the VT Options menu can be used to turn this feature on or off.
<b>+j</b>	Indicates that the <b>xterm</b> command should not do jump scrolling.
<b>-ls</b>	Indicates that the shell that is started in the xterm window is a login shell (in other words, the first character of the <i>ArgumentVector</i> parameter is a dash, indicating to the shell that it should read the user's <b>.login</b> or <b>.profile</b> file).
<b>+ls</b>	Indicates that the shell that is started should not be a login shell (in other words, it will be a usual subshell).
<b>-mb</b>	Indicates that the <b>xterm</b> command should ring a margin bell when the user types near the right end of a line. This option can be turned on and off from the VT Options menu.
<b>+mb</b>	Indicates that the margin bell should not be rung.
<b>-mc</b> <i>Milliseconds</i>	Specifies the maximum time between multiclick selections.
<b>-ms</b> <i>Color</i>	Specifies the color to be used for the pointer cursor. The default is to use the foreground color.
<b>-nb</b> <i>Number</i>	Specifies the number of characters from the right end of a line at which the margin bell, if enabled, will ring. The default is 10.
<b>-rw</b>	Indicates that reverse wraparound should be allowed. This allows the cursor to back up from the leftmost column of one line to the rightmost column of the previous line. This is very useful for editing long shell command lines and is encouraged. This option can be turned on and off from the VT Options menu.
<b>+rw</b>	Indicates that reverse wraparound should not be allowed.

Item	Description
<b>-aw</b>	Indicates that auto wraparound should be allowed. This allows the cursor to automatically wrap to the beginning of the next line when it is at the rightmost position of a line and text is output.
<b>+aw</b>	Indicates that auto wraparound should not be allowed.
<b>-s</b>	Indicates that the <b>xterm</b> command may scroll asynchronously, meaning that the screen does not have to be kept completely up to date while scrolling. This allows the <b>xterm</b> command to run faster when network latencies are high and is typically useful when running across a large Internet or many gateways.
<b>+s</b>	Indicates that the <b>xterm</b> command should scroll synchronously.
<b>-sb</b>	Indicates that some number of lines that are scrolled off the top of the window should be saved and that a scrollbar should be displayed so that those lines can be viewed. This option can be turned on and off from the VT Options menu.
<b>+sb</b>	Indicates that a scrollbar should not be displayed.
<b>-sf</b>	Indicates that Sun Function Key escape codes should be generated for function keys.

Item	Description
<b>+sf</b>	Indicates that the standard escape codes should be generated for function keys.
<b>-si</b>	Indicates that output to a window should not automatically reposition the screen to the bottom of the scrolling region. This option can be turned on and off from the VT Options menu.
<b>+si</b>	Indicates that output to a window should cause it to scroll to the bottom.
<b>-sk</b>	Indicates that pressing a key while using the scrollbar to review previous lines of text should cause the window to be repositioned automatically in the usual position at the bottom of the scroll region.
<b>+sk</b>	Indicates that pressing a key while using the scrollbar should not cause the window to be repositioned.
<b>-sl</b> <i>Number</i>	Specifies the number of lines to save that have been scrolled off the top of the screen. The default is 64.
<b>-t</b>	Indicates that the <b>xterm</b> command should start in Tektronix mode, rather than in VT102 mode. Switching between the two windows is done using the Options menus.
<b>+t</b>	Indicates that the <b>xterm</b> command should start in VT102 mode.
<b>-tm</b> <i>String</i>	Specifies a series of terminal-setting keywords followed by the characters that should be bound to those functions, similar to the <b>stty</b> program. Allowable keywords include: <b>intr</b> , <b>quit</b> , <b>erase</b> , <b>kill</b> , <b>eof</b> , <b>eol</b> , <b>swtch</b> , <b>start</b> , <b>stop</b> , <b>brk</b> , <b>susp</b> , <b>dsusp</b> , <b>rprnt</b> , <b>flush</b> , <b>weras</b> , and <b>Inext</b> . Control characters might be specified as <b>^Character</b> (for example, <b>^c</b> or <b>^u</b> ), and <b>^?</b> may be used to indicate Delete.
<b>-tn</b> <i>Name</i>	Specifies the name of the terminal type to be set in the <b>TERM</b> environment variable. This terminal type must exist in the <b>termcap</b> database and should have <b>li#</b> and <b>co#</b> entries.
<b>-ut</b>	Indicates that the <b>xterm</b> command should not write a record into the <b>/etc/utmp</b> system log file.
<b>+ut</b>	Indicates that the <b>xterm</b> command should write a record into the <b>/etc/utmp</b> system log file.
<b>-vb</b>	Indicates that a visual bell is preferred over an audible one. Instead of ringing the terminal bell whenever the Ctrl+G key sequence signal is received, the window will flash.
<b>+vb</b>	Indicates that a visual bell should not be used.
<b>-wf</b>	Indicates that the <b>xterm</b> command should wait for the window to be mapped the first time before starting the subprocess so that the initial terminal size settings and environment variables are correct. It is the application's responsibility to catch subsequent terminal size changes.
<b>+wf</b>	Indicates that the <b>xterm</b> command should not wait before starting the subprocess.
<b>-C</b>	Indicates that this window should receive console output. This is not supported on all systems. To obtain console output, you must be the owner of the console device, and you must have read and write permission for it. If you are running X windows under <b>xdm</b> on the console screen, you may need to have the session startup and reset programs explicitly change the ownership of the console device in order to get this option to work.
<b>-Sccn</b>	Specifies the last two letters of the name of a pseudoterminal to use in slave mode, plus the number of the inherited file descriptor. The option is parsed <b>``%c%c%d"</b> . This allows the <b>xterm</b> command to be used as an input and output channel for an existing program and is sometimes used in specialized applications.

The following command-line arguments are provided for compatibility with older versions. They may not be supported in the next release as the X Toolkit provides standard options that accomplish the same task.

<b>Item</b>	<b>Description</b>
<code>%geom</code>	Specifies the preferred size and position of the Tektronix window. It is shorthand for specifying the <b>*tekGeometry</b> resource.
<code>#geom</code>	Specifies the preferred position of the icon window. It is shorthand for specifying the <b>*iconGeometry</b> resource.
<code>-T String</code>	Specifies the title for the <b>xterm</b> program's windows. It is equivalent to <b>-title</b> .
<code>-n String</code>	Specifies the icon name for the <b>xterm</b> program's windows. It is shorthand for specifying the <b>*iconName</b> resource. Note that this is not the same as the Toolkit option <b>-name</b> (see the following). The default icon name is the application name.
<code>-r</code>	Indicates that reverse video should be simulated by swapping the foreground and background colors. It is equivalent to <b>-rv</b> .
<code>-w Number</code>	Specifies the width in pixels of the border surrounding the window. It is equivalent to <b>-borderwidth</b> or <b>-bw</b> .

The following standard X Toolkit command-line arguments are commonly used with the **xterm** command:

<b>Item</b>	<b>Description</b>
<code>-bg Color</code>	Specifies the color to use for the background of the window. The default is white.
<code>-bd Color</code>	Specifies the color to use for the border of the window. The default is black.
<code>-bw Number</code>	Specifies the width in pixels of the border surrounding the window.
<code>-fg Color</code>	Specifies the color to use for displaying text. The default is black.
<code>-fn Font</code>	Specifies the font to be used for displaying usual text. The default is fixed.
<code>-name Name</code>	Specifies the application name under which resources are to be obtained, rather than the default executable file name. The <i>Name</i> parameter should not contain . (dot) or * (asterisk) characters.
<code>-title String</code>	Specifies the window title string, which may be displayed by window managers if the user so chooses. The default title is the command line specified after the <b>-e</b> option, if any; otherwise, the application name.
<code>-rv</code>	Indicates that reverse video should be simulated by swapping the foreground and background colors.
<code>-geometry Geometry</code>	Specifies the preferred size and position of the VT102 window; see the <b>X</b> command.
<code>-display Display</code>	Specifies the X server to contact; see the <b>X</b> command.
<code>-xrm ResourceString</code>	Specifies a resource string to be used. This is especially useful for setting resources that do not have separate command-line options.
<code>-iconic</code>	Indicates that the <b>xterm</b> command should ask the window manager to start it as an icon rather than as the usual window.

## Resources

The program understands all of the core X Toolkit resource names and classes as well as:

<b>Item</b>	<b>Description</b>
<code>iconGeometry</code> (class <code>IconGeometry</code> )	Specifies the preferred size and position of the application when iconified. It is not necessarily obeyed by all window managers.
<code>termName</code> (class <code>TermName</code> )	Specifies the terminal type name to be set in the <b>TERM</b> environment variable.
<code>title</code> (class <code>Title</code> )	Specifies a string that may be used by the window manager when displaying this application.
<code>ttyModes</code> (class <code>TtyModes</code> )	Specifies a string containing terminal-setting keywords and the characters to which they may be bound. Allowable keywords include: <b>intr, quit, erase, kill, eof, eol, swtch, start, stop, brk, susp, dsusp, rprnt, flush, weras,</b> and <b>Inext</b> . Control characters may be specified as <i>^Character</i> (for example, <b>^c</b> or <b>^u</b> ) and <b>^?</b> may be used to indicate Delete. This is very useful for overriding the default terminal settings without having run an <b>stty</b> program every time an <b>xterm</b> window is started.

Item	Description
<b>useInsertMode</b> (class <b>useInsertMode</b> )	Forces the use of insert mode by adding appropriate entries to the <b>TERMCAP</b> environment variable. This is useful if the system termcap is broken. The default is <b>false</b> .
<b>utmpInhibit</b> (class <b>UtmpInhibit</b> )	Specifies whether <b>xterm</b> should try to record the user's terminal in <b>/etc/utmp</b> .
<b>sunFunctionKeys</b> (class <b>SunFunctionKeys</b> )	Specifies whether Sun Function Key escape codes should be generated for function keys instead of standard escape sequences.
<b>waitForMap</b> (class <b>WaitForMap</b> )	Specifies whether the <b>xterm</b> command should wait for the initial window map before starting the subprocess. The default is <b>False</b> .

The following resources are specified as part of the **vt100** widget (class **VT100**):

Item	Description
<b>allowSendEvents</b> (class <b>AllowSendEvents</b> )	Specifies whether synthetic key and button events (generated using the X protocol <b>SendEvent</b> request) should be interpreted or discarded. The default is <b>False</b> , meaning they are discarded. Note that allowing such events creates a large security hole.
<b>alwaysHighlight</b> (class <b>AlwaysHighlight</b> )	Specifies whether <b>xterm</b> should always display a highlighted text cursor. By default, a hollow text cursor is displayed whenever the pointer moves out of the window or the window loses the input focus.
<b>appcursorDefault</b> (class <b>AppcursorDefault</b> )	If <b>True</b> , the cursor keys are initially in application mode. The default is <b>False</b> .
<b>appkeypadDefault</b> (class <b>AppkeypadDefault</b> )	If <b>True</b> , the keypad keys are initially in application mode. The default is <b>False</b> .
<b>autoWrap</b> (class <b>AutoWrap</b> )	Specifies whether auto wraparound should be enabled. The default is <b>True</b> .
<b>bellSuppressTime</b> (class <b>BellSuppressTime</b> )	Specifies the number of milliseconds after a bell command is sent during which additional bells will be suppressed. The default is 200. If set to nonzero, additional bells will also be suppressed until the server reports that processing of the first bell has been completed; this feature is most useful with the visible bell.
<b>boldFont</b> (class <b>BoldFont</b> )	Specifies the name of the bold font to use instead of overstriking.
<b>c132</b> (class <b>C132</b> )	Specifies whether the VT102 DECCOLM escape sequence should be honored. The default is <b>False</b> .
<b>charClass</b> (class <b>CharClass</b> )	Specifies comma-separated lists of character class bindings of the form <i>[low-high:val</i> ue. These are used in determining which sets of characters should be treated the same when doing cut and paste. See "Character Classes" on page 230.
<b>curses</b> (class <b>Curses</b> )	Specifies whether the last column bug in the <b>curses</b> function should be worked around. The default is <b>False</b> .
<b>cutNewline</b> (class <b>cutNewline</b> )	If <b>false</b> , triple clicking to select a line does not include the Newline at the end of the line. If <b>true</b> , the Newline is selected. The default is <b>true</b> .
<b>cutToBeginningofLines</b> (class <b>CutToBeginningOfLine</b> )	If <b>false</b> , triple clicking to select a line selects only from the current word forward. If <b>true</b> , the entire line is selected. The default is <b>true</b> .
<b>background</b> (class <b>Background</b> )	Specifies the color to use for the background of the window. The default is white.
<b>foreground</b> (class <b>Foreground</b> )	Specifies the color to use for displaying text in the window. Setting the class name instead of the instance name is an easy way to have everything that would usually be displayed in the text color to change color. The default is black.
<b>cursorColor</b> (class <b>Foreground</b> )	Specifies the color to use for the text cursor. The default is black.



## Item

**eightBitInput** (class **EightBitInput**)

**eightBitOutput** (class **EightBitOutput**)

**font** (class **Font**)

**font1** (class **Font1**)

**font2** (class **Font2**)

**font3** (class **Font3**)

**font4** (class **Font4**)

**font5** (class **Font5**)

**font6** (class **Font6**)

**geometry** (class **Geometry**)

**hpLowerleftBugCompat** (class **hpLowerleftBugCompat**)

**internalBorder** (class **BorderWidth**)

**jumpScroll** (class **JumpScroll**)

## Item

**loginShell** (class **LoginShell**)

**marginBell** (class **MarginBell**)

**multiClickTime** (class **MultiClickTime**)

**multiScroll** (class **MultiScroll**)

**nMarginBell** (class **Column**)

**pointerColor** (class **Foreground**)

**pointerColorBackground** (class **Background**)

**pointerShape** (class **Cursor**)

**resizeGravity** (class **ResizeGravity**)

**reverseVideo** (class **ReverseVideo**)

**reverseWrap** (class **ReverseWrap**)

**saveLines** (class **SaveLines**)

**scrollBar** (class **ScrollBar**)

## Description

If True, meta characters input from the keyboard are presented as a single character with the eighth bit turned on. If False, meta characters are converted into a 2-character sequence with the character itself preceded by **ESC**. The default is True.

Specifies whether 8-bit characters sent from the host should be accepted as is or stripped when printed. The default is True.

Specifies the name of the normal font. The default is fixed.

Specifies the name of the first alternative font.

Specifies the name of the second alternative font.

Specifies the name of the third alternative font.

Specifies the name of the fourth alternative font.

Specifies the name of the fifth alternative font.

Specifies the name of the sixth alternative font.

Specifies the preferred size and position of the VT102 window.

Specifies whether to work around a bug in **xdb**, which ignores termcap and always sends ESC F to move to the lower left corner. **true** causes **xterm** in interpret ESC F as a request to move to the lower left corner of the screen. The default is **false**.

Specifies the number of pixels between the characters and the window border. The default is 2.

Specifies whether jump scrolling should be used. The default is True.

## Description

Specifies whether the shell to be run in the window should be started as a login shell. The default is False.

Specifies whether the bell should be rung when the user types near the right margin. The default is False.

Specifies the maximum time in milliseconds between multiclick select events. The default is 250 milliseconds.

Specifies whether scrolling should be done asynchronously. The default is False.

Specifies the number of characters from the right margin at which the margin bell should be rung, when enabled.

Specifies the foreground color of the pointer. The default is **XtDefaultForeground**.

Specifies the background color of the pointer. The default is **XtDefaultBackground**.

Specifies the name of the shape of the pointer. The default is **xterm**.

Affects the behavior when the window is resized to be taller or shorter. **NorthWest** specifies that the top line of text on the screen stays fixed. If the window is made shorter, lines are dropped from the bottom; if the window is made taller, blank lines are added at the bottom.

This is compatible with the behavior in MIT version X11R4.

**SouthWest** (the default) specifies that the bottom line of text on the screen stays fixed. If the window is made taller, additional saved lines will be scrolled down onto the screen; if the window is made shorter, lines will be scrolled off the top of the screen, and the top saved lines will be dropped.

Specifies whether reverse video should be simulated. The default is False.

Specifies whether reverse wraparound should be enabled. The default is False.

Specifies the number of lines to save beyond the top of the screen when a scrollbar is turned on. The default is 64.

Specifies whether the scrollbar should be displayed. The default is False.

<b>Item</b>	<b>Description</b>
<b>scrollTTYOutput</b> (class <b>ScrollCond</b> )	Specifies whether output to the terminal should automatically cause the scrollbar to go to the bottom of the scrolling region. The default is True.
<b>scrollKey</b> (class <b>ScrollCond</b> )	Specifies whether pressing a key should automatically cause the scrollbar to go to the bottom of the scrolling region. The default is False.
<b>scrollLines</b> (class <b>ScrollLines</b> )	Specifies the number of lines that the <b>scroll-back</b> and <b>scroll-forw</b> actions should use as a default. The default value is 1.
<b>signalInhibit</b> (class <b>SignalInhibit</b> )	Specifies whether the entries in the Main Options menu for sending signals to <b>xterm</b> should be disallowed. The default is False.
<b>tekGeometry</b> (class <b>Geometry</b> )	Specifies the preferred size and position of the Tektronix window.
<b>tekInhibit</b> (class <b>TekInhibit</b> )	Specifies whether the escape sequence to enter Tektronix mode should be ignored. The default is False.
<b>tekSmall</b> (class <b>TekSmall</b> )	Specifies whether the Tektronix mode window should start in its smallest size if no explicit geometry is given. This is useful when running the <b>xterm</b> command on displays with small screens. The default is False.
<b>tekStartup</b> (class <b>TekStartup</b> )	Specifies whether <b>xterm</b> should start up in Tektronix mode. The default is False.
<b>titeInhibit</b> (class <b>TiteInhibit</b> )	Specifies whether <b>xterm</b> should remove <b>ti</b> and <b>te termcap</b> file entries (used to switch between alternate screens during startup of many screen-oriented programs) from the <b>TERMCAP</b> string. If set, the <b>xterm</b> command also ignores the escape sequence to switch to the alternate screen.
<b>translations</b> (class <b>Translations</b> )	Specifies the key and button bindings for menus, selections, programmed strings, and so forth. For more information, see <b>Actions</b> .
<b>visualBell</b> (class <b>VisualBell</b> )	Specifies whether a visible bell (flashing) should be used instead of an audible bell when the Ctrl+G key sequence signal is received. The default is False.

The following resources are specified as part of the **tek4014** widget (class **Tek4014**):

<b>Item</b>	<b>Description</b>
<b>width</b> (class <b>Width</b> )	Specifies the width of the Tektronix window in pixels.
<b>height</b> (class <b>Height</b> )	Specifies the height of the Tektronix window in pixels.
<b>fontLarge</b> (class <b>Font</b> )	Specifies the large font to use in the Tektronix window.
<b>font2</b> (class <b>Font</b> )	Specifies font number 2 to use in the Tektronix window.
<b>font3</b> (class <b>Font</b> )	Specifies font number 3 to use in the Tektronix window.
<b>fontSmall</b> (class <b>Font</b> )	Specifies the small font to use in the Tektronix window.
<b>initialFont</b> (class <b>InitialFont</b> )	Specifies which of the four Tektronix fonts to use initially. Values are the same as for the <b>set-tek-text</b> action. The default is large.
<b>ginTerminator</b> (class <b>GinTerminator</b> )	Specifies what characters should follow a GIN report or status report. The possibilities are `none,' which sends no terminating characters; CRonly, which sends CR; and CR&EOT, which sends both CR and EOT. The default is none.

The resources that may be specified for the various menus are described in the documentation for the **Athena SimpleMenu** widget. Following is a list of the names and classes of the entries in each of the menus.

The mainMenu has the following entries:

Item	Description
<b>securekbd</b> (class <b>SmeBSB</b> )	Invokes the <b>secure()</b> action.
<b>allowsends</b> (class <b>SmeBSB</b> )	Invokes the <b>allow-send-events(toggle)</b> action.
<b>redraw</b> (class <b>SmeBSB</b> )	Invokes the <b>redraw()</b> action.
<b>line1</b> (class <b>SmeLine</b> )	This is a separator.
<b>suspend</b> (class <b>SmeBSB</b> )	Invokes the <b>send-signal(tstp)</b> action on systems that support job control.
<b>continue</b> (class <b>SmeBSB</b> )	Invokes the <b>send-signal(cont)</b> action on systems that support job control.
<b>interrupt</b> (class <b>SmeBSB</b> )	Invokes the <b>send-signal(int)</b> action.
<b>hangup</b> (class <b>SmeBSB</b> )	Invokes the <b>send-signal(hup)</b> action.
<b>terminate</b> (class <b>SmeBSB</b> )	Invokes the <b>send-signal(term)</b> action.
<b>kill</b> (class <b>SmeBSB</b> )	Invokes the <b>send-signal(kill)</b> action.
<b>line2</b> (class <b>SmeLine</b> )	This is a separator.
<b>quit</b> (class <b>SmeBSB</b> )	Invokes the <b>quit()</b> action.

The vtMenu has the following entries:

Item	Description
<b>scrollbar</b> (class <b>SmeBSB</b> )	Invokes the <b>set-scrollbar(toggle)</b> action.
<b>jumpscroll</b> (class <b>SmeBSB</b> )	Invokes the <b>set-jumpscroll(toggle)</b> action.
<b>reversevideo</b> (class <b>SmeBSB</b> )	Invokes the <b>set-reverse-video(toggle)</b> action.
<b>autowrap</b> (class <b>SmeBSB</b> )	Invokes the <b>set-autowrap(toggle)</b> action.
<b>reversewrap</b> (class <b>SmeBSB</b> )	Invokes the <b>set-reversewrap(toggle)</b> action.
<b>autolinefeed</b> (class <b>SmeBSB</b> )	Invokes the <b>set-autolinefeed(toggle)</b> action.
<b>appcursor</b> (class <b>SmeBSB</b> )	Invokes the <b>set-appcursor(toggle)</b> action.
<b>appkeypad</b> (class <b>SmeBSB</b> )	Invokes the <b>set-appkeypad(toggle)</b> action.
<b>scrollkey</b> (class <b>SmeBSB</b> )	Invokes the <b>set-scroll-on-key(toggle)</b> action.
<b>scrollttyoutput</b> (class <b>SmeBSB</b> )	Invokes the <b>set-scroll-on-tty-output(toggle)</b> action.
<b>allow132</b> (class <b>SmeBSB</b> )	Invokes the <b>set-allow132(toggle)</b> action.
<b>cursesemul</b> (class <b>SmeBSB</b> )	Invokes the <b>set-cursesemul(toggle)</b> action.
<b>visualbell</b> (class <b>SmeBSB</b> )	Invokes the <b>set-visualbell(toggle)</b> action.
<b>marginbell</b> (class <b>SmeBSB</b> )	Invokes the <b>set-marginbell(toggle)</b> action.
<b>altscreen</b> (class <b>SmeBSB</b> )	This entry is currently disabled.
<b>line1</b> (class <b>SmeLine</b> )	This is a separator.
<b>softreset</b> (class <b>SmeBSB</b> )	Invokes the <b>soft-reset()</b> action.
<b>hardreset</b> (class <b>SmeBSB</b> )	Invokes the <b>hard-reset()</b> action.
<b>clearsavedlines</b> (class <b>SmeBSB</b> )	Invokes the <b>clear-saved-lines()</b> action.
<b>line2</b> (class <b>SmeLine</b> )	This is a separator.
<b>tekshow</b> (class <b>SmeBSB</b> )	Invokes the <b>set-visibility(tek,toggle)</b> action.
<b>tekmode</b> (class <b>SmeBSB</b> )	Invokes the <b>set-terminal-type(tek)</b> action.
<b>vthide</b> (class <b>SmeBSB</b> )	Invokes the <b>set-visibility(vt,off)</b> action.

The fontMenu has the following entries:

Item	Description
<b>fontdefault</b> (class <b>SmeBSB</b> )	Invokes the <b>set-vt-font(d)</b> action.
<b>font1</b> (class <b>SmeBSB</b> )	Invokes the <b>set-vt-font(1)</b> action.
<b>font2</b> (class <b>SmeBSB</b> )	Invokes the <b>set-vt-font(2)</b> action.
<b>font3</b> (class <b>SmeBSB</b> )	Invokes the <b>set-vt-font(3)</b> action.
<b>font4</b> (class <b>SmeBSB</b> )	Invokes the <b>set-vt-font(4)</b> action.
<b>font5</b> (class <b>SmeBSB</b> )	Invokes the <b>set-vt-font(5)</b> action.
<b>font6</b> (class <b>SmeBSB</b> )	Invokes the <b>set-vt-font(6)</b> action.
<b>fontescape</b> (class <b>SmeBSB</b> )	Invokes the <b>set-vt-font(e)</b> action.
<b>fontsel</b> (class <b>SmeBSB</b> )	Invokes the <b>set-vt-font(s)</b> action.

The tekMenu has the following entries:

Item	Description
<b>tektextlarge</b> (class <b>SmeBSB</b> )	Invokes the <b>set-tek-text(l)</b> action.
<b>tektext2</b> (class <b>SmeBSB</b> )	Invokes the <b>set-tek-text(2)</b> action.
<b>tektext3</b> (class <b>SmeBSB</b> )	Invokes the <b>set-tek-text(3)</b> action.
<b>tektextsmall</b> (class <b>SmeBSB</b> )	Invokes the <b>set-tek-text(s)</b> action.
<b>line1</b> (class <b>SmeLine</b> )	This is a separator.
<b>tekpage</b> (class <b>SmeBSB</b> )	Invokes the <b>tek-page()</b> action.
<b>tekreset</b> (class <b>SmeBSB</b> )	Invokes the <b>tek-reset()</b> action.
<b>tekcopy</b> (class <b>SmeBSB</b> )	Invokes the <b>tek-copy()</b> action.
<b>line2</b> (class <b>SmeLine</b> )	This is a separator.
<b>vtshow</b> (class <b>SmeBSB</b> )	Invokes the <b>set-visibility(vt,toggle)</b> action.
<b>vtmode</b> (class <b>SmeBSB</b> )	Invokes the <b>set-terminal-type(vt)</b> action.
<b>tekhide</b> (class <b>SmeBSB</b> )	Invokes the <b>set-visibility(tek,toggle)</b> action.

The following resources are useful when specified for the **Athena Scrollbar** widget:

Item	Description
<b>thickness</b> (class <b>Thickness</b> )	Specifies the width in pixels of the scrollbar.
<b>background</b> (class <b>Background</b> )	Specifies the color to use for the background of the scrollbar.
<b>foreground</b> (class <b>Foreground</b> )	Specifies the color to use for the foreground of the scrollbar. The <i>thumb</i> of the scrollbar is a simple checkerboard pattern with alternating pixels for foreground and background colors.

## Pointer Usage

After the VT102 window is created, the **xterm** command allows you to select text and copy it within the same or other windows.

The selection functions are invoked when the pointer buttons are used with no modifiers, and when they are used with the Shift key. The assignment of the functions to keys and buttons may be changed through the resource database.

Pointer button 1 (usually left) is used to save text into the cut buffer. Move the cursor to beginning of the text, and then hold the button down while moving the cursor to the end of the region and releasing the button. The selected text is highlighted and is saved in the global cut buffer and made the PRIMARY selection when the button is released.

Double-clicking selects by words, triple-clicking selects by lines, and quadruple-clicking goes back to characters. Multiple-click is determined by the amount of time from button up to button down, so you can change the selection unit in the middle of a selection. If the key or button bindings specify that an X selection is to be made, the **xterm** command will leave the selected text highlighted for as long as it is the selection owner.

Pointer button 2 (usually middle) "types" (pastes) the text from the PRIMARY selection, if any, otherwise from the cut buffer, inserting it as keyboard input.

Pointer button 3 (usually right) extends the current selection. If pressed while closer to the right edge of the selection than the left, it extends or contracts the right edge of the selection. If you contract the selection past the left edge of the selection, the **xterm** command assumes you really meant the left edge, restores the original selection, and then extends or contracts the left edge of the selection.

And the opposite also applies: if pressed while closer to the left edge of the selection than the right, it extends/contracts the left edge of the selection. If you contract the selection past the right edge of the selection, the **xterm** command assumes you really meant the right edge, restores the original selection, and then extends/contracts the right edge of the selection. Extension starts in the selection unit mode that the last selection or extension was performed in; you can multiple-click to cycle through them.

By cutting and pasting pieces of text without trailing new lines, you can take text from several places in different windows and form a command to the shell, for example, or take output from a program and insert it into your favorite editor. Because the cut buffer is globally shared among different applications, regard it as a “file” whose contents you know. The terminal emulator and other text programs should be treating it as if it were a text file; in other words, the text is delimited by new lines.

The scroll region displays the position and amount of text currently showing in the window (highlighted) relative to the amount of text actually saved. As more text is saved (up to the maximum), the size of the highlighted area decreases.

Clicking button 1 with the pointer in the scroll region moves the adjacent line to the top of the display window.

Clicking button 3 moves the top line of the display window down to the pointer position.

Clicking button 2 moves the display to a position in the saved text that corresponds to the pointer's position in the scrollbar.

Unlike the VT102 window, the Tektronix window does not allow the copying of text. It does allow Tektronix GIN mode, and in this mode the cursor will change from an arrow to a cross. Pressing any key will send that key and the current coordinates of the cross cursor. Pressing button one, two, or three will return the letters l, m, and r, respectively.

If the Shift key is pressed when a pointer button is pressed, the corresponding uppercase letter is sent. To distinguish a pointer button from a key, the high bit of the character is set (but this bit is usually stripped unless the terminal mode is RAW; see the `tty` command for details).

## Menus

The `xterm` command has four menus, named `mainMenu`, `vtMenu`, `fontMenu`, and `tekMenu`. Each menu opens under the correct combinations of key and button presses. Most menus are divided into two sections, separated by a horizontal line. The top portion contains various modes that can be altered. A check mark is displayed next to a mode that is currently active. Selecting one of these modes toggles its state. The bottom portion of the menu lists command entries; selecting one of these performs the indicated function.

The `xterm` menu opens when the control key and pointer button one are pressed in a window. The `mainMenu` contains items that apply to both the VT102 and Tektronix windows. The **Secure Keyboard** mode is used when typing in passwords or other sensitive data in an unsecure environment.

Notable entries in the command section of the menu are **Continue**, **Suspend**, **Interrupt**, **Hangup**, **Terminate**, and **Kill**, which send the `SIGCONT`, `SIGTSTP`, `SIGINT`, `SIGHUP`, `SIGTERM`, and `SIGKILL` signals, respectively, to the process group of the process running under `xterm` (usually the shell). The **Continue** function is especially useful if the user has accidentally pressed Ctrl+Z, suspending the process.

The `vtMenu` sets various modes in the VT102 emulation, and is opened when the control key and pointer button two are pressed in the VT102 window. In the command section of this menu, the soft reset entry will reset scroll regions. This can be convenient when some program has left the scroll regions set incorrectly (often a problem when using VMS or TOPS-20).

The full reset entry will clear the screen, reset tabs to every eight columns, and reset the terminal modes (such as wrap and smooth scroll) to their initial states just after the `xterm` command has finished processing the command-line options.

The fontMenu sets the font used in the VT102 window. In addition to the default font and a number of alternatives that are set with resources, the menu offers the font last specified by the Set Font escape sequence (See " Control Sequences" ) and the current selection as a font name (if the PRIMARY selection is owned).

The tekMenu sets various modes in the Tektronix emulation, and is opened when the control key and pointer button two are pressed in the Tektronix window. The current font size is checked in the Modes section of the menu. The **PAGE** entry in the command section clears the Tektronix window.

## Security

X windows environments differ in their security consciousness. MIT servers, run under **xdm**, are capable of using a *magic cookie* authorization scheme that can provide a reasonable level of security for many people. If your server is only using a host-based mechanism to control access to the server (see the **xhost** command), and if you enable access for a host and other users are also permitted to run clients on that same host, there is every possibility that someone can run an application that will use the basic services of the X protocol to snoop on your activities, potentially capturing a transcript of everything you type at the keyboard.

This is of particular concern when you want to type in a password or other sensitive data. The best solution to this problem is to use a better authorization mechanism than host-based control, but a simple mechanism exists for protecting keyboard input in the **xterm** command.

The xterm menu contains a **Secure Keyboard** entry that, when enabled, ensures that all keyboard input is directed *only* to the **xterm** command (using the **GrabKeyboard** protocol request). When an application prompts you for a password (or other sensitive data), you can enable **Secure Keyboard** using the menu, type in the data, and then disable **Secure Keyboard** using the menu again.

Only one X client at a time can secure the keyboard, so when you attempt to enable **Secure Keyboard** it may fail. In this case, the bell will sound. If the **Secure Keyboard** succeeds, the foreground and background colors will be exchanged (as if you selected the **Reverse Video** entry in the Modes menu); they will be exchanged again when you exit secure mode. If the colors do *not* switch, be *very* suspicious that you are being spoofed.

If the application you are running displays a prompt before asking for the password, it is safest to enter secure mode *before* the prompt gets displayed, and to make sure that the prompt gets displayed correctly (in the new colors), to minimize the probability of spoofing. You can also bring up the menu again and make sure that a check mark is displayed next to the entry.

**Secure Keyboard** mode will be disabled automatically if your xterm window becomes iconified (or otherwise unmapped), or if you start up a reparenting window manager (that places a title bar or other decoration around the window) while in **Secure Keyboard** mode. (This is a feature of the X protocol not easily overcome.) When this happens, the foreground and background colors will be switched back and the bell will sound in warning.

## Character Classes

Clicking the middle mouse button twice in rapid succession will cause all characters of the same class (such as letters, white space, punctuation) to be selected. Because different people have different preferences for what should be selected (for example, should file names be selected as a whole or only the separate subnames), the default mapping can be overridden through the use of the **charClass** (class **CharClass**) resource.

This resource is a series of comma-separated *range:value* pairs. The *range* is either a single number or *low-high* in the range of 0 to 127, corresponding to the ASCII code for the character or characters to be set. The *value* is arbitrary, although the default table uses the character number of the first character occurring in the set.

The default table is:

```
static int charClass[128] = {
/* NUL  SOH  STX  ETX  EOT  ENQ  ACK  BEL */
   32,  1,  1,  1,  1,  1,  1,  1,
/* BS   HT   NL   VT   NP   CR   SO   SI */
   1,  32,  1,  1,  1,  1,  1,  1,
/* DLE  DC1  DC2  DC3  DC4  NAK  SYN  ETB */
   1,  1,  1,  1,  1,  1,  1,  1,
/* CAN  EM   SUB  ESC  FS   GS   RS   US */
   1,  1,  1,  1,  1,  1,  1,  1,
/* SP   !   "   #   $   %   &   ' */
   32,  33,  34,  35,  36,  37,  38,  39,
/* (   )   *   +   ,   -   .   / */
   40,  41,  42,  43,  44,  45,  46,  47,
/* 0   1   2   3   4   5   6   7 */
   48,  48,  48,  48,  48,  48,  48,  48,
/* 8   9   :   ;   <   =   >   ? */
   48,  48,  58,  59,  60,  61,  62,  63,
/* @   A   B   C   D   E   F   G */
   64,  48,  48,  48,  48,  48,  48,  48,
/* H   I   J   K   L   M   N   O */
   48,  48,  48,  48,  48,  48,  48,  48,
/* P   Q   R   S   T   U   V   W */
   48,  48,  48,  48,  48,  48,  48,  48,
/* X   Y   Z   [   \   ]   ^   _ */
   48,  48,  48,  91,  92,  93,  94,  48,
/* `   a   b   c   d   e   f   g */
   96,  48,  48,  48,  48,  48,  48,  48,
/* h   i   j   k   l   m   n   o */
   48,  48,  48,  48,  48,  48,  48,  48,
/* p   q   r   s   t   u   v   w */
   48,  48,  48,  48,  48,  48,  48,  48,
/* x   y   z   {   |   }   ~  DEL */
   48,  48,  48, 123, 124, 125, 126,  1};
```

For example, the string 33:48,37:48,45-47:48,64:48 indicates that the exclamation mark, percent sign, dash, period, slash, and & characters should be treated the same way as characters and numbers. This is useful for cutting and pasting electronic mailing addresses and file names.

## Actions

It is possible to rebind keys (or sequences of keys) to arbitrary strings for input by changing the translations for the **vt100** or **tek4014** widgets. Changing the translations for events other than key and button events is not expected, and will cause unpredictable behavior. The following actions are provided for using within the vt100 or tek4014 translations resources:

Item	Description
<b>bell</b> ([Percent])	Rings the keyboard bell at the specified percentage above or below the base volume.
<b>ignore</b> ()	Ignores the event but checks for special pointer position escape sequences.
<b>insert</b> ()	Inserts the character or string associated with the key that was pressed.
<b>insert-seven-bit</b> ()	Is a synonym for <b>insert</b> ().
<b>insert-eight-bit</b> ()	Inserts an 8-bit (meta) version of the character or string associated with the key that was pressed. The exact action depends on the value of the <b>eightBitInput</b> resource.

## Item

**insert-selection**(*SourceName* [, ...])

**keymap**(*Name*)

**pop-up menu**(*MenuName*)

**secure**()

**select-start**()

**select-extend**()

**select-end**(*DestName* [, ...])

**select-cursor-start**()

**select-cursor-end**(*DestName* [, ...])

**set-vt-font**(*d*/1/2/3/4/5/6/*e*/*s* [, *NormalFont* [, *BoldFont*]])

**start-extend**()

**start-cursor-extend**()

**string**(*String*)

**scroll-back**(*Count* [, *Units*])

**scroll-forw**(*Count* [, *Units*])

**allow-send-events**(*On/Off/Toggle*)

**redraw**()

## Description

Inserts the string found in the selection or cutbuffer indicated by the *SourceName* parameter. Sources are checked in the order given (case is significant) until one is found. Commonly used selections include PRIMARY, SECONDARY, and CLIPBOARD. Cut buffers are typically named CUT\_BUFFER0 through CUT\_BUFFER7.

Dynamically defines a new translation table whose resource name is *Name* with the suffix *Keymap* (case is significant). The name None restores the original translation table.

Displays the specified popup menu. Valid names (case is significant) include mainMenu, vtMenu, fontMenu, and tekMenu.

Toggles the **Secure Keyboard** mode described in the section named "Security", and is invoked from the **securekbd** entry in mainMenu.

Begins text selection at the current pointer location. See the section entitled "Pointer Usage" for information on making selections.

Tracks the pointer and extends the selection. Only bind this to Motion events.

Puts the currently selected text into all of the selections or cutbuffers specified by *DestName*.

Is similar to **select-start** except that it begins the selection at the current text cursor position.

Is similar to **select-end** except that it should be used with **select-cursor-start**.

Sets the font or fonts currently being used in the VT102 window. The first argument is a single character that specifies the font to be used:

*d* or *D* indicates the default font (the font initially used when the **xterm** command was started),

*1* through *6* indicate the fonts specified by the *font1* through *font6* resources,

*e* or *E* indicates the normal and bold fonts that have been set through escape codes (or specified as the second and third action arguments, respectively), and

*s* or *S* indicates the font selection (as made by programs such as the **xfonssel** program) specified by the second action argument.

Is similar to **select-start** except that the selection is extended to the current pointer location.

Is similar to **select-extend** except that the selection is extended to the current text cursor position.

Inserts the specified text string as if it had been typed. Quotation is necessary if the string contains white space or nonalphanumeric characters. If the string argument begins with the characters ``0x," it is interpreted as a hex character constant.

Scrolls the text window backward so that text that had previously scrolled off the top of the screen is now visible. The *Count* argument indicates the number of *Units* (which may be *page*, *halfpage*, *pixel*, or *line*) by which to scroll.

Scrolls is similar to **scroll-back** except that it scrolls the other direction.

Sets or toggles the **allowSendEvents** resource and is also invoked by the **allowsends** entry in mainMenu.

Redraws the window and is also invoked by the **redraw** entry in mainMenu.



**Item****send-signal**(*SigName*)**Description**

Sends the signal named by *SigName* to the **xterm** subprocess (the shell or program specified with the **-e** command-line option) and is also invoked by the **suspend**, **continue**, **interrupt**, **hangup**, **terminate**, and **kill** entries in **mainMenu**. Allowable signal names are (case is not significant):

**tstp** (if supported by the operating system),

**suspend** (same as **tstp**),

**cont** (if supported by the operating system),

**int**,

**hup**,

**term**,

**quit**,

**alarm**,

**alarm** (same as **alarm**), and

**kill**.

**quit**()

Sends a **SIGHUP** to the subprogram and exits. It is also invoked by the **quit** entry in **mainMenu**.

**Item****set-scrollbar**(*On/Off/Toggle*)**Description**

Toggles the **scrollbar** resource and is also invoked by the **scrollbar** entry in **vtMenu**.

**set-jumpscroll**(*On/Off/Toggle*)

Toggles the **jumpscroll** resource and is also invoked by the **jumpscroll** entry in **vtMenu**.

**set-reverse-video**(*On/Off/Toggle*)

Toggles the **reverseVideo** resource and is also invoked by the **reversevideo** entry in **vtMenu**.

**set-autowrap**(*On/Off/Toggle*)

Toggles automatic wrapping of long lines and is also invoked by the **autowrap** entry in **vtMenu**.

**set-reversewrap**(*On/Off/Toggle*)

Toggles the **reverseWrap** resource and is also invoked by the **reversewrap** entry in **vtMenu**.

**set-autolinefeed**(*On/Off/Toggle*)

Toggles automatic insertion of linefeeds and is also invoked by the **autolinefeed** entry in **vtMenu**.

**set-appcursor**(*On/Off/Toggle*)

Toggles the handling Application Cursor Key mode and is also invoked by the **appcursor** entry in **vtMenu**.

**set-appkeypad**(*On/Off/Toggle*)

Toggles the handling of Application Keypad mode and is also invoked by the **appkeypad** entry in **vtMenu**.

**set-scroll-on-key**(*On/Off/Toggle*)

Toggles the **scrollKey** resource and is also invoked from the **scrollkey** entry in **vtMenu**.

**set-scroll-on-tty-output**(*On/Off/Toggle*)

Toggles the **scrollTtyOutput** resource and is also invoked from the **scrollttyoutput** entry in **vtMenu**.

**set-allow132**(*On/Off/Toggle*)

Toggles the **c132** resource and is also invoked from the **allow132** entry in **vtMenu**.

**set-cursesemul**(*On/Off/Toggle*)

Toggles the **curses** resource and is also invoked from the **cursesemul** entry in **vtMenu**.

**set-visual-bell**(*On/Off/Toggle*)

Toggles the **visualBell** resource and is also invoked by the **visualbell** entry in **vtMenu**.

**set-marginbell**(*On/Off/Toggle*)

Toggles the **marginBell** resource and is also invoked from the **marginbell** entry in **vtMenu**.

**set-altscreen**(*On/Off/Toggle*)

Toggles between the alternate and current screens.

**soft-reset**()

Resets the scrolling region and is also invoked from the **softreset** entry in **vtMenu**.

**hard-reset**()

Resets the scrolling region, tabs, window size, and cursor keys and clears the screen. It is also invoked from the **hardreset** entry in **vtMenu**.

**clear-saved-lines**()

Performs **hard-reset** (see previous entry) and also clears the history of lines saved off the top of the screen. It is also invoked from the **clearsavedlines** entry in **vtMenu**.

Item	Description
<b>set-terminal-type</b> ( <i>Type</i> )	Directs output to either the vt or tek windows, according to the <i>Type</i> string. It is also invoked by the <b>tekmode</b> entry in vtMenu and the <b>vtmode</b> entry in tekMenu.
<b>set-visibility</b> ( <i>vt/tek, On/Off/Toggle</i> )	Controls whether or not the vt or tek windows are visible. It is also invoked from the <b>tekshow</b> and <b>vthide</b> entries in vtMenu and the <b>vtshow</b> and <b>tekhide</b> entries in tekMenu.
<b>set-tek-text</b> ( <i>large/2/3/small</i> )	Sets font used in the Tektronix window to the value of the resources <b>tektextlarge</b> , <b>tektext2</b> , <b>tektext3</b> , and <b>tektextsmall</b> according to the argument. It is also by the entries of the same names as the resources in tekMenu.
<b>tek-page</b> ()	Clears the Tektronix window and is also invoked by the <b>tekpage</b> entry in tekMenu.
<b>tek-reset</b> ()	Resets the Tektronix window and is also invoked by the <b>tekreset</b> entry in tekMenu.
<b>tek-copy</b> ()	Copies the escape codes used to generate the current window contents to a file in the current directory beginning with the name <b>COPY</b> . It is also invoked from the <b>tekcopy</b> entry in tekMenu.
<b>visual-bell</b> ()	Flashes the window quickly.

The Tektronix window also has the following action:

Item	Description
<b>gin-press</b> ( <i>l/L/m/M/r/R</i> )	Sends the indicated graphics input code.

The default bindings in the VT102 window are:

```

Shift <KeyPress> Prior:      scroll-back(1,halpage) \n\
Shift <KeyPress> Next:      scroll-forw(1,halpage) \n\
Shift <KeyPress> Select:    select-cursor-start \
                             select-cursor-end(PRIMARY,
                             CUT_BUFFER0) \n\
Shift <KeyPress> Insert:    insert-selection(PRIMARY,
                             CUT_BUFFER0) \n\
~Meta<KeyPress>:          insert-seven-bit \n\
Meta<KeyPress>:          insert-eight-bit \n\
!Ctrl <Btn1Down>:        pop-up menu(mainMenu) \n\
!Lock Ctrl <Btn1Down>:    pop-up menu(mainMenu) \n\
~Meta <Btn1Down>:        select-start \n\
~Meta <Btn1Motion>:      select-extend \n\
!Ctrl <Btn2Down>:        pop-up menu(vtMenu) \n\
!Lock Ctrl <Btn2Down>:    pop-up menu(vtMenu) \n\
~Ctrl ~Meta <Btn2Down>:  ignore \n\
~Ctrl ~Meta <Btn2Up>:    insert-selection(PRIMARY,
                             CUT_BUFFER0) \n\
!Ctrl <Btn3Down>:        pop-up menu(fontMenu) \n\
!Lock Ctrl <Btn3Down>:    pop-up menu(fontMenu) \n\
~Ctrl ~Meta <Btn3Down>:  start-extend \n\
~Meta <Btn3Motion>:      select-extend \n\
<BtnUp>:                select-end(PRIMARY, CUT_BUFFER0) \n\
<BtnDown>:              bell(0)

```

The default bindings in the Tektronix window are:

```

~Meta<KeyPress>:          insert-seven-bit \n\
Meta<KeyPress>:          insert-eight-bit \n\
!Ctrl <Btn1Down>:        pop-up menu(mainMenu) \n\
!Lock Ctrl <Btn1Down>:    pop-up menu(mainMenu) \n\
!Ctrl <Btn2Down>:        pop-up menu(tekMenu) \n\
!Lock Ctrl <Btn2Down>:    pop-up menu(tekMenu) \n\
Shift ~Meta<Btn1Down>:    gin-press(L) \n\
~Meta<Btn1Down>:        gin-press(l) \n\

```

```

Shift ^Meta<Btn2Down>:   gin-press(M) \n\
^Meta<Btn2Down>:        gin-press(m) \n\
Shift ^Meta<Btn3Down>:   gin-press(R) \n\
^Meta<Btn3Down>:        gin-press(r)

```

The following is an example of how the **keymap** action is used to add special keys for entering commonly typed works:

```

*VT100.Translations:      #override <Key>F13: keymap(dbx)
*VT100.dbxKeymap.translations:
\
    <Key>F14:   keymap(None) \n\
    <Key>F17:   string("next") string(0x0d) \n\
    <Key>F18:   string("step") string(0x0d) \n\
    <Key>F19:   string("continue") string(0x0d) \n\
    <Key>F20:   string("print ")
                insert-selection(PRIMARY,CUT_BUFFER0)

```

## Environment

The **xterm** command sets the environment variables **TERM** and **TERMCAP** properly for the size window you have created. It also uses and sets the **DISPLAY** environment variable to specify which bitmap display terminal to use. The **WINDOWID** environment variable is set to the X window ID number of the **xterm** window.

## Bugs

Large pastes do not work on some systems. This is not a bug in the **xterm** command; it is a bug in the pseudo terminal driver of those systems. The **xterm** command feeds large pastes to the pty only as fast as the pty will accept data, but some pty drivers do not return enough information to know if the write operation has succeeded.

Many of the options are not resettable after the **xterm** command starts.

Only fixed-width, character-cell fonts are supported.

## Control Sequences

This section lists control sequences available for the **xterm** command.

### Definitions

The following information shows how to interpret key sequences in this section.

Item	Description
<i>c</i>	The literal characters <i>c</i> .
<i>C</i>	A single (required) character.
<i>P<sub>s</sub></i>	A single (usually optional) numeric parameter, composed of one or more digits.
<i>P<sub>m</sub></i>	A multiple numeric parameter composed of any number of single numeric parameters, separated by a ; (semi-colon) character or characters.
<i>P<sub>t</sub></i>	A text parameter composed of printable characters.

### VT100 Mode

Most of these control sequences are standard VT102 control sequences, but there are some sequences here from later DEC VT terminals, too. Major VT102 features not supported are smooth scrolling, double-size characters, flashing characters, and VT52 mode.

There are additional control sequences to provide xterm-dependent functions, like the scrollbar or window size. Where the function is specified by DEC or ISO 6429, the code assigned to it is given in parentheses. The escape codes to designate character sets are specified by ISO 2022; see that document for a discussion of character sets.

Control Sequence	Description
<b>BEL</b>	Bell (Ctrl+G)
<b>BS</b>	Backspace (Ctrl+H)
<b>TAB</b>	Horizontal Tab (HT) (Ctrl+l)
<b>LF</b>	Line Feed or New Line (NL) (Ctrl+j)
<b>VT</b>	Vertical Tab (Ctrl+k) same as LF
<b>FF</b>	Form Feed or New Page (NP) (Ctrl+l) same as LF
<b>CR</b>	Carriage return (Ctrl+m)
<b>SO</b>	Shift Out (Ctrl+n) -> Switch to ALternate Character Set: Invokes the G1 character set.
<b>SI</b>	Shift In (Ctrl+o) -> Switch to Standard Character Set: Invokes the G0 character set (the default).
<b>ESC # 8</b>	DEC Screen Test (DCECALN)
<b>ESC ( C</b>	Designate G0 Character Set (ISO 2022) <b>C = 0</b> DEC Special Character and Line Drawing Set <b>C = A</b> United Kingdom (UK) <b>C = B</b> United States (USASCII)
<b>ESC ) C</b>	Designate G1 Character Set (ISO 2022) <b>C = 0</b> DEC Special Character and Line Drawing Set <b>C = A</b> United Kingdom (UK) <b>C = B</b> United States (USASCII)
<b>ESC * C</b>	Designate G2 Character Set (ISO 2022) <b>C = 0</b> DEC Special Character and Line Drawing Set <b>C = A</b> United Kingdom (UK) <b>C = B</b> United States (USASCII)
<b>ESC + C</b>	Designate G3 Character Set (ISO 2022) <b>C = 0</b> DEC Special Character and Line Drawing Set <b>C = A</b> United Kingdom (UK) <b>C = B</b> United States (USASCII)
<b>ESC 7</b>	Save Cursor (DECSC)
<b>ESC 8</b>	Restore Cursor (DECRC)
<b>ESC =</b>	Application Keypad (DECPAM)
<b>ESC &gt;</b>	Normal Keypad (DECNM)
<b>ESC D</b>	Index (IND)
<b>ESC E</b>	Next Line (NEL)
<b>ESC H</b>	Tab Set (HTS)
<b>ESC M</b>	Reverse Index (RI)
<b>ESC N</b>	Single Shift Select of G2 Character Set (SS2): Affects next character only.
<b>ESC P</b>	Single Shift Select of G3 Character Set (SS2): Affects next character only.
<b>ESC O <math>P_t</math> ESC \</b>	Device Control String (DCS). xterm implements no DCS functions; $P_t$ is ignored. $P_t$ need not be printable characters.
<b>ESC Z</b>	Return Terminal ID (DECID). Obsolete form of ESC [ c (DA)

Control Sequence	Description
ESC [ $P_s$ @	Insert $P_s$ (Blank) Character of Characters (default=1) (ICH)
ESC [ $P_s$ A	Cursor Up $P_s$ Times (default=1) (CUU)
ESC [ $P_s$ B	Cursor Down $P_s$ Times (default=1) (CUD)
ESC [ $P_s$ C	Cursor Forward $P_s$ Times (default=1) (CUF)

Item	Description
ESC [ $P_s$ D	Cursor Backward $P_s$ Times (default=1) (CUB)
ESC [ $P_s$ ; $P_s$ H	Cursor Position [row;column] (default=1) (CUP)
ESC [ $P_s$ J	Erase in Display (ED) $P_s = 0$ Clear Below (Default) $P_s = 1$ Clear Above $P_s = 2$ Clear All
ESC [ $P_s$ K	Erase in Line (EL) $P_s = 0$ Clear to Right (Default) $P_s = 1$ Clear to Left $P_s = 2$ Clear All
ESC [ $P_s$ L	Insert $P_s$ Lines (default=1) (IL)
ESC [ $P_s$ M	Delete $P_s$ Lines (default=1) (DL)
ESC [ $P_s$ P	Delete $P_s$ Characters (default=1) (DCH)
ESC [ $P_s$ ; $P_s$ ; $P_s$ ; $P_s$ ; $P_s$ T	Initiate hilite mouse tracking. Parameters are [Func;Startx;Starty;FirstRow;LastRow]. See Mouse Tracking.
ESC [ $P_s$ c	SendDevice Attributes (DA)Delete $P_s$ Characters (default=1) (DCH) $P_s = 0$ or <b>omitted</b> Request attribute from terminal ESC [ ? 1 ; 2 c ("I am a VT100 with Advanced Video Option.")
ESC [ $P_s$ ; $P_s$ f	Horizontal and Vertical Position [row;column] (default = [1,1]) (HVP)
ESC [ $P_s$ g	Tab Clear (TBC) $P_s = 0$ Clear Current Column (default) $P_s = 3$ Clear All
ESC [ $P_m$ h	Set Modes (SM) $P_s = 4$ Insert Mode (IRM) $P_s = 2 0$ Automatic Newline (LNM)
ESC [ $P_m$ l	Reset Modes (RM) $P_s = 4$ Replace Mode (IRM) $P_s = 2 0$ Normal Linefeed (LNM)
ESC [ $P_m$ m	Character Attributes (SGR) $P_s = 0$ Nomal (default) $P_s = 1$ Bold $P_s = 4$ Underscore $P_s = 5$ Blink (displayed as Bold) $P_s = 7$ Inverse

Item	Description
ESC [ <i>P<sub>s</sub></i> n	Device Status Report (DSR) <i>P<sub>s</sub></i> = 5    Status ReportESC [ 0 n ("OK") <i>P<sub>s</sub></i> = 6    Report Cursor Position (CPR)[row;column] asESC [ r ; c R <i>P<sub>s</sub></i> = 2 0    Automatic Newline (LNM)
ESC [ <i>P<sub>s</sub></i> ; <i>P<sub>s</sub></i> r	Set Scroll Region [top;bottom] (default = fullsize of window) (DECSTBM)
ESC [ <i>P<sub>s</sub></i> x	Request Terminal Parameters (DECREQTPARM)
ESC [ ? <i>P<sub>m</sub></i> h	DEC Private Mode (DECSET) <i>P<sub>s</sub></i> = 1    Application Cursor Keys (DECCKM) <i>P<sub>s</sub></i> = 2    Designate USASCII for character sets G0–G3. (In VT102, this selects VT52 mode (DECANM), which <b>xterm</b> does not support.) <i>P<sub>s</sub></i> = 3    132 Column Mode (DECCOLM) <i>P<sub>s</sub></i> = 4    Smooth (Slow) Scroll (DECSCLM) <i>P<sub>s</sub></i> = 5    Reverse Video (DECSCNM) <i>P<sub>s</sub></i> = 6    Origin Mode (DECOM) <i>P<sub>s</sub></i> = 7    Wraparound Mode (DECAWM) <i>P<sub>s</sub></i> = 8    Auto-repeat Keys (DECARM) <i>P<sub>s</sub></i> = 9    Set Mouse X and Y on button press. See "Mouse Tracking" on page 240. <i>P<sub>s</sub></i> = 3 8    Enter Tektronix Mode (DECTEK) <i>P<sub>s</sub></i> = 4 0    Allow 80 <—> 132 Mode <i>P<sub>s</sub></i> = 4 1 <b>curses</b> function fix <i>P<sub>s</sub></i> = 4 4    Turn On Margin Bell <i>P<sub>s</sub></i> = 4 5    Reverse Wraparound Mode <i>P<sub>s</sub></i> = 4 7    Use Alternate Screen Buffer (unless disabled by <b>titelnhibit</b> resource) <i>P<sub>s</sub></i> = 1 0 0 0    Set Mouse X and Y on button press and release. See "Mouse Tracking" on page 240. <i>P<sub>s</sub></i> = 1 0 0 1    Use Hilite Mouse Tracking.

Item	Description
ESC [ ? $P_m$ l	DEC Private Mode Reset (DECRST) $P_s = 1$ Normal Cursor Keys (DECCCKM) $P_s = 3$ 80 Column Mode (DECCOLM) $P_s = 4$ Jump Fast Scroll (DECSCLM) $P_s = 5$ Normal Video (DECSCNM) $P_s = 6$ Normal Cursor Mode (DECOM) $P_s = 7$ No Wraparound Mode (DECAWM) $P_s = 8$ No Auto-repeat Keys (DECARM) $P_s = 9$ Do not Send Mouse X and Y on button press. $P_s = 4 0$ Disallow 80 $\leftrightarrow$ 132 Mode $P_s = 4 1$ No <b>curses</b> function fix $P_s = 4 4$ Turn Off Margin Bell $P_s = 4 5$ No Reverse Wraparound Mode $P_s = 4 7$ Use Normal Screen Buffer $P_s = 1 0 0 0$ Do not Send Mouse X and Y on button press and release. $P_s = 1 0 0 1$ Do not Use Hilite Mouse Tracking. xxx
ESC [ ? $P_m$ r	Restore DEC Private Mode Values. The value of $P_s$ previously saved is restored. $P_s$ values are the same as DECSET.
ESC [ ? $P_m$ s	Save DEC Private Mode Values. $P_s$ values are the same as DECSET.
ESC ]? $P_s$ ; $P_t$ BEL	Set Text Parameters $P_s = 0$ Change Icon Name and Window Title to $P_t$ $P_s = 1$ Change Icon Name to $P_t$ $P_s = 2$ Change Window Title to $P_t$ $P_s = 5 0$ Set Font to $P_t$
ESC $P_t$ ESC \	Private Message (PM). <b>xterm</b> implements no PM functions; $P_t$ need not be printable characters.
ESC _ $P_t$ ESC \	Application Program Command (APC). Private Message (PM). <b>xterm</b> implements no APC functions; $P_t$ is ignored. $P_t$ need not be printable characters.
ESC c	Full Reset (RIS)
ESC n	Select the G2 Character Set (LS2)
ESC o	Select the G3 Character Set (LS3)
ESC l	Invoke the G3 Character Set as GR (LS3R). Has no visible effect in <b>xterm</b> .
ESC }	Invoke the G2 Character Set as GR (LS2R). Has no visible effect in <b>xterm</b> .
ESC	Invoke the G1 Character Set as GR (LS1R). Has no visible effect in <b>xterm</b> .

### XTERM Description Limitation

The **xterm** terminal description in the DEC.TI file on AIX Version 4 provides underline mode by using the SGR attribute. The SMUL and RMUL attributes are not currently defined in the **XTERM** terminal description on AIX Version 4. Use the more generic capability named SGR.

tput sgr x y

Where  $x$  is either a 1 or a 0 to turn standout mode on or off respectively, and  $y$  is either a 1 or a 0 to turn underline mode on or off respectively. See the article "**terminfo** file format" for more details on the SGR capability.

```
tput sgr 0 1    turn off standout; turn on underline
tput sgr 0 0    turn off standout; turn off underline
tput sgr 1 1    turn on standout; turn on underline
tput sgr 1 0    turn on standout; turn off underline
```

## Mouse Tracking

The VT widget can be set to send the mouse position and other information on button presses. These modes are typically used by editors and other full-screen applications that want to make use of the mouse.

There are three mutually exclusive modes, each enabled (or disabled) by a different parameter in the DECSET (or DECRST) escape sequence. Parameters for all mouse tracking escape sequences generated by the **xterm** command encode numeric parameters in a single character as *value*+040. The screen coordinate system is 1-based.

For example ! is 1. The screen coordinate system is 1-based.

X10 compatibility mode sends an escape sequence on button press encoding the location and the mouse button pressed. It is enabled by specifying parameter 9 to DECSET. On button press, the **xterm** command sends the following "6 characters" .  $C_b$  is button-1.  $C_x$  and  $C_y$  are the  $x$  and  $y$  coordinates of the mouse when the button was pressed.

**ESC [ M  $C_bC_xC_y$**

Normal tracking mode sends an escape sequence on both button press and release. Modifier information is also sent. It is enabled by specifying parameter 1000 to DECSET. On button press or release, the **xterm** command sends the following "key sequence" :

**ESC [ M  $C_bC_xC_y$**

The low two bits of  $C_b$  encode button information: 0=MB1 pressed, 1=MB2 pressed, 2=MB3 pressed, 3=release. The upper bits encode what modifiers were down when the button was pressed and are added together. 4=Shift, 8=Meta, 16=Control.  $C_x$  and  $C_y$  are the  $x$  and  $y$  coordinates of the mouse event. The upper left corner is (1,1).

Mouse hilite tracking notifies a program of a button press, receives a range of lines from the program, highlights the region covered by the mouse within that range until button release, and then sends the program the release coordinates. It is enabled by specifying parameter 1001 to DECSET.

**Attention:** Use of this mode requires a cooperating program or it will hang the **xterm** command. On button press, the same information as for normal tracking is generated; the **xterm** command then waits for the program to send mouse tracking information. *All X events are ignored until the following proper escape sequence is received from the pty:*

**ESC [  $P_s ; P_s ; P_s ; P_s ; T$**

The parameters are *Func*, *Startx*, *Starty*, *FirstRow*, and *LastRow*. The *Func* parameter is nonzero to initiate hilite tracking and 0 (zero) to abort. The *Startx* and *Starty* parameters give the starting  $x$  and  $y$  location for the highlighted region. The ending location tracks the mouse, but is never above row *FirstRow* and is always above row *LastRow*. (The top of the screen is row 1.) When the button is released, the **xterm** command reports the ending position one of two ways: if the start and end coordinates are valid text locations, the **xterm** command reports the "ending position" as follows:



ESC [ t C<sub>x</sub>C<sub>y</sub>

If either coordinate is past the end of the line, the **xterm** command reports the “ending position” as follows:

ESC [ T C<sub>x</sub>C<sub>y</sub>C<sub>x</sub>C<sub>y</sub>C<sub>x</sub>C<sub>y</sub>

The parameters are *Startx*, *Starty*, *Endx*, *Endy*, *Mousex*, and *Mousey*. The *Startx*, *Starty*, *Endx*, and *Endy* parameters give the starting and ending character positions of the region. The *Mousex* and *Mousey* parameters give the location of the mouse at button up, which might not be over a character.

### Tektronix 4014 Mode

Most of these sequences are standard Tektronix 4014 control sequences. The major features missing are the write-thru and defocused modes. This document does not describe the commands used in the various Tektronix plotting modes but does describe the commands to switch modes.

#### Related reference:

“vi or vedit Command” on page 13

#### Related information:

aixterm command

resize command

---

## xwd Command

### Purpose

Dumps the image of an Enhanced X-Windows window.

### Syntax

```
xwd [ -add Value ] [ -frame ] [ -display Display ] [ -help ] [ -nobdrs ] [ -xy ] [ -out File ] [ -root | -id id | -name Name ] [ -icmap ] [ -screen ]
```

### Description

The **xwd** command is an Enhanced X-Windows window dumping utility. The **xwd** command allows you to store window images in a specially formatted dump file. This file can then be read by various other X utilities that perform functions such as redisplaying, printing, editing, formatting, archiving, and image processing. Select the target window by clicking the mouse in the desired window. The keyboard bell rings once at the beginning of the dump and twice when the dump is completed.

### Flags

Item	Description
-add <i>Value</i>	Specifies a signed value to add to every pixel. This option is specific to X11R5.
-frame	This option indicates that the window manager frame should be included when manually selecting a window.
-display <i>Display</i>	Specifies the server connection.
-help	Prints the usage command syntax summary.

Item	Description
<b>-nobdrs</b>	Specifies that the window dump does not include the pixels that compose the X window border. This is useful if you want to include the window contents in a document as an illustration. The result of the <b>-nobdrs</b> flag depends on which window manager is running. Many window managers remove all borders from the client. For example, the <b>XGetWindowAttributes</b> function returns the value of 0 for the <code>border_width</code> field regardless of the border width when the client was started. Therefore, any border that is visible on the screen belongs to the window manager; the client has no knowledge of it. In this case, the <b>-nobdrs</b> flag has no effect.
<b>-out</b> <i>File</i>	Specifies the output file on the command line. The default is to output to standard out.
<b>-root</b>	Indicates that the root window should be selected for the window dump, without requiring the user to select a window with the pointer. This option is specific to X11R5.
<b>-id</b> <i>id</i>	Indicates that the window with the specified resource id should be selected for the window dump, without requiring the user to select a window with the pointer. This option is specific to X11R5.
<b>-name</b> <i>Name</i>	Indicates that the window with the specified WM_NAME property should be selected for the window dump, without requiring the user to select a window with the pointer. This option is specific to X11R5.
<b>-icmap</b>	Forces the first installed colormap of the screen to be used to obtain RGB values. By default, the colormap of the chosen window is used. This option is specific to X11R5.
<b>-screen</b>	Indicates that the GetImage request used to obtain the image should be done on the root window, rather than directly on the specified window. In this way, you can obtain pieces of the other windows that overlap the specified window and, more importantly, capture menus or other popups that are independent windows but appear over the specified window. This option is specific to X11R5.
<b>-xy</b>	Selects xy format dumping instead of the default z format. This option applies to color displays only.

## File

Item	Description
XWDFile.h	X Window dump file format definition file.

---

## xwud Command

### Purpose

Retrieves and displays the dumped image of an Enhanced X-Windows window.

### Syntax

```
xwud [ -in FileName ] [ -noclick ] [ -geometry Geometry ] [ -display Display ] [ -new ] [ -std MapType ] [ -raw ] [ -vis visual_type | visual_id ] [ -help ] [ -rv ] [ -plane Number ] [ -fg Color ] [ -bg Color ]
```

### Description

The **xwud** command retrieves the dumped image of an Enhanced X-Windows window. It does so by displaying in a window an image saved in a specially formatted dump file previously produced by the **xwd** command. The dump file format is determined by the **XWDFile.h** file.

You can use flags to specify color display, window size and position, input field, and visual class or identification. You can also select a single bit plane of the image to display.

### Flags

Item	Description
<b>-bg</b> <i>Color</i>	Specifies the color to display for the 0 (zero) bits in the image if a bitmap image (or a single plane of an image) is displayed.
<b>-display</b> <i>Display</i>	Specifies the server to connect to; see the <b>X</b> command.
<b>-fg</b> <i>Color</i>	Specifies the color to display for the 1 bits in the image if a bitmap image (or a single plane of an image) is displayed.
<b>-geometry</b> <i>Geometry</i>	Specifies the size and position of the window. Typically, you will only specify the position and let the size default to the actual size of the image.
<b>-help</b>	Prints a short description of the allowable options.
<b>-in</b> <i>FileName</i>	Specifies the input file on the command line. If the input file is not specified, the standard input is assumed.
<b>-new</b>	Creates a new color map for displaying the image. If the image characteristics match those of the display, this flag can display the image on the screen faster, but at the cost of using a new color map (which on most terminals causes other windows to go technicolor).
<b>-noclick</b>	Prevents the application from ending when a button in the window is clicked. You can end the application by typing a q or Q character, or the Ctrl-C key sequence.
<b>-plane</b> <i>Number</i>	Selects a single bit plane of the image to display. Planes are numbered, with 0 (zero) being the least significant bit. Use this flag to determine which plane to pass to the <b>xpr</b> command for printing.
<b>-raw</b>	Displays the dumped image in whatever color values currently exist on the screen. This flag is useful when undumping an image back onto the same screen that the image originally came from, while the original windows are still on the screen. This results in getting the image on the screen faster.
<b>-rv</b>	Swaps the foreground and background colors if a bitmap image (or a single plane of an image) displays. This flag is useful when displaying a bitmap image that has the color sense of pixel values 0 and 1 reversed from what they are on the display.
<b>-std</b> <i>MapType</i>	Uses the specified Standard Colormap to display the image. You can obtain the map type by converting the type to uppercase letters, prepending <b>RGB_</b> and appending <b>_MAP</b> . Typical map types are <b>best</b> , <b>default</b> , and <b>gray</b> . See the <code>/usr/lpp/X11/Xamples/clients/xstdcmap</code> for information about creating Standard Colormaps.
<b>-vis</b> <i>visual_type</i>   <i>visual_id</i>	Specifies a particular visual type or visual id. The default picks the <b>best</b> one or you can specify <b>default</b> , which is the same class as the colormap of the root window.  You can specify a particular class: <b>StaticGray</b> , <b>GrayScale</b> , <b>StaticColor</b> , <b>PseudoColor</b> , <b>DirectColor</b> , <b>TrueColor</b> . Specify <b>Match</b> to use the same class as the source image.  Specify an exact visual id (specific to the server) as a hexadecimal number (prefixed with 0x) or as a decimal number. This string is not case sensitive.

## Environment Variables

Item	Description
<b>DISPLAY</b>	Gets the default display.

## Example

To retrieve a specific file from the dump window, enter:

```
xwud -in FileName
```

### Related reference:

“**xpr** Command” on page 202

“**xwd** Command” on page 241



---

## y

The following AIX commands begin with the with the letter *y*.

---

### yacc Command

#### Purpose

Generates an LALR(1) parsing program from input consisting of a context-free grammar specification.

#### Syntax

```
yacc [ -b Prefix ] [ -C ] [ -d ] [ -l ] [ -NnNumber ] [ -NmNumber ] [ -NrNumber ] [ -p Prefix ] [ -s ] [ -t ] [ -v ] [ -y Path ] Grammar
```

#### Description

The **yacc** command converts a context-free grammar specification into a set of tables for a simple automaton that executes an LALR(1) parsing algorithm. The grammar can be ambiguous; specified precedence rules are used to break ambiguities.

You must compile the output file, **y.tab.c**, with a C language compiler to produce a **yyparse** function. This function must be loaded with the **yylex** lexical analyzer, as well as with the **main** subroutine and the **yyerror** error-handling subroutine (you must provide these subroutines). The **lex** command is useful for creating lexical analyzers usable by the **yyparse** subroutine. Simple versions of **main** and **yyerror** subroutines are available through the **yacc** library, **liby.a**. Also, **yacc** can be used to generate C++ output.

You can compile the **yacc**-generated C file (**y.tab.c**) with the **-DYACC\_MSG** option to include code necessary to use the Message Facility. When you use this option during compilation, error messages generated by the **yyparse** subroutine and the **YYBACKUP** macro are extracted from the **yacc\_user.cat** catalog.

This allows you to receive error messages in languages other than English in non-English locales. If the catalog cannot be found or opened, the **yyparse** and **YYBACKUP** subroutines display the default English messages.

The **yacc** command is affected by the **LANG**, **LC\_ALL**, **LC\_CTYPE**, and **LC\_MESSAGES** environment variables.

#### Flags

Item	Description
<b>-b Prefix</b>	Use <i>Prefix</i> instead of <b>y</b> as the prefix for all output file names. The code file <b>y.tab.c</b> , the header file <b>y.tab.h</b> (created when <b>-d</b> is specified), and the description file <b>y.output</b> (created when <b>-v</b> is specified) are changed to <i>Prefix.tab.c</i> , <i>Prefix.tab.h</i> , and <i>Prefix.output</i> , respectively.
<b>-C</b>	Produces the <b>y.tab.C</b> file instead of the <b>y.tab.c</b> file for use with a C++ compiler. To use the I/O Stream Library for input and output, define the macro, <b>_CPP_IOSTREAMS</b> .
<b>-d</b>	Produces the file <b>y.tab.h</b> . This contains the <b>#define</b> statements that associate the <b>yacc</b> -assigned token codes with your token names. This allows source files other than <b>y.tab.c</b> to access the token codes by including this header file.
<b>-l</b>	Does not include any <b>#line</b> constructs in <b>y.tab.c</b> . Use this only after the grammar and associated actions are fully debugged.
<b>-NnNumber</b>	Changes the size of the token and nonterminal names array to <i>Number</i> . The default value is 8000. Valid values are only those greater than 8000.

Item	Description
<b>-Nm</b> <i>Number</i>	Changes the size of the memory states array to <i>Number</i> . Default value is 40000. Valid values are only those greater than 40000.
<b>-Nr</b> <i>Number</i>	Changes the internal buffer sizes to handle large grammars. The default value is 2000. Valid values are only those greater than 2000.
<b>-p</b> <i>Prefix</i>	Use <i>Prefix</i> instead of <b>yy</b> as the prefix for all external names created by the <b>yacc</b> command. External names affected include: <b>yychar</b> , <b>yyval</b> , <b>yydebug</b> , <b>yyvsparse( )</b> , <b>yylex( )</b> , and <b>yyerror( )</b> . (Previously, <b>-p</b> was used to specify an alternate parser; now, <b>-yPath</b> can be used to specify an alternate parser.)
<b>-s</b>	Breaks the <b>yyvsparse</b> function into several smaller functions. Since its size is somewhat proportional to that of the grammar, it is possible for the <b>yyvsparse</b> function to become too large to compile, optimize, or execute efficiently.
<b>-t</b>	Compiles run-time debugging code. By default, this code is not included when <b>y.tab.c</b> is compiled. However, the run-time debugging code is under the control of the preprocessor macro, <b>YYDEBUG</b> . If <b>YYDEBUG</b> has a nonzero value, the C compiler ( <b>cc</b> ) includes the debugging code, regardless of whether the <b>-t</b> flag is used. <b>YYDEBUG</b> should have a value of 0 if you don't want the debugging code included by the compiler. Without compiling this code, the <b>yyvsparse</b> subroutine will have a faster operating speed.  The <b>-t</b> flag causes compilation of the debugging code, but it does not actually turn on the debug mode. To get debug output, the <b>yydebug</b> variable must be set either by adding the C language declaration, <code>int yydebug=1</code> to the declaration section of the <b>yacc</b> grammar file or by setting <b>yydebug</b> through <b>dbx</b> .
<b>-v</b>	Prepares the file <b>y.output</b> . It contains a readable description of the parsing tables and a report on conflicts generated by grammar ambiguities.
<b>-y</b> <i>Path</i>	Uses the parser prototype specified by <i>Path</i> instead of the default <b>/usr/lib/yaccpar</b> file. (Previously, <b>-p</b> was used to specify an alternate parser.)

## Exit Status

This command returns the following exit values:

Item	Description
0	Successful completion.
>0	An error occurred.

## Examples

1. The following command:

```
yacc grammar.y
```

draws **yacc** rules from the **grammar.y** file, and places the output in **y.tab.c**.

2. The following command:

```
yacc -d grammar.y
```

functions the same as example 1, but it also produces the **y.tab.h** file which would contain C-style **#define** statements for each of the tokens defined in the **grammar.y** file.

## Files

Item	Description
y.output	Contains a readable description of the parsing tables and a report on conflicts generated by grammar ambiguities.
y.tab.c	Contains an output file.
y.tab.h	Contains definitions for token names.
yacc.tmp	Temporary file.
yacc.debug	Temporary file.
yacc.acts	Temporary file.
/usr/ccs/lib/yaccpar	Contains parser prototype for C programs.
/usr/ccs/lib/liby.a	Contains a run-time library.

#### Related information:

lex command

Generating a Lexical Analyzer with the lex command

Example program for the lex and yacc programs

## yes Command

### Purpose

Outputs an affirmative response repetitively.

### Syntax

```
yes [ charstring ]
```

### Description

The **yes** command outputs an affirmative response repetitively. Use the **yes** command as piped input to another command that requires an affirmative response before it completes the specified action. For example, the **yes** command is useful when deleting multiple files from a directory. The Ctl-C key sequence terminates the continuous affirmative responses.

**Note:** The current locale is determined by the **LC\_MESSAGES** environment variable or the *charstring* parameter, if specified. The *charstring* parameter can be any single character or character stream. If you enter an *charstring* parameter after issuing the **yes** command, the *charstring* parameter displays to the screen until you type the Ctl-C key sequence.

### Example

To display the word `first` to the screen, type:

```
yes first
```

This statement displays the word until you enter the Ctl-C key sequence.

### File

Item	Description
<code>/usr/bin/yes</code>	Contains the <code>yes</code> command.

#### Related information:

environment file

## ypbind Daemon

### Purpose

Enables client processes to bind, or connect, to an NIS server.

### Syntax

```
/usr/lib/netsh/yp/ypbind [ -s -ypset -ypsetme ]
```

### Description

The **ypbind** daemon binds, or connects, processes on a Network Information Services (NIS) client to services on an NIS server. This daemon, which runs on every NIS client, is started and stopped by the following System Resource Controller (SRC) commands:

```
startsrc -s ypbind
stopsrc -s ypbind
```

When a client requests information from a Network Information Services (NIS) map, the **ypbind** daemon broadcasts on the network for a server. When the server responds, it gives the daemon the Internet address and port number of a host. This is the host that provides the information the client is seeking. The **ypbind** daemon stores this address information in the `/var/yp/binding` directory using a file name of **domainname.version**. Then, the next time the client wants to access an NIS map, the client's **ypbind** daemon refers to the addresses in the **domainname.version** file.

The **ypbind** daemon can maintain bindings to several domains and their servers **-ypsetme** simultaneously. The default domain is the one specified by the **domainname** command at startup time.

#### Note:

1. If a domain becomes unbound (usually when the server crashes or is overloaded), the **ypbind** daemon broadcasts again to find another server.
2. To force a client to bind to a specific server, use the **ypset** command.
3. To find out which server a client is bound to, use the **ypwhich** command.
4. If the `/var/yp/binding/domainname/ypservers` file exists, **ypbind** will attempt to contact the servers listed in that file before broadcasting. The file should contain a list of server IP addresses, one per line.
5. By default, the NIS client will wait indefinitely for the NIS server, during which time, logins to the client system are not possible. It is possible, however, to limit the length of this wait. If the **YPBIND\_MAXWAIT** environment variable is set (usually in `/etc/environment`) before the **ypbind** daemon is started, this value (in seconds) will limit the amount of time the NIS client will wait for the NIS server. If this limit is exceeded, the client behaves as if NIS were unavailable and continues using local files. This will allow local logins, such as root.
6. If a domain becomes unbound and it is listed in the `/var/yp/binding/domainname/ypservers` file, by default **ypbind** daemon attempts to contact the server that is currently down; however, if the **YPBIND\_SKIP** environment variable is set to 1 (usually set in the `/etc/environment` file) before the **ypbind** daemon is started, the server that is currently down will not be contacted again.



## Flags

Item	Description
-s	Runs the <b>ypbind</b> daemon in a secure mode on privileged communications ports.
-ypset	Indicates the local host accepts <b>ypset</b> commands from local or remote hosts.
-ypsetme	Indicates that the local host accepts <b>ypset</b> commands only from the local host. This flag overrides the <b>-ypset</b> flag if both are specified.

### Note:

1. If neither the **-ypset** or **-ypsetme** flags are specified, the local host rejects all **ypset** commands from all hosts. This is the most secure mode because the NIS server cannot change.
2. If neither the **-ypset** or **-ypsetme** flags are specified, the local host rejects all **ypset** commands from all hosts. This is the most secure mode because the NIS server cannot change. However, if no NIS servers exist on the networks directly connected to the client machine, then the **-ypsetme** flag must be used and the NIS server should be specified with the **ypset** command.

## Files

Item	Description
<code>/var/yp/binding</code> directory	Contains Internet addresses and port numbers for NIS servers.
<code>/var/yp/binding/domainname/ypservers</code>	Contains a list of internet addresses, one per line, of servers to attempt to contact before broadcasting.
<code>domainname.version</code>	Binary file that contains the address and port number of the current NIS server.

### Related reference:

“ypset Command” on page 258

“ypwhich Command” on page 260

### Related information:

domainname command

System Resource Controller

---

## ypcat Command

### Purpose

Prints out a Network Information Services (NIS) map.

### Syntax

To Display the Network Information Services Database

```
/usr/bin/ypcat [ -k ] [ -t ] [ -d DomainName ] MapName
```

To Display the Nickname Translation Table

```
/usr/bin/ypcat -x
```

### Description

The **ypcat** command prints out the Network Information Services (NIS) map you specify with the *MapName* parameter. You can specify either a map name or a map nickname. Because the **ypcat** command uses the NIS service, you do not need to specify a server.

## Flags

Item	Description
-k	Displays the keys for those maps in which the values are null or for which the key is not part of the value. (None of the maps derived from files that have an ASCII version in the <i>/etc</i> directory fall into this class.)
-t	Indicates that the name specified by the <i>MapName</i> parameter is <i>not</i> a nickname. This flag causes the <b>ypcat</b> command to bypass the nickname translation table and search only for the map specified by the <i>MapName</i> parameter.
-d <i>DomainName</i>	Searches the specified domain for the specified map.
-x	Displays the nickname translation table. This table lists the map nicknames the command knows of and indicates the map name (as specified by the <i>MapName</i> parameter) associated with each nickname.

## Examples

1. To look at the networkwide password map, **passwd.byname**, type:

```
ypcat passwd
```

In this example, `passwd` is the nickname for the **passwd.byname** map.

2. To locate a map, type:

```
ypcat -t passwd
```

In this example, the **ypcat** command bypasses any maps with the nickname of `passwd` and searches for a map with the full name of `passwd`.

3. To display a map in another domain, type:

```
ypcat -d polaris passwd
```

In this example, the **ypcat** command locates the map named `passwd` in the domain named `polaris`.

4. To display the map nickname translation table, type:

```
ypcat -x
```

In this example, the **ypcat** command displays a list of map nicknames and their associated map names.

### Related information:

`domainname` command

Network File System (NFS) Overview for System Management

---

## ypinit Command

### Purpose

Sets up NIS maps on a Network Information Services (NIS) server.

### Syntax

#### To Set up NIS on an NIS Master Server

```
/usr/sbin/ypinit [ -o ] [ -n ] [ -q ] -m [ SlaveName ... ]
```

#### To Set up NIS on an NIS Slave Server

```
/usr/sbin/ypinit -s MasterName
```

## Description

The **ypinit** command sets up NIS maps on a Network Information Services (NIS) master server or NIS slave server. Only users with root user authority can use the **ypinit** command.

By default, the **ypinit** command uses the ASCII system files as input files for the map being created.

## Flags

Item	Description
<b>-m</b> [ <i>SlaveName...</i> ]	Indicates that the local host is to be the NIS master. If the <b>-q</b> flag is used the <b>-m</b> flag can be followed by the names of the machines that will be the NIS slave servers.
<b>-n</b>	Indicates that the <b>ypinit</b> command is not to stop if it finds errors.
<b>-o</b>	Allows any existing maps for the current NIS domain to be overwritten.
<b>-q</b>	Indicates that the <b>ypinit</b> command is to get arguments from the command line instead of prompting for input.
<b>-s</b> <i>MasterName</i>	Copies NIS maps from the server workstation you specify in the <i>MasterName</i> parameter.

## Examples

1. To set up an NIS master server that functions as the master for all NIS maps, type the following command on the command line:

```
ypinit -m
```

This command invokes the **make** procedure, which follows the instructions in the **/var/yp/Makefile** file.

2. To set up an NIS slave server, type:

```
ypinit -s zorro
```

In this example, the **ypinit** command copies the NIS maps onto your workstation from the NIS server named **zorro**, making your workstation an NIS slave server.

3. To set up an NIS master server without being prompted for input, type:

```
ypinit -o -n -q -m slave
```

**Note:** If the system has previously been configured as an NIS master server, ensure that the directory, **/var/yp/binding**, is removed before executing **ypinit**. If old information is stored in **/var/yp/binding**, it may cause errors to occur during configuration of the NIS master server.

## Files

Item	Description
<b>/etc/bootparams</b>	Lists clients that diskless clients can use for booting.
<b>/etc/passwd</b>	Contains an entry for each user that has permission to log on to the machine.
<b>/etc/group</b>	Contains an entry for each user group allowed to log on to the machine.
<b>/etc/hosts</b>	Contains an entry for each host on the network.
<b>/var/yp/Makefile</b>	Contains rules for making NIS maps.
<b>/etc/networks</b>	Contains the name of each network in the DARPA Internet.
<b>/etc/netmasks</b>	Lists network masks used to implement IP standard subnetting.
<b>/etc/netid</b>	Contains identification information for machines, hosts, and groups.
<b>/etc/rpc</b>	Contains map information for RPC programs.
<b>/etc/services</b>	Contains an entry for each server available through the Internet.
<b>/etc/protocols</b>	Defines Internet protocols used on the local host.
<b>/etc/netgroup</b>	Contains information about each user group on the network.
<b>/etc/ethers</b>	Contains the Ethernet addresses of hosts on the Internet network.
<b>/etc/publickey</b>	Contains public or secret keys for NIS maps.

## Related information:

## ypmatch Command

### Purpose

Displays the values of given keys within a Network Information Services (NIS) map.

### Syntax

#### To Display Key Values for an NIS Map

```
/usr/bin/ypmatch [ -d Domain ] [ -k ] [ -t ] Key... MapName
```

#### To Display the NIS Map Nickname Table

```
/usr/bin/ypmatch -x
```

### Description

The **ypmatch** command displays the values associated with one or more keys within a Network Information Services (NIS) map. Use the *MapName* parameter to specify either the name or nickname of the map you want to search.

When you specify multiple keys in the *Key* parameter, the system searches the same map for all of the keys. Because pattern matching is not available, match the capitalization and length of each key exactly. If the system does not find a match for the key or keys you specify, a diagnostic message is displayed.

### Flags

Item	Description
-d <i>Domain</i>	Specifies a domain other than the default domain.
-k	Prints a key followed by a colon before printing the value of the key. This is useful only if the keys are not duplicated in the values or if you have specified so many keys that the output could be confusing.
-t	Inhibits translation of nickname to map name.
-x	Displays the map nickname table. This lists the nicknames (as specified by the <i>MapName</i> parameter) the command knows of and indicates the map name associated with each nickname.

### Examples

To display the value associated with a particular key, type:

```
ypmatch -d ibm -k host1 hosts
```

In this example, the **ypmatch** command displays the value of the *host1* key from the *hosts* map in the *ibm* domain.

#### Related information:

NIS overview for system management  
NIS reference

---

## yppasswd Command

### Purpose

Changes your network password in Network Information Services (NIS).

### Syntax

```
yppasswd [ -f [ Name ] | -s [ Name [ ShellProg ] ]
```

### Description

The **yppasswd** command changes (or installs) a network password and associates it with the name you specify in the *UserName* parameter. To create or change a password, you must be the owner of the password you want to change. The Network Information Services (NIS) password can be different from the one on your own machine. Root users on an NIS server can change the password of another user without knowing the user's original password. To do this, the Root user enters their password in place of the user's original password. Root users on an NIS client, however, do not have this privilege.

When you enter the **yppasswd** command on the command line, the system prompts you to enter the old password. When you do this, the system prompts you to enter the new password. The password you enter can be as small as four characters long if you use a mixture of uppercase and lowercase characters. Otherwise, the password has to be six characters long or longer. These rules are relaxed if you are insistent enough.

If you enter the old password incorrectly, you have to enter the new password before the system will give you an error message. The system requires both passwords because the **update** protocol sends them to the server at the same time. The server catches the error and notifies you that you entered the old password incorrectly.

To verify the new password, the system prompts you to enter it again. For this new password to take effect, the **yppasswdd** daemon must be running on your NIS server.

**Note:** The **yppasswd** command cannot establish rules for passwords as does the **passwd** command.

### Flags

Item	Description
-f [ <i>Name</i> ]	Changes user <i>Name</i> 's gecos information in the NIS maps. Gecos information is general information stored in the <b>/etc/passwd</b> file.
-s [ <i>Name</i> [ <i>ShellProg</i> ] ]	Changes user <i>Name</i> 's login shell in the NIS maps.

### Example

1. To change a user's NIS password, enter:

```
yppasswd Joe
```

This example demonstrates how to change the NIS password for the user named Joe. The system prompts you to enter Joe's old password and then his new password.

2. To change the login shell to **/bin/ksh** for the user named Joe, if the **yppasswdd** daemon has not been started with the **-noshell** flag, enter:

```
yppasswd -s Joe /bin/ksh
```

3. To change the gecos information in the **passwd** file for the user named Joe, if the **yppasswdd** daemon has not been started with the **-nogecos** flag, enter:

```
yppasswd -f Joe
Old NIS password:
Joe's current gecocos:
John Doe Test User Id
Change (yes) or (no)? >y
To?>Joe User Test User Id
```

#### Related information:

NIS overview for system management

NIS reference

“yppasswdd Daemon”

---

## yppasswdd Daemon

### Purpose

Receives and executes requests from the **yppasswd** command.

### Syntax

```
rpc.yppasswdd FileName [ -nogecos ] [ -nopw ] [ -noshell ] [ -r | -m [ Argument... ] ]
```

### Description

The **yppasswdd** daemon is a server that receives and executes requests for new passwords from the **yppasswd** command. These requests require the daemon to verify the user's old password and change it. The daemon changes the password in the file you specify in the *FileName* parameter, which has the same format as the **/etc/passwd** file.

To make it possible to update the Network Information Services (NIS) password map from remote machines, the **yppasswdd** daemon must be running on the master server that contains the NIS password map.

**Note:** The **yppasswdd** daemon is not run by default, nor can it be started up from the **inetd** daemon like other Remote Procedure Call (RPC) daemons.

The **yppasswdd** daemon can be started and stopped with the following System Resource Controller (SRC) commands:

```
startsrc -s yppasswdd
```

```
stopsrc -s yppasswdd
```

### Flags

Item	Description
<b>-m</b>	Runs the <b>make</b> command using the makefile in the <b>/var/yp</b> directory. This adds the new or changed password to the NIS password map. Any arguments that follow the <b>-m</b> flag are passed to the <b>make</b> command.
<b>-nogecos</b>	Indicates the server will not accept changes for gecocos information from the <b>yppasswd</b> command.
<b>-nopw</b>	Indicates that the server will not accept password changes from the <b>yppasswdd</b> command.
<b>-noshell</b>	Indicates the server will not accept changes for user shells from the <b>yppasswd</b> command.
<b>-r</b>	Directly updates the <b>/var/yp/domainname/passwd.byname</b> and <b>/var/yp/domainname/passwd.byuid</b> database files on the Master server as well as any Slave servers with new or changed passwords. This option is faster than the <b>-m</b> flag because the <b>make</b> command is not run. The <b>-r</b> flag is useful when the database files are large (several thousand entries or more).

**Note:** The System Resource Controller (SRC) starts the **yppasswdd** daemon with the **-m** flag specified by default. Use the **chssys** command to change the default to the **-r** flag.

## Example

To propagate updated passwords immediately, invoke the **yppasswdd** daemon as follows:  
startsrc -s yppasswdd

## Files

Item	Description
<code>/etc/inetd.conf</code>	Defines how the <b>inetd</b> daemon handles Internet service requests.
<code>/var/yp/Makefile</code>	Contains rules for making NIS maps.
<code>/etc/rc.nfs</code>	Contains the startup script for the NFS and NIS daemons.
<code>/etc/security/passwd</code>	Stores password information.

### Related reference:

“yppasswd Command” on page 253

### Related information:

domainname command

passwd command

inetd daemon

`/etc/security/passwd` file

System Resource Controller

---

## yppoll Command

### Purpose

Displays the order number (ID number) of the Network Information Services (NIS) map currently in use on the server.

### Syntax

```
/usr/sbin/yppoll [ -h Host ] [ -d Domain ] MapName
```

### Description

The **yppoll** command uses the **ypserv** daemon to display the order number of the map you specify in the *MapName* parameter. An order number is a map's ID number and is assigned by the system. This number changes whenever a map is updated. Use the **yppoll** command whenever you want to make sure your servers are using the most current version of a particular map.

The **yppoll** command can run on systems that have either version 1 or version 2 of the Network Information Services (NIS) protocol installed. Be aware, however, that each version of the protocol has its own set of diagnostic messages.

**Note:** When specifying a *MapName*, be sure to enter the map's full name. The **yppoll** command does not recognize map nicknames.

### Flags

Item	Description
-h <i>Host</i>	Enables you to specify a server other than the default server. To find out which server the command defaults to, use the <b>ypwhich</b> command.
-d <i>Domain</i>	Enables you to specify a domain other than the default domain. To find out which domain the command defaults to, use the <b>domainname</b> command.

## Examples

1. To look at a map located on a particular host, type:

```
/usr/sbin/yppoll -h thor netgroups.byuser
```

In this example, the **yppoll** command displays the order number for the `netgroups.byuser` map located on the host named `thor`.

2. To look at a map on a domain, type:

```
/usr/sbin/yppoll -d atlantis hosts.byname
```

In this example, the **yppoll** command displays the order number for the `hosts.byname` map located in the domain `atlantis`.

### Related reference:

“ypwhich Command” on page 260

### Related information:

“ypserv Daemon” on page 257

`domainname` command

## yppush Command

### Purpose

Prompts the Network Information Services (NIS) slave servers to copy updated NIS maps.

### Syntax

```
/usr/sbin/yppush [ -v ] [ -d Domain ] MapName
```

### Description

The **yppush** command, which is issued from the `/usr/etc/yp` directory, prompts the Network Information Services (NIS) slave servers to make copies of updated NIS maps. The *MapName* variable specifies that map to be transferred to the slave servers of the master servers. To get a list of the servers it needs to prompt, the **yppush** command reads the **ypservers** map, specified by the *Domain* parameter or the current default domain. When prompted, each slave server uses the **ypxfr** command to copy and transfer the map back to its own database.

You can use the System management interface tool (SMIT) to run this command. To use SMIT, type:

```
smit yppush
```

**Note:** If your system uses version 1 of the NIS protocol, the **ypxfr** command is not the transfer agent.

### Flags



Item	Description
<code>-d Domain</code>	Specifies a domain other than the default domain. The maps for the specified domain must exist.
<code>-v</code>	Displays messages as each server is called and then displays one message for each server's response, if you are using the version 2 protocol. If this flag is omitted, the command displays error messages only. <b>Note:</b> Version 1 of the NIS protocol does not display messages. If your system uses version 1, use the <code>yppoll</code> command to verify that the transfer took place.

## Examples

1. To copy a map from another domain to the slave servers, type:

```
/usr/sbin/yppush -d atlantis netgroup
```

In this example, the `yppush` command copies the `netgroup` map from the `atlantis` domain.

2. To display the in-progress status of the `yppush` command as it calls each slave server, type:

```
/usr/sbin/yppush -v -d atlantis netgroup
```

In this example, the `yppush` command displays in-progress messages as it copies the `netgroup` map from the `atlantis` domain onto each of the network's slave servers.

## Files

Item	Description
<code>/var/yp/DomainName/ypservers.{dir, pag}</code>	Lists servers that the <code>yppush</code> command prompts to make copies of updated NIS maps.

### Related reference:

"`yppoll` Command" on page 255

"`ypxfr` Command" on page 261

### Related information:

System Management Interface Tool (SMIT)

---

## ypserv Daemon

### Purpose

Looks up information in local Network Information Services (NIS) maps.

### Syntax

```
/usr/lib/netsvc/yp/ypserv
```

### Description

The `ypserv` daemon looks up information in its local Network Information Services (NIS) maps. The operations performed by the `ypserv` daemon are defined for the implementor by the NIS Protocol Specification and for the programmer by the `/usr/include/rpcsvc/yp_prot.h` header file. Communication with the `ypserv` daemon is by means of Remote Procedure Calls (RPC).

The `ypserv` daemon runs only on server machines. The `ypserv` daemon is started and stopped by the following System Resource Controller (SRC) commands:

```
startsrc -s ypserv
```

```
stopsrc -s ypserv
```

The `ypserv` daemon performs the following operations on a specified map within an NIS domain:

Item	Description
Match	Takes a key and returns the associated value.
Get_first	Returns the first key-value pair from the map.
Get_next	Enumerates the next key-value pair in the map.
Get_all	Ships the entire NIS map to a requestor in response to a single RPC request.
Get_order_number	Supplies information about a map instead of map entries. The order number actually exists in the map as a key-value pair, but the server does not return it through the normal lookup functions. However, the pair will be visible if you examine the map with the <b>makedbm</b> command.
Get_master_name	Supplies information about a map instead of map entries. The master name actually exists in the map as a key-value pair, but the server does not return it through the normal lookup functions. However, the pair will be visible if you examine the map with the <b>makedbm</b> command.

Log information is written to the `/var/yp/ypserv.log` file if it exists when the **ypserv** daemon starts running.

If the `/var/yp/securenets` file exists, the **ypservr** command only responds to hosts within the ip range specified in this file.

## Files

Item	Description
<code>/etc/rc.nfs</code>	Contains the startup script for the NFS and NIS daemons.
<code>/var/yp/ypserv.log</code>	Contains the log for the <b>ypserv</b> daemon.

### Related information:

makedbm command

Remote Procedure Call Overview for Programming

System Resource Controller

---

## ypset Command

### Purpose

Directs a client machine to a specific server.

### Syntax

```
/usr/sbin/ypset [ -V1 ] [ -d Domain ] [ -h Host ] Server
```

### Description

The **ypset** command directs the **ypbind** daemon on the client to the **ypserv** daemon on the server. The **ypbind** daemon goes to the server you specify in the *Server* parameter to get Network Information Services (NIS) services for the domain you specify in the *Domain* parameter. The **ypbind** daemon gets the NIS services from the **ypserv** daemon on the server.

After the binding is set, it is not tested until a client process (such as the **ypcat** command or the **ypwhich** command) tries to get a binding for the domain. If the attempt to bind fails (the specified server is down or is not running the **ypserv** daemon), the **ypbind** daemon makes another attempt to bind for the same domain.

Specify either a name or an Internet Protocol (IP) address in the *Server* parameter. If you specify a name, the **ypset** command attempts to resolve the name to an IP address through the use of the NIS service. This works only if your machine has a current valid binding for the domain in question. In most cases, you should specify the server as an IP address.

In cases where several hosts on the local network are supplying NIS services, the **ypbind** daemon can rebind to another host. If a server is down or is not running the **ypserv** daemon, the **ypbind** daemon rebinds the client to another server. In this way, the network information service balances the load among the available NIS servers.

Use the **ypset** command if the network:

- Does not support broadcasting.
- Supports broadcasting but does not have an NIS server.
- Accesses a map that exists only on a particular NIS server.

An alternative to using **ypset** is to use the `/var/yp/binding/domain_name/ypservers` file. This file, if present, should contain a list of NIS servers to attempt to bind to, one server per line. If the **ypbind** daemon cannot bind to any of the servers in the **ypservers** file, then it will attempt to use the server specified by **ypset**. If that fails, it will broadcast on the subnet for a NIS server.

## Flags

Item	Description
<b>-d</b> <i>Domain</i>	Specifies a domain other than the default domain.
<b>-h</b> <i>Host</i>	Sets the binding for the <b>ypbind</b> daemon on the specified host instead of on the local host. The host can be specified as a name or as an IP address.
<b>-V1</b>	Binds the specified server for the (old) version 1 NIS protocol.

## Example

To set a server to bind on a host in a particular domain, enter:

```
ypset -d ibm -h venus mars
```

In this example, the **ypset** command causes the host named *venus* to bind to the server named *mars*.

### Related reference:

“ypcat Command” on page 249

“ypwhich Command” on page 260

### Related information:

“ypbind Daemon” on page 248

“ypserv Daemon” on page 257

---

## ypupdated Daemon

### Purpose

Updates information in Network Information Services (NIS) maps.

### Syntax

```
/usr/lib/netsvc/yp/rpc.yupdated [ -i | -s ]
```

## Description

The **ypupdated** daemon updates information in Network Information Services (NIS) maps. Before it can update information, however, the daemon consults the **updaters** file in the **/var/yp** directory to determine which NIS maps should be updated and how they should be changed.

By default, the **ypupdated** daemon requires the most secure method of authentication available to it, either DES (secure) or UNIX (insecure).

The **ypupdated** daemon is started and stopped by the following System Resource Controller (SRC) commands:

```
startsrc -s ypupdated
```

```
stopsrc -s ypupdated
```

## Flags

Item	Description
-s	Accepts only calls authenticated using the secure Remote Procedure Call (RPC) mechanism (AUTH_DES authentication). This disables programmatic updating of NIS maps unless the network supports these calls.
-i	Accepts RPC calls with the insecure AUTH_UNIX credentials. This allows programmatic updating of NIS maps in all networks.

## Examples

To start the **ypupdated** daemon from the command line, type:

```
startsrc -s ypupdated
```

## File

Item	Description
<b>/var/yp/updaters</b>	A makefile for updating NIS maps.

### Related information:

startsrc command

Remote Procedure Call Overview for Programming

System Resource Controller

---

## ypwhich Command

### Purpose

Identifies either the Network Information Services (NIS) server or the server that is the master for a given map.

### Syntax

#### To Identify the NIS Server

```
/usr/bin/ypwhich [ -d Domain ] [ -V1 | -V2 ] [ HostName ]
```

#### To Identify the Master NIS Server for a Map

```
/usr/bin/ypwhich [ -t ] [ -d Domain ] [ -m [ MapName ] ]
```

#### To Display the Map Nickname Table

`/usr/bin/ypwhich -x`

## Description

The **ypwhich** command identifies which server supplies Network Information Services (NIS) services or which server is the master for a map, depending on how the **ypwhich** command is invoked. If invoked without arguments, this command displays the name of the NIS server for the local machine. If you specify a host name, the system queries that host to find out which master it is using.

## Flags

Item	Description
<code>-d Domain</code>	Uses the specified domain instead of the default domain.
<code>-V1</code>	Indicates which server is serving the old version 1 NIS protocol client processes.
<code>-V2</code>	Indicates which server is serving the current version 2 NIS protocol client processes. If neither version is specified, the <b>ypwhich</b> command attempts to locate the server that supplies the version 2 services. If there is no version 2 server currently bound, the <b>ypwhich</b> command then attempts to locate the server supplying version 1 services. Because servers and clients are both backward-compatible, the user need seldom be concerned about which version is currently in use.
<code>-t</code>	Inhibits nickname translation, which is useful if there is a map name identical to a nickname.
<code>-m MapName</code>	Finds the master NIS server for a map. No host can be specified with the <code>-m</code> flag. The <i>MapName</i> variable can be a map name or a nickname for a map. When the map name is omitted, the <code>-m</code> flag produces a list of available maps.
<code>-x</code>	Displays the map nickname table. This lists the nicknames ( <i>MapName</i> ) the command knows of and indicates the map name associated with each nickname.

## Examples

1. To find the master server for a map, type:

```
ypwhich -m passwd
```

In this example, the **ypwhich** command displays the name of the server for the passwd map.

2. To find the map named passwd, rather than the map nicknamed passwd, type:

```
ypwhich -t -m passwd
```

In this example, the **ypwhich** command displays the name of the server for the map whose full name is passwd.

3. To find out which server serves clients that run the old version 1 of the NIS protocol, type:

```
ypwhich -V1
```

4. To display a table of map nicknames, type:

```
ypwhich -x
```

### Related information:

NIS overview for system management

NIS reference

---

## ypxfr Command

### Purpose

Transfers a Network Information Services (NIS) map from an NIS server to a local host.

### Syntax

```
/usr/sbin/ypxfr [ -f ] [ -c ] [ -d Domain ] [ -h Host ] [ -s Domain ] [ -C TID Program IPAddress Port ] [ -S ] MapName
```

## Description

The **ypxfr** command transfers a Network Information Services (NIS) map from an NIS server to the local host as follows:

1. Creates a temporary map in the */var/yp/Domain* directory (which must already exist) on the client.
2. Fetches the map entries from the server and fills in the map on the client, one at a time.
3. Gets and loads the map parameters (order number and server).
4. Deletes any old versions of the map.
5. Moves the temporary map to the real map name.

If the */var/yp/securenets* file exists, the **ypxfr** command only responds to hosts that are listed in this file.

The *MapName* variable specifies the name of a map that will be transferred from an NIS server.

If run interactively, the **ypxfr** command sends output to the user's terminal. If invoked without a controlling terminal, the **ypxfr** command appends its output to the */var/yp/ypxfr.log* file (if the file already exists). This file records each transfer attempt and its results. The **ypxfr** command is most often invoked from the root user's **crontab** file or by the **ypserv** daemon.

To maintain consistent information between servers, use the **ypxfr** command to update every map in the NIS database periodically. Be aware though that some maps change more frequently than others and therefore need to be updated more frequently. For instance, maps that change infrequently, such as every few months, should be updated at least once a month. Maps that change frequently, such as several times a day, should be checked hourly for updates. The **services.byname** map, for example, may not change for months at a time, while the **hosts.byname** map may change several times a day.

To perform periodic updates automatically, use a **crontab** entry. To update several maps at one time, group commands together in a shell script. Examples of a shell script can be found in the */usr/etc/yp* directory in the following files: **ypxfr\_1perday**, **ypxfr\_2perday**, **ypxfr\_1perhour**.

You can use the System management interface tool (SMIT) to run this command. To use SMIT, enter:

```
smit ypxfr
```

## Flags

Item	Description
<b>-C</b> <i>TID Program IPAddress Port</i>	Tells the <b>ypxfr</b> command where to find the <b>yppush</b> command. The <b>ypserv</b> daemon invokes the <b>ypxfr</b> command to call back a <b>yppush</b> command to the host. Use the parameters to indicate the following:  <i>TID</i> Specifies the transaction ID of the <b>yppush</b> command.  <i>Program</i> Specifies the program number associated with the <b>yppush</b> command.  <i>IPAddress</i> Specifies the Internet Protocol address of the port where the <b>yppush</b> command resides.  <i>Port</i> Specifies the port that the <b>yppush</b> command is listening on. <b>Note:</b> This option is only for use by the <b>ypserv</b> daemon.
<b>-c</b>	Prevents sending of a request to Clear Current Map to the local <b>ypserv</b> daemon. Use this flag if the <b>ypserv</b> daemon is not running locally at the time you are running the <b>ypxfr</b> command. Otherwise, the <b>ypxfr</b> command displays an error message and the transfer fails.
<b>-d</b> <i>Domain</i>	Specifies a domain other than the default domain. The maps for the specified domain must exist.
<b>-f</b>	Forces the transfer to occur even if the version at the master is not more recent than the local version.

<b>Item</b>	<b>Description</b>
<b>-h</b> <i>Host</i>	Gets the map from host specified, regardless of what the map says the master is. If a host is not specified, the <b>ypxfr</b> command asks the NIS service for the name of the master and tries to get the map from there. The <i>Host</i> variable can contain a name or an Internet address in the form a.b.c.d.
<b>-S</b>	Requires the <b>ypserv</b> server, from which it obtains the maps to be transferred, use <i>privileged</i> IP ports. Because only root user processes are typically allowed to use privileged ports, this feature adds an extra measure of security to the transfer. If the map being transferred is a secure map, the <b>ypxfr</b> command sets the permissions on the map to 0600.
<b>-s</b> <i>Domain</i>	Specifies a source domain from which to transfer a map that should be the same across domains (such as the <b>services.byname</b> map).

## Examples

To get a map from a host in another domain, enter:

```
/usr/sbin/ypxfr -d ibm -h venus passwd.byname
```

In this example, the **ypxfr** command gets the `passwd.byname` map from the host name `venus` in the `ibm` domain.

## Files

<b>Item</b>	<b>Description</b>
<code>/var/yp/ypxfr.log</code>	Contains the log file.
<code>/usr/sbin/ypxfr_1perday</code>	Contains the script to run one transfer each day, for use with the <b>cron</b> daemons.
<code>/usr/sbin/ypxfr_2perday</code>	Contains the script to run two transfers each day.
<code>/usr/sbin/ypxfr_1perhour</code>	Contains the script for hourly transfers of volatile maps.

### Related reference:

“yppush Command” on page 256

### Related information:

System Management Interface Tool (SMIT)

crontab command

cron command

“ypserv Daemon” on page 257





---

## Z

The following AIX commands begin with the with the letter z.

---

### zcat Command

#### Purpose

Expands a compressed file to standard output.

#### Syntax

```
zcat [ -n ] [ -V ] [ File ... ]
```

#### Description

The **zcat** command allows the user to expand and view a compressed file without uncompressing that file. The **zcat** command does not rename the expanded file or remove the **.Z** extension. The **zcat** command writes the expanded output to standard output.

#### Flags

Item	Description
-n	Omits the compressed file header from the compressed file. <b>Note:</b> Use the -n option if the file was compressed using the -n option.
-V	Writes the current version and compile options to standard error.

#### Parameters

Item	Description
<i>File ...</i>	Specifies the compressed files to expand.

#### Return Values

If the **zcat** command exits with a status of 1 if any of the following events occur:

- The input file was not produced by the **compress** command.
- An input file cannot be read or an output file cannot be written.

If no error occurs, the exit status is 0.

#### Exit Status

Item	Description
0	Successful completion.
>0	An error occurred.

## Examples

To view the `foo.Z` file without uncompressing it, enter:

```
zcat foo.Z
```

The uncompressed contents of the `foo.Z` file are written to standard output. The file is not renamed.

### Related information:

compress command

## zdump Command

### Purpose

Displays the time zone information.

### Syntax

```
zdump [--version] [--help] [-v] [-V] [-c [loyear,hiyear] [-t [lotime,hitime] zonename ...
```

### Description

The **zdump** command prints the time zone information for each zone name that is specified in the command line.

### Flags

Item	Description
<b>--version</b>	Displays the version information of the <b>zdump</b> command.
<b>-v</b>	For each zone name that is specified in the command line, the <b>zdump</b> command prints the following items: <ul style="list-style-type: none"> <li>• The lowest possible value of time.</li> <li>• The time at one day after the lowest possible value of time.</li> <li>• The times both one second before and exactly at each detected discontinuity of time.</li> <li>• The highest possible value of time.</li> <li>• The time at one day before the highest possible value of time.</li> </ul> Each line is followed by <code>isdst=<i>D</i></code> , where <i>D</i> is a positive, zero, or negative value that depends on whether the specified time is Daylight Saving Time, standard time, or an unknown time type. Each line is also followed by <code>gmtoff=<i>N</i></code> , where the specified local time is <i>N</i> seconds east of Greenwich.
<b>-V</b>	Prints an output that is similar to the output of the <b>-v</b> flag except that this flag does not display the times that are relative to the extreme time values. This flag generates output that is easier to compare with different time representations.
<b>-c</b> [ <i>loyear</i> , <i>hiyear</i> ]	Reports the verbose output for the specified years. Cutoff time is computed by using the proleptic Gregorian calendar with year 0 and Universal Time (UT) ignoring leap seconds. The lower year attribute ( <i>loyear</i> ) is exclusive and the upper year attribute ( <i>hiyear</i> ) is inclusive. For example, a <i>loyear</i> value of 1970 excludes a transition that occurs at 1970-01-01 00:00:00 Coordinated Universal Time, but a <i>hiyear</i> value of 1970 includes the transition. The default cutoff time is -500,2500.
<b>-t</b> [ <i>lotime</i> , <i>hitime</i> ]	Reports verbose output for a particular time that is specified in the <i>lotime</i> and <i>hitime</i> attributes in the syntax of <code>yyyy-mm-dd hh:mm:ss</code> Coordinated Universal Time (UTC). For example, 1970-01-01 00:00:00. The cutoff time includes the leap seconds depending on the <i>zonename</i> parameter.
	Similar to the <b>-c</b> flag, the lower time limit is exclusive and the upper time limit is inclusive.

### Parameters

Item	Description
<i>zonename</i>	Represents the name of the zone of which the time zone information is displayed.

## Exit Status

- 0 The command completed successfully.
- >0 An error occurred.

## Examples

- To report time zone information for Singapore, enter the following command:  

```
zdump -v Singapore
```
- To display time zone information for Turkey stopping near the start of the year 2035, enter the following command:  

```
zdump -v -c 2035 Turkey
```
- To report the time zone information for New York during the years 2015 - 2017, enter the following command:  

```
zdump -v -c 2015,2017 America/New_York
```

The output will be similar to the following sample:

```
America/New_York Fri Dec 13 20:45:52 1901 UT = Fri Dec 13 15:45:52 1901 EST isdst=0 gmtoff=-18000
America/New_York Sat Dec 14 20:45:52 1901 UT = Sat Dec 14 15:45:52 1901 EST isdst=0 gmtoff=-18000
America/New_York Sun Mar 8 06:59:59 2015 UT = Sun Mar 8 01:59:59 2015 EST isdst=0 gmtoff=-18000
America/New_York Sun Mar 8 07:00:00 2015 UT = Sun Mar 8 03:00:00 2015 EDT isdst=1 gmtoff=-14400
America/New_York Sun Nov 1 05:59:59 2015 UT = Sun Nov 1 01:59:59 2015 EDT isdst=1 gmtoff=-14400
America/New_York Sun Nov 1 06:00:00 2015 UT = Sun Nov 1 01:00:00 2015 EST isdst=0 gmtoff=-18000
America/New_York Sun Mar 13 06:59:59 2016 UT = Sun Mar 13 01:59:59 2016 EST isdst=0 gmtoff=-18000
America/New_York Sun Mar 13 07:00:00 2016 UT = Sun Mar 13 03:00:00 2016 EDT isdst=1 gmtoff=-14400
America/New_York Sun Nov 6 05:59:59 2016 UT = Sun Nov 6 01:59:59 2016 EDT isdst=1 gmtoff=-14400
America/New_York Sun Nov 6 06:00:00 2016 UT = Sun Nov 6 01:00:00 2016 EST isdst=0 gmtoff=-18000
America/New_York Mon Jan 18 03:14:07 2038 UT = Sun Jan 17 22:14:07 2038 EST isdst=0 gmtoff=-18000
America/New_York Tue Jan 19 03:14:07 2038 UT = Mon Jan 18 22:14:07 2038 EST isdst=0 gmtoff=-18000
```

- To report the time zone information for New York during the years 2015 - 2017 without the lowest and highest time values, enter the following command:  

```
zdump -V -c 2015,2017 America/New_York
```

The output will be similar to the following sample:

```
America/New_York Sun Mar 8 06:59:59 2015 UT = Sun Mar 8 01:59:59 2015 EST isdst=0 gmtoff=-18000
America/New_York Sun Mar 8 07:00:00 2015 UT = Sun Mar 8 03:00:00 2015 EDT isdst=1 gmtoff=-14400
America/New_York Sun Nov 1 05:59:59 2015 UT = Sun Nov 1 01:59:59 2015 EDT isdst=1 gmtoff=-14400
America/New_York Sun Nov 1 06:00:00 2015 UT = Sun Nov 1 01:00:00 2015 EST isdst=0 gmtoff=-18000
America/New_York Sun Mar 13 06:59:59 2016 UT = Sun Mar 13 01:59:59 2016 EST isdst=0 gmtoff=-18000
America/New_York Sun Mar 13 07:00:00 2016 UT = Sun Mar 13 03:00:00 2016 EDT isdst=1 gmtoff=-14400
America/New_York Sun Nov 6 05:59:59 2016 UT = Sun Nov 6 01:59:59 2016 EDT isdst=1 gmtoff=-14400
America/New_York Sun Nov 6 06:00:00 2016 UT = Sun Nov 6 01:00:00 2016 EST isdst=0 gmtoff=-18000
```

## Files

Item	Description
<i>/usr/sbin/zdump</i>	Contains the SystemV <b>zdump</b> command.
<i>/usr/share/lib/zoneinfo</i>	Contains the standard time zone directory.

---

## zic Command

### Purpose

Provides a time zone compiler.

## Syntax

```
zic [ -v ] [ -l LocalTime ] [ -p Posixrules ] [ -d Directory ] [ -L Leapseconds ] [ -y YearIsType ] [ FileName ... ]
```

## Description

The **zic** command processes text from the files named on the command line and creates the time conversion binary files specified in this input. The time zone information is processed from the standard input if file name is specified as - (hyphen).

Input lines in the specified files are made up of fields. Field separators are any number of white space characters. A pound sign (#) in the input file indicates a comment that extends to the end of the line on which the pound sign appears. White space characters and pound signs can be enclosed in double quotes (") if they are to be used as part of a field. Blank lines are ignored.

A rule line has the form:

```
Rule NAME FROM TO TYPE IN ON AT SAVE LETTER/S
```

The fields that make up the rule line are as follows:

### NAME

Provides a random name for the set of rules for which this Rule is applicable.

### FROM

Provides the first year in which the rule applies. Specifying **minimum** (min) indicates the minimum year with a representable time value. Specifying **maximum** (max) indicates the maximum year with a representable time value.

### TO

Provides the final year in which the rule applies. This is indicated as a valid time value or by specifying **minimum** and **maximum**. Specifying **only** is used to repeat the value of the **FROM** field.

### TYPE

Provides the type of year in which the rule applies.

The **TYPE** field has the following values:

**'** The rule applies in all years between **FROM** and **TO**, inclusively.

**uspres** The rule applies in U.S. Presidential election years.

### **nonpres**

The rule applies in years other than U.S. Presidential election years.

If the **TYPE** field has a value other than what is mentioned above, the **zic** command runs the **/usr/sbin/yearistype year type** command to check the type of year.

The **yearistype** command accepts two parameters; the year and the type of year. An exit status of 0 is taken to mean that the year is of the given type. Otherwise, 1 is returned as exit status.

### IN

Represents the month in which the rule takes effect. Month names may be abbreviated.

### ON

Represents the day on which the rule takes effect. Recognized forms include:

- **lastFri** represents the last Friday in the month.
- **lastMon** represents the last Monday in the month.
- A number representing the day of the month. For example, 5 represents the fifth of the month.
- **lastSun** represents the last Sunday in the month.
- **lastMon** represents the last Monday in the month.
- **Sun>=8** represents first Sunday on or after the eighth.
- **Sun<=25** represents last Sunday on or before the 25th.

Names of days of the week can be abbreviated or spelled out in full. Note that there must be no spaces within the **ON** field.

- AT** Represents the time of day at which the rule takes effect. Recognized forms include:
- A number representing time in hours. For example, 2 indicates two hours.
  - 2:00 indicates two o'clock in hours and minutes.
  - 15:00 represents 3 o'clock in the afternoon using the 24-hour format time format.
  - 1:28:14 indicates one o'clock, twenty-eight minutes and fourteen seconds, using the hours, minutes, seconds format.

Any of these forms may be followed by the letter *w* if the given time is local *wall-clock* time or *s* if the given time is local *standard* time. In the absence of *w* or *s*, wall-clock time is assumed.

Regions with more than two types of local time are required to use the local standard time in the **AT** field of the earliest transition time's rule to ensure the accuracy of the earliest transition time that is stored in the resulting time-zone binary.

- SAVE** Represents the amount of time to be added to local standard time when the rule is in effect. This field has the same format as the **AT** field. The *w* and *s* suffixes are not valid with this field.

#### **LETTER/S**

Provides the *variable part* of the time-zone abbreviations that are used when this rule is in effect. When this field contains - (hyphen), the variable is null. The *S* character is used to indicate EST and the *D* character is used to indicate EDT.

A zone line has the form:

```
Zone  NAME  GMTOFF  RULES/SAVE  FORMAT  [UNTIL]
```

The fields that make up a zone line are:

#### **NAME**

Indicates the name of the time zone. This is the name used to create the time conversion information file for the zone.

#### **GMTOFF**

Indicates the amount of time to add to GMT to get standard time in this zone. This field has the same format as the **AT** and **SAVE** fields of rule lines. Begin the field with a minus sign if time must be subtracted from GMT.

#### **RULES/SAVE**

Indicates the name of the rules that apply in the time zone or, alternately, an amount of time to add to local standard time. If value of this field is - (hyphen), then standard time always applies in the time zone.

#### **FORMAT**

Indicates the format for time zone abbreviations in this time zone. The *%s* characters are used to show where the variable part of the time zone abbreviation goes.

#### **UNTIL**

Indicates the time at which the GMT offset or the rules change for a location. It is specified as year, month, day, and time of day. If this is specified, the time zone information is generated from the given GMT offset and rule change until the time specified.

The next line must be a *continuation* line. The continuation line places information starting at the time specified in the **UNTIL** field of the previous line into the file used by the previous line. This line has the same format as a zone line, except that the *Zone* string and the name are omitted. Continuation lines can contain an **UNTIL** field, just as zone lines do, indicating that the next line is a further continuation.

A link line has the form:

Link LINK-FROM LINK-TO

The **LINK-FROM** field should appear as the **NAME** field in a zone line; the **LINK-TO** field is used as an alternate name for that zone.

Except for continuation lines, lines can appear in any order in the input.

The **zic** command has a limitation of compiling input containing a date before 14 December 1901 because dates before this time cannot be represented by a 32-bit **time\_t** data type.

## Flags

Item	Description
<b>-d</b> <i>Directory</i>	Creates time conversion information files in the <i>Directory</i> directory, instead of the <b>/usr/share/lib/zoneinfo/</b> standard directory .
<b>-l</b> <i>TimeZone</i>	Use the <i>TimeZone</i> time zone as local time. The <b>zic</b> command acts as if the file contained a link similar to the following: Link <i>timezone</i> <i>localtime</i>
<b>-L</b> <i>Leapseconds</i>	Reads the leap second information from the <b>leapseconds</b> file. If this option is not used, leap second information does not appear in the output.
<b>-p</b> <i>Posixrules</i>	Use the <i>posixrules</i> rules when handling POSIX-format time zone environment variables. The <b>zic</b> command acts as if the file contains a link as in this example: Link <i>timezone</i> <i>posixrules</i>
<b>-v</b>	Provides a message if a year that appears in a data file is outside the range of years representable by system time values (0:00:00 AM GMT, January 1, 1970, to 3:14:07 AM GMT, January 19, 2038).
<b>-y</b> <i>YearIsType</i>	Uses the given <b>yearistype</b> command rather than <b>/usr/sbin/yearistype</b> command when used to check year types.

## Parameters

Item	Description
<i>FileName</i>	A file containing input lines that specify the time conversion information files to be created. If <i>FileName</i> is - (hyphen), then standard input is read.

## Examples

1. A rule line can have the following format:  
Rule USA 1970 max - Sep Sun<=14 3:00 0 S
2. A zone line can have the following format:  
Zone Turkey 3:00 Turkey EET%s
3. A link line can have the following format:  
Link MET CET
4. To compile a **timezone.infile** file containing input time zone information and place the binaries into the standard time zone **/usr/share/lib/zoneinfo/** directory, type:  
zic timezone.infile
5. To compile a **timezone.infile** file containing input time zone information and place the binaries into a directory specified with **-d** option , type:  
zic -d tzdir timezone.infile
6. To report warnings during compilation of the time zone input file when the range of years are incorrect , type:  
zic -v timezone.infile
7. To compile a **timezone.infile** file that contains input time zone information using the **yearistype** file specified with **-y** flag to check year types, type:  
zic -y year timezone.infile

## Exit Status

- 0 The command completed successfully.
- >0 An error occurred.

## Files

Item	Description
<code>/usr/sbin/yearistype</code>	Contains the <code>yearistype</code> command used to check year types.
<code>/usr/sbin/zic</code>	Contains the SystemV <code>zic</code> command.
<code>/usr/share/lib/zoneinfo</code>	Standard directory used for files create by the <code>zic</code> command.





---

## Notices

This information was developed for products and services offered in the US.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:

*IBM Director of Licensing  
IBM Corporation  
North Castle Drive, MD-NC119  
Armonk, NY 10504-1785  
US*

For license inquiries regarding double-byte character set (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

*Intellectual Property Licensing  
Legal and Intellectual Property Law  
IBM Japan Ltd.  
19-21, Nihonbashi-Hakozakicho, Chuo-ku  
Tokyo 103-8510, Japan*

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some jurisdictions do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM websites are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this IBM product and use of those websites is at your own risk.

IBM may use or distribute any of the information you provide in any way it believes appropriate without incurring any obligation to you.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

*IBM Director of Licensing  
IBM Corporation  
North Castle Drive, MD-NC119  
Armonk, NY 10504-1785  
US*

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement or any equivalent agreement between us.

The performance data and client examples cited are presented for illustrative purposes only. Actual performance results may vary depending on specific configurations and operating conditions.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

Statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

All IBM prices shown are IBM's suggested retail prices, are current and are subject to change without notice. Dealer prices may vary.

This information is for planning purposes only. The information herein is subject to change before the products described become available.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to actual people or business enterprises is entirely coincidental.

#### COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. The sample programs are provided "AS IS", without warranty of any kind. IBM shall not be liable for any damages arising out of your use of the sample programs.

Each copy or any portion of these sample programs or any derivative work must include a copyright notice as follows:

© (your company name) (year).

Portions of this code are derived from IBM Corp. Sample Programs.

© Copyright IBM Corp. \_enter the year or years\_.

---

## Privacy policy considerations

IBM Software products, including software as a service solutions, (“Software Offerings”) may use cookies or other technologies to collect product usage information, to help improve the end user experience, to tailor interactions with the end user or for other purposes. In many cases no personally identifiable information is collected by the Software Offerings. Some of our Software Offerings can help enable you to collect personally identifiable information. If this Software Offering uses cookies to collect personally identifiable information, specific information about this offering’s use of cookies is set forth below.

This Software Offering does not use cookies or other technologies to collect personally identifiable information.

If the configurations deployed for this Software Offering provide you as the customer the ability to collect personally identifiable information from end users via cookies and other technologies, you should seek your own legal advice about any laws applicable to such data collection, including any requirements for notice and consent.

For more information about the use of various technologies, including cookies, for these purposes, see IBM’s Privacy Policy at <http://www.ibm.com/privacy> and IBM’s Online Privacy Statement at <http://www.ibm.com/privacy/details> the section entitled “Cookies, Web Beacons and Other Technologies” and the “IBM Software Products and Software-as-a-Service Privacy Statement” at <http://www.ibm.com/software/info/product-privacy>.

---

## Trademarks

IBM, the IBM logo, and [ibm.com](http://www.ibm.com) are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the web at [www.ibm.com/legal/copytrade.shtml](http://www.ibm.com/legal/copytrade.shtml).

Windows is a trademark of Microsoft Corporation in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.



---

# Index

## A

- accounting system
  - correcting format problems 127
- aliases
  - locating 90
- ARPANET
  - user of whois command 97
- ATE program
  - transferring files with xmodem 184
  - xmodem command 184
- authorization information
  - editing and displaying
    - using xauth command 147

## B

- bindings
  - configures virtual key
    - using xmbind command 181

## C

- command
  - wparprnterr 121
  - X 131
- commands
  - event response resource manager (ERRM)
    - wallevent 78
  - vacation 1
  - vpdadd 59
  - vsdchgserver 63
  - vsdelnode 65
  - vsdnode 67
  - vsdsklst 69
  - w 75
  - wall 77
  - wallevent 78
  - who 92
  - wlmassign 100
  - wlmcntrl 105
  - wlmstat 107
  - xargs 144
  - xdm 153
  - xmbind 181
  - xmodmap 186
  - xntpd 195
  - xterm 219
  - ypcat 249
  - ypinit 250
  - yppasswd 253
  - yppoll 255
  - yppush 256
  - ypset 258
  - ypwhich 260
  - ypxfr 261
- communication channel
  - receiving mail in a secure
    - using xget command 169
  - sending mail in a secure
    - using xsend command 211

- configuring
  - virtual key bindings
    - using xmbind command 181

## D

- daemons
  - yppbind 248
  - yppasswdd 254
  - ypserv 257
  - ypupdated 259
- dump file
  - formatting for printer output
    - using xpr command 202

## E

- ERRM commands
  - elogevent 78
- ERRM scripts
  - wallevent 78
- event response resource manager (ERRM)
  - commands
    - wallevent 78
  - scripts
    - wallevent 78

## F

- files
  - compression 266
  - executable
    - locating 90
  - expansion 265
  - locating sections 89
  - SCCS
    - displaying identifying information 83
    - validating 3
- font directories
  - adding 142
- font path element
  - removing 143
- font servers
  - adding 141
- fonts
  - supplying to X Window display servers 167

## G

- games
  - hunt the wumpus 128

## H

- help
  - describing command functions 84
- host machine
  - controlling access
    - using xhost command 171

## I

- image
  - displaying
    - using xwud command 242
  - dumping
    - using xwd command 241
  - retrieving
    - using xwud command 242

## K

- key bindings
  - configures virtual
    - using xmbind command 181
- keywords
  - SCCS
    - substituting values 9

## L

- locking X display until password is entered
  - using xlock command 178
- logged in, users
  - identifying
    - using who command 92

## M

- mail
  - receiving in a securing communication channel
    - using xget command 169
  - sending in a secure communication channel
    - using xsend command 211
- memory management
  - reporting virtual memory statistics 46
- messages
  - listing the addresses of recipients of
    - using whom command 98
  - prompting for the disposition of
    - using whatnow command 85
  - receiving from a remote system
    - using writesrv command 126
  - sending from a remote system
    - using writesrv command 126
  - sending to other users
    - using write command 122
  - verifying the addresses of recipients of
    - using whom command 98
- MH commands
  - invoking a visual interface for use with
    - using vmh command 32

## N

- network information service 248
- Network Time Protocol command
  - xntpd 192
- NIS commands
  - ypcat 249
  - ypinit 250
  - ypmatch 252
  - yppasswd 253
  - yppoll 255
  - yppush 256
  - ypset 258

### NIS commands (*continued*)

- ypwhich 260
- ypxfr 261
- NIS daemons
  - ypbind 248
  - yppasswdd 254
  - ypserv 257
  - ypupdated 259

## P

- parser
  - creating with the yacc command 245
- path names
  - executable files
    - finding 90
- process suspension
  - suspending execution 76
- program
  - monitoring
    - using watch command 80
- program listing
  - formatting
    - using vgrind command 11
- programs
  - creating a Makefile from an Imakefile 182

## R

- realtime messages
  - sending to other users
    - using write command 122
- repetitive responses
  - generating 247
- resource database
  - performing utilities for X server
    - using xrdb command 208

## S

- SCCS
  - files
    - displaying identifying information 83
    - validating 3
  - keywords
    - substituting values 9
- SCCS commands
  - val 3
  - vc 9
  - what 83
- screen color characterization data
  - loading from root window
    - using xcmsdb command 152
  - queries
    - using xcmsdb command 152
  - removing from root window
    - using xcmsdb command 152
- screen lock
  - controlling
    - using xss Command 216
- scripts
  - event response resource manager (ERRM)
    - wallevent 78
  - wallevent 78

- security
  - locking workstation screens
    - using xss Command 216
- shell scripts
  - repetitive responses
    - generating 247
- shells
  - locating executable files in C 90
- strings, shared
  - using xstr command 217

## T

- terminal emulator
  - providing for X Window System
    - using the xterm Command 219
- time
  - displaying the current
    - using xclock command 150

## U

- user name directory
  - searching for ID
    - using whois command 97
  - searching for nickname
    - using whois command 97
- users
  - displaying login name 95
  - identifying those logged in
    - using who command 92

## V

- vacation command 1
- varyoffvg command 4
- varyonvg command 5
- vgrind command 11
- vi command 13
- vi editor
  - starting 32
- view command 32
- Virtual Frame Buffer (VFB)
  - starting X Server with 131
- virtual key bindings
  - configuring
    - using xmbind command 181
- virtual memory
  - reporting statistics 46
- vmh command 32
- vmo command 34
- volume group
  - activating
    - using varyonvg command 5
  - deactivating
    - using varyoffvg command 4
- vpdadd command 59
- vsdata1st command 61
- vsdchgserver command 63
- vsdelnode command 65
- vsdnode command 67
- vsdsk1st command 69
- vsdvgts command 73

## W

- w command 75
- wall command 77
- wallevent command 78
- wallevent script 78
- watch command 80
- wc command 81
- whatis command 84
- whatnow command 85
- which\_fileset command 91
- who command 92
- whodo command 95
- whois command
  - ARPANET
    - use on 97
    - description of 97
    - example of 98
- wlmassign command 100
- wlmcheck command 103
- wlmcntrl command 105
- wlmstat command 107
- wol command 112
- workstation screens
  - locking
    - using xss Command 216
- wpar\_reg\_script 120
- wparerr command 113
- wparexec
  - creates an application workload partition or specification file 114
- wparprnterr command 121
- write command 122
- writesrv command 126
- wtmpfix command 127

## X

- x command 131
- X display
  - displaying troff files on
    - using xpreview command 204
  - locking until password is entered
    - using xlock command 178
- X render extension
  - starting X Server with 131
- X Server
  - starting 131
- x\_add\_fs\_fpe command 141
- x\_add\_nfs\_fpe command 142
- x\_rm\_fpe command 143
- X-Windows
  - setting environment options 212
  - setting root window parameters 215
- xargs command 144
- xauth command 147
- xclock command 150
  - setting the defaults 151
- xcmsdb Command 152
- xdm command 153
- xfs command 167
- xget command 169
- xhost command 171
- xinit command 172
- xlock command 178
- xlsfonts command 180
- xmbind command 181

- xmodem command 184
- xmodem protocol 184
- xmodmap command 186
- xmpeek command 188
- xmscheck command 189
- xmwlm command 183
- xntpd daemon 192
- xntpdc command 195
- xpr command 202
- xpreview command 204
- xprofiler 207
- xrdb command 208
- xsend command 211
- xsetroot command 215
- xss Command 216
- xstr command 217
- xterm command 219
  - providing terminal emulation  
for X Windows system 219
- xwd command 241
- xwud command 242

## Y

- yacc command 245
- ypbind daemon 248
- ypcat command 249
- ypinit command 250
- ypmatch command 252
- yppasswd command 253
- yppasswdd daemon 254
- yppoll command 255
- yppush command 256
- ypserv daemon 257
- ypset command 258
- ypupdated daemon 259
- ypwhich command 260
- ypxfr command 261

## Z

- zcat command 265
- zdump command 266
- zic command 267







Printed in USA