

Virtual I/O Server and Integrated Virtualization Manager commands

ESCALA Power7



REFERENCE
86 A1 39FF 05

ESCALA Power7

Virtual I/O Server and Integrated Virtualization Manager commands

The ESCALA Power7 publications concern the following models:

- Bull Escala E5-700 (Power 750 / 8233-E8B)
- Bull Escala M6-700 (Power 770 / 9117-MMB)
- Bull Escala M6-705 (Power 770 / 9117-MMC)
- Bull Escala M7-700 (Power 780 / 9179-MHB)
- Bull Escala M7-705 (Power 780 / 9179-MHC)
- Bull Escala E1-700 (Power 710 / 8231-E2B)
- Bull Escala E1-705 (Power 710 / 8231-E1C)
- Bull Escala E2-700 / E2-700T (Power 720 / 8202-E4B)
- Bull Escala E2-705 / E2-705T (Power 720 / 8202-E4C)
- Bull Escala E3-700 (Power 730 / 8231-E2B)
- Bull Escala E3-705 (Power 730 / 8231-E2C)
- Bull Escala E4-700 / E4-700T (Power 740 / 8205-E6B)
- Bull Escala E4-705 (Power 740 / 8205-E6C)

References to Power 755 / 8236-E8C models are irrelevant.

Hardware

May 2012

BULL CEDOC
357 AVENUE PATTON
B.P.20845
49008 ANGERS CEDEX 01
FRANCE

REFERENCE
86 A1 39FF 05

The following copyright notice protects this book under Copyright laws which prohibit such actions as, but not limited to, copying, distributing, modifying, and making derivative works.

Copyright © Bull SAS 2012

Printed in France

Suggestions and criticisms concerning the form, content, and presentation of this book are invited. A form is provided at the end of this book for this purpose.

To order additional copies of this book or other Bull Technical Publications, you are invited to use the Ordering Form also provided at the end of this book.

Trademarks and Acknowledgements

We acknowledge the right of proprietors of trademarks mentioned in this book.

The information in this document is subject to change without notice. Bull will not be liable for errors contained herein, or for incidental or consequential damages in connection with the use of this material.

Contents

Virtual I/O Server and Integrated Virtualization Manager commands listed alphabetically 1

What's new in Virtual I/O Server and Integrated Virtualization Manager commands	1
Reading syntax diagrams	2
Exit status for Virtual I/O Server commands	3
activatevg command.	3
alert command	4
alt_root_vg command	5
backup command.	6
backupios command	11
IVM bkprofdata command	12
bootlist command	13
cattracerpt command	16
cfgassist command	16
cfgdev command	19
cfglnagg command	20
cfgnamesrv command	22
cfgsvc command.	23
chauth command	31
chbdsp command	32
chdate command	33
chdev command.	34
chedition command	35
IVM chhwres command	36
chkdev command	41
chlang command	44
IVM chled command	45
IVM chlparutil command.	46
chlv command	47
chpath command	48
chrep command	49
chrole command.	50
chsp command	52
IVM chsvcevent command	54
IVM chsyscfg command	55
IVM chsysstate command	62
chtcpip command	64
chuser command	65
IVM chvet command	67
chvfilt command	67
chvg command	69
chvopt command	70
chvlog command	70
chvrepo command.	72
cleandisk command	74
cluster command	74
cl_snmp command	78
cplv command	80
cpvdi command	81
deactivatevg command	83
diagmenu command	83
dsmc command	84
entstat command	86
errlog command.	90
exportvg command.	90
extendlv command	91
extendvg command.	92

fcstat command	93
fsck command	95
genvfilt command	96
hostmap command	97
hostname command	99
importvg command	99
installios command	100
invscout command	101
ioslevel Command.	102
ldapadd command	103
ldapsearch command	104
ldfware command	105
license command	107
loadopt command	108
loginmsg command	108
IVM lpar_netboot command	109
IVM lpcfgop command	111
lsauth command	112
lsdev command	113
lsfailedlogin command	116
lsfware command	117
lsgcl command	118
IVM lshwres command	119
IVM lsled command	142
lslparinfo command	144
IVM lslparmigr command	144
IVM lslparutil command	147
lslv command	154
lsmap command	159
IVM lsmemdev command	167
lsnetvc command.	170
lsnports command.	171
lspath command	172
lspv command	175
IVM lsrefcode command	180
lsrep command.	183
lsrole command	184
lssecattr command.	185
lssp command	187
lssvc command.	189
IVM lssvcevents command	191
lssw command	200
IVM lssyscfg command	201
IVM lssysconn command	214
lstcpip command	216
lsuser command	217
IVM lsvet command	219
lsvfilt command	220
lsvg command	221
lsvlog command	225
lsvrepo command.	228
lsvopt command	230
migratepv command	231
IVM migrpar command.	232
mirrorios command	236
mkauth command.	237
IVM mkauthkeys command	239
mkbdsp command.	241
IVM mkgencfg command	243
mkkrb5clnt command	245
mkldap command	245

mklv command	247
mklvcopy command	248
mkpath command	249
mkrole command	251
mkrep command	252
mksp command	252
IVM mksvcevent command	254
IVM mksyscfg command	254
mktcpip command	261
mkuser command	263
mkvdev command	264
mkvfilt command	267
mkvg command	268
mkvlog command	269
mkvopt command	271
mkvt command	272
motd command	272
mount command	273
netstat command	274
oem_platform_level command	278
oem_setup_env command	279
optimizenet command	280
IVM os_install Command	285
passwd command	288
pdump command	289
ping command	289
postprocesssvc command	291
prepdev command	294
redefvg command	295
reducevg command	295
remote_management command	296
replphyvol command	297
restore command	299
restorevgstruct command	307
rmauth command	308
rmbdsp command	309
rmdev command	310
rmlv command	311
rmlvcopy command	312
rmpath command	313
rmrep command	314
rmrole command	315
rmsecattr command	316
rmsp command	317
IVM rmsyscfg command	317
rmtcpip command	318
rmuser command	319
rmvdev command	320
rmvfilt command	321
rmvlog command	321
rmvopt command	323
rmvt command	323
rolelist command	324
IVM rsthwres command	325
IVM rstprofdata command	327
save_base command	328
savevgstruct command	330
seastat command	330
setkst command	331
setsecattr command	332
showmount command	336

shutdown command	336
snap command	337
snapshot command	338
snmp_info command	340
snmp_trap command	342
snmpv3_ssw command	343
startnetsh command	343
startsvcs command	345
startsysdump command	346
starttrace command	347
stopnetsh command	348
stopsvcs command	349
stoptrace command	350
svmon command	351
swrole command	354
syncvg command	355
sysstat command	356
topas command	356
topasrec command	366
tracepriv command	367
tracert command	368
uname command	370
unloadopt command	371
unmirrorios command	371
unmount command	372
updateios command	373
vasistat command	374
vfcmap command	382
viosbr command	383
viosecure command	393
viostat command	396
vlantfw command	400
vmstat command	401
wkldagent command	405
wkldmgr command	406
wkldout command	407
Notices	409
Programming interface information	410
Trademarks	410
Terms and conditions	411

Virtual I/O Server and Integrated Virtualization Manager commands listed alphabetically

This topic contains links to information about the Virtual I/O Server and Integrated Virtualization Manager commands, including syntax statements, flag descriptions, and usage examples. Commands are listed alphabetically.

What's new in Virtual I/O Server and Integrated Virtualization Manager commands

Read about new or significantly changed information in Virtual I/O Server and Integrated Virtualization Manager commands since the previous update of this topic collection.

May 2012

- Added new commands to support intervirtual LAN routing. The newly added commands are “chvfilt command” on page 67, “genvfil command” on page 96, “lsvfilt command” on page 220, “mkvfil command” on page 267, “rmvfil command” on page 321, and “vlantfw command” on page 400.
- Updated the “chdev command” on page 34, “cluster command” on page 74, “lsmmap command” on page 159, “lssp command” on page 187, “replphyvol command” on page 297, and “viosecure command” on page 393 topics.

December 2011

- Added the “snapshot command” on page 338 topic, which has the information about creating a cluster that consists of one or more Virtual I/O Server (VIOS) partitions that are connected to the shared storage pool (SSP) and that have access to distributed storage on the VIOS version 2.2.1.3, or later.
- Updated the “chsp command” on page 52, “cluster command” on page 74, “lsmmap command” on page 159, “lspv command” on page 175, “lssp command” on page 187, “mkbdsp command” on page 241, and “rmbdsp command” on page 309 topics with the information about creating a cluster of one or more VIOS partitions connected to the SSP and that have access to distributed storage on the Virtual I/O Server (VIOS) version 2.2.1.3, or later.

October 2011

- Added the following new commands: “chvlog command” on page 70, “chvrepo command” on page 72, “cleandisk command” on page 74, “lsvlog command” on page 225, “lsvrepo command” on page 228, “mkvlog command” on page 269, and “rmvlog command” on page 321.
- Updated the following commands: “cfgassist command” on page 16, “cfgsvc command” on page 23, “cluster command” on page 74, “hostname command” on page 99, “lssp command” on page 187, “lssvc command” on page 189, “mkbdsp command” on page 241, “mktcpip command” on page 261, “postprocesssvc command” on page 291, “startsvc command” on page 345, “stopsvc command” on page 349, and “viosbr command” on page 383.

May 2011

- Added the new command “cleandisk command” on page 74, which removes any storage pool or cluster signatures that are on the physical volume.
- Updated the following commands: “alert command” on page 4, “cfglnagg command” on page 20, “cluster command” on page 74, and “viosbr command” on page 383.

December 2010

- Added IBM® License Metric Tool metric agent information to “cfigsvc command” on page 23, “lssvc command” on page 189, “postprocesssvc command” on page 291, “startsvc command” on page 345, and “stopsvc command” on page 349.
- Added a new command “chedition command” on page 35, which manages the editions on Virtual I/O Server.
- Added 3 new commands and modified existing commands as a part of the Virtual I/O Server (version 2.2 or later). The newly added commands are “alert command” on page 4 and “cluster command” on page 74. The updated commands are “cfgassist command” on page 16, “chuser command” on page 65, “chsp command” on page 52, “lsmmap command” on page 159, “lspv command” on page 175, “lssp command” on page 187, “mkbdsp command” on page 241, “mkuser command” on page 263, “rmbdsp command” on page 309, and “viosbr command” on page 383.

September 2010

- Added new commands as a part of the functional implementation of role-based access control (RBAC) on the Virtual I/O Server. The newly added commands are “chauth command” on page 31, “chrole command” on page 50, “lsauth command” on page 112, “lsrole command” on page 184, “lssecattr command” on page 185, “mkauth command” on page 237, “mkrole command” on page 251, “rmauth command” on page 308, “rmrole command” on page 315, “rmsecattr command” on page 316, “rolelist command” on page 324, “setkst command” on page 331, “setsecattr command” on page 332, “swrole command” on page 354, and “tracepriv command” on page 367.
- Updated the “mkvdev command” on page 264 as a part of Peer-to-Peer Remote Copy (PPRC) VIOS support.

March 2010

- Added information for systems servers that contain the POWER7® processor.

Reading syntax diagrams

Syntax diagrams are a way to represent command syntax and consist of symbols such as brackets ([]), braces ({ }), and vertical bars (|). The following is a sample of a syntax statement for the **lsdev** command:

```
lsdev { -dev DeviceName | -plc PhysicalLocationCode } [ -child ] [ -field FieldName... ] [ -fmt Delimiter ]
```

The following conventions are used in the command syntax statements:

- Items that must be entered literally on the command line are in **bold**. These items include the command name, flags, and literal characters.
- Items representing variables that must be replaced by a name are in *italics*. These items include parameters that follow flags and parameters that the command reads, such as *Files* and *Directories*.
- Parameters enclosed in brackets are optional.
- Parameters enclosed in braces are required.
- Parameters not enclosed in either brackets or braces are required.
- A vertical bar signifies that you choose only one parameter. For example, [a | b] indicates that you can choose a, b, or nothing. Similarly, { a | b } indicates that you must choose either a or b.
- Ellipses (...) signify the parameter can be repeated on the command line.
- The dash (-) represents standard input.

Exit status for Virtual I/O Server commands

The following table defines the standard return codes returned by all of the Virtual I/O Server commands. Additional return codes unique to a specific command are defined within the individual command description page.

Return Code	Description
0	Success
1	Syntax Error
2	Not Found
1	Command requires an option
1	Command requires the specified option
1	Command requires an attribute
3	Invalid access to execute command
1	Invalid command
1	Invalid flag or argument
1	Invalid option flag
1	Invalid attribute
1	Invalid option combination
1	Specified option requires an argument
1	Specified option requires an attribute
1	Specified option also requires another option
1	Specified option is repeated
1	Attributes cannot be repeated
1	Contains an invalid argument
1	Is invalid
1	Too many arguments
1	Too few arguments
1	Unable to acquire permission to execute command
4	Execution of this command did not complete
10	No device found with physical location
11	Too many matches for physical location
12	Too many matches for physical location code
18	Insufficient memory

activatevg command

Purpose

Activates a volume group.

Syntax

activatevg [-f] *VolumeGroup*

Description

The **activatevg** command activates the volume group specified by the *VolumeGroup* parameter and all associated logical volumes. When a volume group is activated, physical partitions are synchronized if they are not current.

Flags

-f Allows a volume group to be made active that does not currently have a quorum of available disks.

Exit Status

See “Exit status for Virtual I/O Server commands” on page 3.

Examples

1. To activate volume group **vg03**, type:
`activatevg vg03`

Related Information

The **mkvg** command, the **chvg** command, the **extendvg** command, the **reducevg** command, the **mirrorios** command, the **unmirrorios** command, the **lsvg** command, the **deactivatevg** command, the **importvg** command, the **exportvg** command, and the **syncvg** command.

alert command

Purpose

Sets, removes, and lists all the alerts for a cluster and storage pool.

Syntax

To set the alert:

```
alert -set -clustername ClusterName -spname StoragePool [-type threshold] -value ThresholdValue
```

To unset the alert:

```
alert -unset -clustername ClusterName -spname StoragePool [-type threshold]
```

To list all the alerts:

```
alert -list -clustername ClusterName -spname StoragePool [-type threshold]
```

Description

The **alert** command sets, removes, and lists all alerts for a cluster and storage pool. If the **-type** flag is included, then it performs the action based on the specified type. In the current release, the **-type** flag supports only threshold. The **-value** flag can be specified in percentage.

Note: Alerts can be set only for a storage pool.

Flags

-clustername	Specifies the cluster name.
-list	Lists alerts.
-set	Sets an alert.
-spname	Specifies the storage pool name.
-type	Specifies the type of alert. The default is threshold.
-unset	Removes an alert.
-value	Specifies the value of the alert. The value depends on the alert type. For the type threshold, the value is in percentage.

Examples

1. To list alerts within the storage pool, type the following command:

```
alert -list -clustername newcluster -spname sp
```

The system displays the following output:

Pool Name	PoolID	Threshold Percentage
sp	4970216569494228437	80

2. To set an alert within the storage pool, type the following command:

```
alert -set -clustername newcluster -spname sp -value 70
```

alt_root_vg command

Purpose

Creates a copy of (clones) the currently running system to an alternative disk.

Syntax

```
alt_root_vg [-g][-z] [-bundle bundle_name -location images_location] -target target_disks...
```

Description

The **alt_root_vg** command allows users to copy the current root volume group to an alternate disk and to update the operating system to the next fix pack level, without taking the machine down for an extended period of time and mitigating outage risk. This can be done by creating a copy of the current rootvg on an alternate disk and simultaneously applying fix pack updates. If needed, the **bootlist** command can be run after the new disk has been booted, and the bootlist can be changed to boot back to the older level of the operating system.

Note: The *target_disk* cannot be a physical volume that is assigned to a shared memory pool (to be used as a paging space device by a shared memory partition).

Cloning the running rootvg, allows the user to create a backup copy of the root volume group. This copy can be used as a back up in case the rootvg failed, or it can be modified by installing additional updates. One scenario might be to clone a 1.3.0.0 system, and then install updates to bring the cloned rootvg to 1.3.0.0-FP8.0. This would update the system while it was still running. Rebooting from the new rootvg would bring the level of the running system to 1.3.0.0-FP8.0. If there was a problem with this level, changing the bootlist back to the 1.3.0.0 disk and rebooting would bring the system back to 1.3.0.0. Other scenarios would include cloning the rootvg and applying individual fixes, rebooting the system and testing those fixes, and rebooting back to the original rootvg if there was a problem.

At the end of the installation, a volume group, *altinst_rootvg*, is left on the target disks in the varied off state as a placeholder. If varied on, it indicates that it owns no logical volumes; however, the volume group does contain logical volumes, but they have been removed from the Object Data Manager (ODM)

because their names now conflict with the names of the logical volumes on the running system. Do not vary on the `altinst_rootvg` volume group; instead, leave the definition there as a placeholder.

After rebooting from the new alternate disk, the former rootvg volume group is displayed in a `lspv` listing as `old_rootvg`, and it includes all disks in the original rootvg. This former rootvg volume group is set to not vary-on at reboot.

If a return to the original rootvg is necessary, the `bootlist` command is used to reboot from the original rootvg.

Flags

-target <i>target_disks</i>	Specifies a space-delimited list of the name or names of the target disks where the alternate rootvg will be created. These disks must not contain any volume group definitions. The <code>lspv</code> command shows these disks as belonging to volume group <i>None</i> .
-bundle <i>bundle_name</i>	Path name of optional file that contains a list of packages or filesets that are installed after a rootvg clone is created. You must also specify the -location flag when you specify the bundle option. To update all currently installed software to the latest level available on the installation media, specify <i>update_all</i> in place of <i>bundle_name</i> .
-location <i>image_location</i>	Location of the install images or updates to apply after a clone of rootvg is created. The location can be a directory and full-path name or device name, such as <i>/dev/rmt0</i> .
-g	Skips checks to ensure the disk can boot.
-z	Specifies that no types of non-rootvg volume groups are to be imported.

Exit Status

See “Exit status for Virtual I/O Server commands” on page 3.

Examples

To create a boot list of logical devices to be used in the next normal boot, enter the following commands:

```
alt_root_vg -target hdisk2 -bundle my_bundle -location /tmp/update
alt_root_vg -target hdisk2 -bundle my_bundle -location /home/padmin
```

Related Information

The `bootlist` command.

backup command

To back up a file system by *i-node*, specify the *-Level* and *FileSystem* parameters. When used in conjunction with the **-u** flag, the *-Level* parameter provides a method of maintaining a hierarchy of incremental backups for each file system. Specify the **-u** flag and set the *-Level* parameter to *n* to back up only those files that have been modified since the *n-1* level back up. The possible back up levels are 0 to 9. A level 0 back up archives all files in the file system.

The *FileSystem* parameter can specify either the physical device name (block or raw name) or the name of the directory on which the file system is mounted. The default file system is the root (*/*) file system.

Users must have read access to the file system device (such as */dev/hd4*) or have back up authorization to perform backups by *i-node*.

Note:

1. You must first unmount a file system before backing it up by i-node. If you attempt to back up a mounted file system, a warning message is displayed. The **backup** command continues, but the created back up might contain inconsistencies because of changes that might have occurred in the file system during the back up operation.
2. Backing up file systems by i-node truncates the **uid** or **gid** of files having a **uid** or **gid** greater than 65535. When restored, these files might have different values for the **uid** and **gid** attributes. To retain the values correctly, always back up by name those files that have a **uid** or **gid** greater than 65535.
3. You can archive only JFS (Journaled File System) file systems when backing up by i-node. Back up any non-JFS file systems by file name.

Purpose

Backs up files and file systems.

Syntax

To back up files by name:

```
backup -i [ -b Number ] [ -p [ -e RegularExpression ] ] [ -E{force|ignore|warn} ] [ -f Device ] [ -l Number ] [ -U ] [ -o ] [ -q ] [ -v ]
```

To back up file systems by i-node:

```
backup [ [ -Level ] [ -b Number ] [ -c ] [ -f Device ] [ -L Length ] [ -U ] [ -u ] ] [ FileSystem ] | [ -w | -W ]
```

Description

The **backup** command creates copies of your files on a back up medium, such as a magnetic tape or diskette. The copies are in one of the two back up formats:

- Specific files backed up by name using the **-i** flag.
- Entire file system backed up by i-node using the *Level* and *FileSystem* parameters.

If you issue the **backup** command without any parameters, the root file system level 9 i-node back ups to the /dev/rfd0 device. The default syntax is:

```
-9uf/dev/rfd0 /dev/rhd4
```

The default back up device is /dev/rfd0. If flags are specified that are not appropriate for the specified back up device, the **backup** command displays an error message and continues with the back up.

A single back up can span multiple volumes.

Note:

1. Running the **backup** command results in the loss of all material previously stored on the selected output medium.
2. Data integrity of the archive might be compromised if a file is modified during system back up. Keep system activity at a minimum during the system back up procedure.
3. If a back up is made to a tape device with the device block size set to 0, it might be difficult to restore data from the tape unless the default write size was used with the **backup** command. The default write size for the **backup** command can be read by the **restore** command when the tape device block size is 0.

In other words, the **-b** flag should not be specified when the tape device block size is 0. If the **-b** flag of the **backup** command is specified and is different from the default size, the same size must be specified with the **-b** flag of the **restore** command when the archived files are restored from the tape.

Backing up files by name

To back up by name, use the **-i** flag. The **backup** command reads standard input for the names of the files to be backed up.

File types can be special files, regular files, or directories. When the file type is a directory, only the directory is backed up. The files under the directory are not backed up, unless they are explicitly specified.

Note:

1. Files are restored using the same path names as the archived files. Therefore, to create a back up that can be restored from any path, use full path names for the files that you want to back up.
2. When backing up files that require multiple volumes, do not enter the list of file names from the keyboard. Instead, pipe or redirect the list from a file to the **backup** command.
When you enter the file names from the keyboard and the back up process needs a new tape or diskette, the command loses any file names already entered but not yet backed up. To avoid this problem, enter each file name only after the archived message for the previous file has been displayed. The archived message consists of the character a followed by the file name.
3. If you specify the **-p** flag, only files of less than 2 GB are packed.

Backing up file systems by i-node

Flags

-b *Number*

For backups by name, specifies the number of 512-byte blocks; for backups by i-node, specifies the number of 1024-byte blocks to write in a single output operation. When the **backup** command writes to tape devices, the default is 100 for backups by name and 32 for backups by i-node.

The write size is the number of blocks multiplied by the block size. The default write size for the **backup** command writing to tape devices is 51200 (100 × 512) for backups by name and 32768 (32 × 1024) for backups by i-node. The write size must be an even multiple of the tape's physical block size.

The value of the **-b** flag is always ignored when the **backup** command writes to diskette. In this case, the command always writes in clusters that occupy a complete track.

-c

Specifies that the tape is a cartridge, not a nine-track.

-e *RegularExpression*

Specifies that the files with names matching the regular expression are not to be packed. A regular expression is a set of characters, meta characters, and operators that define a string or group of strings in a search pattern. It can also be a string containing wildcard characters and operations that define a set of one or more possible strings. The **-e** flag is applied only when the **-p** flag is specified.

-E

For backups by name, the **-E** flag requires one of the following arguments. If you omit the **-E** flag, **warn** is the default behavior.

force Fails the back up operation on a file if the fixed extent size or space reservation of the file cannot be preserved.

ignore Ignores any errors in preserving extent attributes.

warn Issues a warning if the space reservation or the fixed extent size of the file cannot be preserved.

-f Device

Specifies the output device. To send output to a named device, specify the *Device* variable as a path name such as the following example:

```
/dev/rmt0
```

To send output to the standard output device, specify a - (minus sign). The - (minus sign) feature enables you to pipe the output of the **backup** command to the **dd** command.

You can also specify a range of archive devices. The range specification must be in the following format:

```
/dev/deviceXXX-YYY
```

where XXX and YYY are whole numbers, and XXX must always be less than YYY; for example:

```
/dev/rfd0-3
```

All devices in the specified range must be of the same type. For example, you can use a set of 8 mm, 2.3 GB tapes or a set of 1.44 MB diskettes. All tape devices must be set to the same physical tape block size.

If the *Device* variable specifies a range, the **backup** command automatically goes from one device in the range to the next. After exhausting all of the specified devices, the **backup** command stops and requests that new volumes be mounted on the range of devices.

i

Specifies that files be read from standard input and archived by file name. If relative path names are used, files are restored (with the **restore** command) relative to the current directory at restore time. If full path names are used, files are restored to those same names.

-L Length

Specifies the length of the tape in bytes. This flag overrides the **-c**, **-d**, and **-s** flags. You can specify the size with a suffix of b, k, m, or g to represent blocks (512 bytes), kilo (1024 bytes), mega (1024 KB), or giga (1024 MB), respectively. To represent a tape length of 2 GB, enter **-L 2g**.

Note: Use the **-L** flag only for i-node backups.

-I Number

Limits the total number of blocks to use on the diskette device. The value specified must be a nonzero multiple of the number of sectors per diskette track. This option applies to by-name backups only.

-o

Creates a Version 2-compatible back up by name. This flag is required for compatibility with Version 2 systems because backups by name that are created by a version higher than 2 cannot be restored on Version 2 systems. To create a Version 2-compatible back up by name, use the **-o** flag along with other flags required for backups by name.

Files with attributes and values, such as user IDs and group IDs, that are too large for Version 2 systems are not backed up. A message is displayed for each such file and each value that is too large.

-p

Specifies that the files be packed, or compressed, before they are archived. Only files smaller than 2 GB are packed.

Note: This option should only be used when backing up files from an inactive file system. Modifying a file when a back up is in progress might result in corruption of the back up and an inability to recover the data. When backing up to a tape device that performs compression, this option can be omitted.

-q

Indicates that the removable medium is ready to use. When you specify the **-q** flag, the backup command proceeds without prompting you to prepare the back up medium. Press the Enter key to continue. This option applies only to the first volume. You are prompted for subsequent volumes. The **-q** flag applies only to backups by name.

-U

Specifies to back up any access control list (ACL)s or named extended attributes.

Without this option the image will include only AIX Classic (AIXC) ACLs and Printing Color Files (PCLs) in the archive along with the other regular file data. For files containing NFS4 ACLs, conversion to AIXC will happen by default during archival.

-u	Updates the <code>/etc/dumpdates</code> file with the raw device name of the file system and the time, date, and level of the back up. You must specify the -u flag if you are making incremental backups. The -u flag applies only to backups by i-node.
-v	Causes the backup command to display additional information about the back up. When using the -v flag, the size of the file as it exists on the archive is displayed in bytes. Additionally, a total of these file sizes is displayed when all files have been processed. Directories are listed with a size of 0. Symbolic links are listed with the size of the symbolic link. Hard links are listed with the size of the file, which is how hard links are archived. Block and character devices, if they were backed up, are listed with a size of 0. When the -v flag is not specified, the backup command displays only the names of the files being archived. This option is used only when backing up by file name.
-w	Disabled. If the -w flag is specified, no other flags are applied.
-W	Displays, for each file system in the <code>/etc/dumpdates</code> file, the most recent back up date and level. If the -W flag is specified, no other flags are applied.
-Level	Specifies the back up level (0 to 9). The default level is 9.

Exit Status

This command returns the following exit values:

0	Successful completion.
>0	An error occurred.

Examples

1. To back up all the files and subdirectories in the `/home` directory using full path names, type the following command:

```
find /home -print | backup -i -f /dev/rmt0
```

The **-i** flag specifies that files will be read from standard input and archived by the file name. The **find** command generates a list of all the files in the `/home` directory. The files in this list are full path names. The `|` (pipe symbol) causes this list to be read from standard input by the **backup** command. The **-f** flag directs the **backup** command to write the files to the `/dev/rmt0` tape device. Because the files are archived using full path names, they will be written to the same paths when restored.

2. To back up all the files and subdirectories in the `/home/padmin` directory using relative path names, type the following command:

```
find . -print | backup -i -v -q
```

Each file name in the list generated by the **find** command is preceded by `./` (dot, slash). Because the files are backed up using relative path names, they are written to the current directory when restored. The **-v** flag causes the **backup** command to display additional information about the backup. The files are written to the default backup device `/dev/rfd0`.

3. To back up the `/` (root) file system, type the following command:

```
backup -0 -u -f /dev/rmt0 /
```

The 0 level specifies that all the files in the `/` (root) file system be backed up. The **-u** flag causes the backup command to update the `/etc/dumpdates` file for this backup.

4. To back up all the files in the root (`/`) file system that have been modified since the last level 0 backup, type the following command:

```
backup -1 -u -f /dev/rmt0 /
```

If the `/etc/dumpdates` file does not have an entry for a level 0 backup of the `/` (root) system, all the files in the file system are backed up.

5. To back up virtual optical media files from the virtual media repository, complete the following tasks:
 - a. Create a file called `backup_files` with the names of the files to be backed up. Include the full path name.

- ```
/var/vio/VMLibrary/media_file1
/var/vio/VMLibrary/media_file2
/var/vio/VMLibrary/media_filen
```
- b. Send the file using the **backup** command:
- ```
cat backup_files | backup -i -f /dev/rmt0
```

Related Information

The **restore** command.

backupios command

Purpose

Creates an installable image of the root volume group either onto a bootable tape, file system, or DVD.

Syntax

```
backupios -file {Directory} [-nosvg] [-nomedialib]
```

```
backupios -file {File name} [-mksysb] [-nosvg] [-nomedialib]
```

```
backupios -tape Device [-nopak] [-verify] [-nosvg] [-nomedialib]
```

```
backupios -cd Device {-udf | -cdformat} [-accept] [-nosvg] [-nomedialib]
```

Description

The **backupios** command creates a backup of the Virtual I/O Server and places it onto a file system, bootable tape, or DVD. You can use this backup to reinstall a system to its original state after it has been corrupted. If you create the backup on tape, the tape is bootable and includes the installation programs needed to install from the backup.

If the **-cd** flag is specified, the **backupios** command creates a system backup image to DVD-RAM media. If you need to create multi-volume discs because the image does not fit on one disc, the **backupios** command gives instructions for disk replacement and removal until all the volumes have been created.

Note: Vendor disc drives may support burning to additional disc types, such as CD-RW and DVD-R. Refer to the documentation for your drive to determine which disc types are supported.

If the **-file** flag is specified, the **backupios** command creates a system backup image to the path specified. The file system must be mounted and writable by the Virtual I/O Server root user prior to running the **backupios** command (see “mount command” on page 273 for details). Backing up the Virtual I/O Server to a remote file system will create the **nim_resources.tar** image in the directory you specify. The Virtual I/O Server must have root write access to the server on which the backup will be created. This backup can be reinstalled from the HMC using the **installios** command.

The **backupios** command empties the **target_disks_stanza** section of **bosinst.data** (which is part of the **nim_resources.tar** image) and sets **RECOVER_DEVICES=Default**. This allows the **mksysb** file generated by the command to be cloned to another logical partition. If you plan to use the **nim_resources.tar** image to install to a specific disk, then you need to repopulate the **target_disks_stanza** section of **bosinst.data** and replace this file in the **nim_resources.tar** image. All other parts of the **nim_resources.tar** image must remain unchanged.

Prior to backing up the root volume group, the **backupios** command saves the structure of all user-defined volume groups by calling the **savevgstruct** command for each volume group defined. To avoid having all user-defined volume groups backed up, use the **-nosvg** flag.

Note: The **backupios** command backs up only the volume group structures that are activated. The structures of volumes groups that are deactivated are not backed up.

Flags

-accept	Accepts licenses.
-cd	The Virtual I/O Server backup is placed onto DVD-RAM media.
-cdformat	Creates final CD images that are DVD sized (up to 4.38 GB).
-file	The Virtual I/O Server backup is placed in a file.
-mksysb	When the -mksysb flag is used, the resources used by the installios command are not saved in the image.
-nomedialib	The -nomedialib flag optionally excludes the contents of the virtual media repository from the backup. The -nomedialib flag uses the <code>/etc/exclude.rootvg</code> feature of the mksysb command. When the -nomedialib flag is specified, the backupios command copies the original contents of the <code>/etc/exclude.rootvg</code> file to a repository, appends the <code>"/var/vio/VMLibrary</code> string to the <code>/etc/exclude.rootvg</code> file, ensures that the -e flag is passed to the mksysb command, and restores the original contents of the <code>/etc/exclude.rootvg</code> file. If the -nomedialib option is specified in a backup, and the system is restored using the backup, the restored system will have the <code>"/var/vio/VMLibrary</code> string included in the <code>/etc/exclude.rootvg</code> file. To include the contents of the media library in subsequent backups, it might be necessary to edit the <code>/etc/exclude.rootvg</code> file to remove the <code>"/var/vio/VMLibrary</code> string.
-nopak	Disables software packing of the files as they are backed up. Some tape drives use their own packing or compression algorithms.
-nosvg	Prevents the volume groups structure of user defined volume groups from being saved as part of the backupios process.
-tape	The Virtual I/O Server backup is placed onto a tape.
-udf	Creates a UDF (Universal Disk Format) file system on DVD-RAM media. The default format is Rock Ridge (ISO9660).
-verify	Verifies a tape backup. This flag causes the backupios command to verify the file header of each file on the backup tape and report any read errors as they occur.

Exit Status

See “Exit status for Virtual I/O Server commands” on page 3.

Examples

1. To generate a backup to the optical device named `/dev/cd1`, type:

```
backupios -cd /dev/cd1 -cdformat
```
2. To generate a system backup to a tape device named `/dev/rmt0`, type:

```
backupios -tape /dev/rmt0
```

Related Information

The **mount** command and the **savevgstruct** command.

IVM bkprofdata command

Purpose

Backs up profile data. This command is operable only in an Integrated Virtualization Manager environment.

Syntax

```
bkprofdata -o backup [ -f BackupFile ] [ -m ManagedSystem ]
```

Description

The **bkprofdata** command performs a back up of logical partition configuration information to a file. The file can later be used to restore the partition configuration with the **rstprofdata** command.

Flags

-f BackupFile	The name of the file to write the dump to in the current working directory. If not specified, the default file is /var/adm/lpm/profile.bak .
-o backup	The backup parameter backs up logical partition configuration data to a file.
-m ManagedSystem	The name of the managed system. This attribute is optional because there is only one system to manage. The name may either be the user-defined name for the managed system, or be in the form ttt-mmm*sssssss , where ttt is the machine type, mmm is the model, and sssssss is the serial number of the managed system.

Exit Status

This command has a return code of 0 on success.

Security

This command is available to all users.

Examples

1. To back up the partition configuration data to **/var/adm/lpm/profile.bak**, type:

```
bkprofdata -o backup
```
2. To back up the partition configuration data to **lparData.bak**, type:

```
bkprofdata -o backup -f lparData.bak
```

Related Information

The **rstprofdata** command.

bootlist command

Purpose

Displays and alters the list of boot devices available to the system.

Syntax

```
bootlist -mode mode {[ -ls ] [ [ -rm ] | [ Device [ -attr Attribute=Value ... ] ... ] ] }
```

Description

The **bootlist** command allows the user to display and alter the list of possible boot devices from which the system may be booted. When the system is booted, it will scan the devices in the list and attempt to boot from the first device it finds containing a boot image. This command supports the updating of the following:

- Service boot list. The service list designates possible boot devices for when the system is booted in service mode. How a system is booted in service mode is hardware-platform dependent. It may require

a key switch to be turned to the Service position, a particular function key to be pressed during the boot process, or some other mechanism, as defined for the particular hardware platform.

- Previous boot device entry. This entry designates the last device from which the system booted. Some hardware platforms may attempt to boot from the previous boot device before looking for a boot device in one of the other lists.

The **bootlist** command supports the specification of generic device types as well as specific devices for boot candidates. Devices in the boot device list occur in the same order as devices listed on the invocation of this command.

The selection of the boot list to display or alter is made with the **-mode** *mode* option, where the *mode* variable is one of the keywords: **service**, **normal**, **both**, or **prevboot**. If the **both** keyword is specified, then both the normal boot list and the service boot list will be displayed, or if being altered, will be set to the same list of devices. If the **prevboot** keyword is specified, the only alteration allowed is with the **-rm** flag. The **-rm** flag invalidates the boot list specified by the **-mode** flag.

The devices currently in the boot list may be displayed by using the **-ls** flag. The list of devices that make up the specified boot list will be displayed, one device per line. If a device specified in the boot list is no longer present on the system, a '-' is displayed instead of a name.

Note: When you add a hot plug adapter to the system, that adapter and its child devices might not be available for specification as a boot device when you use the **bootlist** command. You may be required to reboot your system to make all potential boot devices known to the operating system.

Device Choices

The device name specified on the command line can occur in one of two different forms:

- It can indicate a specific device by its device logical name.
- It can indicate a generic or special device type by keyword. The following generic device keywords are supported:

fd	Any standard I/O-attached diskette drive
scdisk	Any SCSI-attached disk (including serial-link disk drives)
badisk	Any direct bus-attached disk
cd	Any SCSI-attached CD-ROM
rmt	Any SCSI-attached tape device
ent	Any Ethernet adapter
fddi	Any Fiber Distributed Data Interface adapter

When a specific device is to be included in the device list, the device's logical name (used with system management commands) must be specified. This logical name is made up of a prefix and a suffix. The suffix is generally a number and designates the specific device. The specified device must be in the Available state. If it is not, the update to the device list is rejected and this command fails. The following devices and their associated logical names are supported (where the bold type is the prefix and the *xx* variable is the device-specific suffix):

fd <i>xx</i>	Diskette-drive device logical names
hdisk <i>xx</i>	Physical-volume device logical names
cd <i>xx</i>	SCSI CD-ROM device logical names
rmt <i>xx</i>	Magnetic-tape device logical names
ent <i>xx</i>	Ethernet-adapter logical names
fddi <i>xx</i>	Fiber Distributed Data Interface adapter logical names

Attribute Choices

Attributes are extra pieces of information about a device that the user supplies on the command line. Since this information is specific to a particular device, generic devices do not have attributes. Attributes apply to the device that immediately precedes them on the command line, which allows attributes to be interspersed among devices on the command line. Currently, only network devices have attributes. These are:

bserver	IP address of the BOOTP server
gateway	IP address of the gateway
client	IP address of the client
speed	Network adapter speed
duplex	The mode of the network adapter

Error Handling

If this command returns with an error, the device lists are not altered. The following device list errors are possible:

- If the user attempts to display or alter a boot list that is not supported by the hardware platform, the command fails, indicating the mode is not supported.
- If the user attempts to add too many devices to the boot list, the command fails, indicating that too many devices were requested. The number of devices supported varies depending on the device selection and the hardware platform .
- If an invalid keyword, invalid flag, or unknown device is specified, the command fails with the appropriate error message.
- If a specified device is not in the Available state, the command fails with the appropriate error message.

Flags

<i>Device</i>	Provides the names of the specific or generic devices to include in the boot list.
-attr <i>Attribute=Value</i>	Specifies the device attribute value pairs to be used instead of the defaults. The <i>Attribute=Value</i> variable can be used to specify one attribute value pair or multiple attribute value pairs for one -attr flag. If you use an -attr flag with multiple attribute value pairs, the list of pairs must be enclosed in quotation marks with a blank space between the pairs. For example, entering -attr Attribute=Value lists one attribute value pair per flag, while entering -attr 'Attribute1=Value1 Attribute2=Value2' lists more than one attribute value pair.
-mode <i>Mode</i>	Specifies which boot list to display or alter. Possible values for the <i>mode</i> variable are normal , service , both , or prevboot .
-ls	Indicates that the specified boot list is to be displayed after any specified alteration is performed. The output is a list of device names.
-rm	Indicates that the device list specified by the -mode flag should be invalidated.

Exit Status

See “Exit status for Virtual I/O Server commands” on page 3.

Examples

1. To invalidate the Service mode boot list, type:
`bootlist -mode service -rm`
2. To make a boot list for Normal mode with devices listed on the command line, type:
`bootlist -mode normal hdisk0 hdisk1 rmt0 fd`

- To attempt to boot through a gateway using Ethernet, and then try other devices, type a command similar to the following. Substitute the IP addresses specified in the example with your IP addresses.

```
bootlist -mode normal ent0 -attr gateway=129.35.21.1 bserver=129.12.2.10 \  
client=129.35.9.23 hdisk0 rmt0 tok0 bserver=129.35.10.19 hdisk1
```

Related Information

The **backupios** command.

cattracerpt command

Purpose

Formats a report from the trace log.

Syntax

```
cattracerpt [ -hookid List | -lshid ] [ -outfile FileName ]
```

Description

The **cattracerpt** command reads the trace log, formats the trace entries, and writes a report to standard output.

Flags

-hookid <i>List</i>	Limits report to hook IDs specified with the <i>List</i> variable. The <i>List</i> parameter items must be separated by commas.
-lshid	Displays the list of hook IDs. The cattracerpt -listid command can be used with the starttrace -event command that includes IDs of trace events.
-outfile <i>File</i>	Writes the report to a file instead of to standard output.

Exit Status

See “Exit status for Virtual I/O Server commands” on page 3.

Examples

- To send a trace report to the **newfile** file, enter:

```
cattracerpt -outfile newfile
```
- To display a list of hook IDs, enter:

```
cattracerpt -lshid
```

Related Information

The **starttrace** command, and the **stoptrace** command.

cfgassist command

Purpose

Manages the Virtual I/O Server.

Managing the Virtual I/O Server includes the following tasks:

- Set date and time zone

- Change passwords
- Set system security
- Virtual I/O Server TCP/IP configuration
- Install and update software
- Storage management
- Devices
- Performance
- Role-based access control
- Shared storage pools
- Electronic Service Agent™

The **cfgassist** command is menu-driven and accessible only by the *padmin* user or any user with the authorization `vios.security.cfgassist` for security reasons.

Syntax

cfgassist

Description

Set date and time zone

Set the date and time zone of the Virtual I/O Server.

Change passwords

Change the passwords for the users of the Virtual I/O Server.

Set system security

Set or unset the security settings by enforcing security rules and configuring firewall settings on the Virtual I/O Server.

Virtual I/O Server TCP/IP configuration

Set the TCP/IP configuration parameters that the Virtual I/O Server requires. The parameters include the host name, the IP address, the interface on which the IP is to be configured, the subnet mask, the gateway, the name server address, and the domain name.

Install and update software

Use to install and update software on the Virtual I/O Server.

Storage Management

Manage the storage on the Virtual I/O Server. Management operations include creation, deletion, or modification to the following storage devices: volume groups, logical volumes, physical volumes, storage pools, and virtual media repositories.

Devices

Manage the devices on the Virtual I/O Server, including virtual devices. Management operations include configuring or unconfiguring devices, changing characteristics of devices, removing devices, and configuring shared Ethernet adapters, link aggregation, and virtual local area network (VLAN).

Performance

Manage the performance tools on the Virtual I/O Server by using the following **topas** functions:

1. **Topas** recordings:

The Virtual I/O Server supports local, central electronics process (CEC), and cluster recording capabilities. These recordings can be either persistent or normal. Persistent recordings are recordings that run on the Virtual I/O Server and continue to run after the Virtual I/O Server reboots. Normal recordings are recordings that run for a specified time interval. The recording data files that are generated are stored in `/home/ios/perf/topas`. Local recordings gather data about the Virtual I/O Server, whereas CEC recordings gather data about any AIX® logical

partitions that are running on the same CEC as the Virtual I/O Server. The data collected consists of dedicated and shared logical partition data, as well as a set of aggregated values that provide an overview of the partition set. Cluster recordings gather data from a list of hosts that are specified in a cluster configuration file. This list is considered a single cluster. The cluster configuration file is populated with host names and IP addresses by using the **xmtopas** configuration option for the **cfgassist** command. These hosts can belong to various CECs that are not on the same CEC as the Virtual I/O Server, although the **xmtopas** option, which is responsible for gathering data outside of the Virtual I/O Server, must be enabled to start the CEC and cluster recordings. The recording format can be specified by the user as either binary or nmon.

Persistent recordings are started by the **cfgassist** command with the option to specify the cut and retention of the recording. The user can specify the number of days of recording to be stored per recording file (cut) and the number of days of recording to be retained (retention) before the recording can be deleted. Not more than one instance of persistent recording of the same type (CEC or local recording) can be run in a system. When a persistent recording is started, the recording command is called with user-specified options. The same set of command-line options used by this persistent recording is added to *inittab* entries. This ensures that the recording is started automatically on reboot or restart of the system.

If a system is already running a persistent local recording (binary or nmon recording format), a user might want to start a new persistent recording of a local binary recording. To start a new persistent recording, the existing persistent recording must first be stopped by using the stop persistent recording option available under the stop recording option. Then a new persistent local recording has to be started from the start persistent local recording option. Starting a persistent recording fails if a persistent recording of the same recording format is already running in the system. Recordings are deleted after the specified retention period. By default, local persistent recordings run on the Virtual I/O Server.

2. Performance Manager

This function configures the Performance Manager (**perfmgr**) agent to gather data about performance of the Virtual I/O Server to analyze the data. The generated recordings are picked up by the Electronic Service Agent (ESA) agent, which sends the data to support for analysis.

After the agent is started, it runs various scripts internally to collect various data and generates a single data file named *stats.send*. All the individual data files and the *stats.send* file are stored in */var/perf/pm/daily/<hostname>/*. Normally, the *stats.send* file is generated every midnight from the list of individual data files, which are generated by internally called scripts. A user can generate the *stats.send* file on demand by running the Retransmit Recorded Data menu.

Role-based access control

Role-based access control (RBAC) concepts are a part of the core set of security functions. RBAC provides a mechanism in the VIOS system through which the *padmin* -specific system functions can also be managed through users by using the roles assigned to them. RBAC addresses two aspects of system management: convenience and flexibility through the split of system management functions as roles. RBAC permits the system administrator to define roles based on the job functions within an organization. The administrator assigns authorizations to these roles, which might require detailed operations to manage resources. Users are granted membership in the roles (defining the authority and responsibility for the user) based on their job requirements. In other words, a role confers to a set of permissions or authorizations to the assigned user. Also implementation could allow for hierarchies of roles. RBAC simplifies the system administration due to the following reasons:

- RBAC permits you to implement and enforce companywide security policies consistently with regard to system management and access control.
- It is expected that a role or job function definition within an organization remains the same as compared to resources and users. Hence, a change in the role definition is not required.

- By breaking down system functions into smaller units, RBAC permits you to protect the system to a greater extent. The isolation enforced around smaller units of administration confines the attackers to the smallest unit of power in the system.
- RBAC reduces the likelihood of making mistakes of commission and omission in granting privileges to users.
- RBAC permits you to enforce the traditional least privilege model of security.

Electronic Service Agent

Electronic Service Agent automatically monitors and collects hardware problem information and sends this information to support. It also can collect hardware, software, system configuration, and performance management information which can help support assist in diagnosing problems.

Shared storage pools

Shared storage pools manage the cluster and VIOS nodes, storage pool, and logical units. The following are the submenu options of shared storage pools:

- Manage Cluster and VIOS Node
- Manage Storage Pool
- Manage Logical Unit

Manage Cluster and VIOS Node

Using this menu option, the user can create or delete a cluster, list existing clusters, add or delete VIOS nodes from a cluster, and list nodes in a cluster.

Manage Storage Pool

Using this menu option, the user can list storage pools in a cluster, list, change, or show physical volumes in a storage pool, and set or modify storage pool threshold alert.

Manage Logical Unit

Using this menu option, the user can create and map logical units, unmap logical units, delete logical units, list logical units, list logical unit maps, create logical unit snapshots, list logical unit snapshots, roll back to snapshot, and delete snapshots.

Security

The `cfgassist` command is a privileged command. To run the command successfully, assume a role that has the authorization, `vios.security.cfgassist`.

Exit Status

See “Exit status for Virtual I/O Server commands” on page 3.

Examples

To start the configuration menu, type the following command:

```
cfgassist
```

Related Information

The `bootlist` command.

cfgdev command

Purpose

Configures devices in the Virtual I/O Server.

Syntax

```
cfgdev [ -dev Name ]
```

Description

The **cfgdev** command configures devices. If there are any devices detected that have no device software installed when configuring devices, the **cfgdev** command returns a warning message with the name or a list of possible names for the device package that must be installed. If the specific name of the device package is determined, it is displayed as the only package name, on a line below the warning message. If the specific name cannot be determined, a colon-separated list of possible package names is displayed on a single line. A package name or list of possible package names is displayed for each of the devices, if more than one device is detected without its device software.

The system displays the following warning message when devices without their device software are detected:

```
cfgdev: 0514-621 WARNING: The following device packages are
      required for device support but are not currently
      installed.
devices.pci.22100020
devices.pci.14101800
devices.pci.scsi:devices.pci.00100300:devices.pci.NCR.53C825
```

In this example, two devices were found that had missing software, and the **cfgdev** command displayed the names of the device packages that must be installed. A third device that also has missing software was found, but in this case, the **cfgdev** command displays several possible device package names.

Attention: To protect the Configuration database, the **cfgdev** command is not interruptible. Stopping this command before execution is complete could result in a corrupted database.

Flags

-dev *Name* Specifies the named device to configure along with its children.

Exit Status

See “Exit status for Virtual I/O Server commands” on page 3.

Examples

1. To configure detected devices attached to the **scsi0** adapter, type:

```
cfgdev -dev scsi0
```

Related Information

The **chdev** command, the **chpath** command, the **lsdev** command, the **lsmap** command, the **lspath** command, the **mkpath** command, the **mkvdev** command, the **rmdev** command, and the **rmpath** command.

cfglnagg command

Purpose

Add or remove adapters from a Link Aggregation or change Link Aggregation attributes.

Syntax

```
cfglnagg {-add [-backup] | -rm} [-parent ParentName] LinkAggregation Adapter
```

```
cfglnagg [-f ] [-parent ParentName] -attr Attribute=NewValue ... LinkAggregation
```

```
cfglnagg -failover [-parent ParentName] LinkAggregation
```

Description

This command adds adapters to a Link Aggregation or removes adapters from a Link Aggregation. This command can also be used to modify Link Aggregation attributes. It is not necessary to detach the Link Aggregation's interface to add or remove adapters or modify most Link Aggregation attributes.

Changing the **use_jumbo_frame** attribute requires the Link Aggregation to be detached. If this attribute is specified, the **cfglnagg** command detaches the Link Aggregation before modifying the **use_jumbo_frame** attributes, then brings the Link Aggregation back up. The user is prompted to continue unless the **-f** flag is specified.

Use the **lsdev** command with the **-attr** flag to view all attributes associated with a Link Aggregation device, including adapters.

The failover occurs only if the adapter in the idle channel is up. If the adapter in the idle channel is down, the Link Aggregation operates on the active one and no failover occurs.

Flags

-add	Adds the specified Adapter to the specified Link Aggregation. If the adapter must be added as a backup adapter, the -backup flag must be specified.
-attr	Changes the value of the <i>LinkAggregation</i> attribute to the <i>NewValue</i> attribute.
-backup	Specifies that the <i>Adapter</i> is being added as a backup adapter.
-f	Instructs the command not to prompt you.
-failover	Forces a failover of the specified Link Aggregation.
-parent	Specifies the name of the parent adapter (for example, SEA).
-rm	Deletes the specified Adapter from the specified Link Aggregation. The specified adapter can be either a primary or backup adapter.

Parameters

<i>Adapter</i>	Specifies the adapter to add or delete.
<i>LinkAggregation</i>	Specifies the Link Aggregation.
<i>ParentName</i>	Specifies the parent name of the Link Aggregation adapter.

Exit Status

See "Exit status for Virtual I/O Server commands" on page 3.

Examples

1. To add adapter **ent8** to Link Aggregation **ent3**, type the following command:

```
cfglnagg -add ent3 ent8
```
2. To force a failover of the Link Aggregation adapter **ent7** from the currently active channel to the idle channel, type the following command:

```
cfglnagg -failover ent7
```

Related Information

The `cfgnamesrv` command, the `entstat` command, the `hostmap` command, the `hostname` command, the `mktcpip` command, the `mkvdev` command, the `optimizenet` command, the `ping` command, the `startnetsh` command, the `stopnetsh` command, the `traceroute` command.

cfgnamesrv command

Purpose

Directly manipulates domain name server entries for local resolver routines in the system configuration database.

Syntax

To Add a Name Server Entry

```
cfgnamesrv -add { -ipaddr IPAddress | -dname DomainName | -slist SearchList}
```

To Delete a Name Server Entry

```
cfgnamesrv -rm { -ipaddr IPAddress | -domain | -slist}
```

To Change a Name Server Entry

```
cfgnamesrv -ch DomainName
```

To Display a Name Server Entry

```
cfgnamesrv -ls [ -namesrv | -domain | -slist ]
```

To Create the Configuration Database File

```
cfgnamesrv -mk [ -ipaddr IPAddress [ -dname DomainName ] [ -slist SearchList ] ]
```

To Change a Search List Entry

```
cfgnamesrv -chslist SearchList
```

Description

The `cfgnamesrv` command adds or deletes domain name server entries for local resolver routines in the system configuration database. To add a name server entry, specify an Internet Protocol address and, optionally, a domain name.

The `cfgnamesrv` command can show one or all domain name server entries in the system configuration database. There are three types of domain name server entries:

- A domain entry identifying the name of the local Internet domain.
- A name server entry that identifies the Internet address of a domain name server for the local domain. The address must be in dotted decimal format.
- A search list entry that lists all the domains to search when resolving hostnames. This is a space delimited list.

One domain entry and a maximum of three name server entries can exist in the system configuration database. One search entry can exist.

Flags

-add	Adds an entry to the system configuration database.
-ch <i>DomainName</i>	Changes the domain name in the system configuration database.
-chslst	Changes the search list.
-dname	Indicates that the command deals with the domain name entry.
-domain	Specifies that the operation is on the domain name. Use this flag with the -rm flag and the -ls flag.
-ipaddr <i>IPAddress</i>	Indicates that the command deals with a name server entry. Use dotted decimal format for the given IP address.
-ls	Shows all domain and name server entries in the configuration system database. If you use the -ipaddr flag, the cfgnamesrv command shows all name server entries. If you use the -domain flag, the cfgnamesrv command shows the domain name entry found in the database.
-mk	Creates the system configuration database.
-namesrv	Specifies that the -ls flag should print all name server entries.
-rm	Deletes an entry in the system configuration database. It must be used with the -ipaddr <i>IPAddress</i> flag or the -domain flag. The -ipaddr flag deletes a name server entry. The -domain flag deletes the domain name entry.
-slst	Specifies that the operation is on the search list. Use this flag with the -rm and -ls flag.
-slst <i>SearchList</i>	Changes the search list in the system configuration database.

Exit Status

See “Exit status for Virtual I/O Server commands” on page 3.

Examples

1. To add a domain entry with a domain name of **abc.aus.century.com**, type:

```
cfgnamesrv -add -dname abc.aus.century.com
```
2. To add a name server entry with IP address 192.9.201.1, type:

```
cfgnamesrv -add -ipaddr 192.9.201.1
```
3. To show all system configuration database entries related to domain name server information used by local resolver routines, type:

```
cfgnamesrv -ls
```

The output is given in the following format:

```
domain xyz.aus.century.com
name server 192.9.201.1
```

4. To set the search list **abc.aus.century.com xyz.aus.century.com**, type:

```
cfgnamesrv -chslst abc.aus.century.com xyz.aus.century.com
```

The **cfglnagg** command, the **entstat** command, the **lsnetsvc** command, the **mktcpip** command, and the **optimizenet** command.

cfgsvc command

Purpose

Configures an agent or service specified by the given name.

Syntax

```
cfgsvc AgentName [-attr Attribute=Value... | -ls | -key]
```

Description

The **cfgsvc** command configures the agent or service specified, by the given name. Use the **lssvc** command to obtain a list of agents or services. After configuring the agent or service, you can start and stop the agent or service by using the **startsvc** and **stopsvc** commands. Each time you run the command, it overwrites previous configurations.

To find the values for required attributes for an agent or service, use the **-ls** flag. If you do not specify a required attribute when you specify this command, a message is displayed that indicates the missing required attribute.

Agents might create **ssh** keys as part of their configuration. You can display the **ssh** public key that is generated during configuration by using the **-key** flag.

Parameters

-attr <i>Attribute=Value</i>	Identifies an attribute and value that are associated with the configuration of an agent or service.
-ls	Lists all the attributes that can be changed for a particular agent or service configuration.
-key	Displays the ssh public key that is generated for a particular agent or service configuration.

Note: The **-key** flag is not valid for the .

Agent or service names

The following agents or services can be managed by the Virtual I/O Server.

DIRECTOR_agent	Attribute names for the Systems Director agent View and track hardware configuration details of the system, and monitor performance and use of critical components, such as processors, disks, and memory.
	RESTART_ON_REBOOT This is an optional configuration attribute for the Systems Director agent. It designates whether the Systems Director agent is restarted if the Virtual I/O Server is rebooted. Valid values are: True False

ITM_premium

Attribute names for the Tivoli® Monitoring (ITM) agents

ITM_cec

The **ITM_premium** and **ITM_cec** agents provide system information including I/O virtualization mappings and system utilization.

These agents require that you specify the **HOSTNAME** attribute for configuration. Also, an ssh public key is generated if one does not exist.

HOSTNAME

Provides the monitoring agent with a Tivoli Enterprise Management System (TEMS) host name. This is a required attribute.

Valid values are:

Host name

IP address

RESTART_ON_REBOOT

Designates whether the monitoring agent is started if the Virtual I/O Server is rebooted. Because you can run only one monitoring agent on the Virtual I/O Server at any given time, configuring more than one monitoring agent to restart is only valid for the last agent that is configured to restart. Run the **lssvc** command by using the ITM agent name to verify which agent restarts.

Valid values are:

True

False

managing_system

Provides the monitoring agent with a Hardware Management Console (HMC) host name. This attribute is not valid for the Integrated Virtualization Manager.

Valid values are:

[user_name@]Host name

IP address

second_managing_system

Provides the monitoring agent with a secondary Hardware Management Console (HMC) host name. This attribute is not valid for the Integrated Virtualization Manager. This attribute is available for the ITM_cec agent. This is an optional attribute.

Valid values are:

[user_name@]Host name

IP address

CEC Provides the ITM_cec agent with a CEC host name.

DIRECTOR_HOST_ADDRESS

Provides the ITM_cec agent with a Director host name or IP address. The default value is blank. This is an optional attribute.

DIRECTOR_AUTHENTICATION

Provides the ITM_cec agent with a Director authentication value. The default value is Yes. This is an optional attribute.

DIRECTOR_PORT_NUMBER

Provides the ITM_cec agent with a Director port number. The default value is 8422. This is an optional attribute.

TSM_base

Attribute names for the Tivoli Storage Manager agent

SERVERNAME

The host name of the Tivoli Storage Manager server to which the Tivoli Storage Manager client is associated.

SERVERIP

The IP address or domain name of the Tivoli Storage Manager server to which the Tivoli Storage Manager client is associated.

NODENAME

The name of the machine on which the Tivoli Storage Manager client is installed.

ITUAM_base

Attribute names for the Tivoli Usage and Accounting Manager agent

ACCT_DATA0

The size, in megabytes, of the first data file that holds daily accounting information.

ACCT_DATA1

The size, in megabytes, of the second data file that holds daily accounting information.

ISYSTEM

The time, in minutes, when the agent generates system interval records.

IPROCESS

The time, in minutes, when the system generates aggregate process records.

Attribute names for the TotalStorage Productivity Center agents

The **TPC** agent is a TotalStorage Productivity Center agent. This agent requires that you specify the **S**, **A**, **devAuth**, and **caPass** attributes for configuration. By default, specifying this agent configures both the **TPC_data** and **TPC_fabric** agents.

S Provides the TotalStorage Productivity Center agent with a TotalStorage Productivity Center server host name. This is a required configuration attribute.

The valid values are:

Host name

IP address

A Provides the TotalStorage Productivity Center agent with an agent manager host name. This is a required attribute.

The valid values are:

Host name

IP address

devAuth

Sets the TotalStorage Productivity Center device server authentication password. This is a required attribute.

The valid value is:

Password

caPass Sets the CA authentication password. This is a required attribute.

The valid value is:

Password

caPort Sets the CA port. This is an optional attribute. The default value is 9510.

The valid value is:

Number

amRegPort

Specifies the agent manager registration port. This is an optional attribute. The default value is 9511.

The valid value is:

Number

amPubPort

Specifies the agent manager public port. This is an optional attribute. The default value is 9513.

The valid value is:

Number

dataPort

Specifies the TotalStorage Productivity Center data server port. This is an optional attribute. The default value is 9549.

The valid value is:

Number

- devPort**
Specifies the TotalStorage Productivity Center device server port. This is an optional attribute. The default value is 9550.
The valid value is:
Number
- newCA**
This is an optional attribute. The default value is true.
The valid values are:
True
False
- oldCA** This is an optional attribute. The default value is true.
The valid values are:
True
False
- daScan** This is an optional attribute. The default value is true.
The valid values are:
True
False
- daScript**
This is an optional attribute. The default value is true.
The valid values are:
True
False
- daInstall**
This is an optional attribute. The default value is true.
The valid values are:
True
False
- faInstall**
This is an optional attribute. The default value is true.
The valid values are:
True
False
- U** Specifies to uninstall the agent. This is an optional attribute.
The valid values are:
All | data | fabric

perfmgr

Attribute names for the performance management data collector agent: see the **postprocesssvc** command.

ipsec_tunnel

ipsec_tunnel works as a service on the Virtual I/O Server. Configurable parameters are as follows:

local_ip

The IP address of the source Virtual I/O Server.

remote_ip

The IP address of the destination Virtual I/O Server.

key

The key value consists of alphanumeric characters only.

ILMT

The IBM License Metric Tool (ILMT) agent helps the License Metric Tool to identify the system on the network to determine which IBM software is installed. It requires the *server* attribute to be configured with the host name or IP of the ILMT server.

server The host name or IP address of the message handler server.

security_level

The security level for the agent communication.

key The key value in alphanumeric characters.

0: Agent communication is not secure (HTTP).

1: Agent communication is secure with only server authentication.

2: Agent communication is secure with both agent and server authentication.

port The port number used by the agent if the *security_level* parameter equals 0.

secure_port

The port number used by the agent if the *security_level* parameter equals 1.

client_auth_secure_port

The port number used by the agent if the *security_level* parameter equals 2.

scan_group

The name of the scan group, which is only used for the initial registration with the server.

use_proxy

Whether the agent is protected by a proxy server. Valid values are y and n.

proxy_port

The port number used by the proxy server. The default value is 3128.

proxy The address of the proxy server.

fips_enabled

Whether the agent uses the Federal Information Processing Standard (FIPS). Valid values are y and n.

trace_size

The maximum size of a single iteration of the agent trace file.

trace_level

The level of trace that is logged during operation of the agent. The default level, MIN, is enough to identify the cause of most problems. MID or MAX trace levels can be enabled if a problem occurs whose cause cannot be immediately identified.

trace_files

The maximum number of iterations of the agent trace file. If the trace file reaches the maximum size defined in the *trace_size* parameter, the remaining messages roll over to another iteration of the file. The *trace_files* parameter defines how many times it can occur.

msglog_files

The number of message log files.

msglog_size

The maximum size of the message log.

tivoli_common_dir

The path to Tivoli Common Directory.

userdata1

The custom data about agents.

Examples

Tivoli Monitoring

1. To list all the attributes that are associated with an agent configuration, type the command as follows:

```
cfgsvc -ls ITM_premium
```

2. To configure the ITM_premium agent with several attributes, type the command as follows:

```
cfgsvc ITM_premium -attr Restart_On_Reboot=TRUE hostname=tems_server  
managing_system=hmc_console
```

Note: When the RESTART_ON_REBOOT attribute is set to TRUE, the ITM agent is restarted when the Virtual I/O Server is being rebooted. The hostname attribute specifies the Tivoli Enterprise Monitoring Server (TEMS) IP address, or hostname is set to *tems_server*. The managing_system attribute specifies the HMC IP address, or hostname is set to *hmc_console*.

3. To display the ITM_premium agent ssh public key, type the command as follows:

```
cfgsvc -key ITM_premium
```

Tivoli Storage Manager

1. To list all the attributes associated with an agent configuration, type the command as follows:

```
cfgsvc -ls TSM_base
```

2. To configure the TSM_base agent with several attributes, type the command as follows:

```
cfgsvc TSM_base -attr SERVERNAME=tsm_server SERVERIP=1.127.0.1  
NODENAME=VIOS
```

Tivoli Usage and Accounting Manager

1. To list all the attributes associated with an agent configuration, type the command as follows:

```
cfgsvc -ls ITUAM_base
```

2. To configure the ITUAM_base agent with several attributes, type the command as follows:

```
cfgsvc ITUAM_base -attr ACCT_DATA0=15 ACCT_DATA1=15  
ISYSTEM=60 IPROCESS=60
```

TotalStorage Productivity Center

1. To list all TotalStorage Productivity Center configurable attributes, type the command as follows:

```
cfgsvc TPC -ls
```

2. To configure TotalStorage Productivity Center with mandatory attributes, type the command as follows:

```
cfgsvc TPC -attr S=tpc_server_hostname  
A=agent_manager_hostname devAuth=password caPass=password
```

Systems Director Agent

1. To list all the attributes associated with an agent configuration, type the command as follows:

```
cfgsvc -ls DIRECTOR_agent
```

2. To configure the DIRECTOR_agent with several attributes, type the command as follows:

```
cfgsvc DIRECTOR_agent -attr RESTART_ON_REBOOT=TRUE
```

ipsec tunnels

1. To list all the attributes associated with the service configuration, type the command as follows:

```
cfgsvc -ls ipsec_tunnel
```

2. To configure **ipsec_tunnel**, type the command as follows:

```
cfgsvc ipsec_tunnel -attr local_ip=1.2.3.4 remote_ip=2.1.4.3 key=abcd1234
```

ILMT

1. To list all ILMT attributes, type the command as follows:
`cfgsvc ILMT -ls`
2. To configure ILMT, type the command as follows:
`cfgsvc ILMT -attr server=1.2.3.4`

Related Information

The **Issvc** command, the **postprocesssvc** command, the **startsvc** command, and the **stopsvc** command.

For more information about the various agents, see the following information:

- Tivoli software and the Virtual I/O Server
- Configuring the Tivoli agents and clients on the Virtual I/O Server
- Systems Director software
- Configuring the Systems Director agent

chauth command

Purpose

Changes user-defined authorization attributes.

Syntax

```
chauth [-a] [Attribute=Value...] [newauth]
```

Description

The **chauth** command modifies attributes for the authorization that is identified by the *newauth* parameter. The command only modifies existing user-defined authorizations in the authorization database. System-defined authorizations cannot be modified with the **chauth** command. To change an attribute of a user-defined authorization, specify the attribute name and the new value by using the *Attribute=Value* parameter. If any specified attribute or attribute value is not valid, the **chauth** command does not modify the authorization.

Important: Modifying the ID of an authorization can affect the system security because, the current value of the ID might be used by some processes and files. In general, use the *id* attribute to modify the ID of an authorization when you are sure that the authorization is not used. The **chauth** command only permits the ID to be set to an unused value greater than 15,000. IDs up to 15,000 are reserved for system-defined authorizations.

Note: After creating new authorizations, run the **setkst** command to update the kernel database with the new authorizations modified.

Security

The **chauth** command is a privileged command. To run the command successfully, assume a role that has the authorization **vios.security.auth.change**.

Flags

-a Specifies the list of attributes to be modified for the authorization.

Attributes

id	Specifies a unique integer that is used to identify the authorization. The value is a decimal integer of 10,001 - 32,768.
dfltmmsg	Specifies the default description to use if message catalogs are not in use. The value is a string.
msgcat	Specifies the message catalog file name that contains the description of the authorization. If the <i>msgcat</i> attribute is specified, the <i>msgset</i> and <i>msgnum</i> attributes must also be specified. The value is a string. If the specified string contains a leading forward slash (/), the value is assumed to be an absolute path name. Otherwise, the user environment defines the directory search path as specified by the catopen routine.
msgset	Specifies the message set number in the file name to retrieve the message number. The file name is specified by the <i>msgcat</i> attribute, and the message number is specified by the <i>msgnum</i> attribute. The value is a decimal integer.
msgnum	Specifies the message number for the description of the authorization in the file and the set. The authorization is specified by the <i>msgcat</i> attribute, and the set number is specified by the <i>msgset</i> attribute. The value is a decimal integer.

Exit Status

0 for success.

Nonzero for command error.

Examples

1. To change the message catalog that is used to provide the authorization description for the *custom* authorization, type the following command:

```
chauth -a msgcat="custom_auths.cat" custom
```
2. To change the message set and number that designates the authorization description for the *custom.test* authorization, type the following command:

```
chauth -a msgset=5 msgnum=24 custom.test
```
3. To change the message catalog for the *custom.test* authorization, type the following command:

```
chauth -a msgset=5 custom.test
```

chbdsp command

Purpose

Changes the characteristics of a backing device within a storage pool.

Syntax

Increase the size of a backing device:

```
chbdsp [-sp StoragePool ] [-bd BackingDevice ] [-size Size]
```

Rename a backing device:

```
chbdsp [-sp StoragePool ] [-bd BackingDevice ] [-mv NewName]
```

Description

The **chbdsp** command changes attributes of the specified backing device. If the **-size** flag is specified the backing device's size is increased. Size can be given in megabytes (###M/m) or gigabytes (###G/g).

If the **-mv** flag is specified, the backing device is renamed.

This command does not support shared storage pools.

Note: The backing device that you specified cannot be assigned to a shared memory pool (to be used as a paging space device by a shared memory partition).

Flags

-bd <i>BackingDevice</i>	Specifies the backing device.
-mv <i>NewName</i>	Specifies the new name of the backing device.
-size <i>Size</i>	Specifies the minimum amount of storage to add to the backing device.
-sp <i>StoragePool</i>	Specifies the storage pool to be used.

Exit Status

23	Specified storage pool is not a valid storage pool.
39	The given backing device is not in the specified storage pool.
40	File backing device names cannot exceed "38" characters in length.

Examples

1. To increase the size of a the backing device `client1`, in the default storage pool by 3 gigabytes, type:
`chbdsp -bd client1 -size 3g`
2. To rename the backing device `cl_data`, in the storage pool `partition_3`, to `bank_data`, type:
`chbdsp -sp partition_3 -bd cl_data -mv bank_data`

chdate command

Purpose

Displays or changes the date, time or time zone.

Syntax

```
chdate [ mmddHHMM [YYyy | yy] ] [-timezone TZ]
```

```
chdate [-year YYyy] [-month mm] [-day dd] [-hour HH] [-minute MM] [-timezone TZ]
```

Description

Displays or changes the system date, time, or time zone. Changes made to the time zone will not take effect until the user logs out. For the time zone to take effect for the entire system, the system must be rebooted. All flags are optional, and the current system time information will be used if a flag was unspecified.

If no flags or arguments are specified, the **chdate** command displays the current date and time.

The *mmddHHMM* [*YYyy* | *yy*] parameters correspond to month, day, hour, minute, and optional 4 or 2 digit year.

Note: The user must have **padmin** authority to change the date and time.

Flags

-year	Sets the year to <i>YYyy</i> .
-month	Sets the month to <i>mm</i> .
-day	Sets the day to <i>dd</i> .
-hour	Sets the hour to <i>HH</i> in 24-hour format.
-minute	Sets the minute to <i>MM</i> .
-timezone	Sets the time zone (for example, CST6CDT).

Exit Status

See “Exit status for Virtual I/O Server commands” on page 3.

Examples

1. To display the current date and time, type:
`chdate`
 2. To change the date to Tue Oct 12 16:30:00 CDT 2004 for a system in the US Central time zone, type:
`chdate -day 12 -month 10 -year 2004 -hour 16 -minute 30`
- or
- ```
chdate 101216302004
```

**Note:** The user must have **padmin** authority to change the date and time.

---

## chdev command

### Purpose

Changes the characteristics of a device.

### Syntax

```
chdev -dev Name [-perm] [-attr Attribute=Value [Attribute=Value...]]
```

### Description

The **chdev** command changes the characteristics of the device specified with the given device logical name (the **-dev *Name*** flag). Some changes may not be allowed when the device is in the **Available** state.

When the **-perm** flag is not specified, the **chdev** command applies the changes to the device and updates the database to reflect the changes. If the **-perm** flag is specified, only the database is updated to reflect the changes, and the device itself is left unchanged. This is useful in cases where a device cannot be changed because it is in use; in which case, the changes can be made to the database with the **-perm** flag, and the changes will be applied to the device when the system is restarted. Not all devices support the **-perm** flag.

**Attention:** To protect the Configuration database, the **chdev** command is not interruptible. To stop this command before execution is complete could result in a corrupted database.

The *lldpsvc* attribute is added to the shared Ethernet adapter (SEA) device for automated virtual network port configuration and management in the Systems Director Management Console (SDMC). The *lldpsvc* attribute activates the virtual server network (VSN) capability on the SEA when you set the attribute to *yes*. The VSN capability can be enabled on the SDMC. The default value of the *lldpsvc* attribute is *no*. This

attribute must be set to *no* before you remove the SEA. For more information, see Verifying that the host supports virtual server network.

## Flags

|                                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|-------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>-attr</b> <i>Attribute=Value</i> | Specifies the device attribute value pairs used for changing specific attribute values. The <i>Attribute=Value</i> parameter can use one attribute value pair or multiple attribute value pairs for one <b>-attr</b> flag. If you use an <b>-attr</b> flag with multiple attribute value pairs, the list of pairs must be separated by spaces. For example, entering <b>-attr Attribute=Value</b> lists one attribute value pair per flag, while entering <b>-attr Attribute1=Value1 Attribute2=Value2</b> lists more than one attribute value pair. |
| <b>-dev</b> <i>Name</i>             | Specifies the device logical name, indicated by the <i>Name</i> parameter, whose characteristics are to be changed.                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>-perm</b>                        | Changes the device's characteristics permanently without actually changing the device. This is useful for devices that cannot be made unavailable and cannot be changed while in the available state. By restarting the system, the changes will be applied to the device. Not all devices support the <b>-perm</b> flag.                                                                                                                                                                                                                            |

## Exit Status

See “Exit status for Virtual I/O Server commands” on page 3.

## Examples

1. To change the retention instructions of the 4mm SCSI tape drive **rmt0** so that the drive does not move the tape to the beginning, then to the end, and then back to the beginning each time a tape is inserted or the drive is powered on, type:  

```
chdev -dev rmt0 -attr ret=no
```
2. To change the SCSI ID of the available SCSI adapter **scsi0** that cannot be made unavailable or changed due to available disk drives connected to it, type:  

```
chdev -dev scsi0 -attr id=6 -perm
```

To apply the change to the adapter, shutdown and restart the system.

## Related Information

The **cfgdev** command, the **chpath** command, the **lsdev** command, the **lsmap** command, the **lspath** command, the **mkpath** command, the **mkvdev** command, the **rmdev** command, and the **rmpath** command.

---

## chedition command

### Purpose

Lists the current edition, changes the current edition to another edition, and installs the bundles associated with the edition.

### Syntax

```
chedition -list | { -exp | -std | -ent } [-dev <bundlesPath> [-preview]
```

### Description

On Virtual I/O Server partition, the **chedition** command helps in listing the current edition, changing the current edition to another edition, and installing the bundles associated with the edition.

## Flags

|                       |                                                                            |
|-----------------------|----------------------------------------------------------------------------|
| <code>-list</code>    | Displays the current edition.                                              |
| <code>-exp</code>     | Changes to Express® edition.                                               |
| <code>-std</code>     | Changes to Standard edition.                                               |
| <code>-ent</code>     | Changes to Enterprise edition.                                             |
| <code>-dev</code>     | Installs the bundles associated with the edition, found in the given path. |
| <code>-preview</code> | Performs a preview of the bundle file installation.                        |

## Examples

### ILMT

1. To list the edition of the system, type the following command:  
`$ chedition -list`
2. To change system edition to Enterprise, type the following command:  
`$ chedition -ent`
3. To change system edition to Standard, type the following command:  
`$ chedition -std`

---

## IVM chhwres command

### Purpose

Changes hardware resources. This command is operable only in an Integrated Virtualization Manager environment.

### Syntax

To add, remove, or modify physical or virtual I/O, memory, or processors settings

```
chhwres -r io | mem | proc | virtualio | mempool [--subtype scsi | slot | eth | fc | pgdev] [-m <managed system>] -o a | r | s | rs [-p <partition name> | --id <partition ID>] [-l <DRC index>] [-s <virtual slot number>] [-q <quantity>] [--procs <quantity>] [--procunits <quantity>] [-w <wait time>] [-d <detail level>] [--force] [-a "<attributes>"] [--entitled <size>] [--device] [--help]
```

To add, remove, or modify Host Ethernet Adapter resource assignments and settings

```
chhwres -r hea [--subtype slot | eth] [-m <managed system>] -o a | r | s [-p <partition name> | --id <partition ID>] [-l <HEA DRC index>] [-g <port group ID>] [--physport <physical port ID>] [--logport <logical port ID>] [-w <wait time>] [-d <detail level>] [--force] [-a "<attributes>"] [--help]
```

### Description

The **chhwres** command changes the hardware resource configuration of the managed system. The **chhwres** command is used to perform dynamic LPAR operations.

### Flags

|                       |                                                                                                                                                                                                                                                     |
|-----------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <code>--device</code> | Specifies the paging space device to add or remove. The <code>--device</code> flag is valid only with <code>-r mempool --subtype pgdev</code> .                                                                                                     |
| <code>-r</code>       | Specifies the type of hardware resources to change. Valid values are <b>io</b> for physical I/O, <b>hea</b> for Host Ethernet Adapter resources, <b>virtualio</b> for virtual I/O, <b>mem</b> for memory, and <b>proc</b> for processing resources. |

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>--subtype</b>                | <p>Specifies the subtype of hardware resources to change. Valid values for <b>--subtype</b> depend on the resource type you specify.</p> <ul style="list-style-type: none"> <li>• <b>io</b>: slot, taggedio</li> <li>• <b>virtualio</b>: eth, scsi, fc, hsl, virtualopti</li> <li>• <b>mempool</b>: pgdev</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>-m</b> <i>managed system</i> | <p>Specifies the name of the managed system for which the hardware resource configuration is to be changed.</p> <p>The name may either be the user-defined name for the managed system, or be in the form <i>ttt-mmm*sssssss</i>, where <i>ttt</i> is the machine type, <i>mmm</i> is the model, and <i>sssssss</i> is the serial number of the managed system.</p>                                                                                                                                                                                                                                                                                                                                                                       |
| <b>-o</b>                       | <p>Specifies the operation to perform. Valid values are <b>a</b> to add hardware resources to a logical partition, <b>r</b> to remove hardware resources from a logical partition, and <b>s</b> to set hardware resource related attributes for a logical partition an Host Ethernet Adapter, or the managed-system.</p>                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>-P</b>                       | <p>Specifies the name of the logical partition for which the operation is to be performed.</p> <p>You can either use this option to specify the name of the logical partition for which the operation is to be performed, or use the <b>--id</b> option to specify the logical partition's ID. The <b>-p</b> and the <b>--id</b> options are mutually exclusive.</p>                                                                                                                                                                                                                                                                                                                                                                      |
| <b>--id</b> <i>partition ID</i> | <p>The logical partition name or <b>--id</b> is required for virtual Ethernet.</p> <p>Specifies the ID of the logical partition for which the operation is to be performed. To perform an add or remove operation, the logical partition must be in the running state.</p> <p>You can either use this option to specify the ID of the logical partition for which the operation is to be performed, or use the <b>-p</b> option to specify the logical partition's name. The <b>--id</b> and the <b>-p</b> options are mutually exclusive.</p> <p>A logical partition is required to be specified with this option or the <b>-p</b> option for all operations except a virtual Ethernet, memory, Host Ethernet Adapter set operation.</p> |
| <b>-l</b>                       | <p>When adding, removing, or moving a physical I/O slot, use this option to specify the DRC index of the slot.</p> <p>When performing an Host Ethernet Adapter operation, use this option to specify the adapter ID of the Host Ethernet Adapter for which the operation is to be performed.</p> <p>This option is not valid for any other operation.</p>                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>-s</b>                       | <p>Specifies the virtual slot number of the virtual I/O adapter to add or remove. When adding or removing a virtual I/O adapter, this option is required.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>-q</b>                       | <p>Specifies the quantity of memory to add or remove. The quantity specified must be in megabytes, it must be a multiple of the memory region size for the managed-system, and it must be greater than 0.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>--procs</b> <i>quantity</i>  | <p>When adding or removing processing resources to or from a logical partition using dedicated processors, use this option to specify the quantity of dedicated processors to add or remove.</p> <p>When adding or removing processing resources to or from a logical partition using shared processors, use this option to specify the quantity of virtual processors to add or remove.</p> <p>The quantity of processing resources specified with this option must be a whole number greater than 0.</p>                                                                                                                                                                                                                                |

|                             |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|-----------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>--procunits</b>          | <p>When adding or removing processing resources to or from a logical partition using shared processors, use this option to specify the quantity of processing units to add or remove.</p> <p>Otherwise, the quantity of processing units specified with this option can have up to 2 decimal places. In either case, the quantity specified must be greater than 0.</p>                                                                                                                                                                       |
| <b>--physport</b>           | <p>Specifies the ID of the Host Ethernet Adapter physical port. This option is required when adding an Host Ethernet Adapter logical port to a logical partition. This option is also required when setting Host Ethernet Adapter physical port attributes. This option is not valid for any other operation.</p>                                                                                                                                                                                                                             |
| <b>--logport</b>            | <p>Specifies the ID of the Host Ethernet Adapter logical port to add or remove. This option is required for a Host Ethernet Adapter add or remove operation. This option is not valid for any other operation.</p>                                                                                                                                                                                                                                                                                                                            |
| <b>-g</b>                   | <p>Specifies the Host Ethernet Adapter port group. This option is required for all Host Ethernet Adapter operations, and is not valid for any other operation.</p>                                                                                                                                                                                                                                                                                                                                                                            |
| <b>-w</b>                   | <p>Specifies the elapsed time, in minutes, after which an add, remove, or move operation will be stopped. <i>wait-time</i> must be a whole number. If wait-time is 0, the operation will not be timed out.</p> <p>If this option is not specified, a default value of 5 minutes is used.</p> <p>This option is valid for all add, remove, and move operations for AIX, Linux, and Virtual I/O Server logical partitions.</p>                                                                                                                  |
| <b>-d</b>                   | <p>Specifies the level of detail to be displayed upon return of an add or remove operation. Valid values are 0 (none) through 5 (highest).</p> <p>If this option is not specified, a default value of 0 is used.</p> <p>This option is valid for all add or remove operations for AIX, Linux, and Virtual I/O Server logical partitions.</p>                                                                                                                                                                                                  |
| <b>--entitled</b>           | <p>Dynamically changes the amount of I/O entitled memory that is assigned to the logical partition. Use the <b>-o</b> parameter to specify which operation to perform.</p>                                                                                                                                                                                                                                                                                                                                                                    |
| <b>--force</b>              | <p>Forces an add or remove operation to be performed for an AIX, Linux, and Virtual I/O Server logical partition that does not have an RMC connection to the Integrated Virtualization Manager. If this command completes successfully, you will need to restart your operating system for the change to take effect. Use this option only if you intentionally configured your LAN to isolate the Integrated Virtualization Manager from the operating system of your logical partition.</p>                                                 |
| <b>-a <i>attributes</i></b> | <p>Specifies the configuration data needed to create virtual I/O adapters or set hardware resource related attributes. The configuration data consists of attribute name/value pairs, which are in comma separated value (CSV) format. The configuration data must be enclosed in double quotation marks. Possible values are:</p> <p><b>adapter_type</b><br/>Whether the adapter is a client or server. Valid values are client and server. IVM requires the value to be client. This attribute is optional.</p> <p><b>addl_vlan_ids</b></p> |

**conn\_speed**

The configured connection speed of the port in Mbps. Possible values are:

- auto
- 10
- 100
- 1000
- 10000

**flow\_control**

Receive flow control value of the port

- 1 (enabled)
- 0 (disabled)

**ieee\_virtual\_eth**

Valid values:

- 0 - not IEEE 802.1Q compatible
- 1 - IEEE 802.1Q compatible

**is\_trunk**

Valid values:

- 0 - no
- 1 - yes

**lhea\_capabilities**

Comma-separated list of logical Host Ethernet Adapter capabilities, with each capability having one of the following formats: *adapter-ID/capability* or *adapter-ID/5/ieq/nieq/qp/cq/mr* where *ieq* (interruptible event queues), *nierq* (non-interruptible event queues), *qp* (queue pairs), *cq* (completion queues), and *mr* (memory regions) each specify the resource amount in addition to the base minimum. Valid values are:

- 0 - minimum
- 1 - low
- 2 - medium
- 3 - high
- 4 - dedicated
- 5 - custom

**max\_pool\_mem**

Maximum amount of physical memory that can be assigned to the memory pool without taking the pool offline. The value should be designated in megabytes and must be a multiple of the logical memory block (LMB) size. This attribute is required. The attribute supports the *=/+="/-="* format.

**max\_rcv\_packet\_size**

Maximum transmission unit (MTU) size of the port. Possible values:

- 1500 (default)
- 9000 (jumbo)

**mem\_weight**

The shared memory weight of the shared memory partition. This flag is used for determining the priority of logical partitions in a memory pool for distribution of memory. This attribute is only supported when the value of the *-o* parameter is *s*. This attribute supports the *=* format.

**paging\_storage\_pool**

The name of the default paging storage pool. This is the storage pool from which new paging space devices are created. This attribute is optional, but can only be defined during the creation stage.

**paging\_vios\_ids**

A comma-separated list of logical partition IDs of paging service partitions to associate with this memory pool. This attribute is optional, but if specified, must be equal to 1.

**paging\_vios\_names**

A comma-separated list of logical partition names of paging service partitions to associate with this memory pool. This attribute is optional, but if specified, must be equal to the name of the Virtual I/O Server logical partition.

**pend\_port\_group\_mcs\_value**

A Host Ethernet Adapter port group attribute value. In order to synchronize the pending value with the current value, you need to reboot the system.

**pool\_mem**

The amount of physical memory to assign, to add, or remove from the memory pool in megabytes. This required value must be a multiple of the logical memory block size.

The attribute supports the `=/+=-/=` format.

**port\_vlan\_id**

This option is required if you are adding a virtual Ethernet adapter.

**promisc\_lpar\_id**

The ID assigned to the promiscuous logical partition. Possible values:

- none
- 1 - 254

**promisc\_lpar\_name**

Name assigned to the promiscuous logical partition.

**sharing\_mode**

The sharing mode of the logical partition.

Valid values are:

**keep\_idle\_procs**: Never share processors

**share\_idle\_procs**: Share processors only when logical partition is inactive

**share\_idle\_procs\_always**: Always share processors

**share\_idle\_procs\_active**: Share processors only when the logical partition is active

**cap**: Capped mode. **uncap**: Uncapped mode.



**remote\_lpar\_id**

The ID of the logical partition that has the virtual SCSI or virtual fibre channel server adapter.

A value of any indicates that any client adapter should be allowed to connect to this adapter.

Integrated Virtualization Manager (IVM) requires the value to be 1. This attribute is optional.

**Note:** The **remote\_lpar\_id** attribute and the **remote\_lpar\_name** attribute are mutually exclusive.

**remote\_lpar\_name**

Specifies a user-defined name for the logical partition where the virtual fibre channel or virtual SCSI server adapter will be created.

IVM will require the value of this to be the name of the Virtual I/O Server logical partition. This attribute is optional.

**Note:** The **remote\_lpar\_id** attribute and the **remote\_lpar\_name** attribute are mutually exclusive.

**remote\_slot\_num**

Specifies the slot number where the virtual fibre channel or virtual SCSI server adapter will be created on the remote logical partition.

If the value is not specified, IVM will use the next available slot. This attribute is optional.

**trunk\_priority**

Valid values are integers between 1 and 15, inclusive. Required for a trunk adapter.

**uncap\_weight**

The weighted average of processing priority when in uncapped sharing mode. The smaller the value, the lower the weight. Possible values are 0 - 255.

**vlan\_id\_list**

List of virtual LAN IDs that the logical port has access to.

**wwpns** Designates the worldwide port names to this adapter, which are entered as a comma-separated list. If the worldwide port names are not specified, IVM will request that a pair be allocated. Each worldwide port name must be a 16-character hexadecimal value, and two worldwide port names must be specified.

**--help**

Displays the help text for this command and exit.

**Exit Status**

This command has a return code of 0 on success.

**Examples****chkdev command****Purpose**

Checks devices for virtual device provisioning capability in the Virtual I/O Server.

**Syntax**

**chkdev** [-dev *Name*] [-verbose] [-field *FieldName ...*] [-fmt *delimiter*]

## Description

The **chkdev** command determines whether a device can be used in physical to virtual transitions in the Virtual I/O Server. Devices include all physical volumes and multipath representations of logical units. If the **-dev** option is used, information is shown for the device specified. If the **-verbose** option is used, detailed information is shown. With the **-field** option, the user can specify which fields to display. With the **-fmt** option, the user can format the output to be separated by a given delimiter.

Capability in a virtual environment refers to the capability for a device to move between a physical environment to a virtual environment. The command displays these fields: *Phys2Virt\_Capable*, *Virt2NPIV\_Capable*, and *Virt2Phys\_Capable*.

For a device to be capable of moving from a physical to a virtual environment, it must pass the address compatibility check. If the physical volume can be moved from a physical to a virtual environment, the *Phys2Virt\_Capable* field has a value of YES. If the device is not capable of this move, the field has a value of NO. If the device is already in use by Virtual Small Computer System Interface (VSCSI), the field has a value of NA.

If a device is capable of moving from a VSCSI LUN environment to an N\_Port ID Virtualization (NPIV) environment, it must pass the address-compatibility check. If this device can be moved, the *Virt2NPIV\_Capable* field has a value of YES. If the device is not capable of this move, the field has a value of NO. If the device is not in use by the virtual target device (VTD), this field has a value of NA. This command is not sufficient to guarantee that the client can read the data after the VSCSI logical unit number (LUN) is removed. The storage area network (SAN) configuration must also be checked by the SAN administrator to ensure that it is mapped to a client NPIV.

Determination of whether a device is capable of moving from a virtual environment to a physical environment is applicable for devices that are mapped as VTDs. For devices to be capable of moving from a virtual environment to a physical environment, it must pass the address capability test. If the physical volume is capable of being transferred from a virtual environment to a physical environment, the *Virt2Phys\_Capable* field has a value of YES. If the device is not capable of this move, the field has a value of NO. If the device is not in use by VTD, the field has a value of NA.

## Flags

|                                    |                                                                                      |
|------------------------------------|--------------------------------------------------------------------------------------|
| <b>-dev</b> <i>Name</i>            | Specifies the device logical name for the device for which compatibility is checked. |
| <b>-verbose</b>                    | Displays additional details, including values for IEEE, unique_id, and PVID.         |
| <b>-field</b> <i>FieldName ...</i> | Specifies a list of fields to be displayed.                                          |
| <b>-fmt</b> <i>delimiter</i>       | Divides the output by using a user-specified delimiter character.                    |

## Exit Status

|    |                                    |
|----|------------------------------------|
| 0  | The command completed successfully |
| >0 | An error occurred.                 |

## Examples

1. To check device capability and to display unique identifier information for a particular device, type the following command:

```
chkdev -dev hdisk5 -verbose
```

The system displays a message similar to the following:

```
Name: hdisk5
IDENTIFIER: 210Chp0-c4JqYs9g04N37006NETAPPfcp
PHYS2VIRT_CAPABLE: NA
VIRT2NPIV_CAPABLE: YES
```

```
VIRT2PHYS_CAPABLE: YES
PVID:
UDID: 210Chp0-c4JqYs9g04N37006NETAPPfcp
IEEE:
VTD: vtscsi0
```

2. To check device capability and to display unique identifier information for a particular device, type the following command:

```
chkdev -dev hdisk6 -verbose
```

The system displays a message similar to the following:

```
Name: hdisk6
IDENTIFIER: 210Chp0-c4JqYs9g04N37006NETAPPfcp
PHYS2VIRT_CAPABLE: YES
VIRT2NPIV_CAPABLE: NA
VIRT2PHYS_CAPABLE: NA
PVID:
UDID:
IEEE: 210Chp0-c4JqYs9g04N37006NETAPPfcp
VTD:
```

3. To check device capability and to display unique identifier information for a particular device, type the following command:

```
chkdev -dev hdisk8 -verbose
```

The system displays a message similar to the following:

```
Name: hdisk8
IDENTIFIER:
PHYS2VIRT_CAPABLE: NO
VIRT2NPIV_CAPABLE: NA
VIRT2PHYS_CAPABLE: NA
PVID:
UDID:
IEEE:
VTD:
```

4. To check device capability for a particular device, type the following command:

```
chkdev -dev hdisk5
```

The system displays a message similar to the following:

```
Name: hdisk5
IDENTIFIER: 210Chp0-c4JqYs9g04N37006NETAPPfcp
PHYS2VIRT_CAPABLE: NA
VIRT2NPIV_CAPABLE: YES
VIRT2PHYS_CAPABLE: YES
```

5. To list the PHYS2VIRT\_CAPABLE field for a particular device, type the following command:

```
chkdev -dev hdisk5 -field name phys2virt_capable
```

The system displays a message similar to the following:

```
Name: hdisk5
PHYS2VIRT_CAPABLE: NA
```

6. To check all devices for device-capability information, type the following command:

```
chkdev
```

The system displays a message similar to the following:

```
Name: hdisk5
IDENTIFIER: 210Chp0-c4JqYs9g04N37006NETAPPfcp
PHYS2VIRT_CAPABLE: NA
VIRT2NPIV_CAPABLE: YES
VIRT2PHYS_CAPABLE: YES
```

```
Name: hdisk6
IDENTIFIER:
 PHYS2VIRT_CAPABLE: YES
 VIRT2NPIV_CAPABLE: NA
 VIRT2PHYS_CAPABLE: NA

Name: hdisk7
IDENTIFIER: 0000c2aaec372704
PHYS2VIRT_CAPABLE: NA
VIRT2NPIV_CAPABLE: NO
VIRT2PHYS_CAPABLE: NO

Name: hdisk8
IDENTIFIER:
PHYS2VIRT_CAPABLE: NO
VIRT2NPIV_CAPABLE: NA
VIRT2PHYS_CAPABLE: NA
```

## Related Information

The `lsdev` command and the `rmdev` command.

---

## chlang command

### Purpose

Changes the language settings for the system.

### Syntax

```
chlang { [-msg msgtran] -lang Name [-dev Media | -ls }
```

### Description

The `chlang` command is a high-level command that changes the language and keyboard settings for the entire Virtual I/O Server. The user needs to log out for language changes to take effect. If the language file sets are not already installed on the system, the `-dev` flag is used to specify their location.

When `chlang` is run with the `-ls` option, all available languages are listed.

### Flags

|                                  |                                                                                                                                                                                                                                         |
|----------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <code>-msg <i>msgtran</i></code> | Modifies the <code>NSLPATH</code> environment variable. The <i>msgtran</i> parameter is a colon-separated list of message translations (locale names) that indicates the message translation hierarchy required for the system or user. |
| <code>-dev <i>Media</i></code>   | Specifies the device or directory containing the images to install.                                                                                                                                                                     |
| <code>-lang <i>Name</i></code>   | Specifies the language-territory (locale name) that will become the locale setting for the <code>LANG</code> environment variable.                                                                                                      |
| <code>-ls</code>                 | Lists available languages.                                                                                                                                                                                                              |

### Exit Status

See “Exit status for Virtual I/O Server commands” on page 3.

### Examples

1. To change the language for the entire system to French Canadian, type:  
`chlang -lang fr_CA`
2. To display available languages:

chlang -ls

## Related Information

The `license` command.

---

## IVM chled command

### Purpose

Change states for physical and virtual LEDs. This command is operable only in an Integrated Virtualization Manager environment.

### Syntax

To change virtual partition LEDs:

```
chled -r sa -t virtuallpar -o Operation { -p PartitionName | --id PartitionID } [-m ManagedSystem]
```

To change virtual system LEDs:

```
chled -r sa -t virtualsys -o Operation [-m ManagedSystem]
```

### Description

The `chled` command changes the states of physical and virtual LEDs.

### Flags

|                                      |                                                                                                                                                                                                                                                                                                                           |
|--------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>-r</b> <i>ResourceType</i>        | The type of LED resource to change. The only valid value is sa for System Attention (SA) LED.                                                                                                                                                                                                                             |
| <b>-t</b> <i>SystemAttentionType</i> | The type of System Attention (SA) LEDs to change.                                                                                                                                                                                                                                                                         |
|                                      | <br>                                                                                                                                                                                                                                                                                                                      |
|                                      | <b>-r sa -t virtuallpar</b>                                                                                                                                                                                                                                                                                               |
|                                      | Change attributes for virtual partition system attention LEDs                                                                                                                                                                                                                                                             |
|                                      | <b>Attributes:</b> lpar_id, lpar_name, state                                                                                                                                                                                                                                                                              |
|                                      | <b>Filters:</b> None                                                                                                                                                                                                                                                                                                      |
|                                      | <br>                                                                                                                                                                                                                                                                                                                      |
|                                      | <b>-r sa -t virtualsys</b>                                                                                                                                                                                                                                                                                                |
|                                      | Change attributes for virtual system attention LEDs                                                                                                                                                                                                                                                                       |
|                                      | <b>Attributes:</b> state                                                                                                                                                                                                                                                                                                  |
|                                      | <b>Filters:</b> None                                                                                                                                                                                                                                                                                                      |
| <b>-o</b>                            | The operation to perform on the LED. Valid values are:                                                                                                                                                                                                                                                                    |
|                                      | <b>off</b> deactivate the LED                                                                                                                                                                                                                                                                                             |
|                                      | <b>on</b> activate the LED                                                                                                                                                                                                                                                                                                |
| <b>-m</b> <i>ManagedSystem</i>       | The name of the managed system. This attribute is optional because there is only one system to manage. The name may either be the user-defined name for the managed system, or be in the form tttt-mmm*sssssss, where tttt is the machine type, mmm is the model, and sssssss is the serial number of the managed system. |
| <b>-p</b> <i>PartitionName</i>       | The name of the partition for which the virtual partition system attention LED is to be changed.                                                                                                                                                                                                                          |
| <b>--id</b> <i>PartitionID</i>       | The ID of the partition for with the virtual partition system attention LED is to be changed.                                                                                                                                                                                                                             |

## Exit Status

This command has a return code of 0 on success.

## Security

This command is not accessible by users with the ViewOnly role.

## Examples

1. To deactivate the virtual system attention LED for the system, type:  
`chled -r sa -t virtualsys -o off`
2. To activate the virtual partition system attention LED for partition lpar3, type:  
`chled -r sa -t virtualpar -o on -p lpar3`
3. To deactivate the virtual partition system attention LED for the partition with ID 3, type:  
`chled -r sa -t virtualpar -o off --id 3`

## Related Information

The `Isled` command.

---

## IVM chlparutil command

### Purpose

Change settings for data collection. This command is operable only in an Integrated Virtualization Manager environment.

### Syntax

```
chlparutil -r config -s SampleRate [-m ManagedSystem]
```

### Description

The `chlparutil` command changes the settings for data collection such as the sample rate at which data is collected.

### Flags

|                                      |                                                                                                                                                                                                                                                                                                                                                                              |
|--------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <code>-r <i>ResourceType</i></code>  | The type of resources to change:<br><br><b>config</b> Change configuration settings.                                                                                                                                                                                                                                                                                         |
| <code>-s <i>SampleRate</i></code>    | The interval in seconds to sample utilization data. An interval of zero disables sampling. Valid sample rates are 0, 30, 60, 300, 1800, and 3600.                                                                                                                                                                                                                            |
| <code>-m <i>ManagedSystem</i></code> | The name of the managed system. This attribute is optional because there is only one system to manage. The name may either be the user-defined name for the managed system, or be in the form <code>ttt-mmm*sssssss</code> , where <code>ttt</code> is the machine type, <code>mmm</code> is the model, and <code>sssssss</code> is the serial number of the managed system. |

## Exit Status

This command has a return code of 0 on success.

## Security

This command is not accessible by users with the ViewOnly role.

## Examples

1. Disable the collection of utilization data:

```
chlparutil -r config -s 0
```

## Related Information

The `lslparutil` command.

---

## chlv command

### Purpose

Changes the characteristics of a logical volume.

### Syntax

To change the name of a logical volume:

```
chlv { -lv NewLogicalVolumeName | -bbr value } LogicalVolume
```

### Description

The `chlv` command changes the characteristics of a logical volume according to the command flags. The *LogicalVolume* parameter can be a logical volume name or logical volume ID. The `chlv` command can rename and set the bad-block relocation (BBR) policy. Changing the name of a logical volume and setting the bad-block relocation policy are the only supported options.

**Note:** A logical volume that has been assigned as a backing device or as a paging space device cannot be renamed. Although supported, it is not recommended that the bad-block relocation policy be enabled on the Virtual I/O Server for Virtual Small Computer System Interface (VSCSI) devices. Instead, the bad-block relocation policy needs to be enabled for VSCSI devices on the client. Doing so ensures that the bad-block relocation policy is enabled for VSCSI devices that use logical volumes and physical volumes as storage and provides better performance. Also, enabling the bad-block relocation policy is recommended for paging spaces used by Active Memory™ Sharing.

### Flags

|                   |                                                                                                                                                                                                                           |
|-------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <code>-lv</code>  | Specifies the logical volume.                                                                                                                                                                                             |
| <code>-bbr</code> | Sets the bad-block relocation policy. The value can be either: <ul style="list-style-type: none"><li>• yes (Causes bad-block relocation to occur.)</li><li>• no (Prevents bad block relocation from occurring.)</li></ul> |

## Examples

1. To change the logical volume name from `oldlv` to `newlv`, type:

```
chlv -lv newlv oldlv
```
2. To enable the bad-block relocation policy for the logical volume `testlv`, type:

```
chlv -bbr yes testlv
```
3. To disable the bad-block relocation policy for the logical volume `lv01`, type:

```
chlv -bbr no lv01
```

---

## chpath command

### Purpose

Changes the operational status of paths to a MultiPath I/O (MPIO) capable device, or changes an attribute associated with a path to an MPIO capable device.

### Syntax

```
chpath -dev Name -op OpStatus [-pdev Parent] [-conn Connection]
```

```
chpath -dev Name -pdev Parent [-conn Connection] [-perm] -attr Attribute=Value...
```

### Description

The **chpath** command either changes the operational status of paths to the specified device (the **-dev Name** flag) or it changes one, or more, attributes associated with a specific path to the specified device. The required syntax is slightly different depending upon the change being made.

The first syntax shown above changes the operational status of one or more paths to a specific device. The set of paths to change is obtained by taking the set of paths which match the following criteria:

- The target device matches the specified device.
- The parent device matches the specified parent (**-pdev Parent**), if a parent is specified.
- The connection matches the specified connection (**-conn Connection**), if a connection is specified.
- The path status is **PATH\_AVAILABLE**

The operational status of a path refers to the usage of the path as part of MPIO path selection. The value of **enable** indicates that the path is to be used while **disable** indicates that the path is not to be used. It should be noted that setting a path to **disable** impacts future I/O, not I/O already in progress. As such, a path can be disabled, but still have outstanding I/O until such time that all of the I/O that was already in progress completes. As such, if **-op disable** is specified for a path and I/O is outstanding on the path, this fact will be displayed.

Disabling a path affects path selection at the device driver level. The **path\_status** of the path is not changed in the device configuration database. The **lspath** command must be used to see current operational status of a path.

The second syntax shown above changes one or more path specific attributes associated with a particular path to a particular device. Note that multiple attributes can be changed in a single invocation of the **chpath** command; but all of the attributes must be associated with a single path. In other words, you cannot change attributes across multiple paths in a single invocation of the **chpath** command. To change attributes across multiple paths, separate invocations of **chpath** are required; one for each of the paths that are to be changed.

### Flags

**-attr Attribute=Value**

Identifies the attribute to change as well as the new value for the attribute. The Attribute is the name of a path specific attribute. The Value is the value which is to replace the current value for the Attribute. The *Attribute=Value* parameter can use one attribute value pair or multiple attribute value pairs for one **-attr** flag. If you use an **-attr** flag with multiple attribute value pairs, the list of pairs must be enclosed in quotation marks with spaces between the pairs. For example, entering **-attr Attribute=Value** lists one attribute value pair per flag, while entering **-attr 'Attribute1=Value1 Attribute2=Value2'** lists more than one attribute value pair.



|                                |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|--------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>-dev</b> <i>Name</i>        | Specifies the logical device name of the target device for the path(s) affected by the change. This flag is required in all cases.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>-pdev</b> <i>Parent</i>     | Indicates the logical device name of the parent device to use in qualifying the paths to be changed. This flag is required when changing attributes, but is optional when change operational status.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>-perm</b>                   | Changes the path's characteristics without actually changing the path. The change takes effect on the path the next time the path is unconfigured and then configured (possibly on the next boot).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>-conn</b> <i>Connection</i> | Indicates the connection information to use in qualifying the paths to be changed. This flag is optional when changing operational status. When changing attributes, it is optional if the device has only one path to the indicated parent. If there are multiple paths from the parent to the device, then this flag is required to identify the specific path being changed.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>-op</b> <i>OpStatus</i>     | Indicates the operational status to which the indicated paths should be changed. The operational status of a path is maintained at the device driver level. It determines if the path will be considered when performing path selection. The allowable values for this flag are: <ul style="list-style-type: none"> <li><b>enable</b> Mark the operational status as enabled for MPIO path selection. A path with this status will be considered for use when performing path selection. Note that enabling a path is the only way to recover a path from a failed condition.</li> <li><b>disable</b> Mark the operational status as disabled for MPIO path selection. A path with this status will not be considered for use when performing path selection. This flag is required when changing operational status. When used in conjunction with the <b>-attr Attribute=Value</b> flag, a usage error is generated.</li> </ul> |

## Exit Status

See "Exit status for Virtual I/O Server commands" on page 3.

## Examples

- To disable the paths between **scsi0** and the **hdisk1** disk device, enter:

```
chpath -dev hdisk1 -pdev scsi0 -op disable
```

The system displays a message similar to one of the following:

```
paths disabled
```

or

```
some paths disabled
```

The first message indicates that all **PATH\_AVAILABLE** paths from **scsi0** to **hdisk1** have been successfully enabled. The second message indicates that only some of the **PATH\_AVAILABLE** paths from **scsi0** to **hdisk1** have been successfully disabled.

## Related Information

The **cfgdev** command, the **chdev** command, the **lsdev** command, the **lsmmap** command, the **lspath** command, the **mkpath** command, the **mkvdev** command, the **rmdev** command, and the **rmpath** command.

---

## chrep command

### Purpose

Change the characteristics of the Virtual Media Repository.

## Syntax

`chrep -size Size`

## Description

The **chrep** command increases the size of the Virtual Media Repository by at least the amount specified with the **-size** flag. The actual amount of storage added to the repository is dependent on the allocation size of the parent storage pool. Use the **lssp** command to determine the parent storage pools allocation size. The specified flag cannot be a negative number.

| Size   | Minimum file storage pool size |
|--------|--------------------------------|
| ###M/m | ###MB                          |
| ###G/g | ###GB                          |

## Flags

**-size Size** Specifies the minimum amount of storage to add to the file storage pool.

## Examples

To increase the size of the Virtual Media Repository by at least 64 megabyte, type the following command:

```
chrep -size 64m
```

---

## chrole command

### Purpose

Changes the role attributes.

### Syntax

```
chrole [Attribute=Value...] [role]
```

### Description

The **chrole** command changes attributes for the role that is identified by the *role* parameter. The role name must already exist. To change an attribute, specify the attribute name and the new value by using the *Attribute=Value* parameter.

If you specify a single incorrect attribute or attribute value with the **chrole** command, the command does not change any attribute.

**Note:** After modifying the roles, run the **setkst** command to update the kernel database with those roles that were modified.

### Security

The **chrole** command is a privileged command. To run the command successfully, assume a role that has the authorization **vios.security.role.change**.

## Attributes

You can set the following attributes with enough authority.

|                       |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|-----------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>auth_mode</b>      | Specifies the authentication that is required to assume the role when the <b>swrole</b> command is used. You can specify the following values: <ul style="list-style-type: none"><li>• NONE<br/>No authentication is required</li><li>• INVOKER<br/>The invoker of the <b>swrole</b> command is required to enter a password to assume the role. The <b>INVOKER</b> value is the default value.</li></ul>                                                                                                                                                                                                                                                              |
| <b>authorizations</b> | Specifies a list of additional authorizations required for this role beyond those defined by the roles in the <b>rolelist</b> attribute. The <i>Value</i> parameter is a list of authorization names, separated by commas.                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>dfltmsg</b>        | Contains the default role-description text to use if message catalogs are not in use.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>groups</b>         | Specifies a list of groups to which a user should belong, in order to effectively use this role. This attribute is only for information and does not automatically make the user a member of the list of groups. The <i>Value</i> parameter is a list of group names, separated by commas.                                                                                                                                                                                                                                                                                                                                                                             |
| <b>id</b>             | Specifies the unique numeric ID for the role. You must specify the <b>id</b> attribute.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>msgcat</b>         | <b>Attention:</b> Do not modify the attribute value after the role is assigned to a user. Contains the file name of the message catalog that holds the one-line description of system roles. The <i>Value</i> parameter is a character string.                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>msgnum</b>         | Contains the index into a message catalog for a description of the role. The <i>Value</i> parameter is an integer.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>msgset</b>         | Contains the message set that includes the role description in the message catalog.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>rolelist</b>       | Lists the roles implied by this role. The <i>Value</i> parameter is a list of role names, separated by commas.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>screens</b>        | Lists the SMIT screen identifiers that allow roles to be mapped to various SMIT screens. The <i>Value</i> parameter is a list of SMIT screen identifiers, separated by commas.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>visibility</b>     | Specifies the visibility status of the role to the system. The <i>Value</i> parameter is an integer. Possible values follow: <ul style="list-style-type: none"><li><b>1</b> The role is enabled, displayed, and selectable. Authorizations contained in this role are applied to the user. If the attribute does not exist or has no value, the default value is 1.</li><li><b>0</b> The role is enabled and displayed as existing, but not selectable through a visual interface. Authorizations contained in this role are applied to the user.</li><li><b>-1</b> The role is disabled. Authorizations contained in this role are not applied to the user.</li></ul> |

## Exit Status

0 for success.

Nonzero for command error.

## Examples

1. To change the authorizations of the *ManagePasswds* role to *vios.security.passwd*, type the following command:  

```
chrole authorizations=vios.security.passwd ManagePasswds
```
2. To change the ID of the role *ManagePasswds* to 16000, type the following command:  

```
chrole id=16000 ManagePasswds
```

---

## chsp command

### Purpose

Changes the characteristics of a storage pool.

### Syntax

To increase the size of a file storage pool:

```
chsp -add [-sp StoragePool] -size Size
```

To add physical volume to a logical volume storage pool:

```
chsp -add [-f] [-sp StoragePool] PhysicalVolume...
```

To remove a physical volume from a logical volume storage pool :

```
chsp -rm [-f] [-sp StoragePool] PhysicalVolume...
```

To set storage pool as the default:

```
chsp -default StoragePool
```

To add one or more physical volumes to a cluster storage pool:

```
chsp -add [-f] -clustername ClusterName -sp StoragePool PhysicalVolume...
```

To replace one or more physical volumes to a cluster storage pool:

```
chsp -replace -clustername ClusterName -sp StoragePool -oldpv PhysicalVolume... -newpv PhysicalVolume...
```

### Description

The **chsp** command adds and removes physical volumes represented by the *PhysicalVolume* parameter from a logical volume and shared storage pool. Before adding a physical volume, the **chsp** command checks to verify that it is not already in another volume group or storage pool, or is not assigned to a shared memory pool (to be used as a paging space device by a shared memory partition). If the system detects a description area from a volume group or storage pool, the command fails. If the **-f** flag is specified, the physical volume is added even if it contains a description area, unless it is a member of another storage pool or volume group, or is assigned to a shared memory pool (to be used as a paging space device by a shared memory partition).

Before removing physical volumes from a logical volume the storage pool, the **chsp** command prompts the user to verify that each logical volume should be removed. If the user specifies the **-f** flag, the logical volumes are removed with out prompting for input. If all physical volumes in a storage pool are removed, the storage pool is also removed.

**Note:** A logical volume storage pool which contains file storage pools or the Virtual Media Repository cannot be removed.

If the **-sp** flag is not specified, the default storage pool is assumed.

If the specified, or default, storage pool is a file storage pool and the **-size** flag is specified, the given pool size is increased by at least the amount specified. The specified flag cannot be a negative number and file storage pools cannot be reduced in size.



- To replace physical volumes from the cluster storage pool, type the command as follows:  
`chsp -replace -clustername newcluster -sp viosp -oldpv hdisk1 -newpv hdisk2`

**Note:** The system displays the progress as a percentage.

---

## IVM chsvcevent command

### Purpose

Changes an existing serviceable event. This command is operable only in an Integrated Virtualization Manager environment.

### Syntax

To close an existing serviceable event:

```
chsvcevent -o close -p ProblemNumber -n Name -c CommentText [-m ManagedSystem]
```

### Description

The **chsvcevent** command closes an existing serviceable event.

### Flags

|                                |                                                                                                                                                                                                                                                                                                                                                                              |
|--------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>-o</b> <i>Operation</i>     | The operation to be performed. The only valid value is <code>close</code> .                                                                                                                                                                                                                                                                                                  |
| <b>-p</b> <i>ProblemNumber</i> | The problem number (problem_number) for the event, as displayed by the <code>lssvcevents</code> command.                                                                                                                                                                                                                                                                     |
| <b>-n</b>                      | A free form name string identifying the person who is closing the event.                                                                                                                                                                                                                                                                                                     |
| <b>-c</b>                      | A free form text comment about why the event is being changed.                                                                                                                                                                                                                                                                                                               |
| <b>-m</b> <i>ManagedSystem</i> | The name of the managed system. This attribute is optional because there is only one system to manage. The name may either be the user-defined name for the managed system, or be in the form <code>ttt-mmm*sssssss</code> , where <code>ttt</code> is the machine type, <code>mmm</code> is the model, and <code>sssssss</code> is the serial number of the managed system. |

### Exit Status

This command has a return code of 0 on success.

### Security

This command is not accessible by users with the `ViewOnly` role.

### Examples

- To close a serviceable event, type:

```
chsvcevent -o close -p 6013EFFF-205F3F22-4CC992E5-F8B6270-7540D8A3
-m 9111-520*XXXXXXX -n My Name -c Closing Comment
```

### Related Information

The `lssvcevents` command, and the `mksvcevent` command.

---

## IVM chsyscfg command

### Purpose

Changes attributes of logical partitions, logical partition profiles, or the managed system. This command is operable only in an Integrated Virtualization Manager environment.

### Syntax

To change system attributes:

```
chsyscfg -r sys { -f ConfigurationFile | -i ConfigurationData } [-m ManagedSystem]
```

To change partition attributes

```
chsyscfg -r lpar { -f ConfigurationFile | -i ConfigurationData } [-m ManagedSystem]
```

To change partition profile attributes, the Host Ethernet Adapter logical port assignments, or the logical Host Ethernet Adapter capabilities

```
chsyscfg -r prof { -f ConfigurationFile | -i ConfigurationData } [-m ManagedSystem]
```

### Description

The **chsyscfg** command changes attributes of logical partitions, logical partition profiles, or the managed system.

### Flags

|                                |                                                                                                                                                                                                                                                                                                                           |
|--------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>-r</b> <i>ResourceType</i>  | The type of resources to change:<br><b>sys</b> Managed system resources<br><b>lpar</b> Logical partition resources<br><b>prof</b> Logical partition profile resources                                                                                                                                                     |
| <b>-m</b> <i>ManagedSystem</i> | The name of the managed system. This attribute is optional because there is only one system to manage. The name may either be the user-defined name for the managed system, or be in the form tttt-mmm*sssssss, where tttt is the machine type, mmm is the model, and sssssss is the serial number of the managed system. |

## **-f** *ConfigurationFile*

The name of the file containing the configuration data needed to change the resources. The configuration data consists of attribute name and value pairs, which are in comma separated value (CSV) format. These attribute name and value pairs form a configuration record. A line feed marks the end of a configuration record. The file must contain one configuration record for each resource to be changed, and each configuration record must be for the same resource type. If the resource type is the managed system, then the file must contain only one configuration record.

The format of a configuration record is as follows:

```
attribute-name=value,attribute-name=value,...<LF>
```

Note that certain attributes accept a comma separated list of values, as follows:

```
"attribute-name=value,value,...",...<LF>
```

When a list of values is specified, the attribute name/value pair must be enclosed in double quotation marks. Depending on the shell being used, nested double quotation marks may need to be preceded by an escape character, which is usually a \ (back slash) character.

If '+=' is used in the attribute name/value pair instead of '=', then the specified value is added to the existing value for the attribute if the attribute is numerical. If the attribute is a list, then the specified values are added to the existing list.

If '-=' is used in the attribute name/value pair instead of '=', then the specified value is subtracted from the existing value for the attribute if the attribute is numerical. If the attribute is a list, then the specified values are deleted from the existing list.

### **Attribute names for logical partitions**

#### **allow\_perf\_collection**

Valid values are:

**0** - do not allow authority

**1** - allow authority

#### **name | lpar\_id**

Name or ID of the logical partition to change (required)

#### **new\_name**

New name for the logical partition.

#### **work\_group\_id**

Valid values are:

**none** - do not participate in the workload management group

**1** - participate in the workload management group

### **Attribute names for logical partition profiles**

#### **auto\_start**

Valid values are:

**0** - do not automatically start with system power on

**1** - automatically start with system power on



**boot\_mode**

Logical partition power on mode. Valid values are:

- norm** - normal
- dd** - diagnostic with default boot list
- ds** - diagnostic with stored boot list
- of** - Open Firmware OK prompt
- sms** - System Management Services

**desired\_io\_entitled\_mem**

The amount of I/O entitled memory for a shared memory partition. This is the portion of memory that is reserved for I/O mappings. Valid values follow:

- auto (automatically manage)
- *Number of megabytes*

If the value is auto, the entitlement is calculated based on the virtual I/O configuration of the logical partition. If the virtual I/O configuration is changed, the entitlement is updated automatically. If auto is not used, no automatic adjustments are made. The default value is auto.

**desired\_mem**

Assigned memory in megabytes

**desired\_procs**

Assigned processors. In shared processing mode, this refers to virtual processors.

**desired\_proc\_units**

Assigned shared processing units.

**lhea\_capabilities**

Comma-separated list of logical Host Ethernet Adapter capabilities, with each capability having one of the following formats: *adapter-ID/capability* or *adapter-ID/5/ieq/nieq/qp/cq/mr* where *ieq* (interruptible event queues), *niesq* (non-interruptible event queues), *qp* (queue pairs), *cq* (completion queues), and *mr* (memory regions) each specify the resource amount in addition to the base minimum. Valid values are:

- 0 - minimum
- 1 - low
- 2 - medium
- 3 - high
- 4 - dedicated
- 5 - custom

**lhea\_logical\_ports**

Comma separated list of Logical Host Ethernet Adapter (LHEA) logical ports, with each logical port having the following format:

*adapter-ID/port-group/physical-port-ID/  
logical-port-ID/allowed-VLAN-IDs*

All 4 '/' characters must be present, but optional values may be omitted. Optional values are allowed-VLAN-IDs.

**lpar\_avail\_priority**

The priority of the logical partition to maintain its entitled processors. If a processor failure occurs, processing resources will be removed first from the lowest priority logical partition. Valid values are 0 - 255.

**Note:** The Virtual I/O Server logical partition must have a higher priority than any other logical partition on the system.

**lpar\_proc\_compat\_mode**

The requested compatibility mode. Use `lssyscfg -r sys -F lpar_proc_compat_modes` to retrieve a list of valid values.

**max\_mem**

Maximum memory in megabytes.

**min\_mem**

Minimum memory in megabytes.

**min\_procs**

Minimum processors. In shared processing mode, this refers to virtual processors.

**max\_procs**

Maximum processors. In shared processing mode, this refers to virtual processors.

**max\_proc\_units**

Maximum shared processing units.

**min\_proc\_units**

Minimum shared processing units.

**max\_virtual\_slots**

Maximum number of virtual I/O adapter slots

**mem\_mode**

Logical partition memory mode. Valid values follow:

- **ded:** dedicated processor mode
- **shared:** shared processor mode

If the memory mode is shared, the logical partition cannot be assigned any physical I/O slots or host Ethernet adapter resources, the `proc_mode` attribute on the logical partition must be shared, and a memory pool must exist.

**Note:** The Virtual I/O Server logical partition only supports the dedicated memory mode.

**mem\_weight**

The shared memory weight of the shared memory partition. The memory weight is used to determine the priority of logical partitions in a memory pool for distributing memory. Valid values are 0 - 255. The default value is 128.

**name | lpar\_name | lpar\_id**

Name or ID of the logical partition to change (required)

**Note:** This command uses the profile name and `lpar_name` interchangeably because this environment does not support multiple profiles per logical partition. When using this command on the Hardware Management Console, you must specify the profile name and logical partition name or ID because it supports multiple profiles per logical partition.

**new\_name**

New name for the logical partition.

**paging\_device**

The paging space device to use if a memory pool is used. A paging space device is a block storage device that has been added to the memory pool and is not designated as a paging device for any other logical partition. This attribute is optional. If omitted, an appropriate paging device is selected automatically. If the `paging_device` value is a blank string and a paging space device is currently assigned to this partition, the paging space device is removed from the logical partition.

**proc\_mode**

Valid values follow:

- **ded**: dedicated processor mode
- **shared**: shared processor mode

**sharing\_mode**

The sharing mode of the logical partition. Valid values follow:

- **keep\_idle\_procs**: Never share processors
- **share\_idle\_procs**: Share processors only when the logical partition is inactive
- **share\_idle\_procs\_always**: Always share processors
- **share\_idle\_procs\_active**: Share processors only when the logical partition is active
- **cap**: Capped mode
- **uncap**: Uncapped mode

**uncap\_weight**

A weighted average of processing priority when in uncapped sharing mode. The smaller the value, the lower the weight. Valid values are: 0 - 255

**virtual\_eth\_adapters**

Comma-separated list of virtual Ethernet adapters, with each adapter having the following format: *slot\_number/is\_ieee/port\_vlan\_id/additional\_vlan\_ids/is\_trunk/is\_required* All 5 '/' characters must be present, but optional values may be omitted. Optional values are `is_ieee`, `additional_vlan_ids`, `is_required` and `is_trunk`.

Valid values for `is_ieee`, `is_trunk`, and `is_required`"

- 0 - no
- 1 - yes

For example, `4/0/2//0/0` specifies a virtual Ethernet adapter with a virtual slot number of 4, is not IEEE 802.1Q enabled, has a port virtual LAN ID of 2, no additional virtual LAN IDs, it is not a trunk adapter, and is not required.

### **virtual\_fc\_adapters**

Comma-separated list of virtual fibre channel adapters. Each item in this list has the following format:

```
virtual slot num/adapter_type/remote_lpar_id/
remote_lpar_name/remote_slot_num/wwpn_list/is_required
```

**Required values:** remote\_lpar\_id, remote\_lpar\_name, adapter\_type, virtual\_slot\_num

**Note:** You can specify either remote\_lpar\_id, remote\_lpar\_name, or use them both, but at least one of the values is required.

Valid values for adapter\_type:

- client
- server

**Note:** If you specify a value for adapter type, the Integrated Virtualization Manager (IVM) requires the adapter type to be a client.

**Optional values:** wwpn\_list, is\_required, remote\_slot\_num

When you add a virtual fibre channel adapter, wwpn\_list can be left blank to allow IVM to automatically assign worldwide port names to the client adapter. If you leave wwpn\_list blank and the virtual slot number specified for this adapter already contains a virtual fibre channel adapter, IVM will use the worldwide port names that are already assigned. New worldwide port names will only be generated if it is a new adapter. If you specify a wwpn\_list value, there must be exactly two values. Each worldwide port name must be a 16-character hexadecimal value. These values are not case sensitive.

Valid values for is\_required:

- 0 - no
- 1 - yes

A value of none or an empty string indicates that no virtual fibre channel adapters should be assigned.

**Note:** If the logical partition being changed is the Virtual I/O Server, you cannot change the current configuration. The IVM handles the client and server adapter as a pair; thus, IVM handles modifications automatically.

### **virtual\_scsi\_adapters**

Comma-separated list of virtual SCSI adapters. Each item in this list has the format:

```
slot_num/adapter_type/remote_lpar_id/remote_lpar_name/
remote_slot_num/is_required
```

The attribute names are not present in the list, just their values are present. If an attribute is optional and is not to be included, then no value is specified for that attribute. For example, `2/client//lpar2/3/0` specifies a virtual client SCSI adapter with a virtual slot number of 2, a server partition name of `lpar2`, a server slot number of 3, and is not required. The server partition ID was omitted.

**Required values:** `slot_num`, `adapter_type`, `remote_lpar_id`, `remote_lpar_name`

**Note:** You can specify `remote_lpar_id`, `remote_lpar_name`, or both, but at least one of the values is required.

**Optional values:** `is_required`, `remote_slot_num`

**Note:** IVM requires that virtual slot number 2 always contain a virtual SCSI adapter, so if you specify adapters in any other slots, the default adapter is still created in slot number 2. If you use the `chsyscfg` command with an empty list for the `virtual_scsi_adapters` attribute, all virtual SCSI adapters are removed, except the default adapter.

Valid values for `adapter_type`:

- `client`: client adapter
- `server`: server adapter, valid for Virtual I/O Server logical partitions only

Valid values for `is_required`:

- 0 - no
- 1 - yes

### **Attribute names for the managed system**

#### **lpar\_comm\_default**

Returns the `lpar_comm_ipaddr` to using the default IP address configured on the system as reported by `lstcpip -interfaces`. Valid values:

- 1: Default IP address will be used.

#### **lpar\_comm\_ipaddr**

The IP address through which client logical partitions will communicate with the management logical partition. This is used primarily for dynamic LPAR. It is defaulted to the first IP address available on your system, but can be manually set if desired. If you set this manually and then change your systems IP address, you must update this value.

**Note:** This attribute supports multiple IP addresses using a comma-separated list.

#### **new\_name**

New name for the managed system

#### **pend\_configured\_max\_lpars**

The maximum number of logical partitions which can be created after restarting the managed system.

This option is deprecated. Instead use:

```
chsyscfg -r prof
```

with attribute "max\_virtual\_slots" value for partition 1.

**-i ConfigurationData** This option allows you to enter configuration data on the command line, instead of using a file. Data entered on the command line must follow the same format as data in a file, and must be enclosed in double quotation marks.

When this option is used, only a single resource can be changed.

The **-i** and the **-f** options are mutually exclusive.

## Exit Status

This command has a return code of 0 on success.

## Security

This command is not accessible by users with the ViewOnly role.

## Examples

1. To change the user defined name for the managed system, type:  

```
chsyscfg -r sys -i "new_name=sys1"
```
2. To change logical partitions using the configuration data in the file `/tmp/lparfile`, type:  

```
chsyscfg -r lpar -f /tmp/lparfile
```
3. To reduce a partition profile's assigned and minimum memory by 256 MB, type:  

```
chsyscfg -r prof -i "lpar_name=partition3,min_mem-=256,desired_mem-=256"
```

## Related Information

The `lssyscfg` command, the `mksyscfg` command, and the `rmsyscfg` command.

---

## IVM chsysstate command

### Purpose

Changes the state of a partition. This command is operable only in the Integrated Virtualization Manager environment.

### Syntax

To activate a partition:

```
chsysstate -r lpar -o on { -n Name | --id PartitionID } [-k KeylockPosition] [-b BootMode] [-m ManagedSystem]
```

To perform a partition shut down using the **shutdown** command on the client operating system:

```
chsysstate -r lpar -o osshutdown { -n Name | --id PartitionID } [-m ManagedSystem]
```

To perform a delayed partition shut down (white button shut down):

```
chsysstate -r lpar -o shutdown { -n Name | --id PartitionID } [-m ManagedSystem]
```

To perform an immediate partition shutdown (operator panel function 8):

```
chsysstate -r lpar -o shutdown --immed { -n Name | --id PartitionID } [-m ManagedSystem]
```

To perform an immediate restart of a partition (operator panel function 3):

```
chsysstate -r lpar -o shutdown --immed --restart { -n Name | --id PartitionID } [-m ManagedSystem]
```

To perform a restart of a partition after initiating a dump (operator panel function 22):

```
chsysstate -r lpar -o dumprestart { -n Name | --id PartitionID } [-m ManagedSystem]
```

To change the keylock position for a partition:

```
chsysstate -r lpar -o chkey -k KeylockPosition { -n Name | --id PartitionID } [-m ManagedSystem]
```

## Description

The **chsysstate** command changes the state of a partition. To avoid data loss, use the shutdown facilities provided by the operating system in the partition.

## Flags

|                                  |                                                                                                                                                                                                                                                                                                                                                                                                                             |
|----------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>-r</b> <i>ResourceType</i>    | The type of resources to change: lpar Logical partition resources                                                                                                                                                                                                                                                                                                                                                           |
| <b>-m</b> <i>ManagedSystem</i>   | The name of the managed system. This attribute is optional because there is only one system to manage. The name may either be the user-defined name for the managed system, or be in the form tttt-mmm*sssssss, where tttt is the machine type, mmm is the model, and sssssss is the serial number of the managed system.                                                                                                   |
| <b>-o</b> <i>Operation</i>       | The type of operation to perform: <ul style="list-style-type: none"><li>• chkey: Changes the keylock position</li><li>• dumprestart: Restarts after initiating a dump</li><li>• on: Power on</li><li>• osshutdown: Shuts down using the client partition's <b>shutdown</b> command. This is the safest shutdown option, so should be used if available (RMC state must be active).</li><li>• shutdown: Shuts down</li></ul> |
| <i>IPL Source</i>                |                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>-b</b> <i>BootMode</i>        | Override the current power on mode setting. Valid values are: <ul style="list-style-type: none"><li>• norm: normal</li><li>• dd: diagnostic with default boot list</li><li>• ds: diagnostic with stored boot list</li><li>• of: Open Firmware OK prompt</li><li>• sms: System Management Services</li></ul>                                                                                                                 |
| <b>-k</b> <i>KeylockPosition</i> | Keylock position. Valid values are: <ul style="list-style-type: none"><li>• norm: normal keylock</li><li>• manual: manual keylock</li></ul>                                                                                                                                                                                                                                                                                 |
| <b>--immed</b>                   | Force the state change immediately.                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>--restart</b>                 | Restart the partition. This flag is valid only if the <b>--immed</b> flag is also specified.                                                                                                                                                                                                                                                                                                                                |

## Exit Status

This command has a return code of 0 on success.

## Security

This command is not accessible by users with the ViewOnly role.

## Examples

1. To power on the partition with an ID of 2 and set the boot mode to System Management Services, type the following command:  
`chsysstate -r lpar -o on --id 2 -b sms`
2. To shut down the partition with an ID of 3, type the following command:  
`chsysstate -r lpar -o shutdown --id 3`
3. To immediately restart the partition with an ID of 3 using the client operating system's shutdown command, type the following command:  
`chsysstate -r lpar -o osshutdown --restart --immed --id 3`

## Related Information

The `lssyscfg` command, the `mksyscfg` command, the `rmsyscfg` command, and the `mkvt` command.

---

## chtcip command

### Purpose

Changes the Virtual I/O Server TCP/IP settings and parameters.

### Syntax

To change a network interface:

```
chtcip [-interface Interface -inetaddr Address -netmask SubnetMask]
```

```
chtcip [-interface Interface -gateway -add New_gateway_address -remove Old_gateway_address]
```

```
chtcip -ip6-interface Interface [-inetaddr Address] [-plen prefixlen] [-gateway]
```

### Description

The `chtcip` command changes the TCP/IP settings and configuration on the Virtual I/O Server.

### Flags

|                                                 |                                                                                                                                                                         |
|-------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <code>-add <i>New_Gateway_Address</i></code>    | Specifies the new default gateway address to add.                                                                                                                       |
| <code>-inetaddr <i>Address</i></code>           | Changes the IP address of the host. Specify the address in dotted decimal notation.                                                                                     |
| <code>-interface <i>Interface</i></code>        | Specifies a particular network interface, for example <code>en0</code> .                                                                                                |
| <code>-ip6</code>                               | Specifies to use the IPv6 version of this command. The <code>-ip6</code> flag allows you to change an existing IPv6 address, state, prefix length, and default gateway. |
| <code>-gateway <i>Gateway</i></code>            | Changes the gateway address for a static route. Specify both the current address and new address in dotted decimal notation.                                            |
| <code>-netmask <i>SubnetMask</i></code>         | Specifies the subnet mask of the gateway, which is used to route through the appropriate subnetwork.                                                                    |
| <code>-plen <i>prefixLen</i></code>             | Specifies the prefix length of the IPv6 interface.                                                                                                                      |
| <code>-remove <i>Old_Gateway_Address</i></code> | Specifies the old default gateway address to remove.                                                                                                                    |

### Exit Status

|    |                                     |
|----|-------------------------------------|
| 0  | The command completed successfully. |
| >0 | An error occurred.                  |



## Examples

1. To changes the current network address and mask to the new settings, type:  
`chtcpip -interface en0 -inetaddr 9.1.1.1 -netmask 255.255.255.0`
2. To changes the default gateway from 9.1.2.3 to 9.2.3.4, type:  
`chtcpip -interface en0 -gateway -add 9.2.3.4 -remove 9.1.2.3`

**Note:** If more than one default gateway is defined, an error will occur. Only one default gateway can be defined by the **mktcpip** command.

## Related Information

The **topas** command, and the **mktcpip** command.

---

## chuser command

### Purpose

Changes user attributes.

### Syntax

`chuser [-ldap] -attr Attribute=Value ... Name`

### Description

The **chuser** command changes attributes for the user identified by the *Name* variable. To change an attribute, specify the attribute name and the new value with the **-attr Attribute=Value** variable.

Use the **-ldap** flag if the user is an LDAP user. Specifying the **-ldap** flag authenticates through the LDAP load module and changes the users attributes specified by the **-attr** flag.

### Attributes

The prime administrator (padmin) user can set the following attributes:

|                       |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|-----------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>account_locked</b> | Indicates whether the user account is locked. Possible values are as follows:<br><br><b>true   yes   always</b><br>The user's account is locked. The values <b>true</b> , <b>yes</b> , and <b>always</b> are equivalent. The user is denied access to the system.<br><br><b>false   no   never</b><br>The user's account is not locked. The values <b>false</b> , <b>no</b> , and <b>never</b> are equivalent. The user is allowed access to the system. The default value is false. |
| <b>default_roles</b>  | Specifies the default roles for the user. The <i>Value</i> parameter, a comma-separated list of valid role names, can only contain roles assigned to the user in the <b>roles</b> attribute. You can use the <b>ALL</b> keyword to signify that the <b>default roles</b> for the user are all their assigned roles.                                                                                                                                                                  |
| <b>expires</b>        | Identifies the expiration date of the account. The <i>Value</i> variable is a 10-character string in the <i>MMDDhhmmyy</i> form, where <i>MM</i> = month, <i>DD</i> = day, <i>hh</i> = hour, <i>mm</i> = minute, and <i>yy</i> = last 2 digits of the years 1939 through 2038. All characters are numeric. If the <i>Value</i> variable is 0, the account does not expire. The default is 0.                                                                                         |
| <b>histexpire</b>     | Defines the period of time (in weeks) that a user cannot reuse a password. The value is a decimal integer string. The default is 0, indicating that no time limit is set.                                                                                                                                                                                                                                                                                                            |
| <b>histsize</b>       | Defines the number of previous passwords a user cannot reuse. The value is a decimal integer string. The default is 0. Only an administrative user can change this attribute.                                                                                                                                                                                                                                                                                                        |

|                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|---------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>loginretries</b> | Defines the number of unsuccessful login attempts allowed after the last successful login before the system locks the account. The value is a decimal integer string. A zero or negative value indicates that no limit exists. Once the user's account is locked, the user is not able to log in until the prime administrator resets the user's <b>account_locked</b> attribute.                                                                                                                   |
| <b>maxage</b>       | Defines the maximum age (in weeks) of a password. The password must be changed by this time. The value is a decimal integer string. The default is a value of 0, indicating no maximum age. The value can be from 0 to 52.                                                                                                                                                                                                                                                                          |
| <b>maxexpired</b>   | Defines the maximum time (in weeks) beyond the <b>maxage</b> value that a user can change an expired password. The value is a decimal integer string. The default is -1, indicating restriction is set. If the <b>maxexpired</b> attribute is 0, the password expires when the <b>maxage</b> value is met. If the <b>maxage</b> attribute is 0, the <b>maxexpired</b> attribute is ignored. The value can be from 0 to 52.                                                                          |
| <b>maxrepeats</b>   | Defines the maximum number of times a character can be repeated in a new password. Since a value of 0 is meaningless, the default value of 8 indicates that there is no maximum number. The value is a decimal integer string. The value can be from 0 to 8.                                                                                                                                                                                                                                        |
| <b>minage</b>       | Specifies the minimum age at which a password can be changed. Passwords must be kept for a minimum period. This value is measured in weeks.                                                                                                                                                                                                                                                                                                                                                         |
| <b>minalpha</b>     | Specifies the minimum number of alphabetic characters.                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>mindiff</b>      | Specifies the minimum number of characters in the new password that are not in the old password. Note: This restriction does not consider position. If the new password is abcd and the old password is edcb, the number of different characters is 1.                                                                                                                                                                                                                                              |
| <b>minlen</b>       | Defines the minimum length of a password. The value is a decimal integer string. The default is a value of 0, indicating no minimum length. The maximum value allowed is 8. This attribute is determined by <b>minlen</b> or ' <b>minalpha + minother</b> ', whichever is greater. The values for ' <b>minalpha + minother</b> ' cannot be greater than 8. If ' <b>minalpha + minother</b> ' is greater than 8, then the effective value for <b>minother</b> is reduced to ' <b>8 - minalpha</b> '. |
| <b>minother</b>     | Defines the minimum number of non-alphabetic characters that must be in a new password. The value is a decimal integer string. The default is a value of 0, indicating no minimum number. The value can be from 0 to 8.                                                                                                                                                                                                                                                                             |
| <b>pgrp</b>         | Defines the Primary Group and Groups membership. The valid entries are <b>staff</b> and <b>view</b> . If this attribute is not defined, the default <b>staff</b> is used.                                                                                                                                                                                                                                                                                                                           |
| <b>pwdwarntime</b>  | Defines the number of days before the system issues a warning that a password change is required. The value is a decimal integer string. A zero or negative value indicates that no message is issued. The value must be less than the difference of the <b>maxage</b> and <b>minage</b> attributes. Values greater than this difference are ignored and a message is issued when the <b>minage</b> value is reached.                                                                               |
| <b>roles</b>        | Lists the administrative roles for this user. The <i>Value</i> parameter is a list of role names, separated by commas.                                                                                                                                                                                                                                                                                                                                                                              |
| <b>fsize</b>        | Defines the soft limit for the largest file a user's process can create or extend. The <i>Value</i> parameter is an integer representing the number of 512-byte blocks. To make files greater than 2G, specify -1 or unlimited. The minimum value for this attribute is 8192.                                                                                                                                                                                                                       |

## Exit Status

See “Exit status for Virtual I/O Server commands” on page 3.

## Security

This command can only be run by the prime administrator (**padmin**) user.

## Examples

To change the expiration date for the **davis** user account to 8 a.m., 1 May, 2010, type:

```
chuser -attr expires=0501080010 davis
```

## Related Information

The **lsuser** command, the **mkuser** command, the **rmuser** command, and the **passwd** command.

---

## IVM chvet command

### Purpose

Activates Virtualization Engine systems technologies. This command is operable only in an Integrated Virtualization Manager environment.

### Syntax

```
chvet -o <e> -k <activation code> [-m <managed system>] [--help]
```

### Description

The **chvet** command performs activation of Virtualization Engine systems technologies on the managed-system.

### Flags

|                                 |                                                                                                                                                                                                                                                                                                                                                                      |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>-o</b>                       | The operation to perform. The only valid value is <b>e</b> to enter an activation code.                                                                                                                                                                                                                                                                              |
| <b>-m</b> <i>managed system</i> | The name of the managed system for which the Virtualization Engine systems technologies activation is to be performed. The name may either be the user-defined name for the managed system, or be in the form <i>ttt-mmm*sssssss</i> , where <i>ttt</i> is the machine type, <i>mmm</i> is the model, and <i>sssssss</i> is the serial number of the managed system. |
| <b>-k</b>                       | The activation code (key) to enter. Letters may be entered in either upper case or lower case.                                                                                                                                                                                                                                                                       |
| <b>--help</b>                   | Display the help text for this command and exit.                                                                                                                                                                                                                                                                                                                     |

### Exit Status

This command has a return code of 0 on success.

### Examples

To enter an activation code, type the following command:

```
chvet -m mySystem -o e -k AlphaNumericString12345
```

---

## chvfilt command

### Purpose

Changes the values for the existing virtual LAN-crossing filter rule.

### Syntax

```
chvfilt [-v <4|6>] -n fid [-a <D|P>] [-z <svlan>] [-Z <dvlan>] [-s <s_addr>] [-d <d_addr>] [-o <src_port_op>] [-p <src_port>] [-O <dst_port_op>] [-P <dst_port>] [-c <protocol>]
```

## Description

The **chvfilt** command is used to change the definition of a virtual LAN-crossing filter rule in the filter rule table.

## Flags

|           |                                                                                                                                                                                                              |
|-----------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>-a</b> | Specifies the action. Valid values follow: <ul style="list-style-type: none"><li>• D (Deny): Blocks traffic</li><li>• P (Permit): Allows traffic</li></ul>                                                   |
| <b>-c</b> | Specifies different protocols to which the filter rule is applicable. Valid values follow: <ul style="list-style-type: none"><li>• udp</li><li>• icmp</li><li>• icmpv6</li><li>• tcp</li><li>• any</li></ul> |
| <b>-d</b> | Specifies the destination address in IPv4 or IPv6 format.                                                                                                                                                    |
| <b>-m</b> | Specifies the source address mask.                                                                                                                                                                           |
| <b>-M</b> | Specifies the destination address mask.                                                                                                                                                                      |
| <b>-n</b> | Specifies the filter ID of the filter rule that should be modified.                                                                                                                                          |
| <b>-o</b> | Specifies the source port or the Internet Control Message Protocol (ICMP) type operation. Valid values follow: <ul style="list-style-type: none"><li>• lt</li><li>• gt</li><li>• eq</li><li>• any</li></ul>  |
| <b>-O</b> | Specifies the destination port or the ICMP code operation. Valid values follow: <ul style="list-style-type: none"><li>• lt</li><li>• gt</li><li>• eq</li><li>• any</li></ul>                                 |
| <b>-p</b> | Specifies the source port or the ICMP type.                                                                                                                                                                  |
| <b>-P</b> | Specifies the destination port or the ICMP code.                                                                                                                                                             |
| <b>-s</b> | Specifies the source address in v4 or v6 format.                                                                                                                                                             |
| <b>-v</b> | Specifies the IP version of the filter rule table. Valid values are 4 and 6.                                                                                                                                 |
| <b>-z</b> | Specifies the virtual LAN ID of the source logical partition.                                                                                                                                                |
| <b>-Z</b> | Specifies the virtual LAN ID of the destination logical partition.                                                                                                                                           |

## Exit Status

This command returns the following exit values:

|    |                        |
|----|------------------------|
| 0  | Successful completion. |
| >0 | An error occurred.     |

## Examples

1. To change a valid filter rule that exists in the kernel, type the command as follows:

```
chvfilt -n 1 -v4 -a P -z 100 -Z 300 -o eq -p 23 -O lt -P 345 -c tcp
```

2. When a filter rule (n=2) does not exist in the kernel, the output is as follows:

```
chvfilt -n 2 -v4 -a P -z 100 -Z 300 -o eq -p 23 -O lt -P 345 -c tcp
```

The system displays the output as follows:

```
ioctl(QUERY_FILTER) failed no filter rule err=2
Cannot Change the filter rule.
```

## Related Information

The **genvfilter** command, the **lsvfilter** command, the **mkvfilter** command, the **rmvfilter** command, and the **vlantfw** command.

---

## chvg command

### Purpose

Sets the characteristics of a volume group.

### Syntax

```
chvg [-unlock] [-suspend | -resume] [-factor <num> [-chksize] VolumeGroup
```

### Description

The **chvg** command changes the characteristics of a volume group.

### Flags

|                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|-----------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>-resume</b>  | Resumes normal I/O operations for a volume group.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>-suspend</b> | Drains I/O's for this volume group and suspends future I/O's.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>-unlock</b>  | Unlocks the volume group. This option is provided if the volume group is left in a locked state by abnormal termination of another LVM operation (such as the command core dumping, or the system crashing).<br><b>Note:</b> Before using the <b>-unlock</b> flag, make sure that the volume group is not being used by another LVM command.                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>-factor</b>  | Changes the limit of the number of physical partitions per physical volume, specified by factor. factor should be between 1-16 for 32 disk volume groups and 1-64 for 128 disk volume groups.<br><br>If factor is not supplied, it is set to the lowest value such that the number of physical partitions in the volume group is less than factor x1016. If factor is specified, the maximum number of physical partitions per physical volume for the volume group changes to factor x1016.                                                                                                                                                                                                                                               |
| <b>-chksize</b> | Examines all the disks in the volume group to determine if they have grown in size. This option does not support resizing a volume group while it is activated in classic or enhanced concurrent mode. This option does not support resizing the <i>rootvg</i> volume group.<br><br>If any disks have grown in size, this option attempts to add additional physical partitions to the physical volume. If necessary, this option determines the appropriate 1016 multiplier and conversion to big volume group.<br><b>Note:</b> After running <b>chvg -chksize</b> , it might be necessary to run the <b>deactivatevg</b> , followed by running the <b>activatevg</b> command on the volume group to see the change of size of the disks. |

### Exit Status

See "Exit status for Virtual I/O Server commands" on page 3.

### Examples

1. To suspend the volume group **vg03**, type:

```
chvg
-suspend vg03
```

2. To resume the volume group **vg03**, type:

```
chvg
-resume vg03
```

3. To change the number of physical partitions per physical volume by 4, type:

```
chvg -factor 4 testvg
```

4. To determine whether the *testvg* volume group has grown in size after adding additional physical partitions to the physical volumes in the group, type:

```
chvg -chksize testvg
```

## Related Information

The **mkvg** command, the **lsvg** command, the **extendvg** command, the **reducevg** command, the **mirrorios** command, the **unmirrorios** command, the **activatevg** command, the **deactivatevg** command, the **importvg** command, the **exportvg** command, and the **syncvg** command.

---

## chvopt command

### Purpose

Change the characteristics of a virtual optical media disk within the Virtual Media Repository.

### Syntax

```
chvopt -name FileName {-access Value | -mv NewName}
```

### Description

The **chvopt** command renames or changes the access permission's of a virtual optical media disk within the Virtual Media Repository. If the **-access** flag is specified, then the permissions can be set to read-only (ro) or read-write (rw). If the **-mv** flag is specified the media file name is changed. The file name cannot be changed if the media file is loaded into a virtual optical device.

### Flags

|                              |                                                                                                           |
|------------------------------|-----------------------------------------------------------------------------------------------------------|
| <b>-dev</b> <i>Value</i>     | Specifies the new access permission. Valid values are:<br><br><b>ro</b> read-only<br><b>rw</b> read-write |
| <b>-name</b> <i>FileName</i> | Specifies the file name whose access permission is to be changed.                                         |
| <b>-mv</b> <i>NewName</i>    | Specifies the new file name.                                                                              |

### Examples

To change the access permissions on the file *clientData* to read-only, type the following command:

```
chvopt -name clientData -access ro
```

---

## chvlog command

### Purpose

Changes the configuration of a virtual log.

## Syntax

To change the properties of a virtual log by specifying the device name:

```
chvlog -dev DeviceName [-client ClientName] [-name LogName] [-lf FileCount] [-lfs FileSize] [-sf FileCount] [-sfs FileSize]
```

To change the properties of a virtual log by specifying its UUID:

```
chvlog {-uuid UUID | -u UUID} [-client ClientName] [-name LogName] [-lf FileCount] [-lfs FileSize] [-sf FileCount] [-sfs FileSize] [-state VirtualLogState]
```

## Description

The **chvlog** command changes the properties of a virtual log device.

The virtual log can be specified by its universal unique Identifier (UUID) (by using the *-uuid* parameter), or by its device name (by using the *-dev* parameter) if it is attached to a virtual SCSI (VSCSI) host adapter.

The UUID of virtual logs can be discovered by using the **lsvlog** command to display the properties of virtual logs.

The *state* property of a virtual log cannot be changed if the virtual log is attached to a VSCSI host adapter because virtual logs must be in the *enabled* state to be attached. To change the state of an attached virtual log, first use the **rmvlog -d** command to detach it from the VSCSI host adapter.

If the specified virtual log is shared across a VIOS cluster by using a shared storage pool, no changes can be made to it unless it is disconnected from VSCSI host adapters on all virtual I/O servers (VIOS) in the cluster.

## Flags

|                    |                                                                                              |
|--------------------|----------------------------------------------------------------------------------------------|
| <b>-c, -client</b> | Sets the client name, up to 96 characters.                                                   |
| <b>-dev</b>        | Changes the virtual log with the specified device name.                                      |
| <b>-lf</b>         | Sets the maximum number of log files, in the range 1 - 1000.                                 |
| <b>-lfs</b>        | Sets the maximum size of each log file, in bytes, or suffixed with K, M, or G.               |
| <b>-n, -name</b>   | Sets the log name, up to 12 characters.                                                      |
| <b>-s, -state</b>  | Sets the state of the virtual log to <i>disabled</i> , <i>enabled</i> , or <i>migrated</i> . |
| <b>-sf</b>         | Sets the maximum number of state files, in the range 1 - 1000.                               |
| <b>-sfs</b>        | Sets the maximum size of each state file, in bytes, or suffixed with K, M, or G.             |
| <b>-u, -uuid</b>   | Changes the virtual log with the specified UUID.                                             |

## Exit Status

Table 1. Command specific return codes

| Return code | Description                                             |
|-------------|---------------------------------------------------------|
| 0           | The virtual log configuration was updated successfully. |
| >0          | An error occurred.                                      |

## Examples

1. To change the log file size on a running virtual log device named *vtlog0* to 2 MB, type the command as follows:

```
chvlog -dev vtlog0 -lfs 2M
```

The system displays the output as follows:

Updated device.

2. To change the state of the virtual log that has UUID 00000000000000003cee6408c885d677 to *disabled*, type the command as follows:

```
chvlog -uuid 00000000000000003cee6408c885d677 -state disabled
```

The system displays the output as follows:

Updated device.

3. To attempt to change the state of the virtual log that has UUID 00000000000000003cee6408c885d677 to *disabled* when it is connected to a VSCSI host adapter, type the command as follows:

```
chvlog -uuid 00000000000000003cee6408c885d677 -state disabled
```

The system displays the output as follows:

To change the state, the virtual log must not be connected to a device.

## Related Information

The **chvcrepo** command, the **lsvlog** command, the **lsvlrepo** command, the **mkvlog** command, and the **rmvlog** command.

---

## chvcrepo command

### Purpose

Modifies the configuration of a virtual log repository.

### Syntax

To change the configuration of the local virtual log repository:

```
chvcrepo [-lf FileCount] [-lfs FileSize] [-sf FileCount] [-sfs FileSize] [-state RepositoryState] [-root Path]
```

To change the configuration of a virtual log repository in a shared storage pool:

```
chvcrepo {-sp StoragePool} [-lf FileCount] [-lfs FileSize] [-sf FileCount] [-sfs FileSize] [-state RepositoryState]
```

### Description

The **chvcrepo** command changes the configuration of virtual log repositories, including the default property values for new virtual logs created with the **mkvlog** command. Changing the default properties does not affect the properties of existing virtual logs.

By default, the local virtual log repository is changed, but a repository in a shared storage pool can be specified instead by using the *-sp* parameter.

The *state* property of a virtual log cannot be changed if virtual logs exist in the repository.

The *-root* parameter can be used only to change the repository root directory for the local virtual log repository, and only if no virtual logs exist in the repository. If the virtual log root directory is changed, any log data in the old location is preserved, but new virtual logs write data to the new location.



## Flags

|                   |                                                                                                                                                               |
|-------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>-lf</b>        | Sets the default maximum number of log files, in the range 1 - 1000.                                                                                          |
| <b>-lfs</b>       | Sets the default maximum size of each log file, in bytes, or suffixed with K, M, or G.                                                                        |
| <b>-s, -state</b> | Sets the state of the repository to <i>disabled</i> or <i>enabled</i> . Virtual logs can be created only in a repository that is in the <i>enabled</i> state. |
| <b>-sf</b>        | Sets the default maximum number of state files, in the range 1 - 1000.                                                                                        |
| <b>-sfs</b>       | Sets the default maximum size of each state file, in bytes, or suffixed with K, M, or G.                                                                      |
| <b>-sp</b>        | Specifies the name of the shared storage pool to use. If not specified, the local repository is changed.                                                      |
| <b>-root</b>      | Sets the root directory for the virtual log repository. All log data is stored here.                                                                          |

## Exit Status

Table 2. Command specific return codes

| Return code | Description                          |
|-------------|--------------------------------------|
| 0           | All files were written successfully. |
| >0          | An error occurred.                   |

## Examples

1. To set the default log file size for new virtual logs in the local virtual log repository to 2 MB, type the command as follows:

```
chvrepo -lfs 2M
```

The system displays the output as follows:

```
Updated repository.
```

2. To set the default log file size for new virtual logs in the virtual log repository in the shared storage pool *sspool1* to 2 MB, type the command as follows:

```
chvrepo -sp sspool1 -lfs 2M
```

The system displays the output as follows:

```
Updated repository.
```

3. To set the root directory for the local virtual log repository, type the command as follows:

```
chvrepo -root /mnt/logs/
```

The system displays the output as follows:

```
Updated repository.
```

4. To attempt to set the root directory for the virtual log repository in the shared storage pool *sspool1*, type the following command:

```
chvrepo -sp sspool1 -root /mnt/logs
```

The system displays the following output:

```
The root directory of a shared storage pool cannot be changed.
```

## Related Information

The **chvlog** command, the **lsvlog** command, the **lsvlrepo** command, the **mkvlog** command, and the **rmvlog** command.

---

## cleandisk command

### Purpose

Removes any storage pool or cluster signatures that are on the physical volume.

### Syntax

```
cleandisk -r hdiskX
```

```
cleandisk -s hdiskX
```

### Description

The **cleandisk** command removes any storage pool or cluster signatures that were detected on the physical volume. The physical volume specified by the **-s** flag has the storage pool signature removed. The physical volume specified by the **-r** flag has the cluster signature removed. Both the **-s** and **-r** flags are destructive to any data that exists on the physical volume.

**Note:** Ensure that you pass the correct *hdisk*. Commands are destructive to the data on the *hdisk*.

### Flags

|                         |                                                                           |
|-------------------------|---------------------------------------------------------------------------|
| <b>-r</b> <i>hdiskX</i> | Specifies the logical device to be cleared of any cluster signature.      |
| <b>-s</b> <i>hdiskX</i> | Specifies the logical device to be cleared of any storage pool signature. |
| <b>-help</b>            | Displays the help text for this command and exit program.                 |

### Exit Status

See “Exit status for Virtual I/O Server commands” on page 3.

### Examples

1. To clear the cluster signature from *hdisk5* and to use it as a paging space device by a shared memory partition, type the command as follows:  

```
cleandisk -r hdisk5
```
2. To clear the storage pool signature from *hdisk6* and to use it as a paging space device by a shared memory partition, type the command as follows:  

```
cleandisk -s hdisk6
```
3. To display the help text for the **cleandisk** command, type the command as follows:  

```
cleandisk -help
```

### Related Information

The **lspv** command and **prepdev** command.

---

## cluster command

### Purpose

Creates a cluster that uses the VIOS partitions and performs cluster-related operations.

### Syntax

To create a cluster:

**cluster -create -clustername** *ClusterName* **-repopvs** *PhysicalVolume* **-spname** *StoragePool* **-sppvs** *PhysicalVolume* **-hostname** *HostName*

To add a VIOS partition node to the cluster:

**cluster -addnode -clustername** *ClusterName* **-hostname** *HostName*

To remove a VIOS partition node from the cluster:

**cluster -rmnode [-f]-clustername** *ClusterName* **{-hostname** *HostName* **| -uuid** *HostUuid***}**

To remove a cluster from the system:

**cluster -delete -clustername** *ClusterName*

To display the health status of a cluster:

**cluster -status -clustername** *ClusterName* **[-field** *FieldName...* **]** **[-fmt** *Delimiter***]** **[-verbose ]**

To synchronize the cluster and Object Data Manager (ODM):

**cluster -sync -clustername** *ClusterName*

To list all clusters:

**cluster -list**

## Description

The **cluster** command is used to create and delete the cluster.

The **status** subcommand displays a global cluster health status and the status for each node in the cluster. The cluster status can be OK, DEGRADED, or DOWN. If the cluster status is OK, all the nodes of the cluster are working. If the cluster health status is DEGRADED, some nodes are working and some nodes are shut down. If the cluster health status is DOWN, all the nodes are shut down. The node status is either OK or DOWN. A node is in OK state if it is up and running, and a node is in DOWN state if it is not up. A pool status is provided for each node that is not DOWN. The pool status is OK for a node if the node can service the shared storage pool (SSP) virtualized storage to the client logical partitions (LPARs). The pool status is DOWN if it is not usable from the node.

If the *-verbose* option is requested, additional status details such as the number of nodes in each state is provided. For each node that is not DOWN, it displays the node roles and the storage pool status from the point of view of the node.

If the *-fmt* option is requested, status output is displayed in a format of delimiter-separated value report.

With the *-field* option, you can specify which fields to display. If no pool information is requested, the command does not retrieve pool status on each node, thereby accelerating the command response time.

## Flags

|                     |                                    |
|---------------------|------------------------------------|
| <b>-addnode</b>     | Adds a member node to the cluster. |
| <b>-clustername</b> | Specifies the name of the cluster. |
| <b>-create</b>      | Creates a cluster.                 |
| <b>-delete</b>      | Deletes a cluster.                 |

|                  |                                                                                                                                                                                                                                                                         |
|------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>-field</b>    | Specifies a list of fields to be displayed. Some of the fields are as follows:<br><br>cluster_name, cluster_state, cluster_id, node_mtm, node_name, node_number, node_roles, node_state, node_id, nodes_down, nodes_ok, number_of_nodes, pool_name, pool_state, pool_id |
| <b>-fmt</b>      | Divides output by a user-specified delimiter.                                                                                                                                                                                                                           |
| <b>-hostname</b> | Specifies the host name or IP address.                                                                                                                                                                                                                                  |
| <b>-list</b>     | Lists all clusters which have membership of the current node.                                                                                                                                                                                                           |
| <b>-repopvs</b>  | Lists physical volumes on which a repository can be created.                                                                                                                                                                                                            |
| <b>-rmnode</b>   | Removes the member node from the cluster.                                                                                                                                                                                                                               |
| <b>-spname</b>   | Specifies the name of the storage pool.                                                                                                                                                                                                                                 |
| <b>-sppvs</b>    | Lists the physical volumes on which a shared storage pool can be created.                                                                                                                                                                                               |
| <b>-status</b>   | Displays the health status of the cluster.                                                                                                                                                                                                                              |
| <b>-sync</b>     | Synchronizes the cluster and the ODM.                                                                                                                                                                                                                                   |
| <b>-verbose</b>  | Displays additional detailed information.                                                                                                                                                                                                                               |

## Examples

1. To create a cluster, type the command as follows:

```
cluster -create -clustername newcluster -repopvs repo_hdisk
-spname viosng -sppvs sp_hdisk -hostname myhost1
```

The system displays the output as follows:

Cluster newcluster has been created successfully.

2. To delete a cluster from the system, type the command as follows:

```
cluster -delete -clustername newcluster
```

The system displays the output as follows:

Cluster newcluster has been removed successfully.

3. To remove a member node from a cluster, type the command as follows:

```
cluster -rmnode -clustername newcluster -hostname myhost2
```

The system displays the output as follows:

myhost2 has been removed from the newcluster cluster.

4. To add a member node in a cluster, type the command as follows:

```
cluster -addnode -clustername newcluster -hostname myhost2
```

The system displays the output as follows:

myhost2 has been added to the newcluster cluster.

5. To display the health status summary, type the command as follows:

```
cluster -status -clustername mycluster1
```

The system displays the output as follows:

| Cluster Name | State |  |
|--------------|-------|--|
| mycluster1   | OK    |  |

  

| Node Name | Pool State | MTM               | Partition # | State |
|-----------|------------|-------------------|-------------|-------|
| node1     | OK         | 9115-505031006A2A | 1           | OK    |
| node2     | DOWN       | 9115-5050310069FA | 1           | OK    |

**Note:** The term *node* is the standard term that is used in the command instead of the term *partition*.

6. To display a detailed health status report, type the command as follows:

```
cluster -status -verbose -clustername mycluster2
```

The system displays the output as follows:

```
Cluster Name: mycluster2
Cluster Id: d23624fe335f11e0bd510011257e1447
Cluster State: DEGRADED
Number of Nodes: 2
Nodes OK: 1
Nodes DOWN: 1
```

```
Node Name: node1
Node Id: da18d52a32b211e0a5850011257e1447
Node MTM: 9115-505031006A2A
Node Number: 1
Node State: OK
Node Roles: PNN
Pool Name: poo11
Pool Id: b118d73c42b322d0a5850011257e1447
Pool State: OK
```

```
Node Name: node2
Node Id: 4d5012722f9f11e0813c0011257e1447
Node MTM: 9115-5050310069FA
Node Number: 1
Node State: DOWN
```

7. To display a formatted health status report, type the command as follows:

```
cluster -status -fmt ',' -clustername mycluster2
```

The system displays the output as follows:

```
mycluster2,OK,myhost1,8233-E8B02061AAFP,55,OK,OK
mycluster2,OK,myhost2,8233-E8B02061AAFP,11,OK,OK
```

8. To display a formatted, detailed health status report, type the command as follows:

```
cluster -status -fmt ',' -clustername mycluster2 -verbose
```

The system displays the output as follows:

```
mycluster2, fbc48dc88de711e0843c00c0dd040850,OK,2,2,0,myhost1,
00eee89c8d6211e0be6a00c0dd040850,8233-E8B02061AAFP,55,OK,DBN,test_pool,
e8b194ca65d22d18000000000000000,OK
mycluster2, fbc48dc88de711e0843c00c0dd040850,OK,2,2,0,myhost2,
2aed54c68de911e0825400c0dd040850, 8233-E8B02061AAFP,11,OK,DBN,test_pool,
e8b194ca65d22d18000000000000000,OK
```

9. To display a formatted health status report with some specific fields, type the command as follows:

```
cluster -status -fmt ':' -field cluster_name node_name pool_label pool_state -clustername
mycluster -verbose
```

The system displays the output as follows:

```
mycluster:myhost1:test_pool:OK
mycluster:myhost2:test_pool:OK
```

10. To display a specific health status report with specific fields, type the command as follows:

```
cluster -status -field node_name node_state cluster_name cluster_state -clustername mycluster2
```

The system displays the output as follows:

```
Node Name: myhost1
Node State: OK
Cluster Name: mycluster2
Cluster State: OK
Node Name: myhost2
Node State: OK
Cluster Name: mycluster2
Cluster State: OK
```

11. To display all clusters that have membership in the current node, type the command as follows:

```
cluster -list
```

The system displays the output as follows:

```
CLUSTER_NAME: mycluster
CLUSTER_ID: a64c8c725bfc11e1993500215e188da0
```

---

## cl\_snmp command

### Purpose

Issues Simple Network Management Protocol (SNMP) requests to agents and processes SNMP responses returned by agents.

### Syntax

```
cl_snmp [-com][-debug Level] [-host TargetHost] [-timeout TimeoutValue] [-retry RetryNumber] [-max MaxRepetitions] [-file ConfigurationFile] [-port PortNumber] [-verbose] [-non NonRepeaters] Function [MIBvariable][VariableType][Value][...]
```

### Description

Use the command to issue SNMP requests to agents and to process SNMP responses returned by agents. The Virtual I/O Server **cl\_snmp** command can be used for SNMPv1, SNMPv2c, and SNMPv3 requests.

### Flags

|                                       |                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|---------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>-com</b>                           | Specifies the community name used to access the specified variables at the destination SNMP agent. If you do not specify a community name, the default name is public. Community names are not required when using the user-based security model.<br><b>Note:</b> Community names are case sensitive.                                                                                                                                   |
| <b>-debug</b> <i>Level</i>            | Specifies a debugging level during run time. The default is 0. The higher the debugging level, the greater the number of messages that are displayed. The levels can be from 0 through 4.                                                                                                                                                                                                                                               |
| <b>-host</b> <i>TargetHost</i>        | Specifies the target host to which you want to send a request. This can be an Internet protocol address, a host name, or a winSNMP name in the clsnmp.conf configuration file. If you do not specify a host, the default is the local host.                                                                                                                                                                                             |
| <b>-timeout</b> <i>TimeoutValue</i>   | Specifies the amount of time (in seconds) that the <b>cl_snmp</b> command waits for a reply from the SNMP agent. The default is 3.                                                                                                                                                                                                                                                                                                      |
| <b>-retry</b> <i>RetryNumber</i>      | Specifies the maximum number of times to retry the command if it timed out. The default is 2.                                                                                                                                                                                                                                                                                                                                           |
| <b>-max</b> <i>MaxRepetitions</i>     | Specifies the number of lexicographic successors to be returned for each variable binding pair after the first <i>-non number</i> successors. The parameter applies only to the <b>getbulk</b> request. This is ignored if the function request is not a <b>getbulk</b> request. For example, starting with successor <i>-non number+1</i> , return <i>-max number</i> of successors for each variable binding pair. The default is 10. |
| <b>-file</b> <i>ConfigurationFile</i> | Specifies the full path and file name of the configuration file.                                                                                                                                                                                                                                                                                                                                                                        |
| <b>-port</b> <i>PortNumber</i>        | Specifies the number of the port that listens for traps. If a port number is not specified, the <b>cl_snmp</b> trap function listens on the well-known port 162, which is the default port for <b>cl_snmp</b> traps.                                                                                                                                                                                                                    |
| <b>-verbose</b>                       | Specifies that the output from a request should be displayed using verbose output, for example, using the textual name instead of the MIB object identifier.                                                                                                                                                                                                                                                                            |
| <b>-non</b> <i>NonRepeaters</i>       | Specifies the number of variable binding pairs (name/value), starting with the first, for which only a single successor is returned. This parameter applies only to <b>getbulk</b> requests. This is ignored if the function request is not a <b>getbulk</b> request. The default is 0.                                                                                                                                                 |

**Function** [*MIBvariable*]  
[*VariableType*] [*Value*] [...]]

Specifies the SNMP function or operation to perform, which is one of the following:

- get
- getnext
- getbulk
- set
- walk
- trap
- findname

**MIBVariable**

Specifies the Management Information Base (MIB) object, using its object descriptor (textual name), object identifier in ASN.1 notation, or a combination of the two. When used with the walk function, this is the MIB object prefix. A prefix can be any leading portion of the complete object identifier. When used with the **findname** command, this object identifier is in the ASN.1 notation.

**Value** Specifies the value to be set by the SET function. If white space is needed in the value, enclose the value in quotation marks. To set a variable to a value that is also a type, you must specify the type.

**VariableType**

Specifies the type of value being set. To complete an SNMP SET request, the SMI\_type must be known. If no type is specified, **cl\_snmp** command searches first the /etc/mib.defs file and then the compiled MIB to determine the type. If the variable is not found, an error is returned. If a *VariableType* is specified, the *VariableType* takes precedence over any type that may be assigned in the MIB. The *VariableType* and value must be compatible. For example, if you specify a type of *number* and a value of *foo*, an error is returned because *foo* is not a number. *VariableType* is not case sensitive. Valid variable types are:

- bitstring
- counter
- counter32
- counter64
- display or displaystring
- gauge
- gauge32
- integer
- integer32
- ipaddress
- nsapaddress
- null
- objectidentifier or OID
- octetstring
- opaque
- opaqueascii
- timeticks
- uinteger

## Request Types

|          |                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|----------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| findname | Sends a request that a search be done to obtain the textual name, for a given <i>MIBVariable</i> input, whose internal ASN.1 value best matches the input ASN.1 value. The search first checks the <code>/etc/mib.defs</code> file, and if a matching textual name is not found, it continues with the compiled management information base (MIB). Only one <i>MIBVariable</i> is allowed per <code>cl_snmp findname</code> invocation. |
| get      | Sends a request to an SNMP agent for a specific MIB variable. The <code>cl_snmp</code> command then waits for a response or times out.                                                                                                                                                                                                                                                                                                  |
| getbulk  | Obtains the value of the variables in the MIB tree specified by the object identifier (OID) or MIB variable name. A single <code>getbulk</code> performs the same function as a series of <code>getnext</code> commands, with fewer data exchanges between the <code>cl_snmp</code> command and the SNMP agent.                                                                                                                         |
| getnext  | Sends a request to an SNMP agent for the next MIB variable that lexicographically follows the <i>MIBVariable</i> specified. The <code>cl_snmp</code> command then waits for a response or times out.                                                                                                                                                                                                                                    |
| set      | Sends a request to an SNMP agent to set a specific <i>MIBVariable</i> . The <code>cl_snmp</code> command then waits for a response or times out.                                                                                                                                                                                                                                                                                        |
| trap     | Listens for SNMP traps and displays trap information when they occur. Uses the default, well-known port 162 or the port number specified on the <code>-port</code> option. The <code>cl_snmp trap</code> function continues to listen for traps until the process is killed or canceled.                                                                                                                                                |
| walk     | Issues a <code>getnext</code> request for a specified prefix, then continues to issue <code>getnext</code> requests for as long as there are variables that match the specified prefix. A prefix can be any leading portion of the complete object identifier.                                                                                                                                                                          |

## Exit Status

See “Exit status for Virtual I/O Server commands” on page 3.

## Examples

To issue an SNMP request, run the following command:

```
cl_snmp -host hostname get sysName.0
```

## Related Information

The `snmp_info` command and `snmp_trap` command.

---

## cplv command

### Purpose

Copies the contents of a logical volume to a new logical volume.

### Syntax

To Copy to a New Logical Volume

```
cplv [-vg VolumeGroup] [-lv NewLogicalVolume | -prefix Prefix] SourceLogicalVolume
```

To Copy to an Existing Logical Volume

```
cplv [-f] SourceLogicalVolume DestinationLogicalVolume
```



## Description

**Attention:** Do not copy from a larger logical volume containing data to a smaller one. Doing so results in a corrupted file system because some data is not copied.

The **cplv** command copies the contents of *SourceLogicalVolume* to a new or existing logical volume. The *SourceLogicalVolume* parameter can be a logical volume name or a logical volume ID. The **cplv** command creates a new logical volume with a system-generated name by using the default syntax. The system-generated name is displayed.

**Note:** The **cplv** command cannot copy logical volumes which are in the open state, including logical volumes that are being used as backing devices for virtual storage.

## Flags

|                                    |                                                                                                                                                                                                           |
|------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>-f</b>                          | Copies to an existing logical volume without requesting user confirmation.                                                                                                                                |
| <b>-lv</b> <i>NewLogicalVolume</i> | Specifies the name to use, in place of a system-generated name, for the new logical volume. Logical volume names must be unique systemwide names, and can range from 1 to 15 characters.                  |
| <b>-prefix</b> <i>Prefix</i>       | Specifies a prefix to use in building a system-generated name for the new logical volume. The prefix must be less than or equal to 13 characters. A name cannot be a name already used by another device. |
| <b>-vg</b> <i>VolumeGroup</i>      | Specifies the volume group where the new logical volume resides. If this is not specified, the new logical volume resides in the same volume group as the <i>SourceLogicalVolume</i> .                    |

## Examples

1. To copy the contents of logical volume fslv03 to a new logical volume, type:

```
cplv fslv03
```

The new logical volume is created, placed in the same volume group as fslv03, and named by the system.

2. To copy the contents of logical volume fslv03 to a new logical volume in volume group vg02, type:

```
cplv -vg vg02 fslv03
```

The new logical volume is created, named, and added to volume group vg02.

3. To copy the contents of logical volume lv02 to a smaller, existing logical volume, lvtest, without requiring user confirmation, type:

```
cplv -f lv02 lvtest
```

---

## cpvdi command

### Purpose

Copies a block device.

### Syntax

```
cpvdi -src input_disk_image -dst output_disk_image [-isp input_storage_pool] [-osp output_storage_pool]
[-progress] [-overwrite] [-unconfigure] [-f]
```

### Description

The **cpvdi** command copies a block device image, which can be either a logical or physical volume, a file-backed device, or a file on another existing disk. The parameters of the **-src** and **-dst** flags should be

fully qualified names when specifying a file or device. Specifying the **-progress** flag prints a number sign (#) on the **stderr** for each block (1 M) of data copied on to the destination.

**Note:** The *output\_disk\_image* device that you specify with the **-dst** flag cannot be a physical volume that is assigned to a shared memory pool (to be used as a paging space device by a shared memory partition).

## Flags

|                                        |                                                                                                                                                                                                                                                                                                           |
|----------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>-src</b> <i>input_disk_image</i>    | Specifies the input backing device. The input backing device might either be a physical or logical volume or a file is located in the caller's home directory, any of its subdirectories or in the /mnt directory. The root user will not have any file location restrictions.                            |
| <b>-dst</b> <i>output_disk_image</i>   | Specifies the output backing device. The output backing device might either be a physical or logical volume or a file is located in the caller's home directory, any of its subdirectories or in the /mnt directory. The root user will not have any file location restrictions.                          |
| <b>-f</b>                              | Overwrites the destination file. However, if the source or destination file is configured as backing device, it copies the file without unconfiguring the device or devices. If the <b>-unconfigure</b> flag is used in addition to the <b>-f</b> flag, then the <b>-unconfigure</b> flag takes priority. |
| <b>-progress</b>                       | Specifies the process status of the command will be expressed by printing a number sign (#) to <b>stderr</b> for every 1024 KB block of data that has been copied.                                                                                                                                        |
| <b>-isp</b> <i>input_storage_pool</i>  | Specifies the name of the storage pool where the input file-backed device exists. Required if the argument of the <b>-src</b> flag is a file-backed device.                                                                                                                                               |
| <b>-osp</b> <i>output_storage_pool</i> | Specifies the name of the storage pool where the input file-backed device exists. Required if the argument of the <b>-dst</b> flag is a file-backed device.                                                                                                                                               |
| <b>-overwrite</b>                      | Overwrites the destination file, if the file is present. If you do not use this flag, an error message displays with a non-zero exit code.                                                                                                                                                                |
| <b>-unconfigure</b>                    | Specifies to unconfigure the virtual target device (VTD) and reconfigure the VTD after a copy operation. Only use this flag if the source or destination file is configured as backing device. If you do not use this flag, an error message displays with a non-zero exit code.                          |

## Exit Status

See “Exit status for Virtual I/O Server commands” on page 3.

## Examples

1. To copy a physical volume disk image to a file, type the following command:  
`cpvdi -src hdisk1 -dst /home/my_home/image1_file`
2. To copy a disk image that is a file to a physical volume, type the following command:  
`cpvdi -src /home/my_home/image2_file -dst hdisk2`
3. To output a hash symbol for every block of data that is copied, type the following command:  
`cpvdi -src hdisk1 -dst hdisk2 -progress`
4. To copy a file-backed device to another file-backed device, type the following command:  
`cpvdi -src image1_file -dst image2_file -isp sp01 -osp sp02`
5. To copy a file1 to file2, when file2 already exists, type the following command:  
`cpvdi -src /home/padmin/file1 -dst /home/padmin/file2 -overwrite`
6. To copy a file-backed device fbd1 as fbd2, when fbd1 is configured as the backing device, type the following command:  
`cpvdi -src fbd1 -dst fbd2 -isp sp01 -osp sp01 -f`

## Related Information

The `backupios` command.

---

## deactivatevg command

### Purpose

Deactivates a volume group.

### Syntax

```
deactivatevg VolumeGroup
```

### Description

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

### Exit Status

See “Exit status for Virtual I/O Server commands” on page 3.

### Examples

1. To deactivate volume group `vg03`, type:  

```
deactivatevg vg03
```

## Related Information

The `mkvg` command, the `chvg` command, the `extendvg` command, the `reducevg` command, the `mirrorios` command, the `unmirrorios` command, the `lsvg` command, the `activatevg` command, the `importvg` command, the `exportvg` command, and the `syncvg` command.

---

## diagmenu command

### Purpose

Places the user into the diagnostic menus.

### Syntax

```
diagmenu
```

### Description

The `diagmenu` command performs hardware problem determination and maintenance. When the user suspects there is a problem, `diagmenu` assists in finding it. Through the diagnostic menus, users can perform tasks such as running diagnostics, hot plugging devices, formatting and certifying media, and managing RAID devices.

Once inside the diagnostic menus, task specific help can be obtained by pressing the **F1** key.

## Related Information

The `errlog` command.

---

## **dsmc command**

### **Purpose**

Backups and stores data generated while using the Tivoli Storage Manager Server.

### **Syntax**

```
dsmc [-incremental | -schedule]
```

```
dsmc [-query | -restore] argv
```

### **Description**

The `dsmc` command allows the user to back up and restore data from the Tivoli Storage Manager Server.

### **Flag**

|                           |                                                                                                                                                                                      |
|---------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <code>-incremental</code> | Backs up all new or changed files or directories in the default client domain or from file systems, directories, or files you specify, unless you exclude them from backup services. |
|---------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

**-query [argument]**

Query functions to query backups on the Tivoli Storage Manager Server

**access** Displays a list of current authorization rules.

**archive** Displays a list of archived files

**backup** Displays a list of back up versions.

**backupset**

Queries a backup set from a local file, tape device, or the Tivoli Storage Manager Server.

**filespace**

Displays a list of file spaces in Tivoli Storage Manager storage. You can also specify a single file space name to query.

**group** Displays information about group backups and their members.

**image** Displays information about image backups.

**inclexcl**

Displays a list of include-exclude statements in the order in which they are processed during backup and archive operations.

**mgmtclass**

Displays information about available management classes.

**node** Displays all the nodes for which an administrative user ID has authority to perform operations.

**options**

Displays all or part of your options and their current settings.

**restore** Displays a list of your restartable restore sessions in the server database.

**schedule**

Displays information about scheduled events for your node.

**session** Displays information about your session, including the current node name, when the session was shed, server information, and server connection information.

**systeminfo**

Gathers Tivoli Storage Manager system information and outputs this information to a file or the console.

**was** Displays backups of the WebSphere® Application Server (WAS) Network Deployment Manager (contains setup, application files, and configuration information) or the Application Server that match the node name and type of the WAS group backup that you specify.

**-restore [argument]**

Restores copies of backup versions of your files from an Tivoli Storage Manager server.

**backupset**

Restores a backup set from the Tivoli Storage Manager server or a local file. You can also restore a backup from a tape device.

**group** Restores specific members or all members of a group backup.

**image** Restores a file system or raw volume image backup.

**nas** Restores the image of a file system belonging to a Network Attached Storage (NAS) file server.

**was** Restores the WebSphere Application Server (WAS) Network Deployment Manager (contains setup, application files, and configuration information) or the Application Server from the Tivoli Storage Manager server.

**-schedule**

Starts the client scheduler on the workstation.

## Tivoli Storage Manager

The Tivoli Storage Manager can help protect computers running a variety of different operating environments, including the Virtual I/O Server, on a variety of different hardware. Configuring the Tivoli Storage Manager client on the Virtual I/O Server enables you to include the Virtual I/O Server in your standard backup framework.

| Attributes  | Description                                                                   | Value                           |
|-------------|-------------------------------------------------------------------------------|---------------------------------|
| DSMC_CONFIG | Specifies the location of the dsm.sys configuration file.                     | /home/padmin/tivoli/tsm/dsm.sys |
| DSMC_DIR    | Specifies the location of the Tivoli Storage Manager configuration directory. | /home/padmin/tivoli/tsm/        |

## Exit Status

The following exit values are returned:

|    |                        |
|----|------------------------|
| 0  | Successful completion. |
| >0 | An error occurred.     |

## Examples

To backup an image to the Tivoli Storage Manager server, type:

```
dsmc -incremental <mksysb_file_to_backup>
```

To view backup images on the Tivoli Storage Manager server, type:

```
dsmc -query backup <mksysb_file_to_backup>
```

---

## entstat command

### Purpose

Shows Ethernet device driver and device statistics.

### Syntax

```
entstat [-all] [-reset] Device_Name
```

### Description

The **entstat** command displays the statistics gathered by the specified Ethernet device driver. The user can optionally specify that the device-specific statistics be displayed in addition to the device generic statistics. If no flags are specified, only the device generic statistics are displayed.

### Flags

|        |                                                                        |
|--------|------------------------------------------------------------------------|
| -all   | Displays all the statistics, including the device-specific statistics. |
| -reset | Resets all the statistics back to their initial values.                |

### Parameters

*Device\_name*                      The name of the Ethernet device, for example, ent0.

## Exit Status

The statistic fields displayed in the output of the **entstat** command and their descriptions are as follows.

**Note:** Some adapters may not support a specific statistic. The value of non-supported statistic fields is always 0.

## Title Fields

|                         |                                                                                                                                                                                                                                                                                                                                                                                         |
|-------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Device Type</b>      | Displays the description of the adapter type.                                                                                                                                                                                                                                                                                                                                           |
| <b>Hardware Address</b> | Displays the Ethernet network address currently used by the device.                                                                                                                                                                                                                                                                                                                     |
| <b>Elapsed Time</b>     | Displays the real time period which has elapsed since last time the statistics were reset. Part of the statistics may be reset by the device driver during error recovery when a hardware error is detected. There will be another Elapsed Time displayed in the middle of the output when this situation has occurred in order to reflect the time differences between the statistics. |

## Transmit Statistics Fields

|                                              |                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|----------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Packets</b>                               | The number of packets transmitted successfully by the device.                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Bytes</b>                                 | The number of bytes transmitted successfully by the device.                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Interrupts</b>                            | The number of transmit interrupts received by the driver from the adapter.                                                                                                                                                                                                                                                                                                                                                        |
| <b>Transmit Errors</b>                       | The number of output errors encountered on this device. This is a counter for unsuccessful transmissions due to hardware or network errors.                                                                                                                                                                                                                                                                                       |
| <b>Packets Dropped</b>                       | The number of packets accepted by the device driver for transmission which were not (for any reason) given to the device.                                                                                                                                                                                                                                                                                                         |
| <b>Max Packets on S/W Transmit Queue</b>     | The maximum number of outgoing packets ever queued to the software transmit queue.                                                                                                                                                                                                                                                                                                                                                |
| <b>S/W Transmit Queue Overflow</b>           | The number of outgoing packets which have overflowed the software transmit queue.                                                                                                                                                                                                                                                                                                                                                 |
| <b>Current S/W+H/W Transmit Queue Length</b> | The number of pending outgoing packets on either the software transmit queue or the hardware transmit queue.                                                                                                                                                                                                                                                                                                                      |
| <b>Broadcast Packets</b>                     | The number of broadcast packets transmitted without any error.                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Multicast Packets</b>                     | The number of multicast packets transmitted without any error.                                                                                                                                                                                                                                                                                                                                                                    |
| <b>No Carrier Sense</b>                      | The number of unsuccessful transmissions due to the no carrier sense error.                                                                                                                                                                                                                                                                                                                                                       |
| <b>DMA Underrun</b>                          | The number of unsuccessful transmissions due to the DMA underrun error.                                                                                                                                                                                                                                                                                                                                                           |
| <b>Lost CTS Errors</b>                       | The number of unsuccessful transmissions due to the loss of the Clear-to-Send signal error.                                                                                                                                                                                                                                                                                                                                       |
| <b>Max Collision Errors</b>                  | The number of unsuccessful transmissions due to too many collisions. The number of collisions encountered exceeded the number of retries on the adapter.                                                                                                                                                                                                                                                                          |
| <b>Late Collision Errors</b>                 | The number of unsuccessful transmissions due to the late collision error.                                                                                                                                                                                                                                                                                                                                                         |
| <b>Deferred</b>                              | The number of outgoing packets deferred during transmission. Deferred means that the adapter had to defer while trying to transmit a frame. This condition occurs if the network is busy when the adapter is ready to transmit. The adapter will only defer the first attempt to send a packet. After that the adapter will transmit the packet without checking. If the network is still busy then a collision will be recorded. |
| <b>SQE Test</b>                              | Contains the number of "Signal Quality Error" Tests (i.e. Heartbeat) performed successfully during transmission.                                                                                                                                                                                                                                                                                                                  |
| <b>Timeout Errors</b>                        | The number of unsuccessful transmissions due to adapter reported timeout errors.                                                                                                                                                                                                                                                                                                                                                  |
| <b>Single Collision Count</b>                | The number of outgoing packets with single (only one) collision encountered during transmission.                                                                                                                                                                                                                                                                                                                                  |

|                                         |                                                                                                                                                                                                                                                                                                                                          |
|-----------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Multiple Collision Count</b>         | The number of outgoing packets with multiple (2 - 15) collisions encountered during transmission                                                                                                                                                                                                                                         |
| <b>Current HW Transmit Queue Length</b> | The number of outgoing packets which currently exist on the hardware transmit queue.                                                                                                                                                                                                                                                     |
| <b>CRC Errors</b>                       | The number of incoming packets with the Checksum (FCS) error.                                                                                                                                                                                                                                                                            |
| <b>DMA Overrun</b>                      | The number of incoming packets with the DMA overrun error.                                                                                                                                                                                                                                                                               |
| <b>Alignment Errors</b>                 | The number of incoming packets with the alignment error.                                                                                                                                                                                                                                                                                 |
| <b>No Resource Errors</b>               | The number of incoming packets dropped by the hardware due to the no resource error. This error usually occurs because the receive buffers on the adapter were exhausted. Some adapters may have the size of the receive buffers as a configurable parameter. Check the device configuration attributes for possible tuning information. |
| <b>Receive Collision Errors</b>         | The number of incoming packets with the collision errors during the reception.                                                                                                                                                                                                                                                           |
| <b>Packet Too Short Errors</b>          | The number of incoming packets with the length error indicating that the packet size is less than the Ethernet minimum packet size.                                                                                                                                                                                                      |
| <b>Packet Too Long Errors</b>           | The number of incoming packets with the length error indicating that the packet size is bigger than the Ethernet maximum packet size.                                                                                                                                                                                                    |
| <b>Packets Discarded by Adapter</b>     | The number of incoming packets dropped by the hardware for any other reasons.                                                                                                                                                                                                                                                            |
| <b>Receiver Start Count</b>             | The number of times that the receiver (receive unit) on the adapter has been started.                                                                                                                                                                                                                                                    |

## Examples

1. To display the device generic statistics for **ent0**, type:

```
entstat ent0
```

This produces output similar to the following:

```
ETHERNET STATISTICS (ent0) :
Device Type: Ethernet High Performance LAN Adapter
Hardware Address: 02:60:8c:2e:d0:1d
Elapsed Time: 0 days 0 hours 8 minutes 41 seconds

Transmit Statistics: Receive Statistics:

Packets: 3 Packets: 2
Bytes: 272 Bytes: 146
Interrupts: 3 Interrupts: 2
Transmit Errors: 0 Receive Errors: 0
Packets Dropped: 0 Packets Dropped: 0
Max Packets on S/W Bad Packets: 0
Transmit Queue:0
S/W Transmit Queue
Overflow: 0
Current S/W+H/W Transmit
Queue Length: 0

Broadcast Packets: 2 CRC Errors: 0
Multicast Packets: 0 Broadcast Packets: 1
No Carrier Sense: 0 Multicast Packets: 0
DMA Underrun: 0 DMA Overrun: 0
Lost CTS Errors: 0 Alignment Errors: 0
Max Collision Errors: 0 No Resource Errors: 0
Late Collision Errors: 0 Receive Collision Errors: 0
Deferred: 0 Packet Too Short Errors: 0
SQE Test: 0 Packet Too Long Errors: 0
Timeout Errors: 0 Packets Discarded by Adapter: 0
Single Collision Receiver Start Count: 1
Count: 0
Multiple Collision Count: 0
```



Current HW Transmit Queue  
Length: 0

General Statistics:

-----  
No mbuf Errors: 0  
Adapter Reset Count: 0  
Driver Flags: Up Broadcast Running Simplex

2. To display the Ethernet device generic statistics and the Ethernet device-specific statistics for **ent0**, type:

entstat -all

Results similar to the following will be displayed:

ETHERNET STATISTICS (ent0) :  
Device Type: Ethernet High Performance LAN Adapter  
Hardware Address: 02:60:8c:2e:d0:1d  
Elapsed Time: 0 days 2 hours 6 minutes 30 seconds

| Transmit Statistics:                     | Receive Statistics:     |
|------------------------------------------|-------------------------|
| -----                                    | -----                   |
| Packets: 3                               | Packets: 2              |
| Bytes: 272                               | Bytes: 146              |
| Interrupts: 3                            | Interrupts: 2           |
| Transmit Errors: 0                       | Receive Errors: 0       |
| Packets Dropped: 0                       | Packets Dropped: 0      |
| Max Packets on S/W                       | Receiver Start Count: 1 |
| Transmit Queue:0                         |                         |
| Bad Packets: 0                           |                         |
| S/W Transmit Queue Overflow: 0           |                         |
| Current S/W+H/W Transmit Queue Length: 0 |                         |

|                                     |                                 |
|-------------------------------------|---------------------------------|
| Broadcast Packets: 0                | Broadcast Packets: 0            |
| Multicast Packets: 0                | Multicast Packets: 0            |
| No Carrier Sense: 0                 | CRC Errors: 0                   |
| DMA Underrun: 0                     | DMA Overrun: 0                  |
| Lost CTS Errors: 0                  | Alignment Errors: 0             |
| Max Collision Errors: 0             | No Resource Errors: 0           |
| Late Collision Errors: 0            | Receive Collision Errors: 0     |
| Deferred: 0                         | Packet Too Short Errors: 0      |
| SQE Test: 0                         | Packet Too Long Errors: 0       |
| Timeout Errors: 0                   | Packets Discarded by Adapter: 0 |
| Single Collision Count: 0           | Receiver Start Count: 1         |
| Multiple Collision Count: 0         |                                 |
| Current HW Transmit Queue Length: 0 |                                 |

General Statistics:

-----  
No mbuf Errors: 0  
Adapter Reset Count: 0  
Driver Flags: Up Broadcast Running Simplex

Ethernet High Performance LAN Adapter Specific Statistics:

-----  
Receive Buffer Pool Size: 37  
Transmit Buffer Pool Size: 39  
In Promiscuous Mode for IP Multicast: No  
Packets Uploaded from Adapter: 0  
Host End-of-List Encountered: 0  
82586 End-of-List Encountered: 0  
Receive DMA Timeouts: 0  
Adapter Internal Data: 0x0 0x0 0x0 0x0 0x0

## Related Information

The `optimizenet` command.

---

## errlog command

### Purpose

Displays or clears the error log.

### Syntax

```
errlog [-ls] [-seq SequenceNumber] | -rm Days]
```

### Description

The `errlog` command generates an error report from entries in the Virtual I/O Server error log or deletes all entries from the error log older the number of days specified by the *Days* parameter. To delete all error-log entries, specify a value of 0 for the *Days* parameter. If no flags are specified a summary report is displayed.

### Flags

|                                         |                                                                                                 |
|-----------------------------------------|-------------------------------------------------------------------------------------------------|
| <code>-ls</code>                        | Displays detailed information about errors in the error log file.                               |
| <code>-rm <i>Days</i></code>            | Deletes error-log entries older than the number of days specified by the <i>Days</i> parameter. |
| <code>-seq <i>SequenceNumber</i></code> | Displays information about a specific error in the error log file by the sequence number.       |

### Exit Status

See “Exit status for Virtual I/O Server commands” on page 3.

### Examples

1. To display a complete detailed report, enter:  
`errlog -ls`
2. To delete error-log entries older than 5 days, enter:  
`errlog -rm 5`
3. To delete all error-log entries, enter:  
`errlog -rm 0`

## Related Information

The `diagmenu` command.

---

## exportvg command

### Purpose

Exports the definition of a volume group from a set of physical volumes.

### Syntax

```
exportvg VolumeGroup
```

## Description

The **exportvg** command removes the definition of the volume group specified by the *VolumeGroup* parameter from the system. The primary use of the **exportvg** command, coupled with the **importvg** command, is to allow portable volumes to be exchanged between systems. Only a complete volume group can be exported, not individual physical volumes.

Using the **exportvg** command and the **importvg** command, you can also switch ownership of data on physical volumes shared between two processors.

Mount points longer than 128 characters will not automatically be re-mounted when the volume group is imported using the **importvg** command and should not be used.

**Note:** A volume group that has a paging space volume on it cannot be exported.

## Exit Status

See “Exit status for Virtual I/O Server commands” on page 3.

## Examples

1. To remove volume group **vg02** from the system, type:  
`exportvg vg02`

**Note:** The volume group must be deactivated before exporting.  
The definition of **vg02** is removed from the system and the volume group cannot be accessed.

## Restrictions

Mount points longer than 128 characters should not be used.

A volume group that has a paging space volume on it cannot be exported.

## Related Information

The **mkvg** command, the **chvg** command, the **lsvg** command, the **reducevg** command, the **mirrorios** command, the **unmirrorios** command, the **activatevg** command, the **deactivatevg** command, the **importvg** command, the **extendvg** command, and the **syncvg** command.

---

## extendlv command

### Purpose

Increases the size of a logical volume.

### Syntax

```
extendlv LogicalVolume Size [PhysicalVolume ...]
```

### Description

The **extendlv** command increases the size of the logical volume. The *LogicalVolume* parameter can be a logical volume name or a logical volume ID. To limit the allocation to specific physical volumes, use the names of one or more physical volumes in the *PhysicalVolume* parameter; otherwise, all the physical volumes in a volume group are available for allocating new physical partitions.

The *Size* parameter specifies the minimum size the logical volume should be increased by. When specifying *Size* the following conventions must be used:

| Size   | Logical volume size |
|--------|---------------------|
| ###M/m | ### MB              |
| ###G/g | ### GB              |

## Exit Status

See “Exit status for Virtual I/O Server commands” on page 3

## Examples

1. To increase the size of the logical volume **lv05** by three megabytes, type:  

```
extendlv lv05 3M
```
2. To request a logical volume named **lv05** with a minimum size of 10MB, type:  

```
extendlv lv05 10M
```

The **extendlv** command will determine the number of partitions needed to create a logical volume of at least that size.

## Related Information

The **mkiv** command, the **lsiv** command, and the **rmlv** command.

---

## extendvg command

### Purpose

Adds physical volumes to a volume group.

### Syntax

```
extendvg [-f] VolumeGroup PhysicalVolume ...
```

### Description

The **extendvg** command increases the size of volume group by adding one or more physical volumes.

The physical volume is checked to verify that it is not already in another volume group. If the system determines the physical volume belongs to a volume group that is activated, it exits. But if the system detects a description area from a volume group that is not activated, it prompts you for confirmation while continuing with the command. The previous contents of the physical volume are lost, so you must be cautious when using the override function.

### Flags

- f** Forces the physical volume to be added to the specified volume group unless the following conditions are true:
- The physical volume is a member of another volume group in the device configuration database.
  - The physical volume is a member of the another volume group that is active.
  - The physical volume is assigned to a shared memory pool (to be used as a paging space device by a shared memory partition).

## Exit Status

See “Exit status for Virtual I/O Server commands” on page 3.

## Examples

To add physical volumes **hdisk3** and **hdisk8** to volume group **vg3**, type the following command:

```
extendvg vg3 hdisk3 hdisk8
```

## Related Information

The **lsvg** command, the **chvg** command, the **mkvg** command, the **reducevg** command, the **mirrorios** command, the **unmirrorios** command, the **activatevg** command, the **deactivatevg** command, the **importvg** command, the **exportvg** command, and the **syncvg** command.

---

## fcstat command

### Purpose

Displays the statistics gathered by the specified fibre channel device driver.

### Syntax

To display statistics gathered by the specified fibre channel device driver:

```
fcstat -Device_Name
```

### Description

The **fcstat** command displays the statistics gathered by the specified fibre channel device driver. It collects the statistics using the following process:

1. Opens the message catalog of **fcstat** and checks the parameter list.
2. Accesses the ODM database for information that relates to the selected adapter.
3. Accesses the ODM database for information that relates to ports of the selected adapter.
4. Opens and accesses adapter statistics.
5. Reports statistics and exits program.

### Flags

|                    |                                                          |
|--------------------|----------------------------------------------------------|
| <i>Device_Name</i> | The name of the fibre channel device, for example, fcs0. |
|--------------------|----------------------------------------------------------|

## Exit Status

The following exit values are returned:

|    |                        |
|----|------------------------|
| 0  | Successful completion. |
| >0 | An error occurred.     |

## Examples

To display the statistics for fibre channel device driver *fcs0*, enter:

```
fcstat fcs0
```



---

## fsck command

### Purpose

Checks file system consistency and interactively repairs the file system.

### Syntax

```
fsck [FileSystem ...]
```

### Description

The **fsck** command checks and interactively repairs inconsistent file systems. Normally, the file system is consistent, and the **fsck** command merely reports on the number of files, used blocks, and free blocks in the file system. If the file system is inconsistent, the **fsck** command displays information about the inconsistencies found and prompts you for permission to repair them. If no *FileSystem* is specified, all file systems are checked.

The **fsck** command checks for the following inconsistencies:

- Blocks or fragments allocated to multiple files.
- i-nodes containing block or fragment numbers that overlap.
- i-nodes containing block or fragment numbers out of range.
- Discrepancies between the number of directory references to a file and the link count of the file.
- Illegally allocated blocks or fragments.
- i-nodes containing block or fragment numbers that are marked free in the disk map.
- i-nodes containing corrupt block or fragment numbers.
- A fragment that is not the last disk address in an i-node. This check does not apply to compressed file systems.
- Files larger than 32KB containing a fragment. This check does not apply to compressed file systems.
- Size checks:
  - Incorrect number of blocks.
  - Directory size not a multiple of 512 bytes.

**Note:** These checks do not apply to compressed file systems.

- Directory checks:
  - Directory entry containing an i-node number marked free in the i-node map.
  - i-node number out of range.
  - Dot (.) link missing or not pointing to itself.
  - Dot dot (..) link missing or not pointing to the parent directory.
  - Files that are not referenced or directories that are not reachable.
- Inconsistent disk map.
- Inconsistent i-node map.

In addition to its messages, the **fsck** command records the outcome of its checks and repairs through its exit value. This exit value can be any sum of the following conditions:

|   |                                                                                                |
|---|------------------------------------------------------------------------------------------------|
| 0 | All checked file systems are now okay.                                                         |
| 2 | The <b>fsck</b> command was interrupted before it could complete checks or repairs.            |
| 4 | The <b>fsck</b> command changed the file system; the user must restart the system immediately. |

## Exit Status

See “Exit status for Virtual I/O Server commands” on page 3.

## Examples

1. To check a file system, enter:

```
fsck /dev/hd1
```

This command checks the unmounted file system located on the `/dev/hd1` device.

---

## genvfilt command

### Purpose

Adds a filter rule for the virtual LAN (VLAN) crossing between logical partitions on the same systems server.

### Syntax

```
genvfilt -v <4|6> -a <D|P> -z <svlan> -Z <dvlan> [-s <s_addr>] [-d <d_addr>] [-o <src_port_op>] [-p <src_port>] [-O <dst_port_op>] [-P <dst_port>] [-c <protocol>]
```

### Description

The **genvfilt** command adds a filter rule for the virtual LAN (VLAN) crossing between logical partitions (LPARs) on the same systems server.

### Flags

|           |                                                                                                                                                                                                                    |
|-----------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>-a</b> | Specifies the action. Valid values follow: <ul style="list-style-type: none"> <li>• D (Deny): Blocks traffic</li> <li>• P (Permit): Allows traffic</li> </ul>                                                      |
| <b>-c</b> | Specifies different protocols to which the filter rule is applicable. Valid values follow: <ul style="list-style-type: none"> <li>• udp</li> <li>• icmp</li> <li>• icmpv6</li> <li>• tcp</li> <li>• any</li> </ul> |
| <b>-d</b> | Specifies the destination address in v4 or v6 format.                                                                                                                                                              |
| <b>-m</b> | Specifies the source address mask.                                                                                                                                                                                 |
| <b>-M</b> | Specifies the destination address mask.                                                                                                                                                                            |
| <b>-o</b> | Specifies the source port or the Internet Control Message Protocol (ICMP) type operation. Valid values follow: <ul style="list-style-type: none"> <li>• lt</li> <li>• gt</li> <li>• eq</li> <li>• any</li> </ul>   |



|           |                                                                                                                                                                                   |
|-----------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>-O</b> | Specifies the destination port or the ICMP code operation. Valid values follow: <ul style="list-style-type: none"> <li>• lt</li> <li>• gt</li> <li>• eq</li> <li>• any</li> </ul> |
| <b>-p</b> | Specifies the source port or the ICMP type.                                                                                                                                       |
| <b>-P</b> | Specifies the destination port or the ICMP code.                                                                                                                                  |
| <b>-s</b> | Specifies the source address in IPv4 or IPv6 format.                                                                                                                              |
| <b>-v</b> | Specifies the IP version of the filter rule table. Valid values are 4 and 6.                                                                                                      |
| <b>-z</b> | Specifies the virtual LAN ID of the source LPAR.                                                                                                                                  |
| <b>-Z</b> | Specifies the virtual LAN ID of the destination LPAR.                                                                                                                             |

## Exit Status

This command returns the following exit values:

|    |                        |
|----|------------------------|
| 0  | Successful completion. |
| >0 | An error occurred.     |

## Examples

1. To add a filter rule to permit TCP data from a source VLAN ID of 100 to a destination VLAN ID of 200 on specific ports, type the command as follows:

```
genvfilt -v4 -a P -z 100 -Z 200 -o lt -p 345 -O lt -P 345 -c tcp
```

## Related Information

The **chvfilt** command, the **lsvfilt** command, the **mkvfilt** command, the **rmvfilt** command, and the **vlanfw** command.

---

## hostmap command

### Purpose

Directly manipulates address-mapping entries in the system configuration database.

### Syntax

To Add an Address-to-Host Name Mapping

```
hostmap -addr IPAddress -host HostName...
```

To Delete an Address-to-Host Name Mapping

```
hostmap -rm IPAddress
```

To Show all Address-to-Host Name Mappings

```
hostmap -ls
```

## Description

The **hostmap** low-level command adds, deletes, or lists address-mapping entries in the system configuration database. Entries in the database are used to map an Internet Protocol (IP) address (local or remote) to its equivalent host names.

An Internet Protocol (IP) address of a given local or remote host may be associated with one or more host names. Represent an IP address in dotted decimal format. Represent a host name as a string with a maximum length of 255 characters, and do not use any blank characters.

### Note:

1. Valid host names or alias host names must contain at least one alphabetic character. If you choose to specify a host name or alias that begins with an x followed by any hexadecimal digit (0-f), the host name or alias must also contain at least one additional letter that cannot be expressed as a hexadecimal digit. The system interprets a leading x followed by a hexadecimal digit as the base 16 representation of an address unless there is at least one character in the host name or alias that is not a hexadecimal digit. Thus, xdeer would be a valid host name, whereas xdee would not.
2. The **hostmap** command does not recognize the following addresses: .08, .008, .09, and .009. Addresses with leading zeros are interpreted as octal, and numerals in octal cannot contain 8s or 9s.

## Flags

|                                 |                                                                                                                                                           |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>-addr</b> <i>IPAddress</i>   | Adds an IP address-to-host name mapping entry for the given Internet Protocol address in the database. Specify the host names with the <b>-host</b> flag. |
| <b>-host</b> <i>HostName...</i> | Specifies a list of host names. Entries in the list should be separated by blanks.                                                                        |
| <b>-ls</b>                      | Shows all entries in the database.                                                                                                                        |
| <b>-rm</b> <i>IPAddress</i>     | Deletes the IP address-to-host name mapping entry in the database that corresponds to the given address specified by the <i>IPAddress</i> variable.       |

## Exit Status

See "Exit status for Virtual I/O Server commands" on page 3.

## Examples

1. To add an entry in the database associating an address with a series of host names, enter the command in the following format:

```
hostmap -addr 192.100.201.7 -host alpha bravo charlie
```

The IP address 192.100.201.7 is specified as the address of the host that has a primary host name of alpha with synonyms of **bravo** and **charlie**.

**Note:** If you attempt to use .08, .008, .09, or .009 in an address to add, you will get an error message that states "IP Address already exists," although the address is not in the database.

2. To list all entries in the database, enter the command in the following format:

```
hostmap -ls
```

## Related Information

The **hostname** command, and the **mktcpip** command.

---

## hostname command

### Purpose

Sets or displays the name of the current host system.

### Syntax

```
hostname [HostName]
```

### Description

The **hostname** command sets or displays the name of the current host system.

**Note:** If the shared storage pool cluster is present on the machine, you cannot change the host name.

### Parameters

*Hostname*                      Sets the primary name of the host.

### Exit Status

See “Exit status for Virtual I/O Server commands” on page 3.

### Examples

1. To set the host name to **rotterdam**, type the command as follows:

```
hostname rotterdam
```

### Related Information

The **mkcftpip** command, the **startnetsvc** command, the **stopnetsvc** command, the **cfglnagg** command, the **entstat** command, the **cfgnamesrv** command, the **hostmap** command, the **traceroute** command, the **ping** command, the **optimizenet** command.

---

## importvg command

### Purpose

Imports a new volume group definition from a set of physical volumes.

### Syntax

```
importvg [-vg VolumeGroup] PhysicalVolume
```

### Description

The **importvg** command makes the previously exported volume group known to the system. The *PhysicalVolume* parameter specifies only one physical volume to identify the volume group; any remaining physical volumes (those belonging to the same volume group) are found by the **importvg** command and included in the import. An imported volume group is automatically activated. When a volume group with file systems is imported, the **/etc/filesystems** file is updated with values for the new logical volumes and mount points.

After importing the volume group, you must run the **fsck** command before the file systems can be mounted. Care should be taken to avoid using mount point longer than 128 characters as the mount

point information would be missing from the LVCB (logical volume control block) if it is longer than 128 characters. In this case, the **importvg** command will not be able to update the `/etc/filesystems` file with the stanza for the newly imported logical volume.

The **importvg** command changes the name of a logical volume if the name already exists in the system. It prints a message and the new name to standard error, and updates the `/etc/filesystems` file to include the new logical volume name.

## Flags

**-vg** *VolumeGroup* Specifies the name to use for the new volume group. If this flag is not used, the system automatically generates a new name.

The volume group name can only contain the following characters: "A" through "Z," "a" through "z," "0" through "9," or "\_" (the underscore), "-" (the minus sign), or "." (the period). All other characters are considered invalid.

## Exit Status

See "Exit status for Virtual I/O Server commands" on page 3.

## Examples

1. To import the volume group **bkgv** from physical volume **hdisk07**, type:  
`importvg -vg bkgv hdisk07`

The volume group **bkgv** is made known to the system.

## Restrictions

Mount points cannot be longer than 128 characters.

## Related Information

The **activatevg** command, the **chvg** commands, the **deactivatevg** command, the **exportvg** command, the **extendvg** command, the **lsvg** command, the **mirrorios** command, the **mkvg** command, the **syncvg** command, and the **unmirrorios** command.

---

## installios command

### Purpose

Installs the Virtual I/O Server. This command is run from the HMC.

### Syntax

```
installios [-p partition_name -i ipaddr or hostname -S subnet_mask -g gateway -d path -s system_name -m mac_address -r profile [-n] [-P speed] [-D duplex] [-l language]] [-t] | -u
```

### Description

The **installios** command installs the Virtual I/O Server. It must be run from the HMC. All of the flags are optional. If no flags are provided, the **installios** wizard will be invoked and the user will be prompted to interactively enter the information contained in the flags.

## Flags

|    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|----|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| -s | Specifies the managed system. The name of the managed system maintained by the HMC. This name must match the name shown on the HMC, not a host name.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| -p | Specifies the partition name. The name of the LPAR that will be installed with Virtual I/O Server. This partition must be of type Virtual I/O Server and the name given for it must match the name shown on the HMC, not a host name.                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| -r | Specifies the profile name. The name of the profile that contains the hardware resources being installed to.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| -d | Specifies the path to installation images. Either <b>/dev/cdrom</b> or the path to a system backup of the Virtual I/O Server created by the <b>backupios</b> command. The path may also specify a remote location mountable by NFS such as <b>hostname:path_to_backup</b>                                                                                                                                                                                                                                                                                                                                                                                                            |
| -i | Specifies the client IP address. The IP address with which the client's network interface will be configured for network installation of the Virtual I/O Server operating system.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| -S | Specifies the client subnet mask. The subnet mask with which the client's network interface will be configured for network installation of the Virtual I/O Server operating system.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| -g | Specifies the client gateway. The default gateway that the client will use during network installation of the Virtual I/O Server operating system.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| -m | Specifies the client MAC address. The MAC address of the client network interface through which the network installation of the Virtual I/O Server will take place.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| -P | Specifies speed (optional) The communication speed with which to configure the client's network interface. This value can be 10, 100, or 1000, and is 100 by default if this flag is not specified.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| -D | Specifies duplex (optional). The duplex setting with which to configure the client's network interface. This value can be full or half and is set to full by default if this flag is not specified.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| -n | Specifies not to configure the client's network interface (optional): If this flag is specified, then the client's network interface will not be configured with the IP settings specified in the flags given to this command after the installation has completed.                                                                                                                                                                                                                                                                                                                                                                                                                  |
| -l | Specifies language (optional): The language in which the license agreement will be displayed before the installation. Upon viewing the license, a prompt will be shown asking if the license is to be accepted. If the prompt is answered with y, then the installation will proceed and the Virtual I/O Server license will be automatically accepted after the installation. If the prompt is answered with n, then the <b>installios</b> command will exit and the installation will not proceed. If this flag is not specified, then the installation will proceed, but the Virtual I/O Server will not be usable until the license is manually accepted after the installation. |
| -t | Specifies to migrate the Virtual I/O Server. This option creates network installation management (NIM) resources to be used for the migration and requires a Virtual I/O Server migration DVD.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| -u | Unconfigure <b>installios</b> (optional). Will manually unconfigure the <b>installios</b> installation resources. This flag is only needed if a problem occurs during the installation and <b>installios</b> does not automatically unconfigure itself.                                                                                                                                                                                                                                                                                                                                                                                                                              |

---

## invscout command

### Purpose

Surveys the host system for currently installed microcode or Vital Product Data (VPD).

### Syntax

```
invscout [-vpd | -report] [-model Type-Model] [-serial SerialNumber]
```

```
invscout -version
```

```
invscout -rpm rpmPackage rpmOption ...
```

**invscout -install** *Device* [-file *FileName*]

## Description

The **invscout** command executes one instance of the stand alone version of the Inventory Scout process. The **invscout** command starts the server daemon side of a client-server version. The Inventory scout process supports two survey types:

- Microcode Survey
- Vital Product Data (VPD) Survey

## Flags

|                                    |                                                                                                                                                                                                                                                                                                                                                                                                                 |
|------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>-file</b> <i>FileName</i>       | Specifies the location of the upgrade.                                                                                                                                                                                                                                                                                                                                                                          |
| <b>-install</b> <i>Device</i>      | Installs microcode for a given device.                                                                                                                                                                                                                                                                                                                                                                          |
| <b>-model</b> <i>Type-Model</i>    | Machine type and model. For a VPD survey, allows input of the host platform machine type and model for hosts that use this information.                                                                                                                                                                                                                                                                         |
| <b>-report</b>                     | For a Microcode Survey, sends a copy of the formatted text report file to the screen from which the command was invoked. This flag is ignored if the <b>-vpd</b> flag is used.                                                                                                                                                                                                                                  |
| <b>-rpm</b> <i>rpmPackage</i>      | RPM options: <ul style="list-style-type: none"><li><b>-i</b> Installs a new package.</li><li><b>-e</b> Uninstalls a package.</li><li><b>-qp</b> Query for an uninstalled package.</li><li><b>--force</b> Used to install a package even if it is already installed.</li><li><b>--ignoreos</b> Allows installation of a package even if the operating system of the host and binary RPM are different.</li></ul> |
| <b>-serial</b> <i>SerialNumber</i> | Serial number. For a VPD survey, allows input of the host serial number for hosts that use this information.                                                                                                                                                                                                                                                                                                    |
| <b>-version</b>                    | Displays the versions of this command and of the logic database currently in use.                                                                                                                                                                                                                                                                                                                               |
| <b>-vpd</b>                        | Sets the survey or concatenation type to VPD (the default is Microcode).                                                                                                                                                                                                                                                                                                                                        |

## Exit Status

The following exit values are returned:

|    |                        |
|----|------------------------|
| 0  | Successful completion. |
| >0 | An error occurred.     |

## Examples

1. To generate report on microcode levels of all the devices, type:  
`invscout -report`
2. To get the vpd survey of the partition. type:  
`invscout -vpd`

---

## ioslevel Command

### Purpose

Reports the latest installed maintenance level of the system.

## Syntax

ioslevel

## Description

The `ioslevel` command displays the Virtual I/O Server level.

## Exit Status

See “Exit status for Virtual I/O Server commands” on page 3

## Related Information

The `lssw` command, the `updateios` command, the `remote_management` command, the `oem_setup_env` command, and the `oem_platform_level` command.

---

## Idapadd command

### Purpose

Opens a connection to an LDAP server, binds, and modifies or adds entries.

### Syntax

```
ldapadd [-a][-b][-c][-C][-d][-D][-f][-g][-G][-h][-i file][-K][-m][-M][-N][-O][-p][-P][-r][-R][-U][-v][-V][-w][-y][-Y][-Z]
```

### Description

The `ldapmodify` command opens a connection to an LDAP server, binds, and modifies or adds entries. The entry information is read from standard input or from file, specified using the `-f` option. `ldapadd` is implemented as a hard link to the `ldapmodify` tool. When invoked as `ldapadd`, the `-a` (add new entry) option is turned on automatically.

### Flags

|                                  |                                                                                                                                                                 |
|----------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <code>-a</code>                  | Forces add operation as default.                                                                                                                                |
| <code>-b</code>                  | Supports binary values from files (old style paths).                                                                                                            |
| <code>-c</code>                  | Specifies continuous operation; do not stop processing on error.                                                                                                |
| <code>-C <i>charset</i></code>   | Sets the character set name to use, as registered with Internet Assigned Numbers Authority (IANA).                                                              |
| <code>-d <i>level</i></code>     | Sets the debugging level in LDAP library.                                                                                                                       |
| <code>-D <i>dn</i></code>        | Binds dn.                                                                                                                                                       |
| <code>-f <i>file</i></code>      | Specifies that entry-modification information should be read from the specified file.<br><b>Note:</b> Standard input is used if the file is not specified.      |
| <code>-g</code>                  | Specifies that trailing spaces should not be stripped on attribute values.                                                                                      |
| <code>-G <i>realm</i></code>     | Specifies that realm is to be used for the DIGEST-MD5 bind mechanism.                                                                                           |
| <code>-h <i>host</i></code>      | Specifies the LDAP server host name.                                                                                                                            |
| <code>-i <i>file</i></code>      | Specifies that read entry-modification information should be read from the specified file.<br><b>Note:</b> Standard input is used if the file is not specified. |
| <code>-K <i>keyfile</i></code>   | Specifies the file to use for keys.                                                                                                                             |
| <code>-m <i>mechanism</i></code> | Performs SASL bind with the given mechanism.                                                                                                                    |
| <code>-M</code>                  | Manages referral objects as normal entries.                                                                                                                     |
| <code>-N <i>key_name</i></code>  | Specifies the private key name to use in the keyfile.                                                                                                           |

|                           |                                                                    |
|---------------------------|--------------------------------------------------------------------|
| <b>-O</b> <i>maxhops</i>  | Specifies the maximum number of referrals to follow in a sequence. |
| <b>-p</b> <i>port</i>     | Specifies the LDAP server port number.                             |
| <b>-P</b> <i>key_pw</i>   | Specifies the keyfile password.                                    |
| <b>-r</b>                 | Forces replace operation as default.                               |
| <b>-R</b>                 | Specifies not to chase referrals.                                  |
| <b>-U</b> <i>username</i> | Specifies the user name for the DIGEST-MD5 bind mechanism.         |
| <b>-v</b>                 | Specifies verbose mode                                             |
| <b>-V</b> <i>version</i>  | Specifies the LDAP protocol version (2 or 3; default is 3).        |
| <b>-w</b> <i>password</i> | Binds password or '?' for non-echoed prompt.                       |
| <b>-y</b>                 | Requests proxied authorization control.                            |
| <b>-Y</b>                 | Uses a secure ldap connection (TLS).                               |
| <b>-Z</b>                 | Uses a secure ldap connection (SSL).                               |

## Exit Status

See “Exit status for Virtual I/O Server commands” on page 3.

## Related Information

The **mkldap** command, and the **ldapsearch** command.

---

## Idapsearch command

### Purpose

Opens a connection to an LDAP server, binds, and performs a search using the filter *filter*.

### Syntax

```
ldapsearch [-a][-A][-b basedn][-B][-C][-d][-D][-e][-f][-F][-G][-h][-i][-k][-K][-l][-L][-m][-M][-n][-N][-o][-O][-p][-P][-q][-R][-s][-t][-T][-U][-v][-V][-w][-y][-Y][-z][-Z][-9] filter [attributes...]
```

where *basedn* equals the base distinguished name for the search, *filter* equals the LDAP search filter, options equals any other flags, and *attributes* equals a whitespace-separated list of attributes to retrieve. If no attribute list is specified, all attributes are retrieved.

**Note:** *basedn* is optional if LDAP\_BASEDN is set in the environment.

### Description

If the **ldapsearch** command finds one or more entries, the attributes specified by *attrs* are retrieved and the entries and values are printed to standard output. If no *attrs* are listed, all attributes are returned.

### Flags

|                          |                                                                                                    |
|--------------------------|----------------------------------------------------------------------------------------------------|
| <b>-a</b> <i>deref</i>   | Specifies how to dereference aliases (never, always, search, or find).                             |
| <b>-A</b>                | Retrieves attribute names only (no values).                                                        |
| <b>-b</b> <i>basedn</i>  | Specifies base distinguished name (dn) for search. LDAP_BASEDN in environment is the default.      |
| <b>-B</b>                | Specifies not to suppress printing of non-ASCII values.                                            |
| <b>-C</b> <i>charset</i> | Sets the character set name to use, as registered with Internet Assigned Numbers Authority (IANA). |
| <b>-d</b> <i>level</i>   | Sets the debugging level in LDAP library.                                                          |
| <b>-D</b> <i>dn</i>      | Binds dn.                                                                                          |
| <b>-e</b>                | Displays LDAP library version information and quit.                                                |



|                            |                                                                                                                                                                                                                                                       |
|----------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>-f</b> <i>file</i>      | Performs sequence of searches using filters in 'file' " must be substituted for the filter.                                                                                                                                                           |
| <b>-F</b> <i>sep</i>       | Prints 'sep' between attribute names and values.                                                                                                                                                                                                      |
| <b>-G</b> <i>realm</i>     | Specifies that realm is to be used for the DIGEST-MD5 bind mechanism.                                                                                                                                                                                 |
| <b>-h</b> <i>host</i>      | Specifies the LDAP server host name.                                                                                                                                                                                                                  |
| <b>-i</b> <i>file</i>      | Performs sequence of searches using filters in 'file' " must be substituted for the filter.                                                                                                                                                           |
| <b>-k</b>                  | Uses server administration control on bind.                                                                                                                                                                                                           |
| <b>-K</b> <i>keyfile</i>   | Specifies the file to use for the keys.                                                                                                                                                                                                               |
| <b>-l</b> <i>time</i>      | Specifies time limit (in seconds) for search.                                                                                                                                                                                                         |
| <b>-L</b>                  | Prints entries in LDIF format (-B is implied).                                                                                                                                                                                                        |
| <b>-m</b> <i>mechanism</i> | Performs SASL bind with the specified mechanism.                                                                                                                                                                                                      |
| <b>-M</b>                  | Manages referral objects as normal entries.                                                                                                                                                                                                           |
| <b>-n</b>                  | Shows what would be done but don't actually perform the action.                                                                                                                                                                                       |
| <b>-N</b> <i>key_name</i>  | Specifies the private key name to use in the keyfile.                                                                                                                                                                                                 |
| <b>-o</b> <i>attr_type</i> | Sorts based on specified attribute type.                                                                                                                                                                                                              |
| <b>-O</b> <i>maxhops</i>   | Specifies the maximum number of referrals to follow in a sequence.                                                                                                                                                                                    |
| <b>-p</b> <i>port</i>      | Specifies the LDAP server port number.                                                                                                                                                                                                                |
| <b>-P</b> <i>key_pw</i>    | Specifies the keyfile password.                                                                                                                                                                                                                       |
| <b>-q</b> <i>pagesize</i>  | Queries page size for paged results.                                                                                                                                                                                                                  |
| <b>-R</b>                  | Specifies not to chase referrals.                                                                                                                                                                                                                     |
| <b>-s</b> <i>scope</i>     | Searches scope. The search scope can be one of the following: base, one, or sub.                                                                                                                                                                      |
| <b>-t</b> <i>scope</i>     | Writes values to files in /tmp.                                                                                                                                                                                                                       |
| <b>-T</b> <i>seconds</i>   | Specifies the number of seconds to wait between pages for paged results.                                                                                                                                                                              |
| <b>-U</b> <i>username</i>  | Specifies the user name for the DIGEST-MD5 bind mechanism.                                                                                                                                                                                            |
| <b>-v</b>                  | Runs in verbose mode.                                                                                                                                                                                                                                 |
| <b>-V</b> <i>version</i>   | Specifies the LDAP protocol version. The version can be 2 or 3. The default is 3.                                                                                                                                                                     |
| <b>-w</b> <i>password</i>  | Binds password or '?' for non-echoed prompt.                                                                                                                                                                                                          |
| <b>-y</b> <i>proxydn</i>   | Sets proxied ID for proxied authorization operation.                                                                                                                                                                                                  |
| <b>-Y</b>                  | Uses a secure LDAP connection (TLS).                                                                                                                                                                                                                  |
| <b>-zsize</b>              | Specifies the size limit (in entries) for the search.                                                                                                                                                                                                 |
| <b>-Z</b>                  | Uses a secure LDAP connection (SSL).                                                                                                                                                                                                                  |
| <b>-9control</b>           | Sets the criticality for control option. The control can be one of the following options: <ul style="list-style-type: none"> <li>• <b>s</b> sets criticality for sorting to false</li> <li>• <b>p</b> sets criticality for paging to false</li> </ul> |

## Exit Status

See "Exit status for Virtual I/O Server commands" on page 3.

## Examples

To open a connection to the LDAP server and search on specific attributes, type:  
`ldapsearch -h vclient.host.com -D cn=admin -w adminpw -b cn=aixdata objectclass=*`

## Related Information

The `mkldap` command and the `ldapadd` command.

---

## ldfware command

### Purpose

Loads the system's flash EPROM with the specified file, which must contain a valid binary flash EPROM image, and then reboots the system.

## Syntax

**ldfware** [ **-dev** *Device* ] **-file** *filename*

**ldfware -commit**

**ldfware -reject**

## Description

Periodically, you need to install fixes for your server firmware. If you do not use an HMC to manage your server, you must get your fixes through your operating system. If your server is not connected to the Internet, you must obtain optical media that contains the server firmware fix and install the server firmware fix from the optical media. The flash update image file is copied to the file system from diskette. The user needs to provide the image on a diskette if the user does not have access to remote file systems or any other files that are on the system. If enough space is not available, an error is reported stating additional system memory is needed. After the file is copied, a warning screen asks for confirmation to continue the update flash. Continuing the update flash reboots the system. The current flash image is not saved.

After you download and install a firmware fix, the fix is temporarily installed until you install it permanently. You might want to use the new level of firmware for a period of time to verify that it works correctly. When you are sure that the new level of firmware works correctly, you can permanently install the firmware fix. Be aware that if you install the firmware fix permanently (copy the temporary firmware level from the temporary side to the permanent side, so that the temporary and permanent sides contain the same level of firmware), you cannot return to the level that was previously on the permanent side. To install it permanently, use the **-commit** flag.

**Attention:** The **ldfware** command reboots the entire system. Do not use this command if more than one user is signed onto the system.

## Flags

|                              |                                                                                                                                                                                                          |
|------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>-commit</b>               | Commits the temporary image when booted from the temporary image. This effectively causes the permanent image to be copied over by the temporary image. This flag is not supported on all systems.       |
| <b>-dev</b> <i>Device</i>    | Specifies that the flash update image file is on diskette. The <i>device_name</i> variable specifies the device. The default <i>device_name</i> is <code>/dev/fd0</code> .                               |
| <b>-file</b> <i>FileName</i> | Specifies the file name of the image file source. The flash update image file source follows this flag. The <i>file_name</i> variable specifies the fully qualified path of the flash update image file. |
| <b>-reject</b>               | Rejects the temporary image when booted from the permanent image. This overwrites the temporary image with the permanent image. This flag is not supported on all systems.                               |

## Exit Status

The following exit codes are returned:

|   |                                                    |
|---|----------------------------------------------------|
| 0 | Completed successfully.                            |
| 1 | Command terminated due to an error.                |
| 2 | Command was aborted by the user.                   |
| 3 | Command was aborted by the user using the F10 key. |

## Examples

1. To update the firmware from a local file, type the following command:  
`ldfware -file /tmp/firmware/latest_flash`
2. To update the firmware that is on a diskette, type the following command:  
`ldfware -device /dev/fd0 -file latest_flash`
3. To update the firmware that is on a CD device, type the following command:  
`mount -cd /mnt ldfware -file /mnt/firmware/latest_flash`

---

## license command

### Purpose

View and accept the license agreement.

### Syntax

```
license { [-view] [-accept] } [-lang Name]
```

```
license [-ls]
```

### Description

The **license** command is used to view and accept the Virtual I/O Server license agreement. If no flags are specified, the current status of the license agreement is displayed. If the license has been accepted, the date and time of acceptance is displayed.

### Flags

|                          |                                                                                                             |
|--------------------------|-------------------------------------------------------------------------------------------------------------|
| <b>-accept</b>           | Accepts the license agreement.                                                                              |
| <b>-lang <i>Name</i></b> | Specifies the language-territory (locale name) that the license will be displayed in. The default is en_US. |
| <b>-ls</b>               | Lists available languages.                                                                                  |
| <b>-view</b>             | Displays the Virtual I/O Server license agreement.                                                          |

### Exit Status

See “Exit status for Virtual I/O Server commands” on page 3.

### Examples

1. To view the license in the en\_US locale, type:  
`license -view`
2. To accept the license in the fr\_FR locale, type:  
`license -accept -lang fr_FR`
3. To view if the license has been accepted, type:  
`license`

### Related Information

The **chlang** command.

---

## loadopt command

### Purpose

Load a virtual optical media disk from the Virtual Media Repository into a virtual optical device.

### Syntax

```
loadopt [-f] [-release] -disk FileName -vtd VirtualTargetDevice
```

### Description

The **loadopt** command loads the specified virtual optical disk (*FileName*) into the specified virtual optical device (*VirtualTargetDevice*).

**Note:** A particular virtual disk can only be loaded into more than one file-backed virtual optical device simultaneously if the virtual DVD is designated as read-only. If the virtual optical device already has a different virtual disk loaded, the command will fail unless the **-f** flag is specified, in which case an implicit unload will be performed first.

### Flags

|                                        |                                                                                                  |
|----------------------------------------|--------------------------------------------------------------------------------------------------|
| <b>-disk</b> <i>FileName</i>           | Specifies the file name of the virtual media to load into the virtual optical device.            |
| <b>-f</b>                              | Force the media to be loaded even if the virtual optical device already has media loaded.        |
| <b>-release</b>                        | Forces the virtual optical device to be unlocked even if the client has a reserve on the device. |
| <b>-vtd</b> <i>VirtualTargetDevice</i> | The name of the virtual target device.                                                           |

### Examples

To load the virtual optical disk *clientData* into the virtual optical device *vopt1*, type the following command:

```
loadopt -disk clientData -vtd vopt1
```

---

## loginmsg command

### Purpose

Modifies the Virtual I/O Server partition's login herald.

### Syntax

```
loginmsg { -reset | "Herald string" }
```

### Description

The **loginmsg** command sets the Virtual I/O Server partition's login herald.

### Flags

|               |                                                     |
|---------------|-----------------------------------------------------|
| <b>-reset</b> | Reset the login message back to the system default. |
|---------------|-----------------------------------------------------|

## Examples

1. To set the login herald to Welcome followed by login: on a separate line, type:  
loginmsg "Welcome\nlogin:"
2. To reset the login herald back to the system default, type:  
loginmsg -reset

---

## IVM lpar\_netboot command

### Purpose

Retrieves the media access control (MAC) address and physical location code from network adapters for a logical partition, or instructs a logical partition to do a network boot. This command is operable only in an Integrated Virtualization Manager environment.

### Syntax

To retrieve a MAC address:

```
lpar_netboot -M -n [-v] [-x] [-f] [-i] [-A] -t ent [-D -s Speed-d Duplex -S Server -G Gateway -C Client -K subnetmask] partition name partition profile manage system
```

To perform a network boot:

```
lpar_netboot [-v [-x] [-f] [-i] [-g args] [-A -D | [-D] -l phys_loc | [-D] -m maddress] -t ent [-D -s Speed-d Duplex -S Server -G Gateway -C Client -K subnetmask] partition name partition profile manage system
```

### Description

The **lpar\_netboot** command instructs a logical partition to do a network boot by having the logical partition send out a bootp request to a server that is specified with the **-S** flag. The server can be a network installation management (NIM) server serving SPOT resources or any server serving network boot images.

If the **-M** and **-n** flags are specified, the **lpar\_netboot** command returns the MAC address and the physical location code for a particular type of network adapter that is specified with the **-t** flag. When the **-m** flag is specified, **lpar\_netboot** boots a partition using a specific network adapter that matches the specified MAC address. When the **-l** flag is specified, **lpar\_netboot** boots a logical partition by using a specific physical location code for the network adapter that matches the specified physical location code. The matching MAC address or physical location code is dependent on the hardware resource allocation in the profile in which the logical partition was booted. The **lpar\_netboot** command also requires arguments for the partition name and the partition profile (which contains the allocated hardware resources), and the name of the managed system in which the logical partition was defined.

### Flags

|                   |                                                                                                                 |
|-------------------|-----------------------------------------------------------------------------------------------------------------|
| <b>-A</b>         | Returns all adapters of the given type.                                                                         |
| <b>-C Client</b>  | Specifies the IP address of the machine to do a network boot.                                                   |
| <b>-D</b>         | Performs a ping test so the adapter can successfully ping the server that is specified with the <b>-S</b> flag. |
| <b>-d Duplex</b>  | Specifies the duplex setting of the machine that is specified with the <b>-C</b> flag.                          |
| <b>-f</b>         | Forces a close virtual terminal session for the logical partition.                                              |
| <b>-G Gateway</b> | Specifies the gateway IP address of the machine that is specified with the <b>-C</b> flag.                      |
| <b>-g args</b>    | Specifies generic arguments for booting.                                                                        |
| <b>-i</b>         | Forces an immediate shutdown of the partition.                                                                  |

|                             |                                                                                                                                                                                                                                                                                                                                                                         |
|-----------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>-K</b> <i>subnetmask</i> | Specifies the mask that the gateway should use in determining the appropriate subnetwork for routing. The subnet mask is a set of 4 bytes, as in the Internet address. The subnet mask consists of high bits (1's) corresponding to the bit positions of the network and subnetwork address, and low bits (0's) corresponding to the bit positions of the host address. |
| <b>-l</b> <i>phys_loc</i>   | Specifies the physical location code of the network adapter to do a network boot.                                                                                                                                                                                                                                                                                       |
| <b>-M</b>                   | Displays the network adapter MAC address and physical location code.                                                                                                                                                                                                                                                                                                    |
| <b>-m</b> <i>maddress</i>   | Specifies the MAC address of the network adapter to do a network boot.                                                                                                                                                                                                                                                                                                  |
| <b>-n</b>                   | Instructs the logical partition to not do a network boot.                                                                                                                                                                                                                                                                                                               |
| <b>-S</b> <i>Server</i>     | Specifies the IP address of the machine to retrieve the network boot image during the network boot.                                                                                                                                                                                                                                                                     |
| <b>-s</b> <i>Speed</i>      | Specifies the speed setting of the machine that is specified with the <b>-C</b> flag.                                                                                                                                                                                                                                                                                   |
| <b>-t</b> <i>ent</i>        | Specifies the type of adapter for MAC address or physical location code discovery, or for a network boot.                                                                                                                                                                                                                                                               |
| <b>-v</b>                   | Displays additional information while the command is running.                                                                                                                                                                                                                                                                                                           |
| <b>-x</b>                   | Displays debug output while the command is running.                                                                                                                                                                                                                                                                                                                     |

## Parameters

| Parameter                | Description                                                                 |
|--------------------------|-----------------------------------------------------------------------------|
| <i>partition_name</i>    | Specifies the name of the partition.                                        |
| <i>partition_profile</i> | Specifies the name of the partition profile to use.                         |
| <i>managed_system</i>    | Specifies the name of the managed system on which the partition is defined. |

## Exit Status

The following exit values are returned:

|    |                                     |
|----|-------------------------------------|
| 0  | The command completed successfully. |
| >0 | An error occurred.                  |

## Security

Access Control: You must have root authority to run the **lpar\_netboot** command.

## Examples

- To retrieve MAC address and physical location code for partition machA with logical partition profile machA\_prof on managed system test\_sys, type:

```
lpar_netboot -M -n -t ent "machA" "machA_prof" "test_sys"
```
- To do a network boot of logical partition machA with a partition profile machA\_prof on managed system test\_sys, type:

```
lpar_netboot -t ent -s auto -d auto -S 9.3.6.49 -G 9.3.6.1 -C 9.3.6.234
 "machA" "machA_prof" "test_sys"
```
- To do a network boot of logical partition machA with the specific MAC address 00:09:6b:dd:02:e8 and logical partition profile machA\_prof on managed system test\_sys, type:

```
OS_install -o allocate -a os_resource=my53resource myclient01
```
- To do a network boot of logical partition machA with the specific physical location code U1234.121.A123456-P1-T6 and logical partition profile machA\_prof on managed system test\_sys, type:

```
lpar_netboot -t ent -l U1234.121.A123456-P1-T6 -s auto -d auto -S 9.3.6.49
 -G 9.3.6.1 -C 9.3.6.234 "machA" "machA_prof" "test_sys"
```

5. To perform a ping test and do a network boot of logical partition machA with logical partition profile machA\_prof on managed system test\_sys, type:

```
lpar_netboot -t ent -D -s auto -d auto -S 9.3.6.49 -G 9.3.6.1 -C 9.3.6.234
"machA" "machA_prof" "test_sys"
```

## Related Information

The `installios` command.

---

## IVM lpcfgop command

### Purpose

Perform a partition configuration image operation. This command is operable only in an Integrated Virtualization Manager environment.

### Syntax

To clear logical partition configuration data at next system restart:

```
lpcfgop -o clear [--force] [-m ManagedSystem]
```

To disable logical partition configuration data at next system restart:

```
lpcfgop -o disable [--force] [-m ManagedSystem]
```

To dump logical partition configuration data to a file:

```
lpcfgop -o dump [-f DumpFile] [-m ManagedSystem]
```

To enable logical partition configuration data at next system restart:

```
lpcfgop -o enable [-m ManagedSystem]
```

### Description

The `lpcfgop` command performs a partition configuration image operation such as clearing, disabling, enabling, or dumping the logical partition configuration data.

### Flags

|                                  |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|----------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <code>-o</code> <i>Operation</i> | The type of operations: <ul style="list-style-type: none"><li>• clear: Marks the logical partition configuration data to be cleared when the managed system is restarted.</li><li>• disable: Marks the logical partition configuration data to be disabled when the managed system is restarted. This will temporarily place the system back into the factory configuration partition mode.</li><li>• dump: Dumps all logical partition configuration data from managed system firmware to a file. This data may be used by service. Use the <code>bkprofddata</code> command to create backups that can be restored.</li><li>• enable: Enables partition configuration data so that it will not be disabled or cleared when the managed system is restarted.</li></ul> |
| <code>-f</code> <i>DumpFile</i>  | The name of the file to write the dump to in the current working directory. If not specified, the default file will be <code>/var/adm/lpm/lparConfig.dump</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <code>--force</code>             | Skip confirmation prompts for clear and disable operations.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |

**-m** *ManagedSystem* The name of the managed system. This attribute is optional because there is only one system to manage. The name may either be the user-defined name for the managed system, or be in the form `ttt-mmm*sssssss`, where `ttt` is the machine type, `mmm` is the model, and `sssssss` is the serial number of the managed system.

## Exit Status

This command has a return code of 0 on success.

## Security

This command is accessible only to the `padmin` user in an IVM environment.

## Examples

1. Clear the partition configuration data (confirmation will be required) during the next reboot:  
`lpcfgop -o clear`
2. Dump the logical partition configuration data to `lparData.dump`:  
`lpcfgop -o dump -f lparData.dump`

## Related Information

The `bkprofdata` command and the `rstprofdata` command.

---

## lsauth command

### Purpose

Displays user and system-defined authorization attributes.

### Syntax

```
lsauth [-C][-f][-a List] {ALL | auth1,[auth2] ...}
```

### Description

The `lsauth` command displays attributes of user-defined and system-defined authorizations from the authorization database. The command can be used to list attributes of all authorizations or specific authorizations. By default, the `lsauth` command displays all authorization attributes. To view selected attributes, use the `-a List` flag. If one or more attributes cannot be read, the `lsauth` command lists the information that is available.

By default, the `lsauth` command lists the attributes of each authorization in one line. It displays attribute information in the form of `Attribute=Value`, each separated by a blank space. To list the authorization attributes in stanza format, use the `-f` flag. To list the information as colon-separated records, use the `-C` flag.

### Security

The `lsauth` command is a privileged command. To run the command successfully, assume a role that has the authorization `vios.security.auth.list`.

### Flags

`-C` Displays the authorization attributes in colon-separated records.



- f** Displays the output in stanzas, with each stanza identified by an authorization name.
- a** Specifies the list of attributes to be listed.

## Exit Status

0 for success.

Nonzero for command error.

## Examples

1. To display all attributes of the *custom* authorization, type the following command:  

```
lsauth custom
```
2. To display the authorization ID and description for the custom authorization in stanza format, type the following command:  

```
lsauth -f -a id description custom
```
3. To display the authorization ID and description for the custom authorization in colon-separated format, type the following command:  

```
lsauth -c -a id description custom
```

---

## Isdev command

### Purpose

Displays Virtual I/O Server devices and their characteristics.

### Syntax

To list devices:

```
lsdev [-type DeviceType...] [-virtual] [-field FieldName...] [-fmt Delimiter] [-state State]
```

To display information about a specific device:

```
lsdev { -dev DeviceName | -plc PhysicalLocationCode } [-child] [-field FieldName...] [-fmt Delimiter]
```

```
lsdev { -dev DeviceName | -plc PhysicalLocationCode } [-attr [Attribute] | -range Attribute | -slot | -vpd | -parent]
```

```
lsdev -vpd
```

```
lsdev -slots
```

### Description

The **lsdev** command displays information about devices in the Virtual I/O Server. If no flags are specified, a list of all devices, both physical and virtual, in the Virtual I/O Server is displayed. To list devices, both physical and virtual, of a specific type use the **-type** *DeviceType* flag. Use the **-virtual** flag to list only virtual devices. Combining both the **-type** and **-virtual** flags will list the virtual devices of the specified type.

To display information about a specific device, use the **-dev** *DeviceName* or **-plc** *PhysicalLocationCode*. Use either the **-child**, **-parent**, **-attr**, **-range**, **-slot**, or **-vpd** flag to specify what type of information is displayed. If none of these flags are used, the name, status, and description of the device will be displayed.

Using the **-vpd** flag, without specifying a device, displays platform-specific information for all devices.

If the **-fmt** *Delimiter* flag is specified the **lsdev** command returns all output in a delimiter separated format. The **-state** *State* flag limits the output to devices in the indicated state. The **-slots** flag produces a list of I/O slot information for built-in adapters that are not hot-pluggable but can have dynamic LPAR operations performed on them.

The **lsdev** output will be truncated to 80 characters per line, unless the **-fmt** flag is used. If the **-fmt** flag is used, the output lines will not be truncated and line wrap may occur, based on what the terminal window line length is set to.

## Flags

|                                         |                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|-----------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>-attr</b> [ <i>Attribute</i> ]       | Displays information about attributes of a given device. If an attribute is specified, its current value is displayed. If no attribute is specified, the following information is displayed about all attributes for the give device:<br><br><b>attribute</b><br>Attribute name<br><br><b>value</b> Current value of the attribute<br><br><b>description</b><br>Description of the attribute<br><br><b>user</b> Whether attribute can be set by the user (TRUE/FALSE) |
| <b>-child</b>                           | Displays the name, status, physical location code, and description for each child of the specified device ( <b>-dev</b> <i>DeviceName</i> or <b>-plc</b> <i>PhysicalLocationCode</i> ).                                                                                                                                                                                                                                                                               |
| <b>-dev</b> <i>DeviceName</i>           | Specifies the device logical name for the device for which information is listed. This flag cannot be used with the <b>-plc</b> flag.                                                                                                                                                                                                                                                                                                                                 |
| <b>-field</b> <i>FieldName</i>          | Specifies the list of fields to be displayed. The following fields are supported:<br><br><b>name</b> Device name<br><br><b>status</b> Device status<br><br><b>physloc</b><br>Physical location code<br><br><b>description</b><br>Description of the device<br><br><b>parent</b><br><b>Note:</b> The <b>-field</b> flag cannot be combined with the <b>-parent</b> , <b>-attr</b> , <b>-range</b> , <b>-slot</b> , or <b>-vpd</b> flags.                               |
| <b>-fmt</b> <i>Delimiter</i>            | Specifies a delimiter character to separate output fields.                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>-parent</b>                          | Displays the name, status, physical location code, and description of the parent device of the given device ( <b>-dev</b> <i>DeviceName</i> or <b>-plc</b> <i>PhysicalLocationCode</i> ).                                                                                                                                                                                                                                                                             |
| <b>-plc</b> <i>PhysicalLocationCode</i> | Specifies the device physical location code for the device for which information is listed. This flag cannot be used with the <b>-dev</b> flag.                                                                                                                                                                                                                                                                                                                       |
| <b>-range</b> <i>Attribute</i>          | Displays the allowed values for the specified attribute.                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>-slot</b>                            | Displays the slot, description, and device name of the specified device ( <b>-dev</b> <i>DeviceName</i> or <b>-plc</b> <i>PhysicalLocationCode</i> ). The device must be in a PCI hot swappable slot.                                                                                                                                                                                                                                                                 |
| <b>-slots</b>                           | Displays a list of I/O slot information for built-in adapters that are not hot-pluggable but can have dynamic LPAR operations performed on them.                                                                                                                                                                                                                                                                                                                      |

**-State** *State* Limits the output to devices in the specified state. The following states are supported:

- 0, defined**  
Server Virtual Adapter.
- 1, available**  
Server Virtual Adapter Physical Location Code
- 2, stopped**  
Client Partition ID

**-type** *DeviceType* Specifies the device type. This flag can be used to restrict output to devices in the specified types. Both physical and virtual devices are listed.

Supported types are as follows:

- adapter** Lists adapters
- disk** Lists disks
- lv** Lists logical volumes and volume groups
- optical** Lists optical devices (cdrom/dvdrom)
- tape** Lists tape devices
- tape4vtd** Lists tape devices available for creating virtual target devices
- tty** Lists tty devices
- ent4sea** Lists all physical Ethernet adapters and Etherchannel adapters available for creating a shared Ethernet adapter
- vent4sea** Lists all virtual Ethernet adapters available for creating shared Ethernet adapter
- ent4ip** Lists all adapters over which interface can be configured.
- sea** Lists all shared Ethernet adapters over which the interface can be configured.

**-virtual** Restricts output to virtual devices only.

**-vpd** Displays platform-specific information for all devices or for a single device when **-dev** *DeviceName* or **-plc** *PhysicalLocationCode* are specified.

## Exit Status

Table 3. Command specific return codes

| Return code | Description                                                                  |
|-------------|------------------------------------------------------------------------------|
| 12          | Specified logical volume belongs to the operating system.                    |
| 13          | Specified physical or logical volume is not valid physical or logical volume |

See “Exit status for Virtual I/O Server commands” on page 3.

## Examples

- To list all virtual adapters and display the **name** and **status** fields, type:

```
lsdev -type adapter -virtual -field name status
```

The system displays a message similar to the following:

```
name status
```

```
vhost0 Available
```

```
vhost1 Available
vhost2 Available
ent6 Available
ent7 Available
ent8 Available
ent9 Available
```

2. To list all devices of type **disk** and display the name and physical location fields, type:

```
lsdev -type disk -field name physloc
```

The system displays a message similar to the following:

```
name physloc

hdisk0 U9111.520.10004BA-T15-L5-L0
hdisk1 U9111.520.10004BA-T15-L8-L0
hdisk2 U9111.520.10004BA-T16-L5-L0
hdisk3 U9111.520.10004BA-T16-L8-L0
hdisk4 UTMP0.02E.00004BA-P1-C4-T1-L8-L0
hdisk5 UTMP0.02E.00004BA-P1-C4-T2-L8-L0
hdisk6 UTMP0.02F.00004BA-P1-C8-T2-L8-L0
hdisk7 UTMP0.02F.00004BA-P1-C4-T2-L8-L0
hdisk8 UTMP0.02F.00004BA-P1-C4-T2-L11-L0
vtscsi0 U9111.520.10004BA-V1-C2-L1
vtscsi1 U9111.520.10004BA-V1-C3-L1
vtscsi2 U9111.520.10004BA-V1-C3-L2
vtscsi3 U9111.520.10004BA-V1-C4-L1
vtscsi4 U9111.520.10004BA-V1-C4-L2
vtscsi5 U9111.520.10004BA-V1-C5-L1
```

3. To display the parent of a devices, type:

```
lsdev -dev hdisk0 -parent
```

The system displays a message similar to the following:

```
parent
```

```
scsi0
```

4. To display all I/O slots that are not hot-pluggable but can have DLPAR operations performed on them, type:

```
lsdev -slots
```

The system displays a message similar to the following:

```
U787A.001.DNZ00Y1-P1-C1 Logical I/O Slot pci4 sisscsia0
U787A.001.DNZ00Y1-P1-T5 Logical I/O Slot pci3 ent0 ent1
U787A.001.DNZ00Y1-P1-T7 Logical I/O Slot pci2 usbhc0 usbhc1
U9111.520.10DFD8C-V2-C0 Virtual I/O Slot vsa0
U9111.520.10DFD8C-V2-C2 Virtual I/O Slot vhost0
U9111.520.10DFD8C-V2-C4 Virtual I/O Slot Unknown
```

## Related Information

The **cfgdev** command, the **chdev** command, the **chpath** command, the **lsmap** command, the **lspath** command, the **mkpath** command, the **mkvdev** command, the **rmdev** command, and the **rmpath** command.

---

## lsfailedlogin command

### Purpose

Lists the contents of the failed login log to the screen.

## Syntax

`lsfailedlogin`

## Description

The `lsfailedlogin` command dumps the contents of the failed login log. The failed login file records unsuccessful login attempts by any user on the Virtual I/O Server.

## Exit Status

See “Exit status for Virtual I/O Server commands” on page 3.

## Security

This command can only be executed by the prime administrator (`padmin`) user.

## Examples

1. To list all failed logins, type:

```
lsfailedlogin
```

## Related Information

The `lsgcl` command.

---

## lsfware command

### Purpose

Displays microcode and firmware levels of the system, adapters and devices.

### Syntax

```
lsfware [-all | -dev Name]
```

### Description

The `lsfware` command displays the platform system firmware microcode level and the service processor microcode levels, if supported. Not all systems contain a service processor, nor do all systems support displaying the system processor level. Information on a specific device is displayed with the `-dev` flag.

The `-all` flag displays system firmware/microcode for all devices. The output from the `-all` flag is always delimiter separated.

### Flags

|                               |                                                                 |
|-------------------------------|-----------------------------------------------------------------|
| <code>-all</code>             | Displays microcode level information for all supported devices. |
| <code>-dev <i>Name</i></code> | Displays microcode level information for the named device.      |

### Examples

1. To display the system firmware level and service processor (if present), type:

```
lsfware
```

The system displays a message similar to the following:

System Firmware level is TCP99256

- To display the microcode level for all supported devices, type:

```
lsfware -all
```

The system displays a message similar to the following:

```
sys0|system:TCP99256
rmt0|C009
scraid0|adapter:4.20.18|adapter-boot:4.00.26
raid-dasd|22:FFC #:DDYS-T0.524D3031.53393446
raid-dasd|26:FFC #:DDYS-T0.524D3031.53393446
raid-dasd|2e:FFC #:DDYS-T0.525A3034.53393243
....
```

---

## Isgcl command

### Purpose

Lists the contents of the global command log to the screen.

### Syntax

```
lsycl
```

### Description

The **lsycl** command lists the contents of the global command log (gcl). This log contains a listing of all commands that have been executed by all Virtual I/O Server users. Each listing contains the date and time of execution as well as the userid the command was executed from.

Global command log file format is as follows:

| Date        | Time     | userid  | Command | Command options                                                  |
|-------------|----------|---------|---------|------------------------------------------------------------------|
| mmm dd yyyy | hh:mm:ss | ccccccc | Command | Command options span to 80 characters then wraps to the next row |

### Exit Status

See "Exit status for Virtual I/O Server commands" on page 3.

### Security

This command can only be executed by the prime administrator (padmin) user.

### Examples

- To list the contents of the global command log, type:  

```
lsycl
```

### Related Information

The **lsfailedlogin** command.

---

## IVM Ishowres command

### Purpose

Lists the hardware resources of a managed system. This command is operable only in an Integrated Virtualization Manager environment.

### Syntax

To list attributes for physical I/O buses

```
Ishowres -r io --subtype bus [--filter "FilterData"] [-F "AttributeNames"] [--header] [-m ManagedSystem]
```

To list attributes for I/O pools

```
Ishowres -r io --subtype iopool --level pool [--filter "FilterData"] [-F "AttributeNames"] [--header] [-m ManagedSystem]
```

To list system attributes for I/O pools

```
Ishowres -r io --subtype iopool --level sys [-F "AttributeNames"] [--header] [-m ManagedSystem]
```

To list attributes for physical I/O buses

```
Ishowres -r io --subtype bus [--filter "FilterData"] [-F "AttributeNames"] [--header] [-m ManagedSystem]
```

To list attributes for physical I/O slots

```
Ishowres -r io --subtype slot [--filter "FilterData"] [-F "AttributeNames"] [--header] [-m ManagedSystem]
```

To list attributes for physical I/O units

```
Ishowres -r io --subtype unit [--filter "FilterData"] [-F "AttributeNames"] [--header] [-m ManagedSystem]
```

To list recoverable I/O resources

```
Ishowres -r io --subtype slot -R [--filter "FilterData"] [-F "AttributeNames"] [--header] [-m ManagedSystem]
```

To list partition attributes for memory

```
Ishowres -r mem --level lpar [--filter "FilterData"] [-F "AttributeNames"] [--header] [-m ManagedSystem]
```

To display memory statistics from a partition

```
Ishowres -r mem --level lpar --stat --filter "FilterData" [-F "AttributeNames"] [--header] [-m ManagedSystem]
```

To list system attributes for memory

```
Ishowres -r mem --level sys [--maxmem] [-F "AttributeNames"] [--header] [-m ManagedSystem]
```

To list recoverable memory resources

```
lshwres -r mem --level lpar -R [--filter "FilterData"] [-F "AttributeNames"] [--header] [-m ManagedSystem]
```

To list information about the memory pool

```
lshwres -r mempool [--filter "FilterData"] [-F "AttributeNames"] [--header] [-m ManagedSystem]
```

To list paging devices associated with the memory pool

```
lshwres -r mempool --subtype pgdev [--filter "FilterData"] [-F "AttributeNames"] [--header] [-m ManagedSystem]
```

To list recoverable memory pools

```
lshwres -r mempool -R [-F "AttributeNames"] [--header] [-m ManagedSystem]
```

To list partition attributes for processors

```
lshwres -r proc --level lpar [--filter "FilterData"] [-F "AttributeNames"] [--header] [-m ManagedSystem]
```

To list shared processor pool attributes

```
lshwres -r proc --level pool [-F "AttributeNames"] [--header] [-m ManagedSystem]
```

To list system attributes for processors

```
lshwres -r proc --level sys [-F "AttributeNames"] [--header] [-m ManagedSystem]
```

To list recoverable processors resources

```
lshwres -r proc --level lpar -R [--filter "FilterData"] [-F "AttributeNames"] [--header] [-m ManagedSystem]
```

To list virtual Ethernet adapter attributes

```
lshwres -r virtualio --subtype eth --level lpar [--filter "FilterData"] [-F "AttributeNames"] [--header] [-m ManagedSystem]
```

To list system attributes for virtual Ethernet adapters

```
lshwres -r virtualio --subtype eth --level sys [-F "AttributeNames"] [--header] [-m ManagedSystem]
```

To list virtual fibre channel adapter attributes

```
lshwres -r virtualio --subtype fc --level lpar [--filter "FilterData"] [-F "AttributeNames"] [--header] [-m ManagedSystem]
```

To list system attributes for virtual fibre channel adapters

```
lshwres -r virtualio --subtype fc --level sys [-F "AttributeNames"] [--header] [-m ManagedSystem]
```

To list virtual SCSI adapter attributes



**lshwres -r virtualio --rsubtype scsi [ --level lpar ] [ --filter "FilterData" ] [ -F "AttributeNames" ] [ --header ] [ -m ManagedSystem ]**

To list virtual serial adapter attributes

**lshwres -r virtualio --rsubtype serial --level lpar [ --filter "FilterData" ] [ -F "AttributeNames" ] [ --header ] [ -m ManagedSystem ]**

To list partition attributes for virtual I/O slots

**lshwres -r virtualio --rsubtype slot --level lpar [ --filter "FilterData" ] [ -F "AttributeNames" ] [ --header ] [ -m ManagedSystem ]**

To list virtual I/O slot attributes

**lshwres -r virtualio --rsubtype slot --level slot [ --filter "FilterData" ] [ -F "AttributeNames" ] [ --header ] [ -m ManagedSystem ]**

To list recoverable virtual I/O resources

**lshwres -r virtualio --rsubtype slot --level slot -R [ --filter "FilterData" ] [ -F "AttributeNames" ] [ --header ] [ -m ManagedSystem ]**

To list the physical Host Ethernet Adapters

**lshwres -r hea --rsubtype phys --level sys [ --filter "FilterData" ] [ -F "AttributeNames" ] [ --header ] [ -m ManagedSystem ]**

To list the physical Host Ethernet Adapters ports

**lshwres -r hea --rsubtype phys --level port [ --filter "FilterData" ] [ -F "AttributeNames" ] [ --header ] [ -m ManagedSystem ]**

To list the physical Host Ethernet Adapters port performance metrics

**lshwres -r hea --rsubtype phys --level port --stat [ --filter "FilterData" ] [ -F "AttributeNames" ] [ --header ] [ -m ManagedSystem ]**

To list the physical Host Ethernet Adapters port groups

**lshwres -r hea --rsubtype phys --level port\_group [ --filter "FilterData" ] [ -F "AttributeNames" ] [ --header ] [ -m ManagedSystem ]**

To list the logical Host Ethernet Adapters and resource settings

**lshwres -r hea --rsubtype logical --level sys [ --filter "FilterData" ] [ -F "AttributeNames" ] [ --header ] [ -m ManagedSystem ]**

To list the logical ports

**lshwres -r hea --rsubtype logical --level port [-R] [ --filter "FilterData" ] [ -F "AttributeNames" ] [ --header ] [ -m ManagedSystem ]**

To list the virtual Ethernet topology

**lshwres -r virtualio --rsubtype eth --level lpar -F device\_name,drc\_name,shared\_adapter,backing\_device**

## Description

The `lshwres` command lists the hardware resources of the managed-system, including physical I/O, virtual I/O, memory, and processing.

## Flags

- r** *ResourceType*  
**--rsubtype** *ResourceSubtype* **--level** *ResourceLevel*
- The combination of `-r`, `--rsubtype`, and `--level` indicates which type of attributes and objects to list. The valid combinations are enumerated below.
- **-r io --rsubtype bus**: List attributes for physical I/O buses
    - Attributes: `unit_phys_loc`, `bus_id`, `backplane_phys_loc`, `bus_drc_index`, `bus_drc_name`
    - Filters: `units`, `buses`
  - **-r io --rsubtype iopool -level sys**: List attributes for physical I/O pools
    - Attributes: `max_io_pools`
    - Filters: `none`
  - **-r io --rsubtype iopool -level pool**: List attributes for physical I/O pools
    - Attributes: `io_pool_id`, `lpar_ids`, `slots`
    - Filters: `{lpar_ids | lpar_names}`, `pools`
  - **-r io --rsubtype slot**: List attributes for physical I/O slots
    - Attributes: `unit_phys_loc`, `bus_id`, `phys_loc`, `drc_index`, `lpar_name`, `lpar_id`, `slot_io_pool_id`, `description`, `feature_codes`, `adapter_feature_codes`, `adapter_descriptions`, `pci_vendor_id`, `pci_device_id`, `pci_subs_vendor_id`, `pci_subs_device_id`, `pci_class`, `pci_revision_id`, `bus_grouping`, `iop`, `iop_info_stale`, `console_capable`, `load_source_capable`, `load_source_attached`, `alt_restart_device_capable`, `alt_restart_device_attached`, `op_console_capable`, `op_console_attached`, `twinas_capable`, `direct_console_capable`, `lan_console_capable`, `vpd_stale`, `vpd_type`, `vpd_model`, `vpd_serial_num`, `parent_slot_drc_index`, `drc_name`
    - Filters: `{lpar_ids | lpar_names}`, `units`, `buses`, `slots`, `pools`
  - **-r io --rsubtype slot -R**: List attributes for physical I/O slots
    - Attributes: `drc_index`, `lpar_name`, `lpar_id`, `drc_name`
    - Filters: `{lpar_ids | lpar_names}`, `units`, `buses`, `slots`, `pools`
  - **-r io --rsubtype taggedio**: List attributes for tagged I/O resources
    - Attributes: `lpar_name`, `lpar_id`, `load_source_slot`, `alt_restart_device_slot`, `recent_alt_restart_device_slot`, `console_slot`, `alt_console_slot`, `op_console_slot`
    - Filters: `{lpar_ids | lpar_names}`
  - **-r io --rsubtype unit**: List attributes for physical I/O units
    - Attributes: `unit_phys_loc`
    - Filters: `units`

- **-r hea --subtype phys --level sys**
  - Attributes: adapter\_id, state, phys\_loc
  - Filters: adapter\_ids
- **-r hea --subtype phys --level port\_group**
  - Attributes: adapter\_id, port\_group, phys\_port\_ids, unassigned\_logical\_port\_ids, curr\_port\_group\_mcs\_value, pend\_port\_group\_mcs\_value, valid\_port\_group\_mcs\_values
  - Filters: adapter\_ids, port\_groups
- **-r hea --subtype phys --level port**
  - Attributes: adapter\_id, port\_group, phys\_port\_id, phys\_port\_type, phys\_port\_state, conn\_state, curr\_conn\_speed, config\_conn\_speed, curr\_duplex, config\_duplex, trans\_flow\_control, rcv\_flow\_control, config\_flow\_control, max\_rcv\_packet\_size, promisc\_lpar\_id, promisc\_lpar\_name, logical\_port\_ids, phys\_port\_loc
  - Filters: adapter\_ids, port\_groups
- **-r hea --subtype phys --level port --stat**
  - Attributes: adapter\_id, port\_group, phys\_port\_id, rcv\_octets, rcv\_packets\_0\_64, rcv\_packets\_65\_127, rcv\_packets\_128\_255, rcv\_packets\_256\_511, rcv\_packets\_512\_1023, rcv\_packets\_1024\_max, rcv\_packets\_dropped\_bad\_FCS, rcv\_packets\_dropped\_internal\_mac\_error, rcv\_packets\_dropped\_in\_range\_length\_error, rcv\_packets\_dropped\_out\_of\_range\_length\_error, rcv\_packets\_dropped\_frame\_too\_long, rcv\_packets\_dropped\_jabber, rcv\_symbol\_error, rcv\_code\_error, rcv\_runt\_frame, rcv\_fragments, rcv\_unsupported\_opcode, rcv\_control\_pause\_frames, rcv\_search\_busy, rcv\_packets\_dropped\_filter, rcv\_packets\_dropped\_other, rcv\_packets\_dropped\_alignment, rcv\_MC\_packets, rcv\_BC\_packets, trans\_MC\_packets, trans\_BC\_packets, trans\_octets, trans\_packets\_length\_0\_64, trans\_packets\_length\_65\_127, trans\_packets\_length\_128\_255, trans\_packets\_length\_255\_511, trans\_packets\_length\_512\_1023, trans\_packets\_length\_1024\_max, trans\_packets\_dropped\_bad\_FCS, trans\_control\_pause\_frames, trans\_tx\_local\_fault\_packets, trans\_tx\_remote\_fault\_packets, trans\_tx\_packets\_dropped\_int\_MAC\_error, trans\_packets\_retried\_single\_collision, trans\_packets\_retried\_multiple\_collision, trans\_packets\_signal\_quality, trans\_packets\_deferred, trans\_packets\_late\_collisions, trans\_packets\_excessive\_collisions, trans\_packets\_no\_carrier, rcv\_overrun\_frames\_dropped
  - Filters: adapter\_ids, port\_groups
- **-r hea --subtype logical --level sys**
  - Attributes: adapter\_id, lpar\_id, lpar\_name, drc\_index, drc\_name, capabilities, ieq, nieq, qp, cq, mr
  - Filters: {lpar\_ids | lpar\_names}, adapter\_ids
- **-r hea --subtype logical --level port**
  - Attributes: adapter\_id, lpar\_id, lpar\_name, state, port\_group, phys\_port\_id, logical\_port\_id, drc\_index, drc\_name, mac\_addr, user\_def\_mac\_addr, vlan\_id\_list
  - Filters: {lpar\_ids | lpar\_names}, adapter\_ids, port\_groups

- **-r mem --level lpar**: List partition attributes for memory
  - Attributes: lpar\_name, lpar\_id, mem\_mode, auto\_io\_entitled\_mem, curr\_io\_entitled\_mem, curr\_min\_mem, curr\_mem, curr\_mem\_weight, curr\_max\_mem, pend\_io\_entitled\_mem, pend\_min\_mem, pend\_mem, pend\_mem\_weight, pend\_max\_mem, run\_io\_entitled\_mem, run\_mem\_weight, run\_min\_mem, run\_mem
  - Filters: {lpar\_ids | lpar\_names}
- **-r mem --level lpar -R**: List partition attributes for memory
  - Attributes: lpar\_name, lpar\_id, mem\_mode, pend\_mem\_weight, pend\_io\_entitled\_mem, run\_mem\_weight, run\_io\_entitled\_mempend\_mem, run\_mem
  - Filters: {lpar\_ids | lpar\_names}
- **-r mem --level lpar -stat**: List partition attributes for memory
  - Attributes: min\_io\_entitled\_mem, optimal\_io\_entitled\_mem, max\_io\_entitled\_mem\_usage
  - Filters: {lpar\_ids | lpar\_names}
- **-r mem --level sys [--maxmem]**: List system attributes for memory
  - Attributes: configurable\_sys\_mem, curr\_avail\_sys\_mem, pend\_avail\_sys\_mem, installed\_sys\_mem, deconfig\_sys\_mem, sys\_firmware\_mem, mem\_region\_size, pend\_mem\_region\_size, possible\_mem\_region\_size, max\_mem\_pools, max\_paging\_vios\_per\_mem\_pool
  - Additional attributes with --maxmem: required\_min\_mem\_aix\_linux, required\_min\_mem\_os400
  - Filters: None
- **-r mempool** : List system attributes for memory pools
  - Attributes: curr\_pool\_mem, pend\_pool\_mem, curr\_max\_pool\_mem, pend\_max\_pool\_mem, sys\_firmware\_pool\_mem, paging\_storage\_pool, paging\_vios\_names, paging\_vios\_ids
  - Filters: None
- **-r mempool -R**: List system attributes for recoverable memory pool resources
  - Attributes: curr\_pool\_mem, pend\_pool\_mem, curr\_max\_pool\_mem, pend\_max\_pool\_mem, sys\_firmware\_pool\_mem, paging\_storage\_pool, paging\_vios\_names, paging\_vios\_ids
  - Filters: None
- **-r mempool --subtype pgdev**: List paging devices associated with the memory pool
  - Attributes: device\_name, size, state, type, phys\_loc, storage\_pool, lpar\_name, lpar\_id, paging\_vios\_name, paging\_vios\_id, redundant\_state, is\_redundant, redundant\_device\_name, redundant\_paging\_vios\_id, redundant\_paging\_vios\_name, redundant\_phys\_loc
  - Filters: {lpar\_ids | lpar\_names}
- **-r proc --level lpar**: List partition attributes for processors
  - Attributes: lpar\_name, lpar\_id, curr\_shared\_proc\_pool\_id, curr\_proc\_mode, curr\_min\_proc\_units, curr\_proc\_units, curr\_max\_proc\_units, curr\_min\_procs, curr\_procs, curr\_max\_procs, curr\_sharing\_mode, curr\_uncap\_weight, pend\_shared\_proc\_pool\_id, pend\_proc\_mode, pend\_min\_proc\_units, pend\_proc\_units, pend\_max\_proc\_units, pend\_min\_procs, pend\_procs, pend\_max\_procs, pend\_sharing\_mode, pend\_uncap\_weight, run\_proc\_units, run\_procs, run\_uncap\_weight
  - Filters: {lpar\_ids | lpar\_names}
- **-r proc --level lpar -R**: List partition attributes for processors
  - Attributes: lpar\_name, lpar\_id, curr\_proc\_mode, curr\_sharing\_mode, pend\_proc\_mode, pend\_proc\_units, pend\_procs, pend\_sharing\_mode, pend\_uncap\_weight, run\_proc\_units, run\_procs, run\_uncap\_weight
  - Filters: {lpar\_ids | lpar\_names}

- **-r proc --level pool:** List shared processor pool attributes
  - Attributes: shared\_proc\_pool\_id, configurable\_pool\_proc\_units, curr\_avail\_pool\_proc\_units, pend\_avail\_pool\_proc\_units
  - Filters: None
- **-r proc --level sys:** List system attributes for processors
  - Attributes: configurable\_sys\_proc\_units, curr\_avail\_sys\_proc\_units, pend\_avail\_sys\_proc\_units, installed\_sys\_proc\_units, deconfig\_sys\_proc\_units, min\_proc\_units\_per\_virtual\_proc, max\_shared\_proc\_pools, max\_virtual\_procs\_per\_lpar, max\_procs\_per\_lpar
  - Filters: None
- **-r virtualio --subtype eth --level lpar:** List virtual Ethernet adapter attributes
  - Attributes: lpar\_name, lpar\_id, slot\_num, state, ieee\_virtual\_eth, port\_vlan\_id, addl\_vlan\_ids, is\_trunk, trunk\_priority, is\_required, mac\_addr
  - Filters: {lpar\_ids | lpar\_names}, vlans, slots
- **-r virtualio --subtype eth --level sys:** List system attributes for virtual Ethernet adapters
  - Attributes: max\_vlans\_per\_port, mac\_prefix
  - Filters: None
- **-r virtualio --subtype eth --level lpar -F device\_name,drc\_name,shared\_adapter,backing\_device:** List virtual Ethernet adapter attributes
 

**Note:** The shared\_adapter and backing\_device are empty for all logical partitions other than the Virtual I/O Server. These attributes might be empty on the Virtual I/O Server if the Integrated Virtualization Manager cannot retrieve the information or if no such adapter exists (shared\_adapter) or if no such shared Ethernet adapter (backing\_device) exists.

  - Attributes: lpar\_name, lpar\_id, slot\_num, state, ieee\_virtual\_eth, port\_vlan\_id, addl\_vlan\_ids, is\_trunk, trunk\_priority, is\_required, mac\_addr
  - Filters: {lpar\_ids | lpar\_names}, vlans, slots
- **-r virtualio --subtype fc --level lpar:** List information about virtual fibre channel adapters for each logical partition
  - Attributes: lpar\_name, lpar\_id, slot\_num, adapter\_type, remote\_lpar\_id, remote\_lpar\_name, remote\_slot\_num, is\_required, wwpns, state
  - Filters: {lpar\_ids | lpar\_names}, slots
- **-r virtualio --subtype fc --level sys:** List system attributes for virtual fibre channel adapters
  - Attributes: num\_wwpns\_remaining, wwpn\_prefix
  - Filters: None
- **-r virtualio --subtype fc --level lpar -F topology:** List information about virtual fibre channel adapters for each logical partition in the form of a slash delimited list (for example, vd\_name/status/fc\_client\_name/fc\_server\_name/phys\_port\_name)
  - Attributes: lpar\_name, lpar\_id, slot\_num, adapter\_type, remote\_lpar\_id, remote\_lpar\_name, remote\_slot\_num, is\_required, wwpns, state
  - Filters: {lpar\_ids | lpar\_names}, slots
- **-r virtualio --subtype hsl:** List system attributes for high-speed link
  - Attributes: hsl\_pool\_id, lpar\_names, lpar\_ids
  - Filters: {lpar\_ids | lpar\_names}, pools

- **-r virtualio --subtype scsi --level lpar**: List virtual SCSI adapter attributes
  - Attributes: lpar\_name, lpar\_id, slot\_num, state, adapter\_type, remote\_lpar\_id, remote\_lpar\_name, remote\_slot\_num, is\_required
  - Filters: {lpar\_ids | lpar\_names}, slots
- **-r virtualio --subtype scsi --level lpar -F topology**: List virtual SCSI adapter attributes in the form of a slash delimited list (for example, vd\_name/status/vscsi\_client\_name/bd\_sp\_name/bd\_name)
 

**Note:** The bd\_sp\_name might be empty if the Virtual I/O Server uses physical volumes rather than a storage pool.

  - Attributes: vd\_name, vscsi\_client\_name, vscsi\_server\_name, bd\_sp\_name, bd\_name
  - Filters: {lpar\_ids | lpar\_names}, slots
- **-r virtualio --subtype slot --level lpar**: List partition attributes for virtual I/O slots
  - Attributes: lpar\_name, lpar\_id, curr\_max\_virtual\_slots, pend\_max\_virtual\_slots
  - Filters: {lpar\_ids | lpar\_names}
- **-r virtualio --subtype slot --level slot**: List virtual I/O slot attributes
  - Attributes: slot\_num, lpar\_name, lpar\_id, config, state, drc\_name
  - Filters: {lpar\_ids | lpar\_names}, slots
- - **-r virtualio --subtype slot --level slot -R**: List virtual I/O slot attributes
    - Attributes: slot\_num, lpar\_name, lpar\_id, drc\_name
    - Filters: {lpar\_ids | lpar\_names}, slots
  - **-r virtualio --subtype virtualopti**: List virtual OptiConnect resource attributes
    - Attributes: virtual\_opti\_pool\_id, lpar\_names, lpar\_ids
    - Filters: {lpar\_ids | lpar\_names}, pools

**--maxmem**

*MaximumMemory*

When this option is specified, the required minimum memory amount needed for partitions to support the maximum memory quantity specified is listed. All memory quantities are in megabytes, and are a multiple of the memory region size for the managed- system.

This information is useful for specifying memory amounts in partition profiles.

**-m ManagedSystem**

This option is only valid when listing system level memory resources.

The name of the managed system. This attribute is optional because there is only one system to manage. The name may either be the user-defined name for the managed system, or be in the form tttt-mmm\*sssssss, where tttt is the machine type, mmm is the model, and sssssss is the serial number of the managed system.

**--filter** *FilterData*

The filters to apply to the resources to be listed. Filters are used to select which resources of the specified resource type are to be listed. If no filters are used, then all of the resources of the specified resource type will be listed. For example, specific partitions can be listed by using a filter to specify the names or IDs of the partitions to list. Otherwise, if no filter is used, then all the partitions in the managed system will be listed.

The filter data consists of filter name/value pairs, which are in comma separated value (CSV) format. The filter data must be enclosed in double quotation marks.

The format of the filter data is as follows:

```
"filter-name=value,filter-name=value,..."
```

Note that certain filters accept a comma separated list of values, as follows:

```
"filter-name=value,value,...",..."
```

When a list of values is specified, the filter name/value pair must be enclosed in double quotation marks. Depending on the shell being used, nested double quotation marks characters might need to be preceded by an escape character, which is usually a backslash character.

Unless otherwise indicated, multiple values can be specified for each filter.

**Valid filter names:**

**adapter\_ids**

The DRC index of the Host Ethernet Adapter in hexadecimal

**buses** The bus ID of the I/O bus to view

**lpar\_ids | lpar\_names**

Name or ID of the logical partition to view

**Note:** The filter value of none for lpar\_ids is supported and displays only paging devices that are not currently assigned to any logical partition. If the memory pool does not exist, and you use the `-r mempool --rsubtype pgdev` attribute, an error is displayed.

**pools** The pool ID of the I/O pool to view

**port\_groups**

The Host Ethernet Adapter port group or groups

**slots** For physical I/O slots, the DRC index of the slot to view. For virtual I/O slots, the virtual slot number of the slot to view

**units** The physical location code of the unit to view

**v lans** The virtual LAN of the virtual Ethernet adapters to view

**-F** *AttributeNames*

A delimiter separated list of attribute names for the desired attribute values to be displayed for each resource. If no attribute names are specified, then values for all of the attributes for the resource will be displayed.

When this option is specified, only attribute values will be displayed. No attribute names will be displayed. The attribute values displayed will be separated by the delimiter which was specified with this option.

This option is useful when only attribute values are desired to be displayed, or when the values of only selected attributes are desired to be displayed.

**Attribute names:**

**adapter\_descriptions**

Indicates a description of the adapter.

**adapter\_feature\_codes**

List of possible feature codes for the I/O adapter sorted by likely match. Each element in the list consists of two feature codes - one for each operating system. This list follows the following format:

```
aix_feature_code1/linux_feature_code1,

/i5_feature_code1,aix_feature_code2/

linux_feature_code2/i5_feature_code2,...
```

**adapter\_id**

Indicates the DRC index of the Host Ethernet Adapter in hexadecimal.

**adapter\_type**

Indicates whether the virtual SCSI, virtual fibre channel, or serial adapter is a client or server. Valid values are client and server.

**addl\_vlan\_ids**

The list of additional IEEE 802.1Q virtual LANs (VLANs) on a virtual Ethernet adapter.

**auto\_io\_entitled\_mem**

Indicates whether the I/O entitled memory of the partition is automatically managed. Valid values follow:

- 0 (Not managed/custom)
- 1 (Automatically managed)

This attribute is only displayed by default if the `mem_mode` attribute is shared. If the value is dedicated and the attribute is requested, the value is null.

**auto\_mem\_region\_size**

Indicates the optimal size of the memory region, as automatically calculated by the firmware.

**Note:** The memory region size applies to the entire system and requires you to reboot the entire system.

**backing\_device**

The name of the physical Ethernet adapter that is connected to the shared Ethernet adapter on the virtual I/O server (VIOS). For example, `ent0`.

**backplane\_phys\_loc**

Physical location code of the backplane on which the bus resides.

**bus\_drc\_index**

DRC index, in hex, of the bus.

**bus\_drc\_name**

DRC name of the bus.

**bus\_grouping**

Indicates whether bus grouping is required. Possible values:

- 0 - not required
- 1 - required

**bus\_id** I/O bus unique ID.



**capabilities**

Values currently applied to the system.

**config** Virtual slot configuration state. Possible values follow:

- empty - no adapter
- ethernet - virtual Ethernet adapter
- fc - virtual fibre channel adapter
- reserved - reserved slot
- scsi - virtual SCSI adapter
- serial - virtual serial adapter
- vasi - virtual asynchronous services interface
- vmc - virtual management channel adapter

**config\_conn\_speed**

The configured connection speed of the port in megabits per second. This value can be set by the user using **chhwres**. Possible values are:

- auto
- 10
- 100
- 1000
- 10000

**config\_duplex**

Configured duplex value of the port. This value can be set by the user using **chhwres**. Possible values are:

- auto
- half
- full

**config\_flow\_control**

Configured flow control value of the port. This value is used for both the receive and transmit flow control. This attribute can also be set using the **chhwres** command. Possible values are:

- 1 (enabled)
- 0 (disabled)

**configurable\_pool\_proc\_units**

Total number of configurable processing units in the shared processing pool.

**configurable\_sys\_mem**

Total amount, in megabytes, of configurable memory on the managed system.

**configurable\_sys\_proc\_units**

Total number of configurable processing units on the managed system.

**conn\_state**

Connection state or link state of the physical port. Possible values are:

- 1 (up)
- 0 (down)
- unavailable

**cq** Total number of completion queues.

**curr\_avail\_pool\_proc\_units**

Current number of configurable processing units in the shared processing pool that are not assigned to partitions.

**curr\_avail\_sys\_mem**

Current amount, in megabytes, of configurable memory on the managed system that is not assigned to partitions.

**curr\_avail\_sys\_proc\_units**

Current number of configurable processing units on the managed system that are not assigned to partitions.

**curr\_conn\_speed**

Speed of the port in Mbps. Possible values are:

- 10
- 100
- 1000
- 10000
- unavailable

**curr\_duplex**

Duplex value of the port. Valid values are:

- half
- full
- unavailable

**curr\_io\_entitled\_mem**

The current I/O entitled memory of the partition in megabytes. This attribute is only displayed, by default, if the value of the mem\_mode attribute is Shared. If the value is Dedicated and the attribute is requested, the value is null.

**curr\_max\_mem**

Maximum amount of memory, in megabytes, that can be dynamically assigned to the partition.

**curr\_max\_pool\_mem**

Maximum amount of physical memory that can be assigned to the memory pool without taking the pool offline. This value is in megabytes.

**curr\_max\_proc\_units**

Maximum number of processing units that can be dynamically assigned to the partition. This attribute is only valid for partitions using shared processors.

**curr\_max\_procs**

Maximum number of processors or virtual processors that can be dynamically assigned to the partition.

**curr\_max\_virtual\_slots**

Maximum number of virtual slots that can be dynamically configured for the partition.

**curr\_mem**

Current amount of memory, in megabytes, assigned to the partition.

**curr\_mem\_weight**

The current shared memory weight of the partition. Possible values are 0 through 255. This attribute is only displayed by default if the mem\_mode attribute is set to the shared value. If the value is dedicated and the attribute is requested, the value is null.

**curr\_min\_mem**

Minimum amount of memory, in megabytes, that can be dynamically assigned to the partition.

**curr\_min\_proc\_units**

Minimum number of processing units that can be dynamically assigned to the partition. This attribute is only valid for partitions using shared processors.

**curr\_min\_procs**

Minimum number of processors or virtual processors that can be dynamically assigned to the partition.

**curr\_pool\_mem**

The amount of physical memory currently assigned to the memory pool in megabytes.

**curr\_port\_group\_mcs\_value**

The current MCS value for the port group.

**curr\_proc\_mode**

Indicates whether the partition is using dedicated or shared processors. The mode cannot change dynamically. Valid values:

- ded - dedicated
- shared - shared

**curr\_proc\_units**

The current number of processing units assigned to the partition. This attribute is only valid for partitions using shared processors.

**curr\_procs**

The current number of processors or virtual processors assigned to the partition.

**curr\_shared\_proc\_pool\_id**

The shared processor pool that this partition is currently participating in. This attribute is only valid for partitions using shared processors.

**curr\_sharing\_mode**

The current sharing mode for a partition. Valid values:

- keep\_idle\_procs - valid with dedicated processor mode
- share\_idle\_procs - valid with dedicated processor mode
- share\_idle\_procs\_always - valid with dedicated processor mode
- share\_idle\_procs\_active - valid with dedicated processor mode
- cap - capped mode. valid with shared processor mode
- uncap - uncapped mode. valid with shared processor mode

**curr\_uncap\_weight**

The current weighted average of processing priority when in uncapped sharing mode. Valid values are 0 - 255.

**deconfig\_sys\_mem**

The amount of memory, in megabytes, on the managed system that has been unconfigured. This includes memory that has been unconfigured by the system due to hardware failure, and memory that has been manually unconfigured.

**deconfig\_sys\_proc\_units**

The number of processing units on the managed system that have been unconfigured. This includes processing units that have been unconfigured by the system due to hardware failure, and processing units that have been manually unconfigured.

**description**

A description of the I/O adapter which is in the slot.

**device\_name**

The Virtual I/O Server device used as a paging space device in the memory pool.

**drc\_index**

The DRC index, in hexadecimal, of the I/O slot.

**drc\_name**

The DRC name of the I/O slot.

**feature\_codes**

The most likely feature code for the I/O adapter. The feature code may be different for each operating system. The **adapter\_feature\_codes** attribute should be used to distinguish the feature code specific to each operating system.

**ieee\_virtual\_eth**

Indicates whether the virtual Ethernet adapter is using IEEE 802.1Q. Valid values follow:

- 0 - No
- 1 - Yes

**ieq** Total number of interruptible event queues.

**installed\_sys\_mem**

Total amount, in megabytes, of memory installed on the managed system.

**installed\_sys\_proc\_units**

Total number of processing units installed on the managed system.

**io\_pool\_id**

A unique identifier for an I/O pool. I/O pools are used in some I/O clustering environments.

**iop** Indicates whether the I/O adapter is an I/O processor. Valid values:

- 0 - no
- 1 - yes

**iop\_info\_stale**

Indicates whether the information about the I/O processor is stale. Valid values:

- 0 - no
- 1 - yes

**is\_required**

Indicates whether the I/O slot or virtual I/O adapter is required for the partition.  
Valid values:

- 0 - no
- 1 - yes

**is\_trunk**

Indicates whether the virtual Ethernet adapter is the trunk or uplink adapter for the virtual LAN. Valid values:

- 0 - no
- 1 - yes

**load\_source\_attached**

Whether this IOP has a load source device. Valid values:

- 0 - not attached
- 1 - attached

**logical\_port\_id**

The ID of the logical port.

**logical\_port\_ids**

Logical port IDs defined on the physical port.

**lpar\_id** Partition ID. Valid values are 1 through the maximum number of partitions supported on the managed system (max\_lpars).

When using the --subtype pgdev flag, this attribute represents the partition ID of the shared memory partition to which this paging device is associated. If no partition is associated, then the value is none.

**lpar\_ids**

List of logical partition IDs of logical partitions using a pool.

**lpar\_name**

User-defined name for the partition.

When using the --subtype pgdev flag, this attribute represents the name of the shared memory partition to which this paging device is associated. If no partition is associated, this attribute is not displayed by default. If no partition is associated but the attribute is requested, the attribute will have a blank value.

**lpar\_names**

List of names of logical partitions using a pool.

**mac\_addr**

MAC address for the virtual Ethernet adapter.

**mac\_prefix**

The first 3 bytes of the MAC address to be assigned to all virtual Ethernet adapters for this managed system. This must be specified as a 3 byte hexadecimal value (for example, 32ab10) and can only be configured in the **mkgenfcfg** command.

**max\_io\_entitled\_mem\_usage**

The highest level of I/O entitled memory used by the specified logical partition since the value was last reset.

- max\_io\_pools**  
Maximum number of I/O pools supported on the managed system.
- max\_mem\_pools**  
Maximum number of memory pools supported. If the system does not support memory pools, then the value is 0; otherwise, it is 1.
- max\_paging\_vios\_per\_mem\_pool**  
Maximum number of paging Virtual I/O Server (VIOS) partitions that can be assigned to a memory pool. If the system supports memory pools, this value is 1.
- max\_procs\_per\_lpar**  
Indicates the largest number of dedicated processors that you can assign as the maximum processors value for a logical partition.
- max\_rcv\_packet\_size**  
Maximum Transmission Unit (MTU) size of the port. This value can also be changed using the **chhwres** command. Possible values:
- 1500 (default)
  - 9000 (jumbo)
- max\_shared\_proc\_pools**  
Maximum number of shared processing pools which are supported on the managed system.
- max\_virtual\_procs\_per\_lpar**  
Indicates the largest number of virtual processors that can be assigned as the maximum processors value for a logical partition.
- max\_vlans\_per\_port**  
Indicates the largest number of VLANs that can be specified as additional IEEE 802.1Q VLANs on a virtual Ethernet adapter.
- mem\_mode**  
Partition memory mode.
- **ded**: Dedicated memory
  - **shared**: Shared memory
- mem\_region\_size**  
The memory region size, in megabytes, for the managed system. The memory regions size represents the granularity of memory allocation to partitions.
- min\_io\_entitled\_mem**  
The minimum I/O entitled memory required for the logical partition to function with the current I/O configuration.
- min\_proc\_units\_per\_virtual\_proc**  
Minimum number of processing units that are required for each virtual processor assigned to a partition.
- mr** Total number of memory regions.
- nieq** Total number of non-interruptible event queues.
- num\_wwpns\_remaining**  
The number of worldwide port names that remain to be allocated on the managed system.
- optimal\_io\_entitled\_mem**  
The amount of I/O entitled memory that would satisfy the requested configuration of all I/O devices.

**paging\_storage\_pool**

The name of default paging storage pool. This attribute is the storage pool from which new paging devices are created. If there is no default paging storage pool, the value is an empty string.

**paging\_vios\_id**

The partition ID of the paging VIOS partition.

**paging\_vios\_ids**

A comma-separated list of partition IDs for paging VIOS partitions associated with this memory pool. This attribute is not displayed by default. If requested, the value displayed is 1.

**paging\_vios\_name**

The name of the paging VIOS partition.

**paging\_vios\_names**

A comma-separated list of partition names for paging VIOS partitions that are associated with this memory pool. This attribute is not displayed by default. If requested, the value displayed is the name of the Virtual I/O Server logical partition.

**parent\_slot\_drc\_index**

The DRC index, in hex, of the parent slot. A value of none indicates there is no parent slot.

**pci\_class**

PCI class code for the I/O adapter. This value is displayed in hexadecimal.

**pci\_device\_id**

ID of the device for this I/O adapter. This value is displayed in hexadecimal.

**pci\_revision\_id**

Vendor-assigned code indicating the revision number of the I/O adapter. This value is displayed in hexadecimal.

**pci\_subs\_device\_id**

Vendor-assigned code identifying the type of driver that is installed for the I/O adapter. This value is displayed in hexadecimal.

**pci\_subs\_vendor\_id**

ID of the vendor that made the driver for the I/O adapter. This value is displayed in hexadecimal.

**pci\_vendor\_id**

ID of the vendor that made this I/O adapter. This value is displayed in hexadecimal.

**pend\_avail\_pool\_proc\_units**

After system restart, the number of configurable processing units in the shared processing pool that are not assigned to partitions.

**pend\_avail\_sys\_mem**

After system restart, the amount, in megabytes, of configurable memory on the managed system that is not assigned to partitions.

**pend\_avail\_sys\_proc\_units**

After system restart, the number of configurable processing units on the managed system that are not assigned to partitions.

**pend\_io\_entitled\_mem**

The pending I/O entitled memory of the partition in megabytes. This attribute is only displayed by default if the mem\_mode attribute has a value of shared. If the value is dedicated and the attribute is requested, the value is null.

**pend\_max\_mem**

After partition restart, the maximum amount of memory, in megabytes, that can be dynamically assigned to the partition.

**pend\_max\_pool\_mem**

The pending value for the maximum amount of physical memory that can be assigned to the memory pool without taking the memory pool offline. This value is in megabytes.

**pend\_max\_proc\_units**

After partition restart, the maximum number of processing units that can be dynamically assigned to the partition. This attribute is only valid for partitions using shared processors.

**pend\_max\_procs**

After partition restart, the maximum number of processors or virtual processors that can be dynamically assigned to the partition.

**pend\_max\_virtual\_slots**

After partition restart, the maximum number of virtual slots that can be dynamically created for the partition.

**pend\_mem**

The target amount of memory, in megabytes, assigned to the partition.

**pend\_mem\_region\_size**

This is the memory region size of the system after the system is restarted.

**pend\_mem\_weight**

The pending shared memory weight of the partition. Possible values are 0 - 255. This attribute is only displayed by default if the mem\_mode attribute is set to the shared value. If the value is dedicated and the attribute is requested, the value is null.

**pend\_min\_mem**

After partition restart, the minimum amount of memory, in megabytes, that can be dynamically assigned to the partition.

**pend\_min\_proc\_units**

After partition restart, the minimum number of processing units that can be dynamically assigned to the partition. This attribute is only valid for partitions using shared processors.

**pend\_min\_procs**

After partition restart, the minimum number of processors or virtual processors that can be dynamically assigned to the partition.

**pend\_pool\_mem**

The amount of physical memory to be assigned to the memory pool in megabytes.

**pend\_port\_group\_mcs\_value**

Pending MCS value of the port group. Setting the MCS value may require a system restart.

**pend\_proc\_mode**

Indicates whether the partition will be using dedicated or shared processors after restart. Valid values:

- **ded:** Dedicated
- **shared:** Shared



**pend\_proc\_units**

The target number of processing units assigned to the partition. This attribute is only valid for partitions using shared processors.

**pend\_procs**

The target number of processors or virtual processors assigned to the partition.

**pend\_shared\_proc\_pool\_id**

The shared processor pool that a partition is will participate in after restart. This attribute is only valid for partitions using shared processors.

**pend\_sharing\_mode**

The target sharing mode for a partition. Valid values:

- keep\_idle\_procs: Valid with pending dedicated processor mode
- share\_idle\_procs: Valid with pending dedicated processor mode
- cap - capped mode: Valid with pending shared processor mode
- uncap - uncapped mode: Valid with pending shared processor mode

**pend\_uncap\_weight**

The target weighted average of processing priority when in uncapped sharing mode. Valid values are: 0 - 255.

**phys\_loc**

If the device has a physical location code for the slot, this attribute is displayed by default with the physical location code as the value. If the device does not have a physical location code, this attribute is not displayed by default and has a blank value.

**phys\_port\_id**

The ID of the physical port in the port group (0 or 1).

**phys\_port\_ids**

List of host Ethernet adapter (HEA) physical port IDs in a port group.

**phys\_port\_loc**

Location code suffix for the physical port.

**phys\_port\_state**

State of the physical port.

**phys\_port\_type**

Speed of the port.

**port\_group**

Indicates the Host Ethernet Adapter port group(s).

**port\_vlan\_id**

Port virtual LAN ID for the virtual ethernet adapter.

**possible\_mem\_region\_size**

The memory region size that is possible of the system.

**promisc\_lpar\_id**

ID assigned to the promiscuous logical partition. Possible values:

- none
- 1 - 254

**promisc\_lpar\_name**

Name assigned to the promiscuous logical partition.

**qp** The total number of queue pairs.

**recv\_flow\_control**

Receive flow control value of the port

- 1 (enabled)
- 0 (disabled)
- unavailable

**recv\_octets, recv\_packets\_0\_64, recv\_packets\_65\_127, recv\_packets\_128\_255, recv\_packets\_256\_511, recv\_packets\_512\_1023, recv\_packets\_1024\_max, recv\_packets\_dropped\_bad\_FCS, recv\_packets\_dropped\_internal\_mac\_error, recv\_packets\_dropped\_in\_range\_length\_error, recv\_packets\_dropped\_out\_of\_range\_length\_error, recv\_packets\_dropped\_frame\_too\_long, recv\_packets\_dropped\_jabber, recv\_symbol\_error, recv\_code\_error, recv\_runt\_frame, recv\_fragments, recv\_unsupported\_opcode, recv\_control\_pause\_frames, recv\_search\_busy, recv\_packets\_dropped\_filter, recv\_packets\_dropped\_other, recv\_packets\_dropped\_alignment, default\_unicast\_QPN, recv\_MC\_packets, recv\_BC\_packets, trans\_MC\_packets, trans\_BC\_packets, trans\_octets, trans\_packets\_length\_0\_64, trans\_packets\_length\_65\_127, trans\_packets\_length\_128\_255, trans\_packets\_length\_255\_511, trans\_packets\_length\_512\_1023, trans\_packets\_length\_1024\_max, trans\_packets\_dropped\_bad\_FCS, trans\_control\_pause\_frames, trans\_tx\_local\_fault\_packets, trans\_tx\_remote\_fault\_packets, trans\_tx\_packets\_dropped\_int\_MAC\_error, trans\_packets\_retried\_single\_collision, trans\_packets\_retried\_multiple\_collision, trans\_packets\_signal\_quality, trans\_packets\_deferred, trans\_packets\_late\_collisions, trans\_packets\_excessive\_collisions, trans\_packets\_no\_carrier, recv\_overrun\_frames\_dropped**

Performance metric values

**remote\_lpar\_id**

For client adapters, this specifies the ID of the partition which has the virtual serial/SCSI server adapter for this adapter. For server adapters, this specifies the ID of the partition which has the only client virtual serial/SCSI adapter allowed to connect to this adapter. A value of any indicates that any client virtual serial/SCSI adapter should be allowed to connect to this adapter.

**remote\_lpar\_name**

The partition name which corresponds to the `remote_lpar_id`.

**remote\_slot\_num**

For client adapters, this specifies the virtual slot number of the virtual serial/SCSI server adapter for this adapter. For server adapters, this specifies the virtual slot number of the only client virtual serial/SCSI adapter allowed to connect to this adapter. A value of any indicates that any client virtual serial/SCSI adapter should be allowed to connect to this adapter.

**required\_min\_mem\_aix\_linux**

For an AIX or Linux partition, this is the required minimum memory amount, in megabytes, for the maximum memory amount specified with the `--maxmem` parameter.

**required\_min\_mem\_os400**

The required minimum memory amount, in megabytes, for the maximum memory amount specified with the **--maxmem** parameter. This attribute only applies for an logical partition.

**run\_io\_entitled\_mem**

The I/O entitled memory of the partition at run time in megabytes. This attribute is only displayed by default if the **mem\_mode** attribute has a value of shared. If the value is dedicated and the attribute is requested, the value is null.

**run\_mem**

Current amount of memory, in megabytes, that the partition has varied on.

**run\_mem\_weight**

The runtime shared memory weight of the logical partition. Possible values are 0 - 255. This attribute is only displayed by default if the **mem\_mode** attribute is set to the shared value. If the value is dedicated and the attribute is requested, the value is null.

**run\_min\_mem**

The amount of memory, in megabytes, returned by a running partition's operating system to represent the smallest memory assignment which is currently supported by the operating system. If the operating system does not support this feature, then the value is 0.

**run\_proc\_units**

Number of processing units that are varied on for the partition.

**run\_procs**

Number of processors or virtual processors that are varied on for the partition.

**run\_uncap\_weight**

The operating system's current setting for the weighted average of processing priority when in uncapped sharing mode. Valid values are: 0 - 255.

**shared\_adapter**

The name of the shared Ethernet adapter on the VIOS. For example, *ent9*.

**shared\_proc\_pool\_id**

A unique decimal identifier for a shared processing pool.

**size**

The size of the paging device in megabytes.

**slot\_io\_pool\_id**

A unique decimal identifier of the I/O pool to which the slot is assigned. A value of none indicates that the slot is not assigned to any I/O pools.

**slot\_num**

Virtual slot number.

- slots** A list of slot DRC indices.
- state** The state of a virtual slot. A virtual slot must be able to transition into the off state before a dynamic reconfiguration of the slot can succeed. Valid states follow:
- 0 - The virtual slot is ready for dynamic reconfiguration.
  - 1 - The virtual slot is not yet ready for dynamic reconfiguration.
- When using paging devices, state indicates the state of the paging device in megabytes. Valid states follow:
- Active - The paging device is in use.
  - Inactive - The paging device is available.
  - Missing backing device - The backing device being used as a paging space device is missing.
  - Defined - The backing device being used as a paging space device is defined.
- storage\_pool**  
If the device comes from a storage pool, this attribute is displayed by default with the storage pool as the value. If the device does not come from a storage pool, this attribute is not displayed by default and has a blank value.
- supports\_hmc**  
Indicates whether or not the virtual serial adapter is enabled for HMC client support. Valid values:
- 0 - no
  - 1 - yes
- sys\_firmware\_mem**  
The amount of memory, in megabytes, on the managed system that is being used by system firmware.
- sys\_firmware\_pool\_mem**  
The amount of memory in the shared memory pool that is reserved for firmware use. This is the greater of the current and pending values.
- trans\_flow\_control**  
Transmit flow control value of the port. Possible values are:
- 1 (enabled)
  - 0 (disabled)
  - unavailable
- trunk\_priority**  
The valid values are integers between 1 and 15, inclusive. This value is required for a trunk adapter.
- unassigned\_logical\_port\_ids**  
List of logical port IDs in the port group that are not assigned to a partition.
- unit\_phys\_loc**  
Physical location code of the I/O unit.
- user\_def\_mac\_addr**  
User defined MAC address for the logical port.
- valid\_port\_group\_mcs\_values**  
Valid MCS values.

|                       |                                                                                                                                                                                                                                                              |
|-----------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>vlan_id_list</b>   | List of virtual LAN IDs that the logical port has access to.                                                                                                                                                                                                 |
| <b>vpd_model</b>      | Model of the I/O adapter.                                                                                                                                                                                                                                    |
| <b>vpd_serial_num</b> | Serial number of the I/O adapter.                                                                                                                                                                                                                            |
| <b>vpd_stale</b>      | Indicates whether or not type, model, and serial number of the adapter is not current. Valid values: <ul style="list-style-type: none"> <li>• 0 - not current</li> <li>• 1 - current</li> </ul>                                                              |
| <b>vpd_type</b>       | Type of the I/O adapter.                                                                                                                                                                                                                                     |
| <b>wwpns</b>          | If the adapter is of type client, the worldwide port names assigned to this adapter are displayed as a comma-separated list. If the adapter is of type server, by default, the attribute is not displayed. If the attribute is requested, the value is null. |
| <b>wwpn_prefix</b>    | A 12-character hexadecimal string that is used by the managed system to generate worldwide port names.                                                                                                                                                       |
| <b>--header</b>       | Display a header record, which is a delimiter separated list of attribute names for the attribute values that will be displayed. This header record will be the first record displayed. This option is only valid when used with the <b>-F</b> option.       |
| <b>-R</b>             | List information for partitions with resources that can be restored due to a dynamic LPAR operation failure.                                                                                                                                                 |

## Exit Status

This command has a return code of 0 on success.

## Security

This command is accessible by all users.

## Examples

1. List system level memory information and include the minimum memory required to support a maximum of 1024 MB:

```
lshwres -r mem --level sys --maxmem 1024
```
2. List all memory information for partitions lpar1 and lpar2, and only display attribute values, following a header of attribute names:

```
lshwres -r mem --level lpar --filter "\lpar_names=lpar1,lpar2\" -F --header
```
3. List all I/O units on the system:

```
lshwres -r io --rsubtype unit
```
4. List all virtual Ethernet adapters on the managed system:

```
lshwres -r virtualio --rsubtype eth --level lpar
```
5. List all virtual slots for partition lpar1:

```
lshwres -r virtualio --rsubtype slot --level slot --filter "lpar_names=lpar1"
```
6. List only the installed and configurable processors on the system:

```
lshwres -r proc --level sys -F installed_sys_proc_units,
configurable_sys_proc_units
```

7. List all physical Host Ethernet Adapter ports on the system:  
lshwres -r hea --subtype phys --level port
8. List all logical Host Ethernet Adapter ports on the system:  
lshwres -r hea --subtype logical --level port
9. List the physical Host Ethernet Adapter adapters on the system:  
lshwres -r hea --subtype phys --level sys
10. List the Host Ethernet Adapter port groups on the system:  
lshwres -r hea --subtype phys --level port\_group

## Related Information

The `lssyscfg` and `chhwres` commands.

---

## IVM Isled command

### Purpose

List states for physical and virtual LEDs. This command is operable only in an Integrated Virtualization Manager environment.

### Syntax

To list physical platform LEDs:

```
Isled -r sa -t phys [-F "AttributeNames"] [--header] [-m ManagedSystem]
```

To list virtual partition LEDs:

```
Isled -r sa -t virtualpar [--filter "FilterData"] [-F "AttributeNames"] [--header] [-m ManagedSystem]
```

To list virtual platform LEDs:

```
Isled -r sa -t virtualsys [-F "AttributeNames"] [--header] [-m ManagedSystem]
```

### Description

The `Isled` command lists the states of physical and virtual LEDs.

### Flags

- |                                      |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|--------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>-r</b> <i>ResourceType</i>        | The type of LED resources to list. The only valid value is sa for System Attention (SA) LEDs.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>-t</b> <i>SystemAttentionType</i> | The type of System Attention (SA) LEDs to list. <ul style="list-style-type: none"> <li>• <b>-r sa -t virtualpar</b>: List attributes for virtual partition system attention LEDs <ul style="list-style-type: none"> <li>– Attributes: lpar_id, lpar_name, state</li> <li>– Filters: { lpar_ids   lpar_names }</li> </ul> </li> <li>• <b>-r sa -t virtualsys</b>: List attributes for virtual system attention LEDs <ul style="list-style-type: none"> <li>– Attributes: state</li> <li>– Filters: None</li> </ul> </li> <li>• <b>-r sa -t phys</b>: List attributes for physical system attention LEDs <ul style="list-style-type: none"> <li>– Attributes: state</li> <li>– Filters: None</li> </ul> </li> </ul> |

|                                   |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|-----------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>-m</b> <i>ManagedSystem</i>    | The name of the managed system. This attribute is optional because there is only one system to manage. The name might either be the user-defined name for the managed system, or be in the form <code>ttt-mmm*sssssss</code> , where <code>ttt</code> is the machine type, <code>mmm</code> is the model, and <code>sssssss</code> is the serial number of the managed system.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>--filter</b> <i>FilterData</i> | <p>The filters to apply to the resources to be listed. Filters are used to select which resources of the specified resource type are to be listed. If no filters are used, then all of the resources of the specified resource type will be listed. For example, specific partitions can be listed by using a filter to specify the names or IDs of the partitions to list. Otherwise, if no filter is used, then all the partitions in the managed system will be listed.</p> <p>The filter data consists of filter name/value pairs, which are in comma separated value (CSV) format. The filter data must be enclosed in double quotation marks.</p> <p>The format of the filter data is as follows:<br/> <code>"filter-name=value,filter-name=value,..."</code></p> <p>Note that certain filters accept a comma separated list of values, as follows:<br/> <code>"filter-name=value,value,...",..."</code></p> <p>When a list of values is specified, the filter name/value pair must be enclosed in double quotation marks. Depending on the shell being used, nested double quotation marks might need to be preceded by an escape character, which is usually a <code>'\'</code> character.</p> <p>Unless otherwise indicated, multiple values can be specified for each filter.</p> <p><b>Valid filter names for -r lpar</b></p> <p><b>lpar_ids</b><br/> ID of the partitions to view</p> <p><b>lpar_names</b><br/> Name of the partitions to view</p> |
| <b>-F</b> <i>AttributeNames</i>   | <p>A delimiter separated list of attribute names for the desired attribute values to be displayed for each resource. If no attribute names are specified, then values for all of the attributes for the resource will be displayed.</p> <p>When this option is specified, only attribute values will be displayed. No attribute names will be displayed. The attribute values displayed will be separated by the delimiter which was specified with this option.</p> <p>This option is useful when only attribute values are desired to be displayed, or when the values of only selected attributes are desired to be displayed.</p> <p><b>Attribute Names</b></p> <ul style="list-style-type: none"> <li>• <b>lpar_id</b>: Unique integer identifier for the partition</li> <li>• <b>lpar_name</b>: name of the partition</li> <li>• <b>state</b>: The current state of the LED. Valid values: <ul style="list-style-type: none"> <li>– off - the LED is off</li> <li>– on - the LED is on</li> </ul> </li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>--header</b>                   | Display a header record, which is a delimiter separated list of attribute names for the attribute values that will be displayed. This header record will be the first record displayed. This option is only valid when used with the -F option.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |

## Exit Status

This command has a return code of 0 on success.

## Security

This command is accessible by all users.

## Examples

1. Display the physical system attention LED for the system:  

```
lsled -r sa -t phys
```
2. Display all of the virtual partition system attention LEDs:  

```
lsled -r sa -t virtualpar
```
3. Display the virtual partition system attention LEDs for partitions lpar1 and lpar2:  

```
lsled -r sa -t virtualpar --filter \"lpar_names=lpar1,lpar2\"
```

## Related Information

The `chled` command.

---

## lsparinfo command

### Purpose

Displays the logical partition number and name.

### Syntax

`lsparinfo`

### Description

The `lsparinfo` command displays LPAR number and LPAR name. If LPAR does not exist, -1 is displayed for LPAR number and NULL for LPAR name.

### Examples

1. To display the logical partition number and name, type:  

```
lsparinfo
```

---

## IVM lsparmigr command

### Purpose

Lists and displays partition migration information. This command is operable only in an Integrated Virtualization Manager environment.

### Syntax

```
lsparmigr -r manager | lpar | msp | procpool | sys | virtualio | mempool | pgdev [-m <managed system>] [-t <managed system>] [--ip <target HMC/IVM IP address>] [-u <target HMC/IVM username>] [--filter "<filter data>"] [-F [<attribute names>]] [--redundantpgvios { 0 | 1 | 2 }] [--header] [--help]
```

### Description

The `lsparmigr` command displays the state of the partition migration and the adapter mappings that might be used during the partition migration.



## Flags

**-r**

The type of resources for which to list partition migration information.

Specify **lpar** to list partition migration information for all of the partitions in managed-system.

Specify **mssp** to list possible source and destination mover service partition (MSP) pairs for an active partition migration of the partition specified with the **lpar\_names** or **lpar\_ids** filter.

Specify **procpool** to list possible shared processor pools in the destination managed system that have enough available processing resources for the partition specified with the **lpar\_names** or **lpar\_ids** filter.

Specify **sys** to list the partition mobility capabilities for managed-system.

Specify **virtualio** to list possible and suggested mappings of the virtual SCSI adapters and virtual fibre channels in the partition specified with the **lpar\_names** or **lpar\_ids** filter to Virtual I/O Server in the destination managed system.

Specify **mempool** to list details for each shared memory pool on the target managed system.

Specify **pgdev** to list paging devices available on the target managed system.

**-m** <managed system>

The name of the source managed system.

**-t** <managed system>

The name of target managed system.

**--ip**

The IP address or hostname of the Integrated Virtualization Manager managing the target system.

**-u**

The user name to use on the Integrated Virtualization Manager that manages the target system. If the **--ip** flag is specified and **-u** flag is not specified, the user name on the source Integrated Virtualization Manager is used on the target Integrated Virtualization Manager.

**--filter** <filter data>

The filter or filters to apply to the resources to be listed. Filters are used to select which resources are to be listed. If no filters are used, then all of the resources will be listed. For example, partition information can be listed by using a filter to specify the names or IDs of the partitions to list. Otherwise, if no filter is used, then information for all of the partitions in the managed system will be listed.

The filter data consists of filter name/value pairs, which are in comma separated value (CSV) format. The filter data must be enclosed in quotation marks.

The format of the filter data is as follows:

```
"filter-name=value,filter-name=value,..."
```

**Note:** Certain filters accept a comma separated list of values, as follows:

```
"filter-name=value,value,...",..."
```

When a list of values is specified, the filter name/value pair must be enclosed in quotation marks. Depending on the shell being used, nested quotation marks may need to be preceded by an escape character, which is usually a `'\'` character.

Valid filter names:

`lpar_names` | `lpar_ids`

Only one of these filters may be specified

**-F** [<attribute names>]

List of the names of the attributes. If no attribute names are specified, then all attributes will be listed.

|                         |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|-------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>-redundantpgvios</b> | Specify the configured level of redundancy for the memory pools to be listed.<br><b>0</b> Do not require redundancy on the target managed system.<br><b>1</b> Only return memory pool information for those that are capable of redundancy.<br><b>2</b> Return all memory pool information regardless of redundancy capabilities.<br>If you do not specify the parameter, memory pools capable of the currently configured level of redundancy are returned.<br><br>Exactly one partition must be specified by either the <code>lpar_ids</code> filter or <code>lpar_names</code> filter.<br><br>If no suitable memory pools are found on the destination managed system, the destination IVM generates an error message that describes the problem. |
| <b>--header</b>         | Displays a header record, which is a delimiter separated list of attribute names for the attribute values that are displayed. This header record is the first record displayed. This option is only valid when used with the <code>-F</code> option.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>--help</b>           | Displays the help text for this command and exit.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |

## Exit Status

This command has a return code of 0 on success.

## Examples

1. To return partition mobility capabilities for the HMC or the Integrated Virtualization Manager, type the following command:

```
lslparmigr -r manager
remote_lpar_mobility_capable=1
```

2. To return partition mobility capabilities for the specified managed system, type the following command:

```
lslparmigr -m zd25ae -r sys
inactive_lpar_mobility_capable=1,num_inactive_migrations_supported=40,
num_inactive_migrations_in_progress=0,
active_lpar_mobility_capable=1,num_active_migrations_supported=40,
num_active_migrations_in_progress=0
```

3. To return a list of partitions on the specified managed system and their migration state, type the following command:

```
$ lslparmigr -r lpar
name=10-0EDDA,lpar_id=1,migration_state=Not Migrating
name=mptest_migration,lpar_id=2,migration_state=Not Migrating
name=lp3,lpar_id=3,migration_state=Not Migrating
name=mp_i51,lpar_id=4,migration_state=Not Migrating
name=mover2,lpar_id=5,migration_state=Migration Starting,
migration_type=active,dest_sys_name=migfspL4,
dest_lpar_id=5,source_msp_name=10-0EDDA,source_msp_id=1,
dest_msp_name=10-0EDEA,dest_msp_id=1,bytes_transmitted=10247078,
bytes_remaining=259547136,
remote_manager=miglp9.yourcity.yourcompany.com,
remote_user=padmin
name=lpar10,lpar_id=10,migration_state=Not Migrating
```

4. To return a list of possible source and destination MSP pairs, type the following command:

```
lslparmigr -r msp -m migfspL2 -t migfspL1
source_msp_name=VIOS,source_msp_id=2,"dest_msp_names=VIOS,VIOS2",
"dest_msp_ids=1,3","ipaddr_mappings=192.168.10.10/en1/1/VIOS/192.
168.10.11/en1,192.168.20.10/en2/3/VIOS2/192.168.20.13/en0"
source_msp_name=VIOS2,source_msp_id=1,"dest_msp_names=VIOS,VIOS2",
"dest_msp_ids=1,3","ipaddr_mappings=192.168.10.12/en0/1/VIOS/192.
168.10.11/en1,192.168.20.12/en1/3/VIOS2/192.168.20.13/en0"
```

5. To return a list of possible and suggested mappings for virtual SCSI and virtual Ethernet adapters, type the following command:

```
lslparmigr -r virtualio -m migfspL2 -t migfspL1 --filter "lpar_ids=5"
"possible_virtual_scsi_mappings=2/VIOS/1,2/VIOS2/3",
"suggested_virtual_scsi_mappings=2/VIOS/1"
```

6. To return a list of existing shared processor pools on the target managed system, type the following command:

```
lslparmigr -r procpool -m migfspL2 -t migfspL1 --filter "lpar_ids=5"
"shared_proc_pool_ids=0,3,6","shared_proc_pool_names=
DefaultPool,ProcPoolA,ProcPoolB"
```

7. To return a list of possible and suggested mappings for virtual fibre channel adapters, type the following command:

```
lslparmigr -r virtualio -m migfspL2 -t migfspL1 --filter "lpar_ids=5"
"possible_virtual_scsi_mappings=2/VIOS/1",
"suggested_virtual_scsi_mappings=2/VIOS/1"
"possible_virtual_fc_mappings=4/VIOS/1,4/VIOS2/2",
"suggested_virtual_fc_mappings=4/VIOS/1"
```

---

## IVM lslparutil command

### Purpose

List utilization metrics for the managed system and partitions. This command is operable only in an Integrated Virtualization Manager environment.

### Syntax

To list utilization metrics:

```
lslparutil {-r sys | pool | lpar | mempool} [--startyear Year] [--startmonth Month] [--startday Day] [
--starthour Hour] [--startminute Minute] [--endyear Year] [--endmonth Month] [--endday Day] [
--endhour Hour] [--endminute Minute] [-n NumberOfEvents] [-spread] [--filter "FilterData"] [-F
"AttributeNames"] [--header] [-m ManagedSystem]
```

To list utilization configuration attributes:

```
lslparutil -r config [-F "AttributeNames"] [--header] [-m ManagedSystem]
```

### Description

The **lslparutil** command lists utilization metrics for the managed system and partitions. The sample rate at which these metrics are gathered can be changed, or disabled. The information gathered from this command allows for the determination of processor utilization on a system and partition level over a given period of time.

## Flags

**-r** *ResourceType*

The type of resources to list:

- **-r config**: List configuration attributes for utilization monitoring
  - **Attributes**: sample\_rate
  - **Filters**: None
- **-r lpar**: List utilization data for logical partitions
  - **Attributes**: time, time\_bin, sys\_time, event\_type, resource\_type, time\_cycles, lpar\_id, uptime, curr\_io\_entitled\_mem, curr\_proc\_mode, curr\_proc\_units, curr\_procs, curr\_sharing\_mode, curr\_uncap\_weight, curr\_5250\_cpw\_percent, curr\_mem, entitled\_cycles, capped\_cycles, mapped\_io\_entitled\_mem, mem\_mode, mem\_overage\_cooperation, phys\_run\_mem, uncapped\_cycles, donated\_cycles, idle\_cycles, run\_latch\_instructions, run\_latch\_cycles, run\_mem\_weight
  - **Filters**: { lpar\_ids | lpar\_names}
- **-r mempool**: List information and statistics about memory pools
  - **Attributes**: time, time\_bin, sys\_time, event\_type, resource\_type, page\_faults, page\_in\_delay, curr\_pool\_mem, lpar\_curr\_io\_entitled\_mem, lpar\_mapped\_io\_entitled\_mem, lpar\_run\_mem, sys\_firmware\_pool\_mem
- **-r pool**: List utilization data for shared processor pools
  - **Attributes**: time, time\_bin, sys\_time, event\_type, resource\_type, time\_cycles, shared\_proc\_pool\_id, curr\_avail\_pool\_proc\_units, configurable\_pool\_proc\_units, borrowed\_pool\_proc\_units, total\_pool\_cycles, utilized\_pool\_cycles
  - **Filters**: None
- **-r sys**: List utilization data for the managed system
  - **Attributes**: time, time\_bin, sys\_time, event\_type, resource\_type, state, configurable\_sys\_proc\_units, configurable\_sys\_mem, curr\_avail\_sys\_proc\_units, curr\_avail\_5250\_cpw\_percent, curr\_avail\_sys\_mem, sys\_firmware\_mem, proc\_cycles\_per\_second
  - **Filters**: None

**--startyear** *Year*

Starting year filter (default 1970)

**--startmonth** *Month*

Starting month filter (default 1) January is 1. December is 12

**--startday** *Day*

Starting day filter (default 1)

**--starthour** *Hour*

Starting hour filter (default 0)

**--startminute** *Minute*

Starting minute filter (default 0)

**--endyear** *Year*

Ending year filter (default now)

**--endmonth** *Month*

Ending month filter (default now)

**--endday** *Day*

Ending day filter (default now)

**--endhour** *Hour*

Ending hour filter (default now)

**--endminute** *Minute*

Ending minute filter (default now)

**--n** *NumberOfEvents*

Starting from the most recent event, the maximum number of samples to return.

**--spread**

Used with -n to distribute displayed samples evenly over time.

**-m** *ManagedSystem*

The name of the managed system. This attribute is optional because there is only one system to manage. The name might either be the user-defined name for the managed system, or be in the form tttt-mmm\*sssssss, where tttt is the machine type, mmm is the model, and sssssss is the serial number of the managed system.

**--filter** *FilterData*

The filters to apply to the resources to be listed. Filters are used to select which resources of the specified resource type are to be listed. If no filters are used, then all of the resources of the specified resource type will be listed. For example, specific partitions can be listed by using a filter to specify the names or IDs of the partitions to list. Otherwise, if no filter is used, then all the partitions in the managed system will be listed.

The filter data consists of filter name/value pairs, which are in comma separated value (CSV) format. The filter data must be enclosed in double quotation marks.

The format of the filter data is as follows:

```
"filter-name=value,filter-name=value,..."
```

Note that certain filters accept a comma separated list of values, as follows:

```
"filter-name=value,value,...",..."
```

When a list of values is specified, the filter name/value pair must be enclosed in double quotation marks. Depending on the shell being used, nested double quotation marks might need to be preceded by an escape character, which is usually a backslash character.

Unless otherwise indicated, multiple values can be specified for each filter.

#### **Valid filter names for -r lpar**

##### **lpar\_ids**

ID of the partitions to view

##### **lpar\_names**

Name of the partitions to view

**-F** *AttributeNames*

A delimiter separated list of attribute names for the desired attribute values to be displayed for each resource. If no attribute names are specified, then values for all of the attributes for the resource will be displayed.

When this option is specified, only attribute values will be displayed. No attribute names will be displayed. The attribute values displayed will be separated by the delimiter which was specified with this option.

This option is useful when only attribute values are desired to be displayed, or when the values of only selected attributes are desired to be displayed.

#### **Attribute Names:**

##### **borrowed\_pool\_proc\_units**

Processing units that are being borrowed from powered off partitions with dedicated processors.

##### **capped\_cycles**

The number of capped processing cycles utilized by this partition since the system started.

##### **configurable\_pool\_proc\_units**

The number of configurable processing units in the shared pool at the time of the sample

##### **configurable\_sys\_mem**

The amount of configurable system memory (in megabytes) at the time of the sample.

**configurable\_sys\_proc\_units**

The number of configurable system processing units at the time of the sample.

**curr\_5250\_cpw\_percent**

The 5250 CPW percent assigned to the partition at the time of the sample.

**curr\_avail\_5250\_cpw\_percent**

The 5250 CPW percent available to be assigned to partitions at the time of the sample.

**curr\_avail\_pool\_proc\_units**

The number of processing units available to be assigned to partitions at the time of the sample.

**curr\_avail\_sys\_mem**

The amount of memory (in megabytes) available to be assigned to partitions at the time of the sample.

**curr\_avail\_sys\_proc\_units**

The number of processing units available to be assigned to partitions at the time of the sample.

**curr\_io\_entitled\_mem**

The size of memory that the logical partition is entitled to map to I/O (in megabytes).

**curr\_mem**

The amount of memory (in megabytes) assigned to the partition at the time of the sample.

**curr\_pool\_mem**

The total amount of physical memory (in megabytes) in the memory pool.

**curr\_proc\_mode**

The processing mode for the partition at the time of the sample. Valid values

- ded: dedicated processor mode
- shared: shared processor mode curr\_proc\_units

**curr\_proc\_units**

Current number of processing units assigned to the partition. This attribute is only valid for partitions using shared processors.

**curr\_procs**

The number of processors or virtual processors assigned to the partition at the time of the sample.

**curr\_sharing\_mode**

The sharing mode of the partition at the time of the sample. Valid values:

- keep\_idle\_procs: valid with dedicated processor mode
- share\_idle\_procs: valid with dedicated processor mode
- cap: capped mode. valid with shared processor mode
- uncap: uncapped mode. valid with shared processor mode

**curr\_uncap\_weight**

The current weighted average of processing priority when in uncapped sharing mode at the time of the sample. The smaller the value, the lower the weight. Valid values are 0-255

**donated\_cycles**

The number of cycles donated by a partition with dedicated processors.

**entitled\_cycles**

The number of processing cycles to which the partition has been entitled since the system started. This value is based on the number of processing units assigned to the partition, and might be greater than, or smaller than the number of cycles actually used.

**event\_type**

The type of event. This will be displayed with a constant value of sample for all samples except the config resource type.

**idle\_cycles**

The number of cycles that the logical partition reported as idle. If the partition does not have the ability to report idle cycles, this value is 0.

**lpar\_curr\_io\_entitled\_mem**

The total entitlement (in megabytes) of all shared memory partitions served by the memory pool.

**lpar\_id** The unique integer identifier for the partition.

**lpar\_mapped\_io\_entitled\_mem**

The total I/O mapped memory (in megabytes) of all active partitions served by the memory pool.

**lpar\_run\_mem**

The total logical real memory in megabytes of all active partitions served by the memory pool.

**mapped\_io\_entitled\_mem**

Size of memory (in megabytes) that the logical partition has mapped to I/O memory.

**mem\_mode**

The partition memory mode. Valid values:

- ded - dedicated memory
- shared - shared memory

If the memory mode is shared, then the logical partition cannot be assigned any physical I/O slots, the `proc_mode` attribute on the partition must be shared, and a memory pool must exist.

**Note:** The Virtual I/O Server logical partition only supports the dedicated memory mode.

**page\_faults**

The total-page fault count of the memory pool since the initialization of the pool, or since the last reboot of the hypervisor, whichever is more recent.

**page\_in\_delay**

The total page in delay (in microseconds) spent waiting for the page faults since the initialization of the memory pool, or since the last reboot of the hypervisor, whichever is more recent.

**phys\_run\_mem**

The size of physical memory (in megabytes) that is backing the logical address space of the logical partition. This value represents the portion of the logical memory assignment that is currently paged in the logical partition.

**proc\_cycles\_per\_second**

The processing cycles per second on one physical processor. This value is static for a particular managed system.

**resource\_type**

The resource type queried. Valid values are sys, pool, and lpar, depending on which value is supplied for the **-r** flag.

**run\_latch\_cycles**

The number of nonidle cycles used by the partition, while the run-latch is set, since the last time the system was started.

**run\_latch\_instructions**

The number of nonidle instructions performed by the partition, while the run-latch is set, since the last time the system was started.

**run\_mem\_weight**

Shared memory weight at run time.

**sample\_rate**

The rate at which samples are obtained. This can be changed with the **chlpoutil** command. Valid values:

- 0: Samples will not be retrieved.
- 30: Samples will be retrieved every 30 seconds. This is the default value.
- 60: Samples will be retrieved every minute.
- 300: Samples will be retrieved every 5 minutes.
- 1800: Samples will be retrieved every 30 minutes.
- 3600: Samples will be retrieved every hour.

**shared\_proc\_pool\_id**

The unique decimal identifier for a shared processing pool.

**state** The state of the managed system at the time of the sample.

**sys\_firmware\_mem**

The amount of memory, in megabytes, on the managed system that is being used by system firmware at the time of the sample.

**sys\_firmware\_pool\_mem**

The amount of memory in the shared memory pool that is reserved for use by the firmware. This is the greater of the current and pending values.

**sys\_time**

The time on the managed system that the sample was taken. The time and sys\_time attributes will have the same value.

**time** The time on the management partition that this sample was taken. The time and sys\_time attributes will have the same value.

**time\_bin**

The time or sys\_time represented as the number of milliseconds since January 1, 1970.

**time\_cycles**

The number of time cycles since the system was started.

**total\_pool\_cycles**

The total number of processing cycles available in the shared pool since the system was started.



|                             |                                                                                                                                                                                                                                                 |
|-----------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>uncapped_cycles</b>      | The number of uncapped processing cycles utilized by this partition since the system was started.                                                                                                                                               |
| <b>uptime</b>               | The amount of time (in seconds) that the partition has been running at the time of the sample.                                                                                                                                                  |
| <b>utilized_pool_cycles</b> | The number of processing cycles in the shared pool that have been utilized since the system was started.                                                                                                                                        |
| <b>--header</b>             | Display a header record, which is a delimiter separated list of attribute names for the attribute values that will be displayed. This header record will be the first record displayed. This option is only valid when used with the -F option. |

## Exit Status

This command has a return code of 0 on success.

## Security

This command is accessible by all users.

## Examples

1. To list the last five shared processing pool utilization metrics, type:

```
lslparutil -r pool -n 5
```

2. To calculate the shared processing pool utilization in percent over a five-minute time period, type:

```
lslparutil -r pool --startyear 2006 --startmonth 10 --startday 4
--starthour 15 --startminute 23 --endyear 2006 --endmonth 10
--endday 4 --endhour 15 --endminute 28
-F time,total_pool_cycles,utilized_pool_cycles
10/04/2006 15:27:56,449504263518104,372389272879
10/04/2006 15:27:26,449442382657200,372195239995
10/04/2006 15:26:56,449382561908822,371998920942
10/04/2006 15:26:26,449320667371346,371797239591
10/04/2006 15:25:56,449258781703684,371592366752
10/04/2006 15:25:26,449196894951060,371394157970
10/04/2006 15:24:56,449135006535822,371192089089
10/04/2006 15:24:27,449075176369863,369243635687
10/04/2006 15:23:56,449013298087726,369040980263
10/04/2006 15:23:26,448951445376558,368850062933

Pool utilization = (utilized_pool_cycle / total_pool_cycles) * 100
Pool utilization = ((372389272879 - 368850062933) /
(449504263518104 - 448951445376558)) * 100
Pool utilization = 0.64%
```

3. To calculate the shared processing pool size over a five-minute time period, type:

```
lslparutil -r pool --startyear 2006 --startmonth 10 --startday 4
--starthour 15 --startminute 23 --endyear 2006 --endmonth 10
--endday 4 --endhour 15 --endminute 28 -F time,time_cycles,total_pool_cycles
10/04/2006 15:27:56,112376065883848,449504263518104
10/04/2006 15:27:26,112360595668767,449442382657200
10/04/2006 15:26:56,112345640481652,449382561908822
10/04/2006 15:26:26,112330166847247,449320667371346
10/04/2006 15:25:56,112314695430447,449258781703684
10/04/2006 15:25:26,112299223741951,449196894951060
10/04/2006 15:24:56,112283751639775,449135006535822
10/04/2006 15:24:27,112268794096846,449075176369863
10/04/2006 15:23:56,112253324526335,449013298087726
10/04/2006 15:23:26,112237861348574,448951445376558
```

```
Pool size = total_pool_cycles / time_cycles
Pool size = (449504263518104 - 448951445376558) /
 (112376065883848 - 112237861348574)
Pool size = 4
```

4. To calculate the processing utilization in percent for partition 1 over the last 11 samples, type:

```
lslparutil -r lpar -F time,lpar_id,entitled_cycles,capped_cycles,uncapped_cycles
--filter lpar_ids=1 -n 11
```

```
06/26/2005 12:13:04,1,13487973395246,353524992184,93964052971
06/26/2005 12:12:33,1,13486720703117,353490258336,93964052971
06/26/2005 12:12:03,1,13485467110700,353456792591,93964052971
06/26/2005 12:11:33,1,13484213859686,353423048854,93964052971
06/26/2005 12:11:03,1,13482961098044,353386674795,93964052971
06/26/2005 12:10:32,1,13481706673802,353350985013,93964052971
06/26/2005 12:10:02,1,13480453156357,353317211748,93964052971
06/26/2005 12:09:32,1,13479199972343,353283141535,93964052971
06/26/2005 12:09:02,1,13477946765207,353248812551,93964052971
06/26/2005 12:08:31,1,13476693184663,353213970760,93964052971
06/26/2005 12:08:01,1,13475439617080,353179654833,93964052971
```

```
Processor utilization = ((capped_cycles + uncapped_cycles)
 / entitled_cycles) * 100
Processor utilization = (((353524992184 - 353179654833)
 + (93964052971 - 93964052971))
 / (13487973395246 - 13475439617080)) * 100
Processor utilization = 2.76%
```

5. To calculate the most recent number of utilized processor units for partition 1, type:

```
lslparutil -r lpar -F time,time_cycles,capped_cycles,uncapped_cycles -n 2 \
--filter lpar_ids=1
```

```
10/06/2006 09:42:58,190122585897822,836322334068,1209599213218
10/06/2006 09:42:28,190107628555119,836215824328,1209507899652
```

```
Processor units used = (capped_cycles + uncapped_cycles) / time_cycles
Processor units used = ((836322334068 - 836215824328)
 + (1209599213218 - 1209507899652)) /
 (190122585897822 - 190107628555119)
Processor units used = 0.01
```

6. To display 5 entries spread over the last day, type:

```
lslparutil -r lpar --startyear 2005 --startmonth 6 --startday 25 \
--starthour 14 --startminute 20 --endyear 2005 --endmonth 6 \
--endday 26 --endhour 14 --endminute 20 -F time,lpar_id,entitled_cycles,\
capped_cycles,uncapped_cycles -n 5 --spread
```

```
06/26/2005 14:20:03,1,64074629124428,2634420796918,3335839807455
06/26/2005 08:20:03,1,63180190141506,2597103712238,3292339240560
06/26/2005 02:20:02,1,62285686629911,2559847748332,3249081303922
06/25/2005 20:20:01,1,61391161857754,2522594738730,3205850397796
06/25/2005 14:20:01,1,60496719757782,2485459490629,3162821474641
```

**Note:** This data is particularly useful for graphing utilization data.

7. To view the number of donated cycles:

```
lslparutil -r lpar -F donated_cycles
```

## Related Information

The `chlparutil` and `lssyscfg` commands.

---

## lslv command

### Purpose

Displays information about a logical volume.

## Syntax

`lslv [ -map | -pv ] LogicalVolume [ -field FieldName ] [ -fmt Delimiter ]`

`lslv -free [ -field Fieldname ] [ -fmt Delimiter ]`

## Description

The `lslv` command displays the characteristics and status of the *LogicalVolume* or lists the logical volume allocation map for the physical partitions on the *PhysicalVolume* in which the logical volume is located. The logical volume can be a name or identifier.

If no flags are specified, the following status is displayed:

|                                  |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|----------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Logical volume</b>            | Name of the logical volume. Logical volume names must be unique systemwide and can range from 1 to 15 characters.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Volume group</b>              | Name of the volume group. Volume group names must be unique systemwide and can range from 1 to 15 characters.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Logical volume identifier</b> | Identifier of the logical volume.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Permission</b>                | Access permission; read-only or read-write.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Volume group state</b>        | State of the volume group. If the volume group is activated with the <code>activatevg</code> command, the state is either <code>active/complete</code> (indicating all physical volumes are active) or <code>active/partial</code> (indicating all physical volumes are not active). If the volume group is not activated with the <code>activatevg</code> command, the state is <code>inactive</code> .                                                                                                                                                                                                                        |
| <b>Logical volume state</b>      | State of the logical volume. The <code>Opened/stale</code> status indicates the logical volume is open but contains physical partitions that are not current. <code>Opened/syncd</code> indicates the logical volume is open and synchronized. <code>Closed</code> indicates the logical volume has not been opened.                                                                                                                                                                                                                                                                                                            |
| <b>Type</b>                      | Logical volume type.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Write verify</b>              | Write verify state of On or Off.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Mirror write consistency</b>  | Mirror write consistency state of Yes or No.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Max LPs</b>                   | Maximum number of logical partitions the logical volume can hold.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>PP size</b>                   | Size of each physical partition.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Copies</b>                    | Number of physical partitions created for each logical partition when allocating.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Schedule policy</b>           | Sequential or parallel scheduling policy.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>LPs</b>                       | Number of logical partitions currently in the logical volume.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>PPs</b>                       | Number of physical partitions currently in the logical volume.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Stale partitions</b>          | Number of physical partitions in the logical volume that are not current.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Bad blocks</b>                | Bad block relocation policy.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Inter-policy</b>              | Inter-physical allocation policy.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Strictness</b>                | Current state of allocation. Possible values are <code>strict</code> , <code>nonstrict</code> , or <code>superstrict</code> . A <code>strict</code> allocation states that no copies for a logical partition are allocated on the same physical volume. If the allocation does not follow the strict criteria, is called <code>nonstrict</code> . A <code>nonstrict</code> allocation states that at least one occurrence of two physical partitions belong to the same logical partition. A <code>superstrict</code> allocation states that no partition from one mirror copy may reside the same disk as another mirror copy. |
| <b>Intra-policy</b>              | Intra-physical allocation policy.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Upper bound</b>               | If the logical volume is super strict, upper bound is the maximum number of disks in a mirror copy.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Relocatable</b>               | Indicates whether the partitions can be relocated if a reorganization of partition allocation takes place.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Mount point</b>               | File system mount point for the logical volume, if applicable.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Label</b>                     | Specifies the label field for the logical volume.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>PV distribution</b>           | The distribution of the logical volume within the volume group. The physical volumes used, the number of logical partitions on each physical volume, and the number of physical partitions on each physical volume are shown.                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>striping width</b>            | The number of physical volumes being striped across.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |

**strip size** The number of bytes per stripe.

The **-free** flag displays logical volumes that are available to be used as backing devices for virtual storage.

Full scripting support is available by using the **-field** *FieldNames* and **-fmt** *Delimiter* flags. The **-field** flag will allow the user to select which output fields to display and in what order, while the **-fmt** flag provides scriptable output. The output fields will be displayed in the order they appear on the command line.

## Flags

**-free** Lists only logical volumes that are available for use as a backing device for virtual SCSI.

**-field** Specifies the list of fields to display. The following fields are supported if no flags are specified:

- lvname** Name of the logical volume. Logical volume names must be unique systemwide and can range from 1 to 15 characters.
- vgname** Name of the volume group. Volume group names must be unique systemwide and can range from 1 to 15 characters.
- lvid** Identifier of the logical volume.
- access** Access permission: read-only or read-write.
- vgstate** State of the volume group. If the volume group is activated with the **activatevg** command, the state is either *active/complete* (indicating all physical volumes are active) or *active/partial* (indicating all physical volumes are not active). If the volume group is not activated with the **deactivatevg** command, the state is *inactive*.
- lvstate** State of the logical volume. The *opened/stale* status indicates the logical volume is open but contains physical partitions that are not current. *Opened/syncd* indicates the logical volume is open and synchronized. *Closed* indicates the logical volume has not been opened.
- type** Logical volume type.
- wverify** Write verify state of on or off.
- mwc** Mirror write consistency state of on or off.
- maxlps** Maximum number of logical partitions the logical volume can hold.
- ppsize** Size of each physical partition.
- copies** Number of physical partitions created for each logical partition when allocating.
- spolicy** Sequential or parallel scheduling policy.
- lps** Number of logical partitions currently in the logical volume.
- pps** Number of physical partitions currently in the logical volume.
- stale** Number of physical partitions in the logical volume that are not current.
- bbpolicy** Bad block relocation policy.

**inter** Inter-physical allocation policy.

**intra** Intra-physical allocation policy.

**ubound**

If the logical volume is super strict, upper bound is the maximum number of disks in a mirror copy.

**relocatable**

Indicates whether the partitions can be relocated if a reorganization of partition allocation takes place.

**mount** File system mount point for the logical volume, if applicable.

**label** Specifies the label field for the logical volume.

**separatepv**

The strictness value. Current state of allocation, strict, nonstrict, or superstrict. A strict allocation states that no copies for a logical partition are allocated on the same physical volume. If the allocation does not follow the strict criteria, it is called nonstrict. A nonstrict allocation states that at least one occurrence of two physical partitions belong to the same logical partition. A superstrict allocation states that no partition from one mirror copy may reside the same disk as another mirror copy.

**serialio** Serialization of overlapping IOs state of yes or no. If serialization is turned on (yes), then overlapping IOs are not allowed on a block range, and only a single IO in a block range is processed at any one time. Most applications, such as file systems and databases, perform serialization; therefore, serialization should be turned off (no). The default setting for new logical volumes is no.

The following fields are supported if the **-pv** flag is specified:

**pvname**

Physical volume disk name

**copies** The following three fields:

- The number of logical partitions containing at least one physical partition (no copies) on the physical volume
- The number of logical partitions containing at least two physical partitions (one copy) on the physical volume
- The number of logical partitions containing three physical partitions (two copies) on the physical volume

**inband**

The percentage of physical partitions on the physical volume that belong to the logical volume and were allocated within the physical volume region specified by Intra-physical allocation policy.

**dist** The number of physical partitions allocated within each section of the physical volume: outer edge, outer middle, center, inner middle, and inner edge of the physical volume.

The following fields are supported if the **-map** flag is specified:

**lpnum** Logical partition number.

**pvname1**

Physical volume name where the logical partition's first physical partition is located.

**ppnum1**

First physical partition number allocated to the logical partition.

**pvname2**

Physical volume name where the logical partition's second physical partition (first copy) is located.

**ppnum2**

Second physical partition number allocated to the logical partition.

The following fields are supported if the **-free** flag is specified:

**lvname**

Logical partition number.

**size**

Physical volume name where the logical partition's first physical partition is located.

**vgname**

Name of the volume group. Volume group names must be unique systemwide and can range from 1 to 15 characters.

**-fmt**  
**-map**

Specifies a delimiter character to separate output fields.

Lists the following fields for each logical partition:

**LPs** Logical partition number.

**PV1**

Physical volume name where the logical partition's first physical partition is located.

**PP1**

First physical partition number allocated to the logical partition.

**PV2**

Physical volume name where the logical partition's second physical partition (first copy) is located.

**PP2**

Second physical partition number allocated to the logical partition.

**-pv**

Lists the following fields for each physical volume in the logical volume:

**PV**

Physical volume name.

**Copies** The following three fields:

- The number of logical partitions containing at least one physical partition (no copies) on the physical volume
- The number of logical partitions containing at least two physical partitions (one copy) on the physical volume
- The number of logical partitions containing three physical partitions (two copies) on the physical volume

**In band**

The percentage of physical partitions on the physical volume that belong to the logical volume and were allocated within the physical volume region specified by Intra-physical allocation policy.

**Distribution**

The number of physical partitions allocated within each section of the physical volume: outer edge, outer middle, center, inner middle, and inner edge of the physical volume.

## Exit Status

See “Exit status for Virtual I/O Server commands” on page 3.

## Examples

1. To display information about logical volume **lv03**, type:

```
lslv lv03
```

Information about logical volume **lv03**, its logical and physical partitions, and the volume group to which it belongs is displayed.

2. To display information about logical volume **lv03** by physical volume, type:

```
lslv -pv lv03
```

The characteristics and status of **lv03** are displayed, with the output arranged by physical volume.

3. To display a list of logical volumes that can be used as backing devices, type:

```
lslv -free
```

The system displays a message similar to the following:

| LV NAME | SIZE(megabytes) | VOLUME GROUP |
|---------|-----------------|--------------|
| lv00    | 64              | rootvg       |
| lv01    | 64              | rootvg       |

4. To display only the type and volume group of logical volume **hd6** and separate the data by a : (colon) , type:

```
lslv hd6 -field type vgroup -fmt :
```

The system displays a message similar to the following:

```
paging:rootvg
```

## Related Information

The **mklv** command, the **extendlv** command, and the **rmlv** command.

---

## lsmmap command

### Purpose

Displays the mapping between physical, logical, and virtual devices.

### Syntax

```
lsmmap {-vadapter ServerVirtualAdapter | -plc PhysicalLocationCode | -all } [-type BackingDeviceType | -net | -npiv] [-field FieldNames] [-fmt Delimiter]
```

```
lsmmap -ams { -vtd PagingDevice | -all } [-type BackingDeviceType] [-field FieldNames] [-fmt Delimiter]
```

```
lsmmap -clustername ClusterName{-all | -hostnameHostName}
```

```
lsmmap -suspend {-vadapter ServerVirtualAdapter | -all } [-type BackingDeviceType] [-field FieldNames] [-fmt Delimiter]
```

### Description

The **lsmmap** command displays the mapping between the virtual host adapters and the physical devices they are backed to. Given an adapter name (*ServerVirtualAdapter*) or physical location code

(*PhysicalLocationCode*) of a server virtual adapter, the device name of each connected virtual target device (child devices), its logical unit number, backing device(s) and the backing devices physical location code is displayed. If the **-net** flag is specified the supplied device must be a virtual server Ethernet adapter.

The **-fmt** flag divides the output by a user-specified delimiter/character (delimiter). The delimiter can be any non-white space character. This format is provided to facilitate scripting.

The **-type** flag limits the types of backing devices displayed to the user-specified list, *BackingDeviceType*. With the exception of the net type, any combination of device types may be specified. The net type cannot be combined with any other type.

The **-npiv** flag displays the server binding information between the virtual fibre channel and the physical fibre channel adapter. It is also used to display client adapter attributes that are sent to the server adapter.

The **-suspend** flag displays the storage that is used by a suspended client. The **rmdev** command can be used to remove suspended adapters. For more information, see the *Examples* section of the **rmdev** command topic.

## Flags

|                     |                                                                                                                                                                                                                                                        |
|---------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>-all</b>         | Specifies the <b>lsmap</b> output that should be displayed for all virtual SCSI devices. If used with the <b>-net</b> flag, virtual Ethernet adapters are displayed. If used with the <b>-npiv</b> flag, virtual fibre channel adapters are displayed. |
| <b>-ams</b>         | Displays paging space device information used in active memory sharing.                                                                                                                                                                                |
| <b>-clustername</b> | Specifies the cluster name.                                                                                                                                                                                                                            |



**-field** *FieldName*

Specifies the list of fields to display.

**Note:** The order that you specify the attributes is not guaranteed to be the same order as they are returned.

For virtual SCSI (default), the following fields are supported:

**svsa** Server virtual adapter.

**physloc**

The physical location code of the server virtual adapter.

**mirrored**

The backing device is part of a Peer-to-Peer Remote Copy (PPRC) pair.

**clientid**

Client partition ID.

**vtd** Virtual target device.

**lun** Logical unit number.

**backing**

Backing device.

**bdphysloc**

The physical location code of the backing device.

**status** Virtual target device status.

For shared Ethernet adapter (**-net** flag), the following fields are supported:

**svea** Virtual Ethernet adapter.

**physloc**

The physical location code of the server virtual adapter.

**sea** Shared Ethernet adapter.

**backing**

Backing device.

**bdphysloc**

The physical location code of the backing device.

**status** Shared Ethernet adapter status.

For NPIV (**-npiv** flag), the following fields are supported:

**name** Virtual fibre channel adapter name.

**physloc** The physical location code of the server virtual fibre channel adapter.

**clntid** Client logical partition ID.

**clntname** Client logical partition name.

**clntos** The operating system on the client logical partition.

**status** Virtual fibre channel adapter status.

**fc** Physical fibre channel adapter name.

**fcphysloc** The physical location of the fibre channel adapter.

**ports** Physical fibre channel port number.

**flags** Virtual fibre channel adapter flags.

**vfclient** Client virtual fibre channel adapter name.

**vfclientdrc** Client virtual fibre channel adapter used for Dynamic Reconfiguration Connection (DRC).

For active memory sharing (**-ams** flag), the following fields are supported:

**paging** Paging virtual target device.

**streamid**  
Stream ID.

**clientid**  
Client logical partition ID.

**status** Paging virtual target device status.

**redundancy**  
Redundancy usage for the paging VIOS logical partitions.

**backing**  
Backing device.

**poolid** Memory pool ID.

**vasi** Virtual Asynchronous Services Interface (VASI) device name.

**pager** Pager device name.

**vbsd** Virtual Block Storage Device (VBSD) name.

For suspended adapter (**-suspend** flag), the following fields are supported:

**svsa** Suspended virtual adapter.

**state** Suspended virtual adapter state.

**clientname**  
Client name.

**streamid**  
Stream ID.

**clientid**  
Client partition ID.

**vtd** Virtual target device.

**status** Virtual target device status.

**lun** Logical unit number.

**backing**  
Backing device.

**bdphysloc**  
The physical location code of the backing device.

**mirrored**  
The backing device is part of a peer-to-peer remote copy (PPRC) pair.

**-fmt** *delimiter*

Specifies a delimiter character to separate output fields.

**-hostname**

Specifies the host name and IP address of the VIOS partition.

**-net**

Specifies the supplied device is a server virtual Ethernet adapter or if used with the **-all** flag all virtual Ethernet adapters and backing devices are displayed.

**-npiv**

Displays NPIV binding information.

**-plc** *PhysicalLocatoinCode*

Specifies the device physical location code of a server virtual adapter. This flag cannot be used with the **-vadapter** flag.

**-suspend**

Displays suspended virtual adapters.

- type** Specifies the type of devices to display. The following fields are supported:
- disk** List physical backing devices.
  - lv** List logical volume backing devices.
  - optical** List optical backing devices.
  - tape** List tape backed devices.
  - file** List file backed devices.
  - file\_disk** List file backed disk devices.
  - file\_opt** List file backed optical devices.
  - cl\_disk** List clustered backing devices.
  - net** List network devices. (This option cannot be used in combination with disk, lv, or optical)
- vadapter** *VirtualServerAdapter* Specifies the device name of a server virtual adapter. This flag cannot be used with the **-plc** flag.
- vtd** *PagingDevice*

## Output Field Definitions

| Field   | Description                     |
|---------|---------------------------------|
| SVSA    | Server Virtual SCSI Adaper      |
| Physloc | Physical Location Code          |
| VTD     | Virtual Target Device           |
| LUN     | Logical Unit Number             |
| SVEA    | Server Virtual Ethernet Adapter |
| SEA     | Shared Ethernet Adapter         |

## Exit Status

| Return code | Description                                                                            |
|-------------|----------------------------------------------------------------------------------------|
| 11          | No VTDs associated with device                                                         |
| 12          | No SEAs associated with device                                                         |
| 15          | Specified device is not a server virtual SCSI adapter                                  |
| 16          | Specified device is not a server virtual Ethernet adapter                              |
| 17          | Specified device in not in the AVAILABLE state                                         |
| 63          | Specified device is not a virtual fibre channel adapter on a server logical partition. |

## Examples

- To list all virtual target devices and backing devices mapped to the server virtual SCSI adapter **vhost2**, type the command as follows:  

```
lsmapi -vadapter vhost2
```

The system displays the following message:

| SVSA           | Physloc                        | Client Partition ID |
|----------------|--------------------------------|---------------------|
| vhost2         | U9111.520.10004BA-V1-C2        | 0x00000004          |
| VTD            | vtscsi0                        |                     |
| Status         | Available                      |                     |
| LUN            | 0x8100000000000000             |                     |
| Backing device | vtd0-1                         |                     |
| Physloc        |                                |                     |
| VTD            | vtscsi1                        |                     |
| LUN            | 0x8200000000000000             |                     |
| Backing device | vtd0-2                         |                     |
| Physloc        |                                |                     |
| VTD            | vtscsi2                        |                     |
| Status         | Available                      |                     |
| LUN            | 0x8300000000000000             |                     |
| Backing device | hdisk2                         |                     |
| Physloc        | U787A.001.0397658-P1-T16-L5-L0 |                     |
| Mirrored       | false                          |                     |

- To list the shared Ethernet adapter and backing device mapped to the virtual server Ethernet adapter **ent4**, type the command as follows:

```
lsmap -vadapter
ent4 -net
```

The system displays a message similar to the following:

```
SVEA
Physloc

ent4 P2-I1/E1

SEA ent5
Backing device ent1
Status Available
Physloc P2-I4/E1
```

- To list the shared Ethernet adapter and backing device mapped to the virtual server Ethernet adapter **ent5** in script format separated by a : (colon), type the command as follows:

```
lsmap -vadapter ent5 -fmt ":"
```

The system displays a message similar to the following:

```
ent5:ent8:ent2
```

- To list all virtual target devices and backing devices, where the backing devices are of type disk or lv, type the command as follows:

```
lsmap -all -type disk lv
```

The system displays a message similar to the following:

| SVSA           | Physloc                        | Client Partition ID |
|----------------|--------------------------------|---------------------|
| vhost0         | U9117.570.10D1B0E-V4-C3        | 0x00000000          |
| VTD            | vtscsi0                        |                     |
| Status         | Available                      |                     |
| LUN            | 0x8100000000000000             |                     |
| Backing device | hdisk0                         |                     |
| Physloc        | U7879.001.DQD0KN7-P1-T12-L3-L0 |                     |
| VTD            | vtscsi2                        |                     |
| Status         | Available                      |                     |
| LUN            | 0x8200000000000000             |                     |

```

Backing device 1v04
Physloc

SVSA Physloc Client Partition
ID

vhost1 U9117.570.10D1B0E-V4-C4 0x00000000

VTD vtscsi1
Status Available
LUN 0x8100000000000000
Backing device 1v03
Physloc

```

5. To list NPIV mapping information, type the command as follows:

```
lsmap -all -npiv
```

The system displays a message similar to the following:

```

Name Physloc ClntID ClntName ClntOS
=====
vfchost0 U8203.E4A.HV40026-V1-C12 1 HV-40026 AIX

Status:NOT_LOGGED_IN
FC name:fc50 FC loc code:U789C.001.0607088-P1-C5-T1
Ports logged in:3
Flags:1 <not_mapped, not_connected>
VFC client name: VFC client DRC:

```

6. To list VIOS shared storage pool mappings, type the command as follows:

```
lsmap -all
```

The system displays the following output:

| SVSA           | Physloc                                  | Client Partition ID |
|----------------|------------------------------------------|---------------------|
| vhost0         | U8233.E8B.HV32001-V2-C2                  | 0x00000000          |
| VTD            | vtscsi0                                  |                     |
| Status         | Available                                |                     |
| LUN            | 0x8100000000000000                       |                     |
| Backing device | testLU.b1277fffd5f38acb365413b55e51638   |                     |
| Physloc        |                                          |                     |
| Mirrored       | N/A                                      |                     |
| VTD            | vtscsi1                                  |                     |
| Status         | Available                                |                     |
| LUN            | 0x8200000000000000                       |                     |
| Backing device | test_LU.8f5a2c27dce01bf443383a01c7f723d0 |                     |
| Physloc        |                                          |                     |
| Mirrored       | N/A                                      |                     |
| VTD            | vtscsi2                                  |                     |
| Status         | Available                                |                     |
| LUN            | 0x8300000000000000                       |                     |
| Backing device | c1LU.2cc8eadef6df0db3da76e21c79c7377f    |                     |
| Physloc        |                                          |                     |
| Mirrored       | N/A                                      |                     |

7. To list all cluster names, type the command as follows:

```
lsmap -all -clustername debit_cl
```

The system displays the following output:

| Physloc                   | Client Partition ID |
|---------------------------|---------------------|
| U9179.MHB.10A9DAP-V4-C111 | 0x00000065          |

```
VTD vtscsi0
LUN 0x8100000000000000
Backing device debit81_rootvg.ad089db46c2dc3769e71bf7875607c6f
```

```

Physloc Client Partition ID

U9179.MHB.10A9DAP-V1-C31 0x00000015
VTD vtscsi0
LUN 0x8200000000000000
Backing device debit01_rootvg.34cd008de649b8fde83bb16d395d48c9
```

```

Physloc Client Partition ID

U9179.MHB.10A9DAP-V3-C111 0x00000065
VTD vtscsi0
LUN 0x8200000000000000
Backing device debit81_rootvg.ad089db46c2dc3769e71bf7875607c6f
```

8. To list all suspended adapters, type the command as follows:

```
lsmmap -all -suspend
```

The system displays the following output:

```

SVSA State Client Name Stream ID Client Partition ID

susadpt0 suspended 17*9117-MMB*10002EP 0xf70e4facfc46e9a0 0x00000017
```

```
VTD vtscsi0
Status Defined
LUN 0xx81000000000000000
Backing device hdisk5
Physloc U789D.001.DQDWAXR-P1-C4-T1-W205600A0B86E4A5A-L5000000000000
Mirrored false
```

```
VTD vtscsi1
Status Defined
LUN 0x820000000000000000
Backing device hdisk6
Physloc U789D.001.DQDWAXR-P1-C4-T1-W205600A0B86E4A5A-L6000000000000
Mirrored false
```

## Related Information

The **cfgdev** command, the **chdev** command, the **chpath** command, the **lsdev** command, the **lspath** command, the **mkpath** command, the **mkvdev** command, the **rmdev** command, and the **rmpath** command.

---

## IVM lsmemdev command

### Purpose

Lists block storage devices on the Virtual I/O Server that are capable of being added to a shared memory pool for use as paging space devices. This command is operable only in an Integrated Virtualization Manager environment.

### Syntax

```
lsmemdev -r avail [-m ManagedSystem] [-p <VIOS LPAR name> | --id <VIOS LPAR ID>] [--min <minimum size in MB>] [--max <maximum size in MB>] [--filter "FilterData"] [-F "AttributeNames"] [--header] [--help]
```

## Description

The `lsmemdev` command lists block storage devices on the Virtual I/O Server that are capable of being added to a shared memory pool for use as paging space devices. Only devices in the available state are displayed.

## Flags

|                                         |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|-----------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <code>-r</code> <i>ResourceType</i>     | The type of resources to list:<br><br>avail: available resources                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <code>-m</code> <i>ManagedSystem</i>    | The name of the managed system. This attribute is optional because only one system is managed. The name can either be the user-defined name for the managed system, or be in the form <code>ttt-mmm*sssssss</code> , where <code>ttt</code> is the machine type, <code>mmm</code> is the model, and <code>sssssss</code> is the serial number of the managed system.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <code>--filter</code> <i>FilterData</i> | The filters to apply to the resources to be listed. Filters are used to select which resources of the specified resource type are to be listed. If no filters are used, all of the resources of the specified resource type are listed. For example, specific logical partitions can be listed by using a filter to specify the names or IDs of the logical partitions to list. Otherwise, if no filter is used, all the logical partitions in the managed system are listed.<br><br>The filter data consists of filter name/value pairs, which are in comma-separated value (CSV) format. The filter data must be enclosed in quotation marks.<br><br>The format of the filter data is as follows:<br>"filter-name=value,filter-name=value,..."<br><br>Note that certain filters accept a comma-separated list of values, as follows:<br>"filter-name=value,value,...",..."<br><br>When a list of values is specified, the filter name/value pair must be enclosed in quotation marks. Depending on the shell being used, nested quotation marks might need to be preceded by an escape character, which is usually a backslash (\) character.<br><br>Unless otherwise indicated, multiple values can be specified for each filter.<br><br><b>Valid filter names:</b><br><br>types   storage_pools   redundant |



**-F** *AttributeName*

A delimiter separated list of attribute names for the desired attribute values to be displayed for each resource. If no attribute names are specified, values for all of the attributes for the resource are displayed.

When this option is specified, only attribute values are displayed. No attribute names are displayed. The attribute values displayed are separated by the delimiter which was specified with this option.

This option is useful when only attribute values are desired to be displayed, or when the values of only selected attributes are desired to be displayed.

**Attribute names for partitions:**

**device\_name**

The Virtual I/O Server name of a block storage device.

**phys\_loc**

If the device has a physical location code, this attribute is displayed by default with the physical location code as the value. If the device does not have a physical location code, this attribute is not displayed by default and has a blank value.

**redundant\_capable**

This attribute is not displayed by default, but if requested has a value of 0.

**size** Size of the block storage device in megabytes.

**storage\_pool**

If the device comes from a storage pool, this attribute is displayed by default with the storage pool as the value. If the device does not come from a storage pool, this attribute is not displayed by default and has a blank value.

**type** Type of the underlying device. Valid values follow:

- logical
- phys
- storage\_pool

**--header**

Displays a header record, which is a delimiter-separated list of attribute names for the attribute values that are displayed. This header record is the first record displayed. This option is only valid when used with the **-F** option.

## Exit Status

This command has a return code of 0 on success.

## Security

This command is accessible by all users.

## Examples

To list the storage devices that are capable of being added to a memory pool, type:

```
lsmemdev -r avail
```

## Related Information

The **lshwres** and **chhwres** commands.

---

## Isnetsvc command

### Purpose

Displays the status of a network service.

### Syntax

`Isnetsvc NetworkService`

### Description

The `Isnetsvc` command displays the status of a network service. Use the `NetworkService` parameter to specify which service should have its status displayed.

### Parameters

*NetworkService*

Specify on the following values:

**cimserver**

Returns the status of the `cimserver` daemon.

**inetd**

Returns the status of the `inetd` subsystem. The `inetd` subsystem must be in the active state for the `telnet` and `ftp` daemons to be active. If the `inetd` subsystem is in the inoperative state, when you start the `startnetsvc` command with any of the supported network services it reactivates the `inetd` subsystem.

**ssh**

Returns the status of the `ssh` daemon.

**telnet**

Returns the status of the `telnet` daemon.

**ftp**

Returns the status of the `ftp` daemon.

**xntpd**

Returns the status of the `xntpd` daemon.

### Exit Status

9 Invalid network service

### Examples

1. To list the status of the `inetd` subsystem, type:

```
lsnetsvc inetd
```

This command will return either active or not active.

2. To list the status of the `telnet` daemon, type:

```
lsnetsvc telnet
```

This command will return either active or not active.

3. To list the status of the `ftp` daemon, enter:

```
lsnetsvc ftp
```

This command will return either active or not active.

### Related Information

The `entstat` command, the `hostmap` command, the `hostname` command, the `mktcpip` command, the `netstat` command, the `optimizenet` command, the `startnetsvc` command, and the `stopnetsvc` command.

---

## Isnports command

### Purpose

Lists available ports that are capable of N\_Port ID Virtualization (NPIV) and related information.

### Syntax

```
Isnports [-fmt Delimiter] [-field Fieldname]
```

### Description

The **Isnports** command displays information for all the ports capable of NPIV. If you use the name of the adapter driver (for a particular physical port) in the **vfcmap** command, the **map\_port** attribute is set for a virtual fibre channel adapter on the server logical partition. This attribute maps the virtual fibre channel adapter on the server logical partition to a physical fibre channel port.

Without a flag, information is displayed in a column. Specifying the **-fmt** flag formats the output with a specified delimiter. If no NPIV ports are in the Virtual I/O Server logical partition, the error code E\_NO\_NPIV\_PORTS(62) is displayed.

### Flags

|                              |                                                                                 |
|------------------------------|---------------------------------------------------------------------------------|
| <b>-fmt</b> <i>delimiter</i> | Specifies a delimiter character to separate output fields.                      |
| <b>-field</b>                | Specifies a list of fields to be displayed. The following fields are supported: |
| <b>name</b>                  | Physical port name                                                              |
| <b>physloc</b>               | Physical port location code                                                     |
| <b>fabric</b>                | Fabric support                                                                  |
| <b>tports</b>                | Total number of NPIV ports                                                      |
| <b>aports</b>                | Number of available NPIV ports                                                  |
| <b>swwpns</b>                | Total number of target worldwide port names supported                           |
| <b>awwpns</b>                | Number of target worldwide port names available                                 |

### Output Field Definitions

| Field   | Description                                           |
|---------|-------------------------------------------------------|
| name    | Physical port name                                    |
| physloc | Physical port location code                           |
| fabric  | Fabric support                                        |
| tports  | Total number of NPIV ports                            |
| aports  | Number of available NPIV ports                        |
| swwpns  | Total number of target worldwide port names supported |
| awwpns  | Number of target worldwide port names available       |

## Exit Status

| Return code | Description                                 |
|-------------|---------------------------------------------|
| 62          | System does not have ports capable of NPIV. |

## Examples

1. To list all the NPIV-capable ports, type:

```
lsnports
```

The system displays a message similar to the following:

| Name | physloc                    | fabric | tports | aports | swwpns | awwpns |
|------|----------------------------|--------|--------|--------|--------|--------|
| fcs0 | U789D.001.DQDMLWV-P1-C1-T1 | 1      | 64     | 64     | 2048   | 2047   |
| fcs1 | U787A.001.DPM0WVZ-P1-C1-T2 | 1      | 63     | 62     | 504    | 496    |

2. To list all the NPIV-capable ports in script format separated by two colons (::), type:

```
lsnports -fmt "::"
```

The system displays a message similar to the following:

```
fcs0::U789D.001.DQDMLWV-P1-C1-T1::1::64::64::2048::2047
fcs1::U787A.001.DPM0WVZ-P1-C1-T2::1::63::62::504::496
```

3. To list specific fields (name, swwpns, and awwpns) of all the NPIV-capable ports, type:

```
lsnports -field name swwpns awwpns
```

The system displays a message similar to the following:

| name | swwpns | awwpns |
|------|--------|--------|
| fcs0 | 2048   | 2047   |

## Related Information

The **lsmmap** and **vfcmap** commands.

---

## lspath command

### Purpose

Displays information about paths to a MultiPath I/O (MPIO) capable device.

### Syntax

```
lspath [-dev DeviceName] [-pdev Parent] [-status Status] [-conn Connection] [-field FieldName] [-fmt Delimiter]
```

```
lspath -dev DeviceName -pdev Parent [-conn Connection] -lsattr [-attr Attribute...]
```

```
lspath -dev DeviceName -pdev Parent [-conn Connection] -range -attr Attribute
```

### Description

The **lspath** command displays one of three types of information about paths to an MPIO capable device. It either displays the operational status for one or more paths to a single device, or it displays one or more attributes for a single path to a single MPIO capable device. The first syntax shown above displays the operational status for one or more paths to a particular MPIO capable device. The second syntax displays one or more attributes for a single path to a particular MPIO capable device. Finally, the third syntax displays the possible range of values for an attribute for a single path to a particular MPIO capable device.

## Displaying Path Status with the `lspath` Command

When displaying path status, the set of paths to display is obtained by searching the device configuration database for paths that match the following criteria:

- The target device name matches the device specified with the `-dev` flag. If the `-dev` flag is not present, then the target device is not used in the criteria.
- The parent device name matches the device specified with the `-pdev` flag. If the `-pdev` flag is not present, then parent is not used in the criteria.
- The connection matches the connection specified with the `-conn` flag. If the `-conn` flag is not present, then connection is not used in the criteria.
- The path status matches status specified with the `-status` flag. If the `-status` flag is not present, the path status is not used in the criteria.

If none of the `-dev`, `-pdev`, `-conn`, or `-status` flags are specified, then all paths known to the system are displayed.

By default, this command will display the information in columnar form. When no flags are specified that qualify the paths to display, the format of the output is:

```
status device parent
```

Possible values that can appear for the status column are:

### **enabled**

Indicates that the path is configured and operational. It will be considered when paths are selected for IO.

### **disabled**

Indicates that the path is configured, but not currently operational. It has been manually disabled and will not be considered when paths are selected for IO.

**failed** Indicates that the path is configured, but it has had IO failures that have rendered it unusable. It will not be considered when paths are selected for IO.

### **defined**

Indicates that the path has not been configured into the device driver.

### **missing**

Indicates that the path was defined in a previous boot, but it was not detected in the most recent boot of the system.

### **detected**

Indicates that the path was detected in the most recent boot of the system, but for some reason it was not configured. A path should only have this status during boot and so this status should never appear as a result of the `lspath` command.

## Displaying Path Attributes with the `lspath` Command

When displaying attributes for a path, the path must be fully qualified. Multiple attributes for a path can be displayed, but attributes belonging to multiple paths cannot be displayed in a single invocation of the `lspath` command. Therefore, in addition to the `-lattr`, `-dev`, and `-pdev` flags, the `-conn` flags are required to uniquely identify a single path. For example:

- if only one path between a device and a specific parent, the `-conn` flag is not required
- if there are multiple paths between a device and a specific parent, the `-conn` flag is required

Furthermore, the `-status` flag is not allowed.

By default, this command will display the information in columnar form.

attribute    value    description                    user\_settable

## Flags

- attr** *Attribute*                    Identifies the specific attribute to list. The '*Attribute*' is the name of a path specific attribute. When this flag is provided, only the identified attribute is displayed. Multiple instances of this flag may be used to list multiple attributes. If this flag is not specified at all, all attributes associated with the identified path will be listed.
- lsattr**                                Displays the attribute names, current values, descriptions, and user-settable flag values for a specific path.
- dev** *Name*                            Specifies the logical device name of the target device whose path information is to be displayed.
- field** *FieldNames*                Specifies the list of fields to display. The following fields are supported:
- status**    Status of the path
  - name**      Name of the device
  - parent**    Name of the parent device
  - conn**      Path connection.
- fmt** *Delimiter*                    Specifies a delimiter character to separate output fields.
- pdev** *Parent*                        Indicates the logical device name of the parent device of the path(s) to be displayed.
- range**                                Displays the legal values for an attribute name. The **-range** flag displays the list attribute values in a vertical column as follows:
- ```

Value1
Value2
.
.
ValueN

```
- The **-range** flag displays the range attribute values as *x...n(+i)* where *x* is the start of the range, *n* is the end of the range, and *i* is the increment.
- status** *Status* The **-status** *Status* flag indicates the status to use in qualifying the paths to be displayed. When displaying path information, the allowable values for this flag are:
- enabled** Display paths that are **enabled** for MPIIO path selection.
 - disabled** Display paths that are **disabled** from MPIIO path selection.
 - failed** Display paths that are **failed** due to IO errors.
 - available** Display paths whose **path_status** is **PATH_AVAILABLE** (that is, paths that are configured in the system, includes **enabled**, **disabled**, and **failed** paths).
 - defined** Display paths whose **path_status** is **PATH_DEFINED**.
 - missing** Display paths whose **path_status** is **PATH_MISSING**.
- conn** *Connection* Indicates the connection information to use in qualifying the paths to be displayed.

Exit Status

Return code	Description
1	Invalid status value.

Examples

1. To display, without column headers, the set of paths whose operational status is disabled, enter:
`lspath -status disabled`

The system will display a message similar to the following:

```
disabled hdisk1 scsi1
disabled hdisk2 scsi1
disabled hdisk23 scsi8
disabled hdisk25 scsi8
```

2. To display the set of paths whose operational status is failed, enter:
`lspath -status failed`

The system will display a message similar to the following:

```
failed hdisk1 scsi1
failed hdisk2 scsi1
failed hdisk23 scsi8
failed hdisk25 scsi8
```

3. If the target device is a SCSI disk, to display all attributes for the path to parent scsi0 at connection 5,0, use the command:

```
lspath -dev hdisk10 -pdev scsi0 -conn "5,0" -lsattr
```

The system will display a message similar to the following:

```
weight      1      Order of path failover selection true
```

Related Information

The **lsmap** command, the **mkpath** command, the **chpath** command, and the **rmpath** command.

lspv command

Purpose

Displays information about a physical volume within a volume group or a cluster.

Syntax

```
lspv [ -avail | -free | -size ][ -field Fieldname... ] [ -fmt Delimiter ]
```

```
lspv [ -map | -lv | -pv | -size] PhysicalVolume
```

```
lspv [ -map | -lv | -pv] PhysicalVolume [-field Fieldname ] [ -fmt Delimiter ]
```

To list the physical volumes in the cluster storage pool:

```
lspv -clustername ClusterName -sp StoragePool
```

To list all the physical volumes that are capable of being in a cluster storage pool:

```
lspv -clustername ClusterName -capable
```

Description

The **lspv** command displays information about the physical volume if the specific physical volume name is given. If the **lspv** command is run without any arguments, the default is to print every known physical

volume in the system along with its physical disk name, physical volume identifiers (PVIDs), to which volume group, if any, the physical volume belongs, and the state if the volume group is active.

When the *PhysicalVolume* parameter is used, the following characteristics of the specified physical volume are displayed:

Physical volume	Name of the physical volume
Volume group	Name of volume group. Volume group names must be unique systemwide names and can be from 1 to 15 characters long.
PV Identifier	The physical volume identifier for this physical disk.
VG Identifier	The volume group identifier of which this physical disk is a member.
PVstate	State of the physical volume. If the volume group that contains the physical volume is activated with the activatevg command, the state is active, missing, or removed. If the physical volume is deactivated with the deactivatevg command, the state is varied off.
Allocatable	Allocation permission for this physical volume.
Logical volumes	Number of logical volumes using the physical volume.
Stale PPs	Number of physical partitions on the physical volume that are not current.
VG descriptors	Number of volume group descriptors on the physical volume.
PP size	Size of physical partitions on the volume.
Total PPs	Total number of physical partitions on the physical volume.
Free PPs	Number of free physical partitions on the physical volume.
Used PPs	Number of used physical partitions on the physical volume.
Free distribution	Number of free partitions available in each intra-physical volume section.
Used distribution	Number of used partitions in each intra-physical volume section.

Flags

-avail	Lists only physical volumes that are available for use as a backing device for virtual SCSI. If the physical volume is assigned to a shared memory pool (to be used as a paging space device by a shared memory partition), it is not available and is not listed.
-capable	Lists only physical volumes those are shared storage pool capable.
-clustername	Specifies the cluster name to list the physical volumes
-free	Lists only physical volumes that are available for use as a backing device. If the physical volume is already used as a backing device or is assigned to a shared memory pool (to be used as a paging space device by a shared memory partition), it is not available and is not listed.
-field <i>FieldNames</i>	Specifies the list of fields to display. The following fields are supported if no physical volume is specified:

pvname

Physical volume disk name

pvid

Physical volume identifier

vgname

Volume group the physical volume is in

pvstate

Physical volume state (active, missing, removed, varied off)

The following fields are supported if a physical volume is specified:

pvname

Physical volume disk name

vgname

Volume group the physical volume is in

pvid Physical volume identifier

vgid Volume group identifier

pvstate
Physical volume state (active, missing, removed, varied off)

allocatable
Allocation permission for this physical volume.

stale Number of stale partitions on the disk

ppsize Physical partition size

numlv Number of logical volumes

size Number of physical partitions and total disk size

vgds Number of volume group descriptor areas within the volume group.

free Number of free partitions and free space

pvused
Number of used partitions and used space

maxreq
Maximum transfer size of physical volume

freedist
Number of free partitions available in each intra-physical volume section.

usedist Number of used partitions in each intra-physical volume section

hotspare

-field *FieldNames*
(continued)

The following fields are supported if the **-pv** flag is specified:

range A range of consecutive physical partitions contained on a single region of the physical volume.

ppstate
The current state of the physical partitions: free, used, stale, or vgda

region The intra-physical volume region in which the partitions are located.

lvname
The name of the logical volume to which the physical partitions are allocated.

type The type of the logical volume to which the partitions are allocated.

mount File system mount point for the logical volume, if applicable.

The following fields are supported if the **-map** flag is specified:

physical
Physical volume name and physical partition number.

logical Logical volume name and logical partition number. If mirrored the mirror number is also shown. If the partition is stale this is also shown.

The following fields are supported if the **-free** or **-avail** flag is specified:

pvname
Physical volume disk name.

pvid Physical volume identifier

size Size of the physical volume.

The following fields are supported if the **-size** flag and a physical volume are specified:

- pvname**
Physical volume disk name.
- pvid** Physical volume identifier
- size** Size of the physical volume.
- fmt Delimiter** Specifies a delimiter character to separate output fields.
- lv** Lists the following fields for each logical volume on the physical volume:
 - LVname**
Name of the logical volume to which the physical partitions are allocated.
 - LPs** The number of logical partitions within the logical volume that are contained on this physical volume.
 - PPs** The number of physical partitions within the logical volume that are contained on this physical volume.
 - Distribution**
The number of physical partitions, belonging to the logical volume, that are allocated within each of the following sections of the physical volume: outer edge, outer middle, center, inner middle, and inner edge of the physical volume.
 - Mount Point**
File system mount point for the logical volume, if applicable.
- map** Lists the following fields for each logical volume on the physical volume:
PVname:PPnum [LVname: LPnum [:Copynum] [PPstate]] Where:
 - PVname**
Name of the physical volume as specified by the system.
 - PPnum** Physical partition number.
 - LVname**
Name of the logical volume to which the physical partitions are allocated. Logical volume names must be system-wide unique names, and can range from 1 to 64 characters.
 - LPnum** Logical partition number. Logical partition numbers can range from 1 to 64,000.
 - Copynum**
Mirror number.
 - PPstate** Only the physical partitions on the physical volume that are not current are shown as stale.
- pv** Lists the following fields for each physical partition on the physical volume:
 - Range** A range of consecutive physical partitions contained on a single region of the physical volume.
 - State** The current state of the physical partitions: free, used, stale, or vgda.
Note: If a volume group is converted to a big vg format, it may be necessary to use some data partitions for volume group descriptor area. These partitions are marked vgda.
 - Region** The intra-physical volume region in which the partitions are located.
 - LVname**
The name of the logical volume to which the physical partitions are allocated.
 - Type** The type of the logical volume to which the partitions are allocated.
 - Mount Point**
File system mount point for the logical volume, if applicable.
- size** Displays the size of one or all physical volumes in megabytes.

-sp Specifies the storage pool.

Exit Status

See “Exit status for Virtual I/O Server commands” on page 3.

Examples

1. To display the status and characteristics of physical volume **hdisk3**, type the command as follows:

```
lspv hdisk3
```

2. To display all physical volumes in the system, type:

```
lspv
```

The system displays the following output:

```
hdisk0 0000000012345678 rootvg active
hdisk1 10000BC876543258 vg00 active
hdisk2 ABCD000054C23486 None
```

The previous example shows that physical volume **hdisk0** contains the volume group **rootvg**, and it is activated. Physical volume **hdisk1** contains the volume group **vg00**, and it is activated. Physical volume **hdisk2** does not contain an active volume group.

3. To display all physical volumes that can be virtual SCSI backing devices, type the command as follows:

```
lspv -avail
```

The system displays the following output:

```
lspv -avail
NAME          PVID                      SIZE(megabytes)
hdisk2        00c3e35c99c55ebd         7820
hdisk3        00c3e35c99c0a332         7820
hdisk4        00cbe8ddc00fbaad         7820
```

4. To display all physical volumes that can be virtual SCSI backing devices and are not currently a backing device, type the command as follows:

```
lspv -free
```

The system displays the following output:

```
hdisk3        10000BC876543258         None      None
hdisk4        ABCD000054C23486         None      None
```

5. To list all the physical volumes in the cluster storage pool, type the command as follows:

```
lspv -clustername newcluster -sp SP
```

The system displays the following output:

```
PV NAME      SIZE(MB)  PVUID
hdisk1      10240    200B75CXHW1031907210790003IBMfcp
```

6. To list the physical volumes that are shared storage pool capable , type the command as follows:

```
lspv -clustername newcluster -capable
```

The system displays the following output:

```
PV NAME      SIZE(MB)  PVUID
hdisk0      30720    200B75CXHW1025F07210790003IBMfcp
hdisk1      10240    200B75CXHW1031007210790003IBMfcp
hdisk2      10240    200B75CXHW1031107210790003IBMfcp
hdisk3      10240    200B75CXHW1031207210790003IBMfcp
hdisk4      10240    200B75CXHW1031307210790003IBMfcp
cldisk1     10240    200B75CXHW1031907210790003IBMfcp
hdisk7      10240    200B75CXHW1031A07210790003IBMfcp
```

Related Information

The `migratepv` command.

IVM Isrefcode command

Purpose

Lists reference codes for partitions or the managed system. This command is operable only in an Integrated Virtualization Manager environment.

Syntax

To list reference codes for the managed system:

```
Isrefcode -r sys [ -n Number ] [ --filter "FilterData" ] [ -F "AttributeNames" ] [ --header ] [ -m ManagedSystem ]
```

To list reference codes for partitions:

```
Isrefcode -r lpar [ -n Number ] [ --filter "FilterData" ] [ -F "AttributeNames" ] [ --header ] [ -m ManagedSystem ]
```

Description

The `Isrefcode` command lists reference codes for partitions or the managed system.

Flags

<code>-r ResourceType</code>	The type of resources to list: <ul style="list-style-type: none">• -r sys: List reference codes for the managed system.<ul style="list-style-type: none">– Attributes: refcode_num, time_stamp, refcode, word2, word3, word4, word5, word6, word7, word8, word9, fru_call_out_loc_codes– Filters: None• -r lpar: List reference codes for partitions.<ul style="list-style-type: none">– Attributes: lpar_name, lpar_id, time_stamp, refcode, word2, word3, word4, word5, word6, word7, word8, word9, fru_call_out_loc_codes– Filters: { lpar_ids lpar_names}
<code>-n Number</code>	The number of reference codes to list. The default is to list one. Reference codes are listed in order, with the most recent reference code first.
<code>-m ManagedSystem</code>	The name of the managed system. This attribute is optional because there is only one system to manage. The name might either be the user-defined name for the managed system, or be in the form tttt-mmm*sssssss, where tttt is the machine type, mmm is the model, and sssssss is the serial number of the managed system.

--filter *FilterData*

The filters to apply to the resources to be listed. Filters are used to select which resources of the specified resource type are to be listed. If no filters are used, then all of the resources of the specified resource type will be listed. For example, specific partitions can be listed by using a filter to specify the names or IDs of the partitions to list. Otherwise, if no filter is used, then all the partitions in the managed system will be listed.

The filter data consists of filter name/value pairs, which are in comma separated value (CSV) format. The filter data must be enclosed in double quotation marks.

The format of the filter data is as follows:

```
"filter-name=value,filter-name=value,..."
```

Certain filters accept a comma separated list of values, as follows:

```
"filter-name=value,value,...",..."
```

When a list of values is specified, the filter name/value pair must be enclosed in double quotation marks. Depending on the shell being used, nested double quotation marks might need to be preceded by an escape character, which is usually a backslash character.

Unless otherwise indicated, multiple values can be specified for each filter.

Valid filter names for -r lpar:

lpar_ids

ID of the partition to view

lpar_names

Name of the partitions to view

-F *AttributeNames* A delimiter separated list of attribute names for the desired attribute values to be displayed for each resource. If no attribute names are specified, then values for all of the attributes for the resource will be displayed.

When this option is specified, only attribute values will be displayed. No attribute names will be displayed. The attribute values displayed will be separated by the delimiter which was specified with this option.

This option is useful when only attribute values are desired to be displayed, or when the values of only selected attributes are desired to be displayed.

Attribute Names:

fru_call_out_loc_codes
Location codes of the field replaceable unit related to the reference code

lpar_id Unique integer identifier for the partition

lpar_name
Name of the partition

refcode The ASCII reference code string

refcode_num
The sequence number of the reference code.

time_stamp
The time that the reference code was created in the format: MM/DD/YYYY HH:MM:SS where MM is the two digit month, DD is the two digit day, YYYY is the four digit year, HH is the two digit hour, MM is the two digit minute, and SS is the two digit second.

fru_call_out_loc_codes
The Field Replaceable Unit (FRU) numbers for the FRUs whose absence or failure caused this reference code to be received. This field might be used for other values.

--header Display a header record, which is a delimiter separated list of attribute names for the attribute values that will be displayed. This header record will be the first record displayed. This option is only valid when used with the -F option.

Exit Status

This command has a return code of 0 on success.

Security

This command is accessible by all users.

Examples

1. To list the current reference code for the managed system, type:
lsrefcode -r sys
2. To list the current reference code for all partitions, type:
lsrefcode -r lpar
3. To list the last 25 reference codes for partitions p1 and p2, only viewing the lpar_id and refcode attributes, type:
lsrefcode -r lpar -n 25 --filter "\lpar_names=p1,p2\" -F lpar_id,refcode

Isrep command

Purpose

Lists and displays information about the Virtual Media Repository.

Syntax

Isrep [-**field** *Fieldname*] [-**fmt** *Delimiter*]

Description

The **Isrep** command displays information about the Virtual Media Repository. The following information is displayed; the size of the repository and free space, parent storage pool, size and free space, and the name, size, associated virtual target device, and access state of all virtual optical media in the repository.

This command will provide full scripting support through the use of the **-field** and **-fmt** flags.

Flags

-field *FieldName*

size	Total Size
free	Free space
parent pool	Parent storage pool name
parent size	Parent storage pool size
parent free	Parent storage pool free space
name	Number of backing device file
file size	Backing device file size
optical	Virtual Target Device off the virtual optical device media is loaded in
access	Media access, read-only (ro) or read-write (rw)

-fmt *Delimiter*

Specifies a delimiter character to separate output fields.

Examples

1. To display information about the Virtual Media repository and all virtual media with the repository, type the following command:

```
lsrep
```

The system displays output similar to the following:

SIZE(mb)	FREE(mb)	PARENT SIZE	PARENT FREE		
2039	299	18352	16304		
Name		File Size	Optical		Access
clientCD		640	vtopt3		ro
installDVD1		1000	vtopt16		rw
installDVD2		100	None		rw

2. To display only the parent storage pool size, type the following command:

```
lsrep -field "parent size"
```

```
Parent Size  
30624
```

lsrole command

Purpose

Displays role attributes.

Syntax

```
lsrole [-c] | [-f] | [-C] [-a List] {ALL | role1 [, role2] ...}
```

Description

The **lsrole** command displays the role attributes. You can use this command to list all attributes of all the roles or all the attributes of specific roles. Since there is no default parameter, you must enter the **ALL** keyword to see the attributes of all the roles. By default, the **lsrole** command displays all role attributes. To view selected attributes, use the **-aList** flag. If one or more attributes cannot be read, the **lsrole** command lists as much information as possible.

By default, the **lsrole** command lists attributes of each role in one line. It displays attribute information as *Attribute=Value* definitions, each separated by a blank space. To list the role attributes in stanza format, use the **-f** flag. To list the information as colon-separated records, use the **-c** flag.

Security

The **lsrole** command is a privileged command. To run the command successfully, assume a role that has the authorization **vios.security.role.list**.

Flags

-a List Lists the attributes to display. The *List* variable can include any attribute that is defined in the **chrole** command. To specify more than one attribute, use a blank space between attribute names. If an empty list is specified, only the role names are displayed. In addition to the attributes defined in the **chrole** command, the following attributes can also be listed with the **-a** flag:

all_auths

Traverses the role hierarchy of the specified roles and gathers all the authorizations. The *all_auths* attribute differs from the *authorizations* attribute because the **lsrole** command only lists the explicit authorizations of the specified roles for that attribute.

user Displays the users that are granted the specified roles.

description

Displays the text description of the role as indicated by the *dfltmsg*, *msgcat*, *msgset*, and *msgnum* attributes for the role.

-c Displays the role attributes in colon-separated records, as follows:

```
# role: attribute1: attribute2: ...  
Role: value1: value2: ...
```


- C** Displays the role attributes in colon-separated records, which are easier to parse than the output of the **-c** flag:
- ```
role:attribute1:attribute2: ...
 role:value1:value2: ...
 role2:value1:value2: ...
```
- The output is preceded by a comment line that has details about the attribute represented in each colon-separated field. If you specify the **-a** flag, the order of the attributes matches the order specified in the **-a** flag. If a role does not have a value for a given attribute, an empty field is displayed. The last field in each entry is ended by a newline character rather than a colon.
- f** Displays the output in stanzas, with each stanza identified by a role name. Each *Attribute=Value* pair is listed in a separate line:
- ```
Role:
  attribute1=value
  attribute2=value
  attribute3=value
```

Exit Status

0 for success.

Non-zero for command error.

Examples

- To display the attributes *rolelist* and *groups* of the role *ManageAllUsers* in a colon format, type the following command:

```
lsrole -c -a rolelist groups ManageAllUsers
```
- To list all attributes of the *ManageAllUsers* role, type the following command:

```
lsrole ManageAllUsers
```

Issecattr command

Purpose

Lists the security attributes of a command, a device, a privileged file, or a process.

Syntax

```
Issecattr {-c | -d | -p[-h][-A] | -f} [-C | -F] [-a List] { ALL | Name [,Name] ...}
```

Description

The **Issecattr** command lists the security attributes of one or more commands, devices, or processes. The command interprets the *Name* parameter as either a command, a device, a privileged file, or a process based on whether the **-c** (command), **-d** (device), **-f** (privileged file), or **-p** (process) flag is specified. If the **-c** flag is specified, the *Name* parameter must include the full path to the commands. If the **-d** flag is specified, the *Name* parameter must include the full path to the devices. If the **-f** flag is specified, the *Name* parameter must include the full path to the file. If the **-p** flag is specified, the *Name* parameter must be the numeric process identifier (PID) of an active process on the system. Use the **ALL** keyword to list the security attributes for all commands, devices, files, or processes. By default, the **Issecattr** command displays all of the security attributes for the specified object. To view the selected attributes, use the **-a List** flag.

By default, the **lssecattr** command lists the security attributes on one line. It displays the attribute information as the definitions of `Attribute=Value`, each separated by a blank space. To list the attributes in stanza format, use the **-F** flag. To list the attributes as colon-separated records, use the **-C** flag.

Important: Users can list the security attributes for their own commands, scripts, and AIX commands, but cannot list security attributes for VIOS commands.

Note: When security attributes are set for a command or device, user must run the **setkst** command to update the kernel database.

Security

The **lssecattr** command is a privileged command. To run the command successfully, assume a role that has the following authorizations:

vios.security.cmd.list

Required to list the attributes of a command with the **-c** flag.

vios.security.device.list

Required to list the attributes of a device with the **-d** flag.

vios.security.file.list

Required to list the attributes of a device with the **-f** flag.

vios.security.proc.list

Required to list the attributes of a process with the **-p** flag.

Flags

-A	Displays the list of authorizations used by a specified process. This flag can only be used with the -p flag.
-a	Specifies the list of attributes to be listed.
-C	Displays the privileged security attributes in colon-separated records.
-c	The <i>Name</i> parameter specifies the full paths to one or more commands on the system that have entries in the <code>/etc/security/privcmds</code> privileged command database.
-d	The <i>Name</i> parameter specifies the full paths to one or more devices on the system that have entries in the <code>/etc/security/privdevs</code> privileged device database.
-f	The <i>Name</i> parameter specifies the full paths to one or more files on the system that have entries in the <code>/etc/security/privfiles</code> privileged files database.
-F	Displays the output in stanza format, with each stanza identified by an object name. Each pair of <code>Attribute=Value</code> is listed on a separate line.
-h	Displays the full hierarchy of privileges for the process. By default, only the highest level of privilege is listed.
-p	The <i>Name</i> parameter specifies the numeric process identifiers (PID) of one or more active processes on the system.

Exit Status

0 for success.

Non-zero for command error.

Examples

1. To display the access authorization and the innate privileges of the `/home/padmin/mybinary` command, type the following command:

```
lssecattr -c -a accessauths innateprivs /home/padmin/mybinary
```
2. To display all the security attributes of the `/dev/mydev` device, type the following command:

- ```
lssecattr -d /dev/mydev
```
3. 4. To display the privileges for the effective and used privilege sets of two processes in a colon format, type the following command:
 

```
lssecattr -p -C -a eprivs uprivs 38483,57382
```
  4. To display the read authorization list of the `/etc/security/user` file, type the following command:
 

```
lssecattr -f -a readauths /etc/security/user
```
  5. To display the used authorizations for a process in a stanza format, type the following command:
 

```
lssecattr -F -p -A 34890
```

## lssp command

### Purpose

Lists and displays information about storage pools.

### Syntax

To list all available storage pools:

```
lssp [-type PoolType [-field Fieldname] [-fmt Delimiter]
```

To display information about a specific storage pool:

```
lssp -detail | -bd [-sp StoragePool] [-field Fieldname] [-fmt Delimiter]
```

To display the default storage pool:

```
lssp -default
```

To display all the VIOS cluster objects in a shared storage pool:

```
lssp -clustername ClusterName -sp StoragePool -bd
```

To display storage pools in a specific cluster:

```
lssp -clustername ClusterName
```

### Description

The **lssp** command displays information about storage pools in the Virtual I/O Server. If no flags are specified, a list of all defined storage pools, their total size, free space, minimum allocation size, and number of backing devices contained in the pool and the type of pool is displayed. If the **-type** flag is specified only storage pools of the indicated type are displayed. If the **-detail** flag is specified, detailed information about the storage pool is displayed. If the **-bd** flag is specified, a list of all the backing devices in the specified (or default) storage pool is displayed along with their size and associated virtual target device, or None, and virtual host adapter, or None. If the **-default** flag is specified, the default storage pool is displayed.

This command provides full scripting support through the use of the **-field** and **-fmt** flags.

### Flags

- |                                        |                                                       |
|----------------------------------------|-------------------------------------------------------|
| <b>-clustername</b> <i>ClusterName</i> | Specifies the cluster name to list the storage pools. |
| <b>-default</b>                        | Displays the default storage pool.                    |

**-field** *FieldName* The following fields are supported if no flags are specified:

- pool** Storage pool name
- size** Total size
- free** Free space
- alloc** Minimum allocation size
- bds** Number of backing devices
- type** Type of pool

The following fields are supported if the **-detail** flag is specified for a logical volume pool:

- pvname** Name of the physical volume.
- pvid** Physical volume identifier
- size** The size of the physical volume.

The following fields are supported if the **-detail** flag is specified for a file pool:

- name** Name of the parent storage pool.

The following fields are supported if the **-bd** flag is specified:

- bdname** Backing device name.
- size** The size of the logical volume.
- vtd** Virtual target device.
- svsa** Server virtual SCSI adapter.

**-fmt** *Delimiter* Specifies a delimiter character to separate output fields.

**-bd** Displays information about the backing devices or logical units in the storage pool.

**-detail** Displays detailed information about the file that is backed up and the logical volume-based storage pool.

**-sp** *StoragePool* Specifies which storage pool to display information about.

**-type** *PoolType* Specifies the type of pool to list. The following fields are supported:

- lvpool** List only logical volume pools.
- fbpool** List only file pools.

## Examples

- To list all the storage pools, type the command as follows:

```
lssp
```

The system displays output as follows:

| Pool            | Size(mb) | Free(mb) | Alloc | Size(mb) | Bds    | Type |
|-----------------|----------|----------|-------|----------|--------|------|
| PRODClient      | 30624    | 28576    | 32    | 0        | LVPOOL |      |
| rootvg          | 30656    | 14208    | 64    | 0        | LVPOOL |      |
| DEVClient       | 18352    | 18352    | 16    | 0        | LVPOOL |      |
| PRODClient_FBP  | 1016     | 985      | 32    | 3        | FBPOOL |      |
| PRODClient_FBP2 | 1016     | 1005     | 32    | 1        | FBPOOL |      |

- To display the default storage pool, type the command as follows:

```
lssp -default
```
- To display detailed information about the logical volume storage pool `sp_sp00`, type the command as follows:

```
lssp -detail -sp sp_sp00
```

The system displays output as follows:

| NAME   | PVID             | SIZE(megabytes) |
|--------|------------------|-----------------|
| hdisk3 | 00cdfd8c85bd4b2e | 34624           |
| hdisk2 | 00cdfd8c525d94a2 | 34624           |

- To display information about the backing devices in the logical volume storage pool rootvg, type the command as follows:

```
lssp -bd -sp rootvg
```

The system displays output as follows:

| NAME | SIZE(megabytes) | VTD     | SVSA   |
|------|-----------------|---------|--------|
| lv01 | 96              | vtscsi1 | vhost0 |
| lv02 | 64              | vtscsi2 | vhost0 |

- To list all storage pools within a cluster, type the command as follows:

```
lssp -clustername newcluster
```

The system displays output as follows:

```
POOL_NAME: SP
POOL_SIZE: 714240
FREE_SPACE: 702047
TOTAL_LU_SIZE: 10240
TOTAL_LUS: 15
POOL_TYPE: CLPOOL
POOL_ID: 000000000903311C000000004F42A526
```

- To list all the VIOS cluster objects in a shared storage pool, type the command as follows:

```
lssp -clustername newcluster -sp SP -bd
```

**Note:** The system displays *-1* if you get an error while retrieving the number of logical units (LUs) within the shared storage pool.

The system displays output as follows:

| Lu Name   | Size(mb) | ProvisionType | Lu Udid                          |
|-----------|----------|---------------|----------------------------------|
| Test_LU   | 10240    | THIN          | 7b74e7b07a0c4885cf136683861a23b8 |
| TestI2_LU | 20280    | THIN          | 8a2db5dd43b6de0480bff8143d37b61d |

---

## Issvc command

### Purpose

Lists the available agents and services.

### Syntax

```
Issvc AgentName
```

### Description

The **Issvc** command lists all the available agents and services that can be managed by the Virtual I/O Server command-line interface. If an agent name or service name is passed to the **Issvc** command, a list of attributes with their configured values is displayed. These agents and services are managed by the **cfgsvc**, **startsvc**, and **stopsvc** commands.

### Agent or service names

The following agents can be managed by the Virtual I/O Server.

**DIRECTOR\_agent**                Lists the attributes and values for the Systems Director agent.

|                     |                                                                                                                                           |
|---------------------|-------------------------------------------------------------------------------------------------------------------------------------------|
| <b>ITM_premium</b>  | Lists the attributes and values for the Tivoli Monitoring agents.                                                                         |
| <b>ITM_cec</b>      |                                                                                                                                           |
| <b>TSM_base</b>     | Lists the attributes and values for the Tivoli Storage Manager agent.                                                                     |
| <b>ITUAM_base</b>   | Lists the attributes and values for the Tivoli Usage and Accounting Manager agent.                                                        |
| <b>TPC</b>          | Lists the attributes and values for the TotalStorage Productivity Center agent.                                                           |
| <b>perfmgr</b>      | Does not have attributes that the user can set. See the <b>postprocesssvc</b> command for more information about working with this agent. |
| <b>ipsec_tunnel</b> | Lists the secure tunnels that are created on the system. If none are created, the message <i>No tunnels created</i> is displayed.         |
| <b>ILMT</b>         | Lists the IBM License Metric Tool (ILMT) server that is configured and the agent status on Virtual I/O Server.                            |

## Exit Status

See “Exit status for Virtual I/O Server commands” on page 3.

## Examples

1. To list available agents, type the command as follows:

```
lssvc
```

2. To list values of attributes by agent name, type one or more of the commands as follows:

```
lssvc ITM_premium
```

```
lssvc ITUAM_base
```

```
lssvc TSM_base
```

These commands produce the output as follows:

```
$lssvc ITM_premium
HOSTNAME:tems_server
MANAGING_SYSTEM:hmc_console
RESTART_ON_REBOOT:TRUE

$lssvc ITUAM_base
ACCT_DATA0:
ACCT_DATA1:
ISYSTEM:
IPROCESS:

$lssvc TSM_base
SERVERNAME:
SERVERIP:
NODENAME:
```

3. To list the mandatory attributes (for TotalStorage Productivity Center) with their configured values, type the command as follows:

```
$lssvc TPC
```

This command produce the output as follows:

```
$lssvc TPC
A:
S:
devAuth:
caPass:
```

4. To list the status of ipsec tunnels on the Virtual I/O Server, type the command as follows:

```
$lssvc ipsec_tunnel
```

This command produces the output similar as follows:

| Phase | Tun Id | Status  | Local Id | Remote Id |
|-------|--------|---------|----------|-----------|
| 1     | 1      | Dormant | N/A      | 4.3.2.1   |
| 2     | 1      | Dormant | 1.2.3.4  | 4.3.2.1   |

## ILMT

1. To list ILMT details, type the command as follows:

```
$lssvc ILMT
```

This command produces the output as follows:

```
Server: 1.2.3.4
Agent Status: active
```

## Related Information

The **cfgsvc** command, the **startsvc** command, the **stopsvc** command, and the **postprocesssvc** command.

For more information about the various agents, see the following information:

- Tivoli software and the Virtual I/O Server
- Configuring the Tivoli agents and clients on the Virtual I/O Server
- Systems Director software
- Configuring the Systems Director agent

---

## IVM lssvcevents command

### Purpose

List attributes of console or serviceable events. This command is operable only in an Integrated Virtualization Manager environment.

### Syntax

To list console events:

```
lssvcevents -t console [-d NumberDays | -i NumberMinutes] [--filter "FilterData"] [-F "AttributeName"] [--header] [-m ManagedSystem]
```

To list serviceable events:

```
lssvcevents -t hardware [-d NumberDays | -i NumberMinutes] [--filter "FilterData"] [-F "AttributeName"] [--header] [-m ManagedSystem]
```

To list field replaceable units (frus) for a specific serviceable events

```
lssvcevents -t fru --filter "FilterData" [-F "AttributeName"] [--header] [-m ManagedSystem]
```

To list comments for a specific serviceable event

```
lssvcevents -t comment --filter "FilterData" [-F "AttributeName"] [--header] [-m ManagedSystem]
```

To list service objects associated with a specific serviceable event

```
lssvcevents -t service_object -filter "FilterData" [-F "AttributeName"] [--header] [-m ManagedSystem]
```

To list the status of dynamic LPAR events:

lssvcevents -t dlpdr [ --filter "FilterData" ] [ -F "AttributeNames" ] [ --header ] [ -m *ManagedSystem* ]

## Description

The `lssvcevents` command lists attributes of console or serviceable events.

## Flags

**-t** *EventType*

The type of event to list:

### console

Console events - events created by the applications

- **Attributes:** time, userid, pid, name, category, severity, time\_bin, text
- **Filters:** severities, categories, name

### hardware

Serviceable events

- **Attributes:** problem\_num, pmh\_num, refcode, status, first\_time, last\_time, sys\_name, reporting\_name, sys\_mtms, reporting\_mtms, enclosure\_mtms, failing\_mtms, text, firmware\_fix, created\_time, analyzing\_sfp, refcode\_extension, firmware\_pkg\_name, firmware\_pkg\_status, reporting\_sfp\_name, reporting\_sfp\_mtms, failing\_sfp\_mtms, severity, lpar\_id, lpar\_name, lpar\_hostname, lpar\_os\_type, notification\_type, notification\_status, duplicate\_count, analyzing\_sfp\_mtms, analyzing\_sfp\_name, called\_home\_sys\_mtms, sys\_log\_id, platform\_log\_id, subsystem\_id, creator\_id, ipl\_state, symptom, failing\_lpar\_id, failing\_lpar\_name, failing\_lpar\_os\_type, fru\_part\_nums, fru\_phys\_locs, first\_time\_bin, last\_time\_bin, created\_time\_bin
- **Filters:** status, problem\_nums, refcodes, fru\_part\_nums, fru\_phys\_locs, reporting\_mtms, failing\_mtm

**fru** Field replaceable units (frus) for a serviceable event.

- **Attributes:** part\_num, class, description, phys\_loc, prev\_replaced, replaced\_time, serial\_num, replacement\_grp, ccin, logic\_ctl\_mtms, power\_ctl\_mtms, replaced\_time\_bin
- **Filters:** problem\_nums (required)

### comment

Comments for a serviceable event.

- **Attributes:** time,commenter,text
- **Filters:** problem\_nums (required), status

### service\_object

Service objects for a serviceable event. These are not accessible to users without the DEUser or SRUser roles.

- **Attributes:** key, sys\_log\_id, notification\_type, platform\_log\_id, severity, creator\_id, refcode, subsystem\_id, lpar\_os\_type, failing\_mtms, text, lpar\_id, lpar\_name, lpar\_hostname, first\_time, last\_time, duplicate\_count, eed\_ptr, first\_time\_bin, last\_time\_bin
- **Filters:** problem\_nums (required), status



**dlpar** Status of dynamic LPAR events - events used to synchronize resources that can be dynamically configured while the partition is active. These can be used to determine why the runtime and pending values of a resource are not synchronized. Typically, the last two status records are stored per partition per resource type.

- **Attributes:** lpar\_id, resource\_type, sequence\_num, status\_code, time, internal\_rc, drmgr\_cmd, drmgr\_rc, drmgr\_stdout, drmgr\_stderr
- **Filters:** lpar\_ids, resource\_types

**-d** *NumberDays*

The number of previous days for which to view events. This might not be used with the **-i** flag, and is only applicable to **-t console**, and **-t hardware**. If this flag and the **-i** flag are omitted, the default is 7 days.

**-i** *NumberMinutes*

The number of minutes for which to view events. This might not be used with the **-d** flag, and is only applicable to **-t console**, and **-t hardware**.

**-m** *ManagedSystem*

The name of the managed system. This attribute is optional because there is only one system to manage. The name might either be the user-defined name for the managed system, or be in the form tttt-mmm\*sssssss, where tttt is the machine type, mmm is the model, and sssssss is the serial number of the managed system.

**--filter** *FilterData*

The filters to apply to the resources to be listed. Filters are used to select which resources of the specified resource type are to be listed. If no filters are used, then all of the resources of the specified resource type will be listed. For example, specific partitions can be listed by using a filter to specify the names or IDs of the partitions to list. Otherwise, if no filter is used, then all the partitions in the managed system will be listed.

The filter data consists of filter name/value pairs, which are in comma separated value (CSV) format. The filter data must be enclosed in double quotation marks.

The format of the filter data is as follows:

```
"filter-name=value,filter-name=value,..."
```

Note that certain filters accept a comma separated list of values, as follows:

```
"filter-name=value,value,...",..."
```

When a list of values is specified, the filter name/value pair must be enclosed in double quotation marks. Depending on the shell being used, nested double quotation marks might need to be preceded by an escape character, which is usually a `'\'` character.

Unless otherwise indicated, multiple values can be specified for each filter.

**Valid filter names for -t console:**

severities, categories, name

**Valid filter names for -t hardware:**

status, problem\_nums, refcodes, fru\_part\_nums, fru\_phys\_locs, reporting\_mtms, failing\_mtms

**Valid filter names for -t fru:**

problem\_nums (required)

**Valid filter names for -t comments and -t service\_objects:**

problem\_nums (required), status

**Valid filter names for -t dlpar:**

lpar\_ids, resource\_types

**-F AttributeNames**

A delimiter separated list of attribute names for the desired attribute values to be displayed for each resource. If no attribute names are specified, then values for all of the attributes for the resource will be displayed.

When this option is specified, only attribute values will be displayed. No attribute names will be displayed. The attribute values displayed will be separated by the delimiter which was specified with this option.

This option is useful when only attribute values are desired to be displayed, or when the values of only selected attributes are desired to be displayed.

**Attribute names:**

**analyzing\_sfp**

Name or MTMS of the service focal point system analyzing the event.

**analyzing\_sfp\_mtms**

MTMS of the service focal point system analyzing the event.

**analyzing\_sfp\_name**

Name of the service focal point system analyzing the event.

**called\_home\_sys\_mtms**

MTMS of the system to which the event was called home.

**category**

Category or client type of the console event. Valid values:

- GUI: Web interface
- CLI: Command line interface
- AP: Access process

**ccin** CCIN of the FRU

**class** Class or type of the FRU

**commenter**

Name of the person adding a comment to the event

**created\_time**

Time the event was created.

**creator\_id**

The ID of the entity that created the event. Valid values:

- C: Hardware Management Console
- E: Service processor
- H: hypervisor
- W: Power
- L: Partition firmware

**description**

Description of the FRU

**drmgr\_cmd**

The command used for a dynamic LPAR event. The `drmgr` command is run on the client partition's operating system to synchronize a particular resource.

**drmgr\_rc**

The return code for the command used for a dynamic LPAR event. The `drmgr` command is run on the client partition's operating system to synchronize a particular resource.

**drmgr\_stdout**

The standard output for the command used for a dynamic LPAR event. The `drmgr` command is run on the client partition's operating system to synchronize a particular resource.

**drmgr\_stderr**

The standard error for the command used for a dynamic LPAR event. The `drmgr` command is run on the client partition's operating system to synchronize a particular resource.

**duplicate\_count**

Number of duplicates for this event.

**enclosure\_mtms**

Enclosure MTMS

**failing\_lpar\_id**

Unique ID for the failing partition

**failing\_lpar\_name**

Name of the failing partition

**failing\_lpar\_os\_type**

Operating system type of the failing partition

**failing\_mtms**

MTMS of the failing system

**failing\_sfp\_mtms**

MTMS of the failing systems service focal point

**firmware\_fix**

Indicates whether a firmware fix is available for the event

**firmware\_pkg\_name**

Package name of a possible firmware fix

**firmware\_pkg\_status**

Package status of a possible firmware fix

**first\_time**

First time this event was reported

**fru\_part\_nums**

Part numbers of the FRUs

**fru\_phys\_locs**

Unique physical location codes of the FRUs

**internal\_rc**

The return code for a dynamic LPAR event. This will be non zero only if the dynamic LPAR command was unexpectedly not able to be sent to the client partition.

**ipl\_state**

State of the system when this event occurred

**key**

Unique ID for a service object for a particular serviceable event.

**last\_time**

Last time this event was reported

**logic\_ctl\_mtms**

MTMS of the unit that logically controls the unit that the FRU is located in

**lpar\_hostname**

Hostname of the partition that created this event

**lpar\_id** ID of the partition that created this event

**lpar\_name**  
Name of the partition that created this event

**lpar\_os\_type**  
Operating system type of the partition that created this event

**name** Name of the application that created the console event

**notification\_status**  
The status of the notification type

**notification\_type**  
The notification type for the event. Valid values:

- Yes: Call home
- No: Customer notify

**part\_num**  
Part number of the FRU

**phys\_loc**  
Unique physical location code of the FRU

**pid** Process ID of the process generating the console event

**platform\_log\_id**  
Unique ID of the platform log for this event

**pmh\_num**  
PMH or tracking number

**power\_ctl\_mtms**  
MTMS of the unit that power controls the unit that the FRU is located in

**prev\_replaced**  
Indicates if the FRU has been previously replaced

**problem\_num**  
A unique ID for the event

**refcode** Reference code for the event

**refcode\_extension**  
Extended reference code for the event

**replaced\_time**  
Time the FRU was replaced

**replacement\_grp**  
Replacement priority and grouping for the FRU. Valid values:

- H: Multiple high priority FRUs should be acted on as a group
- M: Medium priority FRUs should be acted on, one at a time, in the order given.
- A: Medium priority group A FRUs should be acted on as a group.
- B: Medium priority group B FRUs should be acted on as a group.
- C: Medium priority group C FRUs should be acted on as a group.
- L: Low priority FRUs should be acted on only after all other priority call-outs failed to resolve the problem.

**reporting\_mtms**  
MTMS of the reporting unit

**reporting\_name**  
Name of the reporting unit

**reporting\_sfp\_mtms**  
MTMS of the service focal point reporting the event

**reporting\_sfp\_name**  
Name of the service focal point reporting the event

**resource\_type**  
The resource type for the dynamic LPAR event. These are all the resources that can be dynamically reconfigured. Valid values:

- mem
- proc
- proc\_units
- uncap\_weight
- memory
- io\_entitled\_mem
- mem\_weight

**sequence\_num**  
The sequence number for this dynamic LPAR event. Each resource type might have more than one entry. The sequence number distinguishes the entries, and increments over time.

**serial\_num**  
Serial number of the FRU

**severity**  
Severity of the event. Valid values:

- 10: Recovered error, general (10)
- 20: Predictive error, general (20)
- 21: Predictive error, degraded performance (21)
- 22: Predictive fault might be corrected after platform re-IPL (22)
- 23: Predictive error, fault might be corrected after IPL, degraded performance (23)
- 24: Predictive error, loss of redundancy (24)
- 40: Unrecovered error, general (40)
- 41: Unrecovered error, bypassed with degraded performance (41).
- 44: Unrecovered error, bypassed with loss of redundancy (44)
- 45: Unrecovered error, bypassed with loss of redundancy and performance (45)
- 48: Unrecovered error, bypassed with loss of function (48)
- 60: Error on diagnostic test, general (60)
- 61: Error on diagnostic test, resource might produce incorrect result (61)

**status** Status of the event. Valid values:

- Open: Event is in the open state
- Closed: Event has been closed

### **status\_code**

Status code for the event. Valid values for dynamic LPAR events:

- 0: Synchronization successful
- 1: Synchronization in progress
- 2: Resource will not synchronize because the partition is a workload group participant
- 3: Resource will not synchronize because the partition communication state is not active
- 4: Resource will not synchronize because the partition does not support dynamic LPAR of this resource type
- 5: Resource will not synchronize because the partition is not in the Running state
- 6: Resource will not synchronize because the partition is unable to remove any more memory dynamically
- 7: Resource will not synchronize because the synchronization command failed to run for an unknown reason
- 8: Resource is not synchronized because the RMC command failed. The system will retry. If the partition is in the Running state with an active network connection, check the return code and contact your support representative.
- 9: Resource is not synchronized because the drmgr command on the partition failed. The system will retry. Check the return code, and the command output.
- 10: Resource will not synchronize because the requested assigned value is less than the current minimum. Restart your partition in order to complete the synchronization.
- 11: Resource will not synchronize because the requested assigned value is greater than the current maximum. Restart your partition in order to complete the synchronization.
- 12: Resource will not synchronize because the pending and current processing modes do not match. Restart your partition in order to complete the synchronization.
- 13: Resource will not synchronize because IVM is unable to determine dynamic LPAR capabilities of the logical partition.
- 255: Resource synchronization has not yet been attempted. It might take a few seconds depending on your system utilization before synchronization is attempted.

**subsystem\_id**

The subsystem causing the event. Valid values include:

- 10 - 1F: Processor subsystem including internal cache
- 20 - 2F: Memory subsystem including external cache
- 30 - 3F: I/O subsystem (hub, bridge, bus)
- 40 - 4F: I/O adapter, device and peripheral
- 50 - 5F: CEC hardware
- 60 - 6F: Power/Cooling subsystem
- 70 - 79: Other subsystem
- 7A - 7F: Surveillance error
- 80 - 8F: Platform firmware
- 90 - 9F: Software
- A0 - AF: External environment

**symptom**

Symptom of the event

**sys\_log\_id**

Unique ID of the system log for the event

**sys\_mtms**

MTMS of the system

**text**

Text of the event

**time**

Time of the console or dynamic LPAR event

**time\_bin**

Time in milliseconds since Jan 1, 1970

**userid**

User ID of the user that ran the command that caused the console event

**--header**

Display a header record, which is a delimiter separated list of attribute names for the attribute values that will be displayed. This header record will be the first record displayed. This option is only valid when used with the **-F** option.

## Exit Status

This command has a return code of 0 on success.

## Security

This command is accessible by all users, except for certain types and attributes as noted, which require SR User or DE User roles.

## Examples

1. To list the serviceable events that occurred today, type:  

```
lssvcevents -t hardware -d 0
```
2. To list the console events that occurred within the past 3 days, type:  

```
lssvcevents -t console -d 3
```
3. To list all of the open serviceable events for the system, type:  

```
lssvcevents -t hardware --filter "status=open"
```
4. To list the associated FRUs for a specific serviceable event, type:  

```
lssvcevents -t fru
--filter problem_nums=6013EFFF-205E9F22-4CC931E5-F892358-A0F6C1D6
```

## Related Information

The `chsvcevent` command, and the `mksvcevent` command.

---

## Issw command

### Purpose

Lists installed software products.

### Syntax

```
Issw [-hist]
```

### Description

The `Issw` command displays information about installed filesets or fileset updates. If the `-hist` parameter is not specified, the name, most recent level, state, and description of all filesets is displayed. Part information (usr, root, and share) is consolidated into the same listing. For formatted filesets, it displays the most recent maintenance level. Any interim fixes on the system are also displayed.

If the `-hist` flag is specified, installation and update history information is displayed.

### Output Values

The following sections define terms used in several of the output fields. Note that not all output values are defined here. Only the ones that require explanation are defined.

### State Values

The `state` field in the `Issw` output gives the state of the fileset on your system. It can have the following values:

| State              | Description                                                                                                                                                                                                                                                                                                                     |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>APPLIED</b>     | The specified fileset is installed on the system. The <b>APPLIED</b> state means that the fileset can be removed with the <code>updateios</code> command and the previous level of the fileset restored.                                                                                                                        |
| <b>APPLYING</b>    | An attempt was made to apply the specified fileset, but it did not complete successfully, and cleanup was not performed.                                                                                                                                                                                                        |
| <b>BROKEN</b>      | The specified fileset or fileset update is broken and should be reinstalled before being used.                                                                                                                                                                                                                                  |
| <b>COMMITTED</b>   | The specified fileset is installed on the system. The <b>COMMITTED</b> state means that a commitment has been made to this level of the software. A committed fileset update cannot be rejected, but a committed fileset base level and its updates (regardless of state) can be removed by the <code>updateios</code> command. |
| <b>EFIX LOCKED</b> | The specified fileset was installed successfully and locked.                                                                                                                                                                                                                                                                    |
| <b>OBSOLETE</b>    | The specified fileset was installed with an earlier version of the operating system but has been replaced by a repackaged (renamed) newer version. Some of the files that belonged to this fileset have been replaced by versions from the repackaged fileset.                                                                  |
| <b>COMMITTING</b>  | An attempt was made to commit the specified fileset, but it did not complete successfully, and cleanup was not performed.                                                                                                                                                                                                       |
| <b>REJECTING</b>   | An attempt was made to reject the specified fileset, but it did not complete successfully, and cleanup was not performed.                                                                                                                                                                                                       |



## Action Values

The **action** field in the **lssw** output identifies the installation action that was taken for the fileset. The following values may be found in this field:

| Action  | Definition                                                        |
|---------|-------------------------------------------------------------------|
| APPLY   | An attempt was made to apply the specified fileset.               |
| CLEANUP | An attempt was made to perform cleanup for the specified fileset. |
| COMMIT  | An attempt was made to commit the specified fileset.              |
| REJECT  | An attempt was made to reject the specified fileset.              |

## Status Values

The **status** field in the **lssw** output identifies the resultant status in the history of installation actions. The following values may be found in this field:

| Status   | Description                                                        |
|----------|--------------------------------------------------------------------|
| BROKEN   | The fileset was left in a broken state after the specified action. |
| CANCELED | The specified action was canceled before it completed.             |
| COMPLETE | The commitment of the fileset has completed successfully.          |

## Flags

**-hist** Displays the installation and update history information.

## Exit Status

See “Exit status for Virtual I/O Server commands” on page 3

## Examples

1. To list all installed software, type:  
`lssw`
2. To list installation and update history, type:  
`lssw -hist`

## Related Information

The **updateios** command, the **ioslevel** command, the **remote\_management** command, the **oem\_setup\_env** command, and the **oem\_platform\_level** command.

---

## IVM lssyscfg command

### Purpose

List attributes of partitions, partition profiles, or the managed system. This command is operable only in an Integrated Virtualization Manager environment.

### Syntax

To list partition attributes:

```
lssyscfg -r lpar [--filter "FilterData"] [-F "AttributeNames"] [--header] [-m ManagedSystem]
```

To list partition profile attributes:

```
lssyscfg -r prof [--filter "FilterData"] [-F "AttributeNames"] [--header] [-m ManagedSystem]
```

To list system attributes:

```
lssyscfg -r sys [-F "AttributeNames"] [--header] [-m ManagedSystem]
```

## Description

The `lssyscfg` command lists attributes of partitions, partition profiles, or the managed system.

## Flags

|                                   |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|-----------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>-r</b> <i>ResourceType</i>     | The type of resources to list:<br><br>lpar: Logical partition resources<br>prof: Logical partition profile resources<br>sys: Managed system resources                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>-m</b> <i>ManagedSystem</i>    | The name of the managed system. This attribute is optional because there is only one system to manage. The name may either be the user-defined name for the managed system, or be in the form <code>ttt-mmm*sssssss</code> , where <code>ttt</code> is the machine type, <code>mmm</code> is the model, and <code>sssssss</code> is the serial number of the managed system.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>--filter</b> <i>FilterData</i> | The filters to apply to the resources to be listed. Filters are used to select which resources of the specified resource type are to be listed. If no filters are used, then all of the resources of the specified resource type will be listed. For example, specific partitions can be listed by using a filter to specify the names or IDs of the partitions to list. Otherwise, if no filter is used, then all the partitions in the managed system will be listed.<br><br>The filter data consists of filter name/value pairs, which are in comma separated value (CSV) format. The filter data must be enclosed in double quotes.<br><br>The format of the filter data is as follows:<br>"filter-name=value,filter-name=value,..."<br><br>Note that certain filters accept a comma separated list of values, as follows:<br>"filter-name=value,value,...",..."<br><br>When a list of values is specified, the filter name/value pair must be enclosed in double quotes. Depending on the shell being used, nested double quote characters may need to be preceded by an escape character, which is usually a <code>'\'</code> character.<br><br>Unless otherwise indicated, multiple values can be specified for each filter.<br><br><b>Valid filter names for partitions:</b><br><br>lpar_names   lpar_ids - name or ID of the partitions to view<br>work_groups - work groups to which the partitions belong<br><br><b>Valid filter names for partition profiles:</b><br><br>lpar_names   lpar_ids: name or ID of the partition profiles<br>profile_names: profile names for the partitions<br><b>Note:</b> This option is not valid when listing managed systems. |

**-F AttributeNames**

A delimiter separated list of attribute names for the desired attribute values to be displayed for each resource. If no attribute names are specified, then values for all of the attributes for the resource will be displayed.

When this option is specified, only attribute values will be displayed. No attribute names will be displayed. The attribute values displayed will be separated by the delimiter which was specified with this option.

This option is useful when only attribute values are desired to be displayed, or when the values of only selected attributes are desired to be displayed.

**Attribute names for partitions:**

**allow\_perf\_collection**

Permission for the partition to retrieve shared processor pool utilization information. Valid values are:

- 0: do not allow authority
- 1: allow authority

**auto\_start**

Valid values are:

- 0 - do not automatically start with system power on
- 1 - automatically start with system power on

**boot\_mode**

Partition power on mode. This attribute is only available for AIX and Linux logical partitions. Valid values are:

- norm: normal
- dd: diagnostic with default boot list
- ds: diagnostic with stored boot list
- of: Open Firmware OK prompt
- sms: System Management Services
- null: not applicable

**curr\_lpar\_proc\_compat\_mode**

Displays the current negotiated compatibility mode. This is the mode that the partition operating system has negotiated, and will be less than or equal to the desired mode at logical partition boot time.

**Note:** It is possible that the mode will be greater than the desired mode, if the partition is powered off. The current value only changes when the logical partition restarts.

**curr\_profile**

The current profile is always equal to the partition name

**default\_profile**

The default profile is always equal to the partition name

**desired\_lpar\_proc\_compat\_mode**

Displays the requested compatibility mode.

**dlpar\_mem\_capable**

Indicates if the partition supports dynamic LPAR of memory. Valid values:

- 0 - no
- 1 - yes

**dlpar\_proc\_capable**

Indicates if the partition supports dynamic LPAR of processing resources. Valid values:

- 0 - no
- 1 - yes

**logical\_serial\_num**

A globally unique string for this partition

**lpar\_env**

The operating environment for this partition Valid values:

- aixlinux: a partition type that supports AIX or Linux.
- vioserver: a Virtual I/O Server partition

**lpar\_id** unique integer identifier for the partition

**lpar\_keylock**

Partition keylock position. Valid values:

- norm: normal keylock
- manual: manual keylock

**mem\_synchronized**

The current and pending memory values for this partition are synchronized.

**name** Name of the partition

**os\_version**

The version of the operating system running that is in the logical partition.

**power\_ctrl\_lpar\_ids**

A list of partitions which have power control over this partition. Valid values:

- none: No partitions

**proc\_synchronized**

The current and pending processing values for this partition are synchronized.

**resource\_config**

Valid values:

- 0 - resources are not available to power on with system
- 1 - resources are available to power on with system

**rmc\_ipaddr**

The IP address of the client partition. This IP address is used by RMC to connect to the client partition for dynamic LPAR.

**rmc\_state**

The state of the RMC connection between the management partition and the client partition. The RMC connection is used primarily for dynamic LPAR.

Valid values:

- inactive
- active
- unknown
- none - RMC not configured. This partition has never been registered with RMC.

**rmc\_osshutdown\_capable**

Indicates if the partition supports shutdown via the RMC connection. This allows the management partition to safely shutdown the client partition via `chsysstate -o osshutdown` Valid values follow:

- 0 - no
- 1 - yes

**state** The current runtime state of the partition. Valid values follow:

- Not Activated
- Starting
- Running
- Shutting Down
- Error
- Open Firmware
- Not Available

**uptime** Partition uptime in seconds.

**work\_group\_id**

Valid values follow:

- none: do not participate in the workload management group
- 1: participate in the workload management group

**Attribute names for partition profiles:****all\_resources**

Valid values:

- 0: This partition will not own all physical resources on the system

**auto\_start**

Valid values:

- 0 - do not automatically start with system power on
- 1 - automatically start with system power on

**boot\_mode**

Partition power on mode. Valid values are:

- norm - normal
- dd - diagnostic with default boot list
- ds - diagnostic with stored boot list
- of - Open Firmware OK prompt
- sms - System Management Services

**conn\_monitoring**

Valid values:

- 0 - connection monitoring is disabled
- 1 - connection monitoring is enabled

**desired\_io\_entitled\_mem**

The amount of I/O entitled memory for a shared memory partition. This is the portion of memory that is reserved for I/O mappings.

- auto (automatically manage)
- *Number of megabytes*

If the value is auto, the entitlement is calculated based on the virtual I/O configuration of the partition. If the virtual I/O configuration is changed, the entitlement is updated automatically. If auto is not used, no automatic adjustments are made. The default value is auto.

**desired\_mem**

The assigned megabytes of memory for this partition.

**desired\_procs**

The assigned number of processors for this partition. In shared processing mode, this refers to virtual processors.

**desired\_proc\_units**

The assigned number of processing units for this partition.

**io\_slots**

Comma separated list of I/O slots for the partition. Each item in this list has the format:

drc\_index/slot\_io\_pool\_id/is\_required

Valid values for is\_required:

- 0 - no
- 1 - yes

**lhea\_logical\_ports**

Comma-separated list of logical Host Ethernet Adapter (LHEA) logical ports, with each logical port having the following format:

adapter-ID/port-group/physical-port-ID  
/logical-port-ID/allowed-VLAN-IDs

**lhea\_capabilities**

Comma-separated list of LHEA capabilities, with each capability having one of the following formats:

adapter-ID/capability

adapter-ID/5/ieq/nieq/qp/cq/mr

The values for *ieq* (interruptible event queues), *niesq* (non-interruptible event queues), *qp* (queue pairs), *cq* (completion queues), and *mr* (memory regions) specify the resource amount in addition to the base minimum. Valid values are:

- 0 - minimum
- 1 - low
- 2 - medium
- 3 - high
- 4 - dedicated
- 5 - custom

**lpar\_id** Unique integer identifier for the partition.

**lpar\_io\_pool\_ids**

Valid values are:

- none - this partition is not part of an I/O pool.

**lpar\_name**

Name of the partition.

**lpar\_proc\_compat\_mode**

The currently requested compatibility mode of the processor for the logical partition.

**max\_mem**

The maximum megabytes of memory for this partition.

**max\_procs**

The maximum number of processors for this partition. In shared processing mode, this refers to virtual processors. Valid values are:

- keep\_idle\_procs - never share processors
- share\_idle\_procs - share processors only when the partition is inactive
- share\_idle\_procs\_active - share processors only when partition is active
- share\_idle\_procs\_always - always share processors
- cap - capped mode
- uncap - uncapped mode

**max\_proc\_units**

The maximum number of processing units for this partition.

**max\_virtual\_slots**

Maximum number of virtual I/O adapter slots.

**mem\_mode**

The partition memory mode.

- ded - dedicated memory
- shared - shared memory

If the memory mode is shared, then the logical partition cannot be assigned any physical I/O slots, the `proc_mode` attribute on the logical partition must be shared, and a memory pool must exist.

**Note:** The Virtual I/O Server logical partition only supports the dedicated memory mode.

**mem\_weight**

The shared memory weight of the logical partition with shared memory. This flag is used for determining priority of logical partitions in a memory pool for distribution of memory.

**min\_mem**

The minimum megabytes of memory for this logical partition.

**min\_procs**

The minimum number of processors for this logical partition. In shared processing mode, this refers to virtual processors.

**min\_proc\_units**

The minimum number of processing units for this logical partition.

**name** Name of the profile.

**os\_type**

The operating system environment for this logical partition. Valid values:

- aixlinux: an RPA partition type which supports AIX or Linux.
- vioserver: a Virtual I/O Server partition

**paging\_device**

The paging space device to use if a memory pool is used. A paging device is a block storage device that has been added to the memory pool and is not designated as a paging device for any other logical partition. If the `paging_device` value is a blank string, no paging device is currently assigned.

**power\_ctrl\_lpar\_ids**

A list of logical partitions which have power control over this logical partition. Valid values are:

- none - No partitions

**power\_ctrl\_lpar\_names**

A list of logical partitions which have power control over this logical partition. Valid values are:

- none - No partitions



**primary\_paging\_vios\_id**

The ID of the primary paging Virtual I/O Server (VIOS) partition that provides access to the paging space devices for the shared memory partitions. A paging VIOS partition is a VIOS logical partition that is assigned to the shared memory pool.

**primary\_paging\_vios\_name**

The name of the primary paging VIOS partition that provides access to the paging space devices for the shared memory partitions. A paging VIOS partition is a VIOS logical partition that is assigned to the shared memory pool.

**proc\_mode**

Valid values:

- ded: dedicated processor mode
- shared: shared processor mode

**secondary\_paging\_vios\_id**

The ID of the secondary paging Virtual I/O Server (VIOS) partition that provides access to the paging space devices for the shared memory partitions. A paging VIOS partition is a VIOS logical partition that is assigned to the shared memory pool.

**secondary\_paging\_vios\_name**

The name of the secondary paging VIOS partition that provides access to the paging space devices for the shared memory partitions. A paging VIOS partition is a VIOS logical partition that is assigned to the shared memory pool.

**sharing\_mode**

Valid values:

- share\_idle\_procs
- proc\_mode=ded
- share\_idle\_procs\_active
- share\_idle\_procs\_always
- proc\_mode=shared
- cap
- uncap

**uncap\_weight**

A weighted average of processing priority when in uncapped sharing mode. The smaller the value, the lower the weight. Valid values are: 0 - 255.

**virtual\_eth\_adapters**

Comma-separated list of virtual Ethernet adapters, with each adapter having the following format: slot\_number/is\_ieee/port\_vlan\_id/additional\_vlan\_ids/is\_trunk/is\_required

All 5 '/' characters must be present, but optional values may be omitted. Optional values are additional-vlan-IDs, and is-trunk. Valid values for is\_ieee, is\_trunk, and is\_required:

- 0 - no
- 1 - yes

For example, 4/0/2//0/0 specifies a virtual Ethernet adapter with a virtual slot number of 4, is not IEEE 802.1Q enabled, has a port virtual LAN ID of 2, no additional virtual LAN IDs, it is not a trunk adapter, and is not required.

A value of none indicates that there are no virtual Ethernet adapters.

**virtual\_fc\_adapters**

Comma-separated list of virtual fibre channel adapters. Each item in the list has the following format:

```
virtual slot num/adapter_type/remote_lpar_id/
remote_lpar_name/remote_slot_num/wwpn_list/is_required
```

**virtual\_scsi\_adapters**

Comma-separated list of virtual SCSI adapters. Each item in the list has the following format:

```
slot_num/adapter_type/remote_lpar_id/remote_lpar_name/
remote_slot_num/is_required
```

**virtual\_serial\_adapters**

Comma-separated list of virtual serial adapters. Each item in the list has the following format:

```
slot_num/adapter_type/supports_hmc/remote_lpar_id/
remote_lpar_name/remote_slot_num/is_required
```

The attribute names are not present in the list, just their values are present. If an attribute is optional and is not to be included, then no value would be specified for that attribute. For example, `0/server/1/any//any/1` specifies a virtual server serial device that has a virtual slot number of 0, supports HMC, supports any remote partition, supports any remote slot, and is required.

A value of none indicates that there are no virtual serial adapters.

Valid values for `adapter_type`:

- client: client adapter
- server: server adapter, valid for Virtual I/O Server partitions only

Valid values for `supports_hmc`:

- 0 - no
- 1 - yes

Valid values for `is_required`:

- 0 - no
- 1 - yes

**Attribute names for the managed system:****active\_lpar\_share\_idle\_procs\_capable**

Indicates if the platform supports setting the share while active sharing mode.

Valid values:

- 0 - not capable
- 1 - capable

**active\_mem\_sharing\_capable**

The capability of the managed system to use a shared memory pool. Valid values follow:

- 0 - Not capable
- 1 - Capable

**capabilities**

Displays a comma-separated list of capabilities.

**cod\_mem\_capable**

Valid values:

- 0- not capable of memory Capacity on Demand
- 1 - capable of memory Capacity on Demand

**cod\_proc\_capable**

Valid values:

- 0- not capable of processor Capacity on Demand
- 1 - capable of processor Capacity on Demand

**config\_version**

The version of the configuration data in platform firmware of the current logical partition .

**curr\_configured\_max\_lpars**

The current maximum number of logical partitions supported by the management logical partition

This option is deprecated. Instead use:

```
lshwres -r virtualio --subtype slot --level lpar
```

with attribute "curr\_max\_virtual\_slots" value for partition 1

**dlpar\_mem\_capable**

Valid values:

- 0: Changes do not take effect until the next reboot of the logical partition or platform.
- 1: Changes take effect immediately.

**lhea\_capable**

Indicates if the managed system has one or more Host Ethernet Adapters.

**lpar\_avail\_priority**

The priority of the partition to maintain its entitled processors. If a processor failure occurs, processing resources will be removed first from the lowest priority partition.

**Note:** The Virtual I/O Server partition must have a higher priority than any other partition on the system.

**lpar\_avail\_priority\_capable**

Specifies if the platform supports the lpar\_avail\_priority attribute. Valid values:

- 0 - not capable
- 1 - capable

**lpar\_comm\_default**

Indicates if the **lpar\_comm\_ipaddr** is using the default IP address, or if the user has manually set this using **chsyscfg**. Valid values:

- 0: User has manually set the IP address
- 1: Default IP address is used. This is the first IP address configured on your system as reported by `lstcpip -interfaces`

**lpar\_comm\_ipaddr**

The IP address through which client partitions will communicate with the management partition. This is used primarily for dynamic LPAR. It is defaulted to the first IP address configured on your system, but can be manually set if desired.

**Note:** This attribute supports multiple IP addresses using a comma-separated list.

**lpar\_proc\_compat\_mode\_capable**

The supported compatibility modes of the processor for the logical partition. Valid values:

- 0: The managed system does not support setting the processor compatibility mode for a partition.
- 1: The managed system supports setting the processor compatibility mode for a partition.

**lpar\_proc\_compat\_modes**

A comma separated list of compatibility modes that the managed system supports.

**max\_lpars**

Maximum number of partitions supported by firmware.

**max\_power\_ctrl\_lpars**

Maximum number of power controlling partitions per controlled partition.

**max\_vscsi\_remote\_lpar\_id**

Indicates the largest partition ID of the remote partition with a virtual SCSI server adapter.

**max\_micro\_lpar\_id**

Indicates largest partition ID of a partition using the MicroPartition technology.

**micro\_lpar\_capable**

Valid values:

- 0: not capable of creating shared processor partitions
- 1: capable of creating shared processor partitions

**mfg\_default\_config**

Specifies whether or not the system is in the manufacturing default partition configuration.

Valid values:

- 0: No
- 1: Yes

**name** Name for the managed system.

**pend\_configured\_max\_lpars**

The maximum number of partitions supported by the management partition after the next restart.

This option is deprecated. Instead use:

```
lshwres -r virtualio --subtype slot --level lpar
```

with attribute "pend\_max\_virtual\_slots" value for partition 1.

**pend\_lpar\_config\_state**

Valid values:

enabled: The logical partition configuration data will be enabled during the next restart operation.

disabled: The logical partition configuration data will be disabled during the next restart operation.

cleared: The logical partition configuration data will be cleared to manufacturing defaults during the next restart operation.

**power\_off\_policy**

Valid values:

- 0- Power off the managed system after all partitions are powered off.

**serial\_num**

Serial number for the managed system.

**service\_lpar\_id**

ID of the partition with platform service authority.

**service\_lpar\_name**

Name of the partition with platform service authority.

**state** Valid values:

- Operating - The managed system is running.

**sys\_time**

The UTC time of system firmware in the format *month/day/year hour:minute:second*.

**type\_model**

Type and model for the managed system

**vet\_activation\_capable**

Whether the platform supports PowerVM® Editions system technologies activation. Valid values:

- 0 - not capable
- 1 - capable

**virtual\_fc\_capable**

A value indicating whether the managed system supports a virtual fibre channel. Valid values follow:

- 0: The managed system does not support a virtual fibre channel.
- 1: The managed system supports a virtual fibre channel.

**--header**

Display a header record, which is a delimiter separated list of attribute names for the attribute values that will be displayed. This header record will be the first record displayed. This option is only valid when used with the **-F** option.

## Exit Status

This command has a return code of 0 on success.

## Security

This command is accessible by all users.

## Examples

1. To list the attributes for the managed system, type:  

```
lssyscfg -r sys
```
2. To list only the user-defined name, machine type and model, and serial number for the managed system, type:  

```
lssyscfg -r sys -F name,type_model,serial_num
```
3. To list all partitions and only display attribute values for each partition following a header of attribute names, type:  

```
lssyscfg -r lpar -F --header
```
4. To list the partitions named lpar1, lpar2, and lpar3, type:  

```
lssyscfg -r lpar --filter \"lpar_names=lpar1,lpar2,lpar3\"
```
5. To list the partition profile for partition lpar2, type:  

```
lssyscfg -r prof --filter lpar_names=lpar2
```
6. To view the desired compatibility mode in the profile, type:  

```
lssyscfg -r prof -F lpar_proc_compat_mode
```

## Related Information

The **chsyscfg** command, the **mksyscfg** command, and the **rmsyscfg** command.

---

## IVM lssysconn command

### Purpose

List connection information for systems. This command is operable only in an Integrated Virtualization Manager environment.

### Syntax

To list the service processor network connection information:

```
lssysconn -r all [-F "AttributeName"] [--header] [-m ManagedSystem]
```

### Description

The **lssysconn** command lists the service processor network connection information.

## Flags

- r** *ResourceType*      The type of resources to list:
- all - Lists all connections  
Attributes: resource\_type, type\_model\_serial\_num, sp, side, ipaddr, alt\_ipaddr, state, eth\_loc\_code, alt\_eth\_loc\_code  
Filters: None
- m** *ManagedSystem*      The name of the managed system. This attribute is optional because there is only one system to manage. The name may either be the user-defined name for the managed system, or be in the form tttt-mmm\*sssssss, where tttt is the machine type, mmm is the model, and ssssssss is the serial number of the managed system.
- F** *AttributeNames*      A delimiter separated list of attribute names for the desired attribute values to be displayed for each resource. If no attribute names are specified, then values for all of the attributes for the resource will be displayed.
- When this option is specified, only attribute values will be displayed. No attribute names will be displayed. The attribute values displayed will be separated by the delimiter which was specified with this option.
- This option is useful when only attribute values are desired to be displayed, or when the values of only selected attributes are desired to be displayed.
- The following attributes are available:
- resource\_type**  
Indicates the resource type. This attribute always reads 0.
- type\_model\_serial\_num**  
The type-model and serial number of the system in the form ttt-mmm\*sssssss, where tttt is the machine type, mmm is the model, and ssssssss is the serial number of the managed system.
- sp**      The type of service processor. This attribute always reads unavailable.
- side**      The current side of the service processor. This attribute always reads unavailable.
- ipaddr**      The IP Address of the first Ethernet device on the service processor.
- alt\_ipaddr**  
The IP Address of the second Ethernet device on the service processor.
- state**      The connection state to the service processor. This attribute always reads No Connection.
- eth\_loc\_code**  
The physical location code of the first Ethernet device on the service processor.
- alt\_eth\_loc\_code**  
The physical location code of the second Ethernet device on the service processor.
- header**      Display a header record, which is a delimiter separated list of attribute names for the attribute values that will be displayed. This header record will be the first record displayed. This option is only valid when used with the -F option.

## Exit Status

This command has a return code of 0 on success.

## Security

This command is accessible by all users.

## Examples

1. To list all system connections, type:

```
lssysconn -r all
```

## Related Information

The `lssyscfg` command.

---

## Istcpip command

### Purpose

Displays the Virtual I/O Server TCP/IP settings and parameters.

### Syntax

```
lstcpip -interfaces [-fmt delimiter]
```

```
lstcpip [-state] [-num] [-field] [-fmt delimiter]
```

```
lstcpip [-num] [-routtable] [-routinfo] [-state] [-arp]
```

```
lstcpip -stored
```

```
lstcpip -adapters
```

```
lstcpip [-sockets] [-family {inet | inet6 | unix}]
```

```
lstcpip -namesrv
```

```
lstcpip [-state [-field FieldName ...]] | [-routtable [-field FieldName ...]] [-fmt delimiter]
```

```
lstcpip -hostname
```

**Note:** If IPv6 is configured on the Virtual I/O Server, the `lstcpip` command lists an IPv6 address. Due to its size, the IPv6 address spans both the Network and Address fields of the output screen.

### Description

The `lstcpip` command displays the current and stored TCP/IP setting such as IP address, routing table, sockets, name server settings, and so forth.

### Flags

|                        |                                                   |
|------------------------|---------------------------------------------------|
| <code>-adapters</code> | Lists Ethernet adapters on the system.            |
| <code>-arp</code>      | Displays the current ARP table entries.           |
| <code>-family</code>   | Specifies the INET, INET6, or UNIX socket family. |
| <code>-field</code>    | Specifies a list of fields to display.            |
| <code>-fmt</code>      | Divides output by a user-specified delimiter.     |
| <code>-hostname</code> | Displays the system hostname.                     |



|                    |                                                                                                                                                                               |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>-interfaces</b> | Displays all of the interfaces configured on the system. Specifically, it displays their addresses, network masks, states, and mac addresses.                                 |
| <b>-namesrv</b>    | Lists DNS name servers in search order and domain name.                                                                                                                       |
| <b>-num</b>        | Displays numeric output, rather than trying to resolve host names.                                                                                                            |
| <b>-routtable</b>  | Displays the routing tables.                                                                                                                                                  |
| <b>-routinfo</b>   | Displays the routing tables, including the user-configured and current costs of each route.                                                                                   |
| <b>-sockets</b>    | Displays information about currently open sockets.                                                                                                                            |
| <b>-state</b>      | Displays the current state of all configured interfaces.                                                                                                                      |
| <b>-stored</b>     | Displays stored TCP/IP configuration, which will be applied when the system starts. It will list interface IP addresses, any defined static routes, host names, and DNS info. |

## Examples

1. To list the Virtual I/O Server TCP/IP configuration, type:  

```
lstcpip -stored
```
2. To list the current routing table, type:  

```
lstcpip -routtable
```
3. To list open inet sockets, type:  

```
lstcpip -sockets -family inet
```
4. To display the states of all interfaces using a delimiter, type:  

```
lstcpip -states -fmt "/"
```

This command produces output similar to the following:

```
sit0/1480/link#2/9.3.126.60/0/0/0/0/0
sit0/1480/commo126060./austin.ixx.com/0/0/0/0/0
en2/1500/link#3/0.9.6b.6e.e3.72/871825/0/16305/1/0
en2/1500/9.3.126/commo126060.austi/871825/0/16305/1/0
en2/1500/fe80::209:6bff:fe6e:e372/871825/0/16305/1/0
lo0/16896/link#1/5013/0/5019/0/0
lo0/16896/127/localhost/5013/0/5019/0/0
lo0/16896/::1/5013/0/5019/0/0
```

5. To display all interface information, type:  

```
lstcpip -interfaces
```

This command produces output similar to the following:

| Name | Address                  | mask/Pfx | State | MAC               |
|------|--------------------------|----------|-------|-------------------|
| en2  | fe80::209:6bff:fe6e:e372 | 64       | up    | 00:09:6b:6e:e3:72 |
| en3  | -                        | -        | down  | 00:09:6b:6e:e3:73 |
| et2  | -                        | -        | down  | 00:09:6b:6e:e3:72 |
| et3  | -                        | -        | down  | 00:09:6b:6e:e3:73 |
| en4  | -                        | -        | down  | 0e:f0:c0:00:30:04 |

---

## lsuser command

### Purpose

Displays user account attributes.

### Syntax

```
lsuser [ALL | Name[, Name] ...]
```

## Description

The **lsuser** command displays the user account attributes. You can use this command to list all attributes of all the system user accounts or all the attributes of specific user accounts. If you specify more than one user account, each user account is separated by a comma. If you do not specify any user accounts, the attributes of all user accounts are displayed.

The **lsuser** command lists each user's attributes on one line. It displays attribute information as *Attribute=Value* definitions, each separated by a blank space.

**Note:** The **lsuser** command only displays the users that were created on the Virtual I/O Server. It filters out all users that do not use the *rsh* shell or have the following role: `roles=RunDiagnostics`. The **lsuser** command does not display all LDAP users. For example, the command does not display an LDAP user that was created on another LDAP client, that does not exist on the Virtual I/O Server.

## Exit Status

See "Exit status for Virtual I/O Server commands" on page 3.

## Security

This command can be run by any users. However, user attributes are only displayed for the **padmin** user.

## Examples

1. To list all users on the system, type:

```
lsuser
```

The system displays output similar to the following for the **padmin** user:

```
padmin roles=PAdmin account_locked=false expires=0 histexpire=0
histsize=0 loginretries=0 maxage=0 maxexpired=-1 maxrepeats=8 minage=0
minalpha=0 mindiff=0 minlen=0 minother=0 pwdwarntime=0
sally roles=DEUser account_locked=false expires=0 histexpire=0
histsize=0 loginretries=0 maxage=0 maxexpired=-1 maxrepeats=8 minage=0
minalpha=0 mindiff=0 minlen=0 minother=0 pwdwarntime=330
henry roles=DEUser account_locked=false expires=0 histexpire=0
histsize=0 loginretries=0 maxage=0 maxexpired=-1 maxrepeats=8 minage=0
minalpha=0 mindiff=0 minlen=0 minother=0 pwdwarntime=330
admin1 roles=Admin account_locked=false expires=0 histexpire=0
histsize=0 loginretries=0 maxage=0 maxexpired=-1 maxrepeats=8 minage=0
minalpha=0 mindiff=0 minlen=0 minother=0 pwdwarntime=330
deuser1 roles=DEUser account_locked=false expires=0 histexpire=0
histsize=0 loginretries=0 maxage=0 maxexpired=-1 maxrepeats=8 minage=0
minalpha=0 mindiff=0 minlen=0 minother=0 pwdwarntime=330
sadan roles=Admin account_locked=false expires=0 histexpire=0
histsize=0 loginretries=0 maxage=0 maxexpired=-1 maxrepeats=8 minage=0
minalpha=0 mindiff=0 minlen=0 minother=0 pwdwarntime=330 registry=LDAP
SYSTEM=LDAP sruer1 roles=SRUser,RunDiagnostics
account_locked=false expires=0 histexpire=0
histsize=0 loginretries=0 maxage=0 maxexpired=-1 maxrepeats=8 minage=0
minalpha=0 mindiff=0 minlen=0 minother=0 pwdwarntime=330
view1 roles=ViewOnly account_locked=false expires=0 histexpire=0
histsize=0 loginretries=0 maxage=0 maxexpired=-1 maxrepeats=8 minage=0
minalpha=0 mindiff=0 minlen=0 minother=0 pwdwarntime=330
```

The system displays the following information for other users:

```
padmin roles=PAdmin
sally roles=DEUser
henry roles=DEUser
admin1 roles=Admin
```

```
deuser1 roles=DEUser
sadan roles=Admin
sruser1 roles=SRUser
view1 roles=ViewOnly
```

2. To display the attributes of user admin1, type the following command:

```
lsuser admin1
```

3. To display the attributes of user admin1 and user admin2, type the following command:

```
lsuser admin1, admin2
```

## Related Information

The **chuser** command, the **mkuser** command, the **rmuser** command, and the **passwd** command.

---

## IVM lsvet command

### Purpose

List Capacity on Demand advanced functions activation information. This command is operable only in an Integrated Virtualization Manager environment.

### Syntax

```
lsvet -t {code | hist} -m managed-system [-F [attribute-names] [--header]] [--help]
```

### Description

The **lsvet** command lists activation information for Capacity on Demand advanced functions. Capacity on Demand advanced functions include PowerVM Editions and Enterprise Enablement. Capacity on Demand advanced functions are sometimes referred to as Virtualization Engine systems technologies.

### Flags

**-t**

The IP address or host name of the managed system. Use **code** for the Virtualization Engine systems technologies activation codes, and **hist** for the Virtualization Engine systems technologies activation history log.

#### Attribute name

time\_stamp=08/30/2007 00:16:28,entry=[VIOSI0500040A-0336] processor capacity limit enabled.

**-m**

The name of the managed system for which information is to be listed. The name may either be the user-defined name for the managed system, or be in the form *ttt-mmm\*sssssss*, where *ttt* is the machine type, *mmm* is the model, and *sssssss* is the serial number of the managed system. The *ttt-mmm\*sssssss* form must be used if there are multiple managed systems with the same user-defined name.

**-F**

A delimiter separated list of attribute names representing the desired attribute values to display. If this option is specified without any attribute names, then all of the attributes will be displayed.

When this option is specified, only attribute values will be displayed. No attribute names will be displayed. The attribute values displayed will be separated by the delimiter which was specified with this option.

This option is useful when only attribute values are desired to be displayed, or when the values of only selected attributes are desired to be displayed.

|                 |                                                                                                                                                                                                                                                        |
|-----------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>--header</b> | Display a header record, which is a delimiter separated list of attribute names for the attribute values that will be displayed. This header record will be the first record displayed. This option is only valid when used with the <b>-F</b> option. |
| <b>--help</b>   | Display the help text for this command and exit.                                                                                                                                                                                                       |

## Exit Status

This command has a return code of 0 on success.

## Examples

1. To display the activation code generation information, type the following command:

```
lsvet -m sys1 -t code
```

2. To display the activation history log, type the following command:

```
lsvet -m 9117-570*1001213 -t hist
time_stamp=02/07/2006 19:52:03,entry=HSCL0421 POWER Hypervisor code entered.
time_stamp=02/07/2006 19:52:03,entry=HSCL0403 Virtual I/O server
capability enabled.
time_stamp=02/07/2006 19:52:03,entry=HSCL0405 Micro-partitioning
capability enabled.
time_stamp=02/07/2006 19:52:03,entry=HSCL0406 Multiple partitions enabled.
```

---

## lsvfilt command

### Purpose

Lists virtual LAN-crossing filter rules from the filter table.

### Syntax

```
lsvfilt [-a]
```

### Description

The **lsvfilt** command is used to list the virtual LAN-crossing filter rules and their status.

### Flags

|           |                                     |
|-----------|-------------------------------------|
| <b>-a</b> | Lists only the active filter rules. |
|-----------|-------------------------------------|

### Exit Status

This command returns the following exit values:

|    |                        |
|----|------------------------|
| 0  | Successful completion. |
| >0 | An error occurred.     |

### Examples

1. To list all the active filter rules in the kernel, type the command as follows:

```
lsvfilt -a
```

### Related Information

The **chvfilt** command, the **genvfilt** command, the **mkvfilt** command, the **rmvfilt** command, and the **vlantfw** command.

---

## lsvg command

### Purpose

Displays information about volume groups.

### Syntax

```
lsvg [-map | -lv | -pv] [-field FieldName] [-fmt Delimiter] VolumeGroup...
```

### Description

The **lsvg** command displays information about volume groups. If you use the *VolumeGroup* parameter, only the information for that volume group is displayed. If you do not use the *VolumeGroup* parameter, a list of the names of all defined volume groups is displayed.

When information from the Device Configuration database is unavailable, some of the fields will contain a question mark (?) in place of the missing data. The **lsvg** command attempts to obtain as much information as possible from the description area when the command is given a logical volume identifier.

Full scripting support is provided to the **lsvg** command by using the **-field** *FieldNames* and **-fmt** *Delimiter* flags. The **-field** flag will allow the user to select which output fields to display and in what order, while the **-fmt** flag provides scriptable output. The output fields will be displayed in the order they appear on the command line.

If you do not specify any flags, the following information will be displayed:

|                           |                                                                                                                                                                                                                                                                           |
|---------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Volume group</b>       | Name of the volume group. Volume group names must be unique and can range from 1 to 15 characters.                                                                                                                                                                        |
| <b>Volume group state</b> | State of the volume group. If the volume group is active, the state is either active/complete (indicating all physical volumes are active) or active/partial (indicating some physical volumes are not active). If the volume group is not active, the state is inactive. |
| <b>Permission</b>         | Access permission: read-only or read-write.                                                                                                                                                                                                                               |
| <b>Max LVs</b>            | Maximum number of logical volumes allowed in the volume group.                                                                                                                                                                                                            |
| <b>LVs</b>                | Number of logical volumes currently in the volume group.                                                                                                                                                                                                                  |
| <b>Open LVs</b>           | Number of logical volumes within the volume group that are currently open.                                                                                                                                                                                                |
| <b>Total PVs</b>          | Total number of physical volumes within the volume group.                                                                                                                                                                                                                 |
| <b>Active PVs</b>         | Number of physical volumes that are currently active.                                                                                                                                                                                                                     |
| <b>VG identifier</b>      | The volume group identifier.                                                                                                                                                                                                                                              |
| <b>PP size</b>            | Size of each physical partition.                                                                                                                                                                                                                                          |
| <b>Total PPs</b>          | Total number of physical partitions within the volume group.                                                                                                                                                                                                              |
| <b>Free PPs</b>           | Number of physical partitions not allocated.                                                                                                                                                                                                                              |
| <b>Alloc PPs</b>          | Number of physical partitions currently allocated to logical volumes.                                                                                                                                                                                                     |
| <b>Quorum</b>             | Number of physical volumes needed for a majority.                                                                                                                                                                                                                         |
| <b>VGDS</b>               | Number of volume group descriptor areas within the volume group.                                                                                                                                                                                                          |
| <b>Auto-on</b>            | Automatic activation at IPL (yes or no).                                                                                                                                                                                                                                  |
| <b>Concurrent</b>         | States whether the volume group is Concurrent Capable or Non-Concurrent Capable.                                                                                                                                                                                          |
| <b>Auto-Concurrent</b>    | States whether you should auto activate the Concurrent Capable volume group in concurrent or non-concurrent mode. For volume groups that are Non-Concurrent Capable, this value defaults to Disabled.                                                                     |
| <b>VG Mode</b>            | The mode of the volume group: Concurrent or Non-Concurrent.                                                                                                                                                                                                               |
| <b>Node ID</b>            | Node id of this node if volume group is in concurrent node.                                                                                                                                                                                                               |
| <b>Active Nodes</b>       | Node ids of other concurrent nodes that have this volume group active.                                                                                                                                                                                                    |
| <b>Max PPs Per PV</b>     | Maximum number of physical partitions per physical volume allowed for this volume group.                                                                                                                                                                                  |

|                    |                                                                                                                                     |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------|
| <b>Max PVs</b>     | Maximum number of physical volumes allowed in this volume group.                                                                    |
| <b>LTG size</b>    | Logical track group size, in number of kilobytes, of the volume group.                                                              |
| <b>BB POLICY</b>   | Bad block relocation policy of the volume group.                                                                                    |
| <b>SNAPSHOT VG</b> | Snapshot volume group name if the snapshot volume group is active else snapshot volume group identifier.                            |
| <b>PRIMARY VG</b>  | Original volume group name of a snapshot volume group if the original volume group is active else original volume group identifier. |

## Flags

|                 |                                                                                                                                                                                                                                                                                                                                                                             |
|-----------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>-field</b>   | Specifies the list of fields to display. The following fields are supported if no flags are specified:                                                                                                                                                                                                                                                                      |
| <b>vgname</b>   | Name of the volume group. Volume group names must be unique systemwide and can range from 1 to 15 characters.                                                                                                                                                                                                                                                               |
| <b>vgstate</b>  | State of the volume group. If the volume group is activated with the <b>activatevg</b> command, the state is either <b>active/complete</b> (indicating all physical volumes are active) or <b>active/partial</b> (indicating all physical volumes are not active). If the volume group is de-activated with the <b>deactivatevg</b> command, the state is <b>inactive</b> . |
| <b>access</b>   | Access permission: read-only or read-write.                                                                                                                                                                                                                                                                                                                                 |
| <b>maxlvs</b>   | Maximum number of logical volumes allowed in the volume group.                                                                                                                                                                                                                                                                                                              |
| <b>numlvs</b>   | Number of logical volumes currently in the volume group.                                                                                                                                                                                                                                                                                                                    |
| <b>openlvs</b>  | Number of logical volumes within the volume group that are currently open.                                                                                                                                                                                                                                                                                                  |
| <b>totalpvs</b> | Total number of physical volumes within the volume group.                                                                                                                                                                                                                                                                                                                   |
| <b>stalepvs</b> | Number of PVs which are not current. The data is stale.                                                                                                                                                                                                                                                                                                                     |
| <b>stalepps</b> | Number of PPs which are not current. The data is stale.                                                                                                                                                                                                                                                                                                                     |
| <b>totalpps</b> | Total number of physical partitions within the volume group.                                                                                                                                                                                                                                                                                                                |
| <b>freepps</b>  | Number of physical partitions not allocated.                                                                                                                                                                                                                                                                                                                                |
| <b>usedpps</b>  | Number of physical partitions currently allocated to logical volumes.                                                                                                                                                                                                                                                                                                       |
| <b>quorum</b>   | Number of physical volumes needed for a majority.                                                                                                                                                                                                                                                                                                                           |
| <b>vgds</b>     | Number of volume group descriptor areas within the volume group.                                                                                                                                                                                                                                                                                                            |
| <b>auton</b>    | Automatic activation at IPL (yes or no).                                                                                                                                                                                                                                                                                                                                    |
| <b>pppervg</b>  | Maximum number of physical partitions allowed in this volume group.                                                                                                                                                                                                                                                                                                         |
| <b>ppperpv</b>  | Maximum number of physical partitions per physical volume allowed for this volume group.                                                                                                                                                                                                                                                                                    |
| <b>maxpvs</b>   | Maximum number of physical volumes allowed in this volume group. This information is displayed only for 32 and 128 PV volume groups.                                                                                                                                                                                                                                        |

**ltgsize** Logical track group size of the volume group. The maximum amount of data that can be transferred in one I/O request to the disks of the volume group. The LTG size will be displayed in kilobytes unless the LTG size is greater than 1 MB, in which case megabytes will be used. It is capable of dynamically determining the LTG size based on the disk topology and it is listed as Dynamic. If that capability is disabled by the user with the option, then it will be listed as Static.

**bbpolicy**

Bad block relocation policy of the volume group.

**hotspare**

**autosync**

The following fields are supported if the **-lv** flag is specified:

**lvname** A logical volume within the volume group.

**type** Logical volume type.

**lps** Number of logical partitions in the logical volume.

**pvs** Number of physical partitions used by the logical volume.

**lvstate** State of the logical volume. `Opened/stale` indicates the logical volume is open but contains partitions that are not current. `Opened/syncd` indicates the logical volume is open and synchronized. `Closed` indicates the logical volume has not been opened.

**mount** File system mount point for the logical volume, if applicable.

The following fields are supported if the **-pv** flag is specified:

**pvname**

A physical volume within the volume group.

**pvstate** State of the physical volume.

**totalpps**

Number of physical partitions on the physical volume.

**freepps**

Number of free physical partitions on the physical volume.

**dist** The number of physical partitions allocated within each section of the physical volume: outer edge, outer middle, center, inner middle, and inner edge of the physical volume.

**-fmt**

Specifies a delimiter character to separate output fields.

**-pv**

Lists the following information for each physical volume within the group specified by the *VolumeGroup* parameter:

**Physical volume**

A physical volume within the group.

**PVstate**

State of the physical volume.

**Total PPs**

Total number of physical partitions on the physical volume.

**Free PPs**

Number of free physical partitions on the physical volume.

**Distribution**

The number of physical partitions allocated within each section of the physical volume: outer edge, outer middle, center, inner middle, and inner edge of the physical volume.

- lv** Lists the following information for each logical volume within the group specified by the `VolumeGroup` parameter:
- LV** A logical volume within the volume group.
  - Type** Logical volume type.
  - LPs** Number of logical partitions in the logical volume.
  - PPs** Number of physical partitions used by the logical volume.
  - PVs** Number of physical volumes used by the logical volume.
- Logical volume state**  
State of the logical volume. `Opened/stale` indicates the logical volume is open but contains partitions that are not current. `Opened/syncd` indicates the logical volume is open and synchronized. `Closed` indicates the logical volume has not been opened.
- Mount Point**  
File system mount point for the logical volume, if applicable.
- map** Lists the following fields for each logical volume on the physical volume:  
*PVname:PPnum [LVname: LPnum [:Copynum] [PPstate]]*
- PVname** Name of the physical volume as specified by the system.
  - PPnum** Physical partition number. Physical partition numbers can range from 1 to 1016.
  - LVname** Name of the logical volume to which the physical partitions are allocated. Logical volume names must be system-wide unique names, and can range from 1 to 64 characters.
  - LPnum** Logical partition number. Logical partition numbers can range from 1 to 64,000.
  - Copynum** Mirror number.
  - PPstate** Only the physical partitions on the physical volume that are not current are shown as stale.

## Exit Status

See “Exit status for Virtual I/O Server commands” on page 3.

## Examples

1. To display the names of all volume groups within the system, type:

```
lsvg
```

2. To display information about volume group **vg02**, type:

```
lsvg vg02
```

The characteristics and status of both the logical and physical partitions of volume group **vg02** are displayed.

3. To display the names, characteristics, and status of all the logical volumes in volume group **vg02**, type:

```
lsvg -lv vg02
```



## Related Information

The **mkvg** command, the **chvg** command, the **extendvg** command, the **reducevg** command, the **mirrorios** command, the **unmirrorios** command, the **activatevg** command, the **deactivatevg** command, the **importvg** command, the **exportvg** command, and the **syncvg** command.

---

## lsvlog command

### Purpose

Lists virtual logs.

### Syntax

To list all virtual logs:

```
lsvlog [-detail | -field FieldList]
```

To list all virtual logs in the local repository:

```
lsvlog -local [-detail | -field FieldList]
```

To list all virtual logs in the specified shared storage pool:

```
lsvlog -sp StoragePool [-detail | -field FieldList]
```

To list the virtual log with the specified UUID:

```
lsvlog -u UUID | -uuid UUID [-detail | -field FieldList]
```

To list the virtual log with the specified device name:

```
lsvlog -dev DeviceName [-detail | -field FieldList]
```

To list all virtual logs that are currently attached to VSCSI host adapters:

```
lsvlog -vtd [-detail | -field FieldList]
```

The client name is an arbitrary string that is either specified or automatically copied from the host name of a client LPAR when the virtual log is created. To list all virtual logs with the specified client name:

```
lsvlog -client ClientName [-detail | -field FieldList]
```

To list all the virtual logs attached to the specified virtual SCSI (VSCSI) host adapter:

```
lsvlog -vadapter Adapter [-detail | -field FieldList]
```

To list all the virtual logs in the specified state:

```
lsvlog -state VirtualLogState [-detail | -field FieldList]
```

### Description

The **lsvlog** command displays the properties of the specified virtual log or set of virtual logs. Various parameters can be used to list all virtual logs; to filter by universal unique Identifier (UUID), device name, virtual SCSI (VSCSI) host adapter, client name, and state; or to show only connected virtual logs.

Three output formats are available. The default output format shows one row per virtual log, with the following columns:

- Client name
- Log name
- UUID
- Device name and attached VSCSI host adapter, if the virtual log is connected.

The detailed format (specified with the **-detail** parameter) displays the virtual logs grouped by Client name, with each of the properties listed on a separate line as follows:

- Log name
- UUID
- Virtual target device
- VSCSI host adapter
- Virtual log state
- Status of the target device (if connected)
- Logical unit address on the VSCSI host adapter
- Name of the shared storage pool in which the log data is stored
- Directory in which the log data is stored
- Maximum number of log files
- Maximum size of each log file
- Maximum number of state files
- Maximum size of each state file

The field format (specified with the **-field** parameter) displays the virtual logs based on the given fields. Each field can be separated by a character that is not alphanumeric. The possible fields are as follows:

|                  |                                                                 |
|------------------|-----------------------------------------------------------------|
| <b>client</b>    | Client name                                                     |
| <b>devstatus</b> | Status of the target device (if connected)                      |
| <b>dir</b>       | Directory in which the log data is stored                       |
| <b>lf</b>        | Maximum size of each log file                                   |
| <b>lfs</b>       | Maximum number of log files                                     |
| <b>lua</b>       | Logical unit address on the VSCSI host adapter                  |
| <b>name</b>      | Log name                                                        |
| <b>parent</b>    | VSCSI host adapter                                              |
| <b>sf</b>        | Maximum size of each state file                                 |
| <b>sfs</b>       | Maximum number of state files                                   |
| <b>sp</b>        | Name of the shared storage pool in which the log data is stored |
| <b>state</b>     | Virtual log state                                               |
| <b>uuid</b>      | UUID                                                            |
| <b>vtd</b>       | Virtual target device                                           |

## Flags

|                    |                                                                   |
|--------------------|-------------------------------------------------------------------|
| <b>-c, -client</b> | Shows only virtual logs for the specified client name.            |
| <b>-d, -detail</b> | Shows extended output for the virtual logs.                       |
| <b>-dev</b>        | Shows only the virtual log with the specified device name.        |
| <b>-field</b>      | Shows output by using the given list of fields.                   |
| <b>-l, -local</b>  | Shows only the virtual logs in the local repository.              |
| <b>-s, -state</b>  | Shows only virtual logs with the specified state.                 |
| <b>-sp</b>         | Shows only the virtual logs in the specified shared storage pool. |
| <b>-u, -uuid</b>   | Shows only the virtual log with the specified UUID.               |

- v, -vadapter** Shows only virtual logs connected to the specified adapter.
- vtd** Shows only virtual logs with virtual target devices.

## Exit Status

Table 4. Command specific return codes

| Return code | Description                          |
|-------------|--------------------------------------|
| 0           | All files were written successfully. |
| >0          | An error occurred.                   |

## Examples

- To list all virtual logs, type the command as follows:

```
lsvlog
```

The system displays the output as follows:

| Client Name | Log Name | UUID                             | VTD           |
|-------------|----------|----------------------------------|---------------|
| lpar-03     | syslog   | 02392437473b6c552680a9ddd2fd8d06 | vhost1/vtlog1 |
| lpar-02     | syslog   | 956f8c1c25208091495c721e0796f456 | vhost0/vtlog0 |
| lpar-01     | audit    | 9705340b31a7883573a1cd04b2254efd |               |
| lpar-01     | syslog   | b27a94a8e187ee5c917577c2a2df0268 |               |

- To list the virtual log with the UUID 02392437473b6c552680a9ddd2fd8d06, type the command as follows:

```
lsvlog -uuid 02392437473b6c552680a9ddd2fd8d06
```

The system displays the output as follows:

| Client Name | Log Name | UUID                             | VTD           |
|-------------|----------|----------------------------------|---------------|
| lpar-03     | syslog   | 02392437473b6c552680a9ddd2fd8d06 | vhost1/vtlog1 |

- To list the details of the virtual log with the UUID 0000000000000000f5e421165dfcc619, type the command as follows:

```
lsvlog -uuid 02392437473b6c552680a9ddd2fd8d06 -detail
```

The system displays the output as follows:

```
Client Name: lpar-03
```

```
Log Name: syslog
UUID: 0000000000000000f5e421165dfcc619
Virtual Target Device: vtlog1
Parent Adapter: vhost1
State: enabled
Logical Unit Address: 8100000000000000
Storage Pool:
Log Directory: /var/vio/vlogs/lpar-03/syslog
Maximum Log Files: 10
Maximum Log File Size: 1048576
Maximum State Files: 10
Maximum State File Size: 1048576
```

- To list all virtual logs by using a custom field string, displaying the UUID, a tab character, the size of each state file, the colon (:), and the number of state files, type the command as follows:

```
lsvlog -field "uuid\tsize:sf"
```

The system displays the output as follows:

```
02392437473b6c552680a9ddd2fd8d06 1048576:10
956f8c1c25208091495c721e0796f456 1048576:10
9705340b31a7883573a1cd04b2254efd 1048576:5
b27a94a8e187ee5c917577c2a2df0268 65536:20
```

## Related Information

The `chvlog` command, `chvrepo` command, the `lsvrepo` command, the `mkvlog` command, and the `rmvlog` command.

---

## lsvrepo command

### Purpose

Lists the virtual log repositories.

### Syntax

To display the properties of all virtual log repositories:

```
lsvrepo [-detail | -field FieldList]
```

To display the properties of the local virtual log repository:

```
lsvrepo -local [-detail | -field FieldList]
```

To display the properties of the virtual log repository in the specified shared storage pool:

```
lsvrepo -sp StoragePool [-detail | -field FieldList]
```

To display the properties of the virtual log repositories in the specified state:

```
lsvrepo -state RepositoryState [-detail | -field FieldList]
```

### Description

The `lsvrepo` command lists and displays the configuration of virtual log repositories. You can use the `-local` and `-sp` flags to view a specific virtual log repository. Several output formats are available. The default output format displays one row for every virtual log repository and consists of columns as follows:

- Shared storage pool name
- Repository state
- Root directory of the repository

You can view the detailed format by using the `-detail` flag. In this format, the virtual log repositories are displayed with each of the following properties:

- Shared storage pool name
- Repository state
- Root directory of the repository
- Maximum number of log files
- Maximum size of each log file
- Maximum number of state files
- Maximum size of each state file

You can view the field format by using the `-field` flag. In this format, the virtual log repositories are displayed based on the fields that you specify. Each field can be separated by a character that is not alphanumeric. You can specify the fields as follows:

|              |                                  |
|--------------|----------------------------------|
| <b>lf</b>    | Maximum size of each log file    |
| <b>lfs</b>   | Maximum number of log files      |
| <b>path</b>  | Root directory of the repository |
| <b>sf</b>    | Maximum size of each state file  |
| <b>sfs</b>   | Maximum number of state files    |
| <b>sp</b>    | Shared storage pool name         |
| <b>state</b> | Repository state                 |

## Flags

|                |                                                                                             |
|----------------|---------------------------------------------------------------------------------------------|
| <b>-detail</b> | Displays the extended output for the virtual log repository.                                |
| <b>-field</b>  | Displays the output with the specified list of fields.                                      |
| <b>-local</b>  | Displays the properties of the local virtual log repository.                                |
| <b>-state</b>  | Displays the properties of the virtual log repositories in the specified state.             |
| <b>-sp</b>     | Displays the properties of the virtual log repository in the specified shared storage pool. |

## Exit Status

Table 5. Command-specific return codes

| Return code | Description                          |
|-------------|--------------------------------------|
| 0           | All files were written successfully. |
| >0          | An error occurred.                   |

## Examples

1. To view the properties of all virtual log repositories, type the command as follows:

```
lsvlrepo
```

The system displays the output as follows:

Local Virtual Log Repository:

```
Repository State: enabled
Path: /var/vio/vlogs
Maximum Log Files: 2
Maximum Log File Size: 1048576
Maximum State Files: 2
Maximum State File Size: 1048576
```

Virtual Log Repository for Shared Storage Pool spool1:

```
Repository State: enabled
Path: /var/vio/SSP/cluster1/D_E_F_A_U_L_T_061310/vlogs/
Maximum Log Files: 2
Maximum Log File Size: 1048576
Maximum State Files: 2
Maximum State File Size: 1048576
```

2. To view the properties of the local virtual log repository, type the command as follows:

```
lsvlrepo -local
```

The system displays the output as follows:

Local Virtual Log Repository:

```
Repository State: enabled
Path: /var/vio/vlogs
Maximum Log Files: 2
Maximum Log File Size: 1048576
Maximum State Files: 2
Maximum State File Size: 1048576
```

3. To view the properties of the virtual log repository in the *sspool1* shared storage pool, type the command as follows:

```
lsvlrepo -sp sspool1
```

The system displays the output as follows:

```
Virtual Log Repository for Shared Storage Pool sspool1:
Repository State: enabled
Path: /var/vio/SSP/cluster1/D_E_F_A_U_L_T_061310/vlogs/
Maximum Log Files: 2
Maximum Log File Size: 1048576
Maximum State Files: 2
Maximum State File Size: 1048576
```

4. To view the properties of all virtual log repositories by using a custom field string, type the command as follows:

```
lsvlrepo -field "state-path lf"
```

The system displays the output as follows:

```
enabled-/var/vio/vlogs 2
enabled-/var/vio/SSP/cta1/D_E_F_A_U_L_T_061310/vlogs/ 2
```

## Related Information

The **chvlog** command, the **chvlrepo** command, the **lsvlog** command, the **mkvlog** command, and the **rmvlog** command.

---

## Isvopt command

### Purpose

Lists and displays information about the systems virtual optical devices.

### Syntax

```
lsvopt [-vtd VirtualTargetDevice] [-field Fieldnames] [-fmt Delimiter]
```

### Description

The **lsvopt** command displays information about file backed virtual optical devices in the Virtual I/O Server. If no flags are specified, a list of all defined file backed virtual optical devices, loaded media, and the size of the media is displayed. If the **-vtd** flag is specified, the media loaded and its size is displayed for the given virtual target device.

This command will provide full scripting support through the use of the **-field** and **-fmt** flags.

### Flags

|                                |                                                                                                                                                                                                                                          |
|--------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>-vtd</b>                    | Specifies the Virtual target device of the file backed virtual optical media                                                                                                                                                             |
| <b>-field</b> <i>FieldName</i> | The following fields are supported if no flags are specified:<br><b>vtd</b> Virtual target device of file backed virtual optical device<br><b>media</b> File name of loaded media or No Media<br><b>size</b> Size of loaded media or n/a |
| <b>-fmt</b> <i>Delimiter</i>   | Specifies a delimiter character to separate output fields.                                                                                                                                                                               |

## Examples

To display information about all file backed virtual optical devices in the Virtual I/O Server, type the following command:

```
lsvopt
```

The system displays output similar to the following:

| VTD     | Media       | Size(mb) |
|---------|-------------|----------|
| testopt | No Media    | n/a      |
| vtopt2  | No Media    | n/a      |
| vtopt3  | clientCD    | 640      |
| vtopt4  | No Media    | n/a      |
| vtopt5  | No Media    | n/a      |
| vtopt6  | No Media    | n/a      |
| vtopt7  | No Media    | n/a      |
| vtopt8  | No Media    | n/a      |
| vtopt9  | No Media    | n/a      |
| vtopt10 | No Media    | n/a      |
| vtopt11 | No Media    | n/a      |
| vtopt12 | No Media    | n/a      |
| vtopt13 | clientCD    | 640      |
| vtopt14 | No Media    | n/a      |
| vtopt15 | No Media    | n/a      |
| vtopt16 | installDVD1 | 1000     |
| vtopt17 | installDVD2 | 100      |

---

## migratepv command

### Purpose

Moves allocated physical partitions from one physical volume to one or more other physical volumes.

### Syntax

```
migratepv [-lv LogicalVolume] SourcePhysicalVolume DestinationPhysicalVolume ...
```

### Description

The **migratepv** command moves allocated physical partitions and the data they contain from the *SourcePhysicalVolume* to one or more other physical volumes, *DestinationPhysicalVolume*. All physical volumes must be within the same volume group. The specified source physical volume cannot be included in the list of *DestinationPhysicalVolume* parameters.

The allocation of the new physical partitions follows the policies defined for the logical volumes that contain the physical partitions being moved.

If you specify a logical volume that contains the boot image, the **migratepv -lv** command attempts to find enough contiguous partitions on one of the target physical volumes. If the migration is successful, the **migratepv** command will indicate a change in the boot device as well as the new boot physical volume. The attempted migration fails if the **migratepv -lv** command is unable to find enough contiguous space to satisfy the request.

**Note:** All logical volume manager migration functions work by creating a mirror of the logical volumes involved, then resynchronizing the logical volumes. The original logical volume is then removed. If the **migratepv** command is used to move a logical volume containing the primary dump device, the system will not have an accessible primary dump device during the execution of the command. Therefore, a dump taken during this execution may fail.

## Flags

**-lv** Moves only the physical partitions allocated to the specified logical volume and located on the specified source physical volume.

## Exit Status

| Return code | Description                                           |
|-------------|-------------------------------------------------------|
| 8           | The physical volume is not assigned to a volume group |

## Examples

1. To move physical partitions from **hdisk1** to **hdisk6** and **hdisk7**, type:

```
migratepv hdisk1 hdisk6 hdisk7
```

Physical partitions are moved from one physical volume to two others within the same volume group.

2. To move physical partitions in logical volume **lv02** from **hdisk1** to **hdisk6**, type:

```
migratepv -lv lv02 hdisk1 hdisk6
```

Only those physical partitions contained in **lv02** are moved from one physical volume to another.

## Related Information

The **lspv** command.

---

## IVM migrpar command

### Purpose

Moves active or inactive logical partitions from one physical system to another physical system. This command is operable only in an Integrated Virtualization Manager environment.

### Syntax

To validate a migration

```
migrpar [-o v -m <managed system> -t <managed system> --ip <target HMC/IVM IP address> [-u <target HMC/IVM username>]] -p <partition name> | --id <partition ID> [-n <profile name>] [-f <input data file> | -i "<input data>"] [-w <wait time>] [-d <detail level>]
```

To migrate a logical partition

```
migrpar [-o m -m <managed system> -t <managed system> --ip <target HMC/IVM IP address> [-u <target HMC/IVM username>]] -p <partition name> | --id <partition ID> [-n <profile name>] [-f <input data file> | -i "<input data>"] [-w <wait time>] [-d <detail level>] [--async] [-v] | [-redundantpgvios { 0 | 1 | 2 }]
```

To stop a migration

```
migrpar [-o s -m <managed system> {-p <partition name> | --id <partition ID> } [--help]
```

To recover from a failed partition migration

```
migrpar [-o r -m <managed system> [--ip <target HMC/IVM IP address>] [-u <target HMC/IVM username>]] [-p <partition name> | --id <partition ID>] [--force] [--help]
```



## Description

The `migrpar` command validates, starts, stops, and recovers a partition migration. The Integrated Virtualization Manager determines which type of migration to perform based on the state of the partition referenced in the command.

## Flags

|                                             |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|---------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <code>-o</code>                             | The partition migration operation. Possible values: <ul style="list-style-type: none"><li>• <code>s</code> - stops a partition migration</li><li>• <code>m</code> - validates and migrates a partition if the validation succeeds</li><li>• <code>r</code> - recover from a failed partition migration</li><li>• <code>v</code> - validates a partition migration</li></ul>                                                                                                                                                                                                                                                |
| <code>-m managed system</code>              | The name of the source managed system for the partition migration operation. The name may be the user-defined name for the managed system, or may be in the form <code>ttt-mmm*sssssss</code> , where <code>ttt</code> is the machine type, <code>mmm</code> is the model, and <code>sssssss</code> is the serial number of the managed system.                                                                                                                                                                                                                                                                            |
| <code>-t managed system</code>              | The name of the target system for the partition migration operation. The name may be the user-defined name for the managed system, or may be in the form <code>tttmmm*sssssss</code> , where <code>ttt</code> is the machine type, <code>mmm</code> is the model, and <code>sssssss</code> is the serial number of the managed system. The <code>-t</code> flag is required when using the <code>-o m</code> or <code>-o v</code> flags.                                                                                                                                                                                   |
| <code>--ip target HMC/IVM IP address</code> | The IP address or hostname of the target system managing the Hardware Management Console or Integrated Virtualization Manager.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <code>-u target HMC/IVM username</code>     | The user name to use on the Integrated Virtualization Manager that manages the target system or the HMC. If <code>--ip</code> flag is specified and <code>-u</code> flag is not specified, the user name on the source HMC or Integrated Virtualization Manager is used on the target HMC or Integrated Virtualization Manager.                                                                                                                                                                                                                                                                                            |
| <code>-p partition name</code>              | The name of the partition on which to perform the migration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <code>--id partition ID</code>              | The ID of the partition on which to perform the migration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <code>-n profile name</code>                | The name of the partition profile to be created for the migrated partition on the destination managed system. If this option is omitted when migrating a partition, then the last activated profile for the partition will be replaced with the current partition configuration on the destination managed system.<br><br>This option is only valid when migrating a partition or validating a partition migration. <b>Note:</b> This option is only valid if you are migrating to a system managed by the HMC. If you are migrating to a system managed by the Integrated Virtualization Manager, the option is not used. |

**-f** *input data file*

The name of the file containing the input data for the **migr1par** command. The data given in the file specified with the **-f** flag, or the data specified with **-i**, must be in comma-separated (CSV) format. These switches can be used with the migrate (**-o m**) and the validate (**-o v**) operations. The following attributes are supported:

```
virtual_scsi_mappings, virtual_fc_mappings, source_msp_name,
source_msp_ipaddr, source_msp_id,
dest_msp_name, dest_msp_ipaddr,
dest_msp_id, shared_proc_pool_id,
shared_proc_pool_name, paging_device
primary_paging_vios_id
primary_paging_vios_name
```

The data specified with the `virtual_scsi_mappings` or `virtual_fc_mappings` attribute consists of one or more source virtual SCSI or virtual fibre channel adapters to destination Virtual I/O Server logical partitions in the following format:

```
client_virtual_slot_num/dest_vios_lpar_name/
dest_vios_lpar_id
```

#### Attribute names

##### **dest\_msp\_id**

Specifies the partition ID to use on the destination managed system.

##### **dest\_msp\_ipaddr**

Specifies the IP address of the mover service partition of the destination managed system.

**Note:** This value is verified by comparing it with the output of the **lstcpip -interfaces** command.

##### **dest\_msp\_name**

Specifies the mover service partition name of the destination managed system.

##### **paging\_device**

The paging space device to use if a memory pool is used. A paging device is a block storage device that has been added to the memory pool and is not designated as a paging device for any other logical partition. If the `paging_device` value is a blank string, there is no paging device currently assigned.

##### **primary\_paging\_vios\_id**

The ID of the primary paging Virtual I/O Server (VIOS) partition that provides access to the paging space devices for the shared memory partitions. A paging VIOS partition is a VIOS logical partition that is assigned to the shared memory pool.

##### **primary\_paging\_vios\_name**

The name of the primary paging VIOS partition that provides access to the paging space devices for the shared memory partitions. A paging VIOS partition is a VIOS logical partition that is assigned to the shared memory pool.

##### **shared\_proc\_pool\_id**

The unique decimal identifier for a shared processing pool in which this logical partition should be on the target system. The default ID is zero. If Integrated Virtualization Manager is the destination managed system, the ID must equal 0.

##### **shared\_proc\_pool\_name**

Specifies the name of the shared processor pool in which this partition should be on the target system. This attribute is only valid for logical partitions using shared processors. The default value is `DefaultPool`.

**source\_msp\_id**

Specifies the partition ID to use on the source managed system. On IVM, this ID must equal the ID of the Virtual I/O Server.

**source\_msp\_ipaddr**

Specifies the IP address of the mover service partition of the source managed system.

**Note:** This value is verified by comparing it with the output of the **lstcpip -interfaces** command.

**source\_msp\_name**

Specifies the mover service partition name of the source managed system. On IVM, this name must equal the name of the Virtual I/O Server logical partition.

**virtual\_fc\_mappings**

Comma-separated list of virtual fibre channel adapters. Each item in this list has the format `slot_num/vios_lpar_name/vios_lpar_id`. For example, `4/vios2/3` specifies a virtual fibre channel adapter on a client logical partition with a virtual slot number of 4, a VIOS partition name of `vios2`, and the ID of the destination VIOS logical partition of 3.

**virtual\_scsi\_mappings**

Comma-separated list of virtual SCSI adapters. Each item in this list has the format `slot_num/vios_lpar_name/vios_lpar_id`. For example, `2/vios/1` specifies a client virtual SCSI adapter with a virtual slot number of 2, a VIOS partition name of `vios`, and the ID of the destination VIOS logical partition of 1.

**-i** *input data*

The input data for the **migr\_lpar** command. The format of the filter data is `attr_name1=value,attr_name2=value,...` or `attr_name1=value1,value2,...`

**-w** *wait time*

The maximum time, in minutes, to wait for the operating system commands issued by the HMC or Integrated Virtualization Manager to the partition to be migrated to complete.

**-d** *detail level*

The level of detail requested from the operating system commands issued by the HMC or Integrated Virtualization Manager to all partitions participating in the migration. Values range from 0 (none) to 5 (highest).

**--async**

The command returns after the validation is completed. This flag does not wait for the migration to complete. This is only valid with the **-o m** flag.

**-v**

Enables verbose mode for the partition migration operation. When verbose mode is enabled, detail messages and warning messages are displayed for a successful partition migration. Detail messages and warning messages are always displayed for a partition migration that fails, regardless of whether this option is specified.

**--force**

Forces a recover operation to proceed when errors are encountered. This option is only valid when recovering from a failed partition migration and can be initiated from either the source VIOS logical partition or the destination managed system. If the operation is initiated from the destination managed system, only the destination managed system is recovered.

**redundantpgvios** {0 | 1 | 2}

Specifies whether to configure the partition for redundancy on the target managed system. Possible values follow:

**0** Do not configure the partition for redundancy on the target managed system.

**1** Configure the partition for redundancy on the target managed system.

**2** Configure the partition for redundancy if possible on the target managed system. If redundancy is not possible, configure as nonredundant.

If you do not specify the parameter, the current redundancy level of the migrating partition is used.

**--help**

Displays the help text for this command and exit.

## Exit Status

This command has a return code of 0 on success.

## Examples

1. To validate a partition migration, type the following command:

```
migr1par -o v -m migfspL1 --ip migivm2 -t migfspL2 --id 5
-i "shared_proc_pool_name=ProcPoolA"
```

2. To perform a partition migration, type the following command:

```
migr1par -o m -m migfspL1 --ip migivm2 -t migfspL2 --id 5 -i "source_msp_id=1,
source_msp_ipaddr=9.3.252.192,dest_msp_id=1,dest_msp_ipaddr=9.3.126.77"
```

3. To perform a partition migration for logical partitions that use virtual fibre channel adapters, type the following command:

```
migr1par -o m -m migfspL1 --ip migivm2 -t migfspL2 --id 5
-i "virtual_fc_mappings=5/VIOS/1,6/VIOS3/3"
```

4. To stop a partition migration, type the following command:

```
migr1par -o s -m migfspL1 --id 5
```

5. To recover from a failed partition migration, type the following command:

```
migr1par -o r -m migfspL1 --id 5
```

---

## mirrorios command

### Purpose

Mirrors all the logical volumes on rootvg. This command may reboot the partition as part of its operation.

### Syntax

```
mirrorios [-f] [-defer] [PhysicalVolume ...]
```

### Description

The **mirrorios** command takes all the logical volumes on the rootvg volume group and mirrors those logical volumes. The target physical drives must already be members of the volume group.

The **mirrorios** command attempts to mirror the logical volumes onto any of the disks in a volume group. To control which drives are used for mirroring, you must include the list of disks in the input parameters, *PhysicalVolume*. Mirror strictness is enforced. The **mirrorios** command mirrors the logical volumes, using the default settings of the logical volume being mirrored.

The Virtual I/O Server will restart when the command is completed and you are prompted to continue. The Virtual I/O Server will be restarted as part of the command's operation. If the user does not want the VIOS to reboot as part of the command's operation, the user may specify the **-defer** option, which is only available for VIOS 1.4 or later. No need to reboot for VIOS 1.5 or later, so you should use the **-defer** option. If the command restarts as part of its operation, then the user can specify not to be prompted to continue. The **-f** option tells the command to run without prompting you.

**Note:** For best results, mirror the rootvg volume group on all Virtual I/O Server partitions.

Only the prime administrator (padmin) can run this command.

## Flags

|                       |                                                                                                                                        |
|-----------------------|----------------------------------------------------------------------------------------------------------------------------------------|
| <b>-f</b>             | Runs the command without prompting you to continue.                                                                                    |
| <b>-defer</b>         | Instructs the command not to reboot the VIOS. Running the command without the <b>-defer</b> option will automatically reboot the VIOS. |
| <i>PhysicalVolume</i> | Specifies the target physical volume name. The volume must already be a member of the volume group.                                    |

## Exit Status

|   |                                                               |
|---|---------------------------------------------------------------|
| 5 | The Virtual I/O Server is already mirrored                    |
| 6 | Boot LV not found                                             |
| 7 | The physical volume appears to belong to another volume group |

## Examples

1. To mirror the Virtual I/O Server root volume group to physical volume **hdisk4**, type the following command:  

```
mirrorios -f hdisk4
```
2. To mirror the Virtual I/O Server root volume group to physical volume **hdisk4**, but defer a system reboot, type the following command:  

```
mirrorios -defer hdisk4
```

## Related Information

The **activatevg** command, the **chvg** command, the **deactivatevg** command, the **exportvg** command, the **importvg** command, the **lsvg** command, the **mkvg** command, the **syncvg** command, the **unmirrorios** command, and the **alt\_root\_vg**.

---

## mkauth command

### Purpose

Creates a new user-defined authorization.

### Syntax

```
mkauth [-a] [Attribute=Value...] [newauth]
```

### Description

The **mkauth** command creates a user-defined authorization in the authorization database. You can create authorization hierarchies by using a dot (.) in the *auth* parameter to create an authorization of the form *ParentAuth.SubParentAuth.SubSubParentAuth.....* All the parent elements in the *auth* parameter must already exist in the authorization database before the new authorization is created. The maximum number of parent elements that you can use to create an authorization is eight.

Authorization attributes can be set at creation time through the *Attribute=Value* parameter. Every authorization that you create must have a value for the ID authorization attribute. If you do not specify the value by using the **mkauth** command, the command automatically generates a unique ID for the authorization. If you specify an ID, the value must be unique and greater than 15000 because IDs 1 to 15000 are reserved for system-defined authorizations.

For the current VIOS commands and their associated authorizations, see command authorization chart, Authorizations corresponding to Virtual I/O Server commands.

## Restrictions

You are not allowed to create any authorizations for VIOS commands, unlike AIX. In AIX, an authorized user can create his own hierarchy of authorizations for all commands. But you can create your own authorizations for your own commands or scripts in VIOS. You cannot create any authorizations that start with “**vios.**” or “**aix.**” The authorizations that start with “**vios.**” and “**aix.**” are considered system-defined authorizations. Hence, you cannot add any further hierarchies to these authorizations.

Authorization names must not begin with a dash (-), a plus sign (+), an at sign (@), or a tilde (~) character. They must not contain any space, tab, or newline characters. You cannot use the keywords **ALL**, **default**, **ALLOW\_OWNER**, **ALLOW\_GROUP**, **ALLOW\_ALL**, or an asterisk (\*) as an authorization name. Additionally, do not use any of the following characters within an authorization string:

- : (colon)
- " (quotation mark)
- # (number sign)
- , (comma)
- = (equal sign)
- \ (backslash)
- / (forward slash)
- ? (question mark)
- ' (single quotation mark)
- ` (grave accent)

**Note:** After creating authorizations, run the **setkst** command to update the kernel database with the new authorizations.

## Security

The **mkauth** command is a privileged command. To run the command successfully, assume a role that has the authorization **vios.security.auth.create**.

## Flags

**-a** Specifies the name of the new virtual optical media file.

## Attributes

|                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|-----------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>id</b>       | Specifies a unique integer that is used to identify the authorization. Valid values are 10,001 - 32,768.                                                                                                                                                                                                                                                                                                                                                               |
| <b>dfltmmsg</b> | Specifies the default description to use if message catalogs are not in use. The value is a string.                                                                                                                                                                                                                                                                                                                                                                    |
| <b>msgcat</b>   | Specifies the message catalog file name that contains the description of the authorization. If the <i>msgcat</i> attribute is specified, the <i>msgset</i> , and <i>msgnum</i> attributes must also be specified. The value is a string. If the specified string contains a leading forward slash (/), the value is assumed to be an absolute path name. Otherwise, the user environment defines the directory search path as specified by the <b>catopen</b> routine. |
| <b>msgset</b>   | Specifies the message set number in the file name to retrieve the message number. The file name is specified by the <i>msgcat</i> attribute, and the message number is specified by the <i>msgnum</i> attribute. The value is a decimal integer.                                                                                                                                                                                                                       |
| <b>msgnum</b>   | Specifies the message number for the description of the authorization in the file and the set. The authorization is specified by the <i>msgcat</i> attribute, and the set number is specified by the <i>msgset</i> attribute. The value is a decimal integer.                                                                                                                                                                                                          |

## Exit Status

0 for success.

Non-zero for command error.

## Examples

1. To create a top-level authorization *custom* and to have the **mkauth** command assign an appropriate ID value, type the following command:

```
mkauth custom
```

2. To create a *custom.test* child authorization and to assign an ID and default description, type the following command:

```
mkauth -a id=11000 dfltmmsg="My authorization" custom.test
```

---

## IVM mkauthkeys command

### Purpose

Allows key-based SSH authentication between two systems. This updates the `~/.ssh/authorized_keys2` file with the specified public key. It can also be used to push the users public key to a remote Integrated Virtualization Manager or HMC system.

### Syntax

To add the SSH key as an authorized key locally:

```
mkauthkeys { -a | --add } <key string>
```

To remove the SSH key locally:

```
mkauthkeys { -r | --remove } [-u <user>] <key string>
```

To exchange public keys with a remote system:

```
mkauthkeys { -a | --add } -- ip <remote system> [-u <user>] <key string>
```

To test the remote non-interactive authentication using the SSH key:

```
mkauthkeys --test -- ip <remote system> [-u <user>]
```

### Description

The **mkauthkeys** command updates the Integrated Virtualization Manager user's `authorized_keys2` file.

### Flags

|                 |                                                                                                                           |
|-----------------|---------------------------------------------------------------------------------------------------------------------------|
| <b>-a</b>       | Adds the <b>ssh</b> command key.                                                                                          |
| <b>-g</b>       | Displays the public key for the specified user and generates the user's public and private key pair if they do not exist. |
| <b>-r</b>       | Removes the key for the specified user ID and host.                                                                       |
| <b>--add</b>    | Adds the <b>ssh</b> command key.                                                                                          |
| <b>--remove</b> | Removes the key for the specified user ID and host.                                                                       |
| <b>--test</b>   | Verifies authentication to the remote host                                                                                |

|                                |                                                                                                                                                                                                                                                                      |
|--------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>--ip</b> <remote server IP> | Allows installing this user's public key on the specified remote HMC or Integrated Virtualization Manager system for the user specified with <b>-u</b> flag. If <b>-u</b> flag is not specified, the remote user's public key will be installed on the local system. |
| <b>-u</b> <i>username</i>      | Specifies the user name to add or remove the key. You must have hmcsuperadmin or PAdmin authority to add or remove key for other users.                                                                                                                              |
| <i>key string</i>              | The <b>ssh</b> command key to add to or the ID to remove.                                                                                                                                                                                                            |

## Exit Status

This command has a return code of 0 on success.

## Examples

1. To add the SSH key generated for user *joe@somehost*, type the following command:  

```
mkauthkeys -a 'adB8fqeZs2d-gg+q joe@somehost'
```
2. To display the current user's public key, type the following command:  

```
mkauthkeys -g
```
3. To display the public key for *fred*, type the following command:  

```
mkauthkeys -g -u fred
```
4. To remove the SSH key generated for user *joe@somehost*, type the following command:  

```
mkauthkeys -r 'adB8fqeZs2d-gg+q joe@somehost'
```
5. To remove all SSH keys generated for user *joe@somehost*, type the following command:  

```
mkauthkeys -r 'joe@somehost'
```
6. To add the SSH key generated for this user to a remote.host as user *fred*, type the following command:  

```
mkauthkeys -a --ip remote.host -u fred
```

**Note:** The user will be prompted for the password on the remote.host.

7. To enable *user* from *somesystem* with public key *ssh-rsa thersakeygoeshere=* to access the Virtual I/O Server without using a password, type the following command:  

```
mkauthkeys -a ssh-rsa thersakeygoeshere= user@somesystem
```
8. To remove the key from the authorized key list, type the following command:  

```
mkauthkeys -r ssh-rsa thersakeygoeshere= user@somesystem
```
9. To remove all keys that end with the string *user@somesystem*, type the following command:  

```
mkauthkeys -r user@somesystem
```
10. To allow the *padmin* user to remove a key for any user, type the following command:  

```
mkauthkeys -r -u user ssh-rsa thersakeygoeshere= user@somesystem
```
11. To add the current user's local public key to the authorized keys list on a remote system, and add the user's remote public key to the authorized keys list on the local system, type the following command:  

```
mkauthkeys -a --ip othersystem.com
```
12. To add the current user's local public key to the authorized keys list for remote user *user* on a remote system, and add the remote public key of *user* to the authorized keys list for the current user on the local system, type the following command:  

```
mkauthkeys -a --ip othersystem.com -u user
```
13. To verify non-interactive authentication using the SSH key for the current user, type the following command:  

```
mkauthkeys --test --ip othersystem.com
```



**Note:** If it returns 0, then non-interactive authentication is working properly. If **mkauthkeys** returns nonzero, then non-interactive authentication is not configured properly and displays the following message: [VIOSE0104200B-0217] Permission denied (publickey,password,keyboard-interactive).

14. To verify non-interactive authentication using the SSH key for the current user on local system to specify the user on a remote system, type the following command:

```
mkauthkeys --test --ip othersystem.com -u user
```

---

## mkbdsp command

### Purpose

Assigns storage from a storage pool, which serves as a backing device for a virtual SCSI (VSCSI) adapter.

### Syntax

To create a backing device file or logical volume:

```
mkbdsp [-sp StoragePool] Size -bd BackingDevice
```

To assign an existing file or logical volume as a backing device:

```
mkbdsp [-sp StoragePool]-bd BackingDevice -vadapter ServerVirtualSCSIAdapter [-tn TargetDeviceName]
```

To create a new file or logical volume as a backing device:

```
mkbdsp [-sp StoragePool] Size [-bd BackingDevice] -vadapter ServerVirtualSCSIAdapter [-tn TargetDeviceName]
```

To create a logical unit in a shared storage pool:

```
mkbdsp -clustername ClusterName -sp StoragePool Size -bd BackingDevice LU [-thick]
```

To assign a logical unit as a backing device in a shared storage pool:

```
mkbdsp -clustername ClusterName -sp StoragePool { -bd LogicalUnit | -luudid LUUIDID } -vadapter ServerVirtualSCSIAdapter [-tn TargetDeviceName]
```

To create a new logical unit as a backing device in a shared storage pool:

```
mkbdsp -clustername ClusterName -sp StoragePool Size -bd BackingDevice LU -vadapter ServerVirtualSCSIAdapter [-tn TargetDeviceName] [-thick]
```

### Description

The **mkbdsp** command assigns a backing device to a VSCSI server adapter. If the **-sp** flag is not specified the default storage pool is used. The storage pool must be specified when working with file-backed devices and logical units. The default storage pool is still used when working with logical volumes. If the storage size is given, the **mkvdev** command creates a backing device of at least the specified size and assign it as the backing device. When working with file backed devices the **-bd** flag must be specified. The system does not generate a name. The type of backing device created is determined by the storage pool type. The size can be given in the number of megabytes (M or m), the number of gigabytes (G or g), or the number of physical partitions.

#### Note:

- Specifying physical partitions works only for logical volume backing devices.

- The backing device specified cannot be assigned to a shared memory pool (to be used as a paging space device by a shared memory partition).

You must specify the name for the newly created backing device by using the **-bd** flag in combination with the size parameter. Naming the backing device is optional when working with logical volumes. You also have the option of assigning the name for the newly created virtual target device by using the **-tn** flag in combination with the **-vadapter** flag.

## Flags

|                     |                                                                                                                                         |
|---------------------|-----------------------------------------------------------------------------------------------------------------------------------------|
| <b>-bd</b>          | Specifies the backing device or logical unit name.                                                                                      |
| <b>-clustername</b> | Specifies the cluster name                                                                                                              |
| <b>-luudid</b>      | Specifies the logical unit (LU) UDID in case specified lu is not unique.                                                                |
| <b>-sp</b>          | Specifies the storage pool to be used.                                                                                                  |
| <b>-thick</b>       | Creates the device as a thick-provisioned device. The default is a thin-provisioned device.                                             |
| <b>-tn</b>          | Specifies the name of the target device.<br><b>Note:</b> The only values accepted are alphanumeric, a dash, an underscore, or a period. |
| <b>-vadapter</b>    | Specifies the VSCSI server adapter.                                                                                                     |

## Exit Status

|    |                                                          |
|----|----------------------------------------------------------|
| 23 | Specified storage pool is not a valid storage pool.      |
| 26 | Specified name is already used. Choose a different name. |
| 34 | Specified name is reserved. Choose a different name.     |

## Examples

1. To create a virtual target device that maps a 3 GB backing device from the default storage pool the virtual SCSI server adapter vhost3, type the command as follows:  

```
mkbdsp -bd bname 3g -vadapter vhost3
```
2. To create an LU in a specific shared storage pool, type the command as follows:  

```
mkbdsp -clustername newcluster -sp viossp 100M -bd LU
```

The system displays the output as follows:

```
Lu Name:LU
Lu Udid:c960d8f854d4064d74b7d0017c4063a2
```
3. To map an LU with a specific virtual adapter, type the command as follows:  

```
mkbdsp -clustername newcluster -sp viossp -bd LU -vadapter vhost0
```

The system displays the output as follows:

```
Assigning file "LU" as a backing device.
VTD:vtscsi0
```
4. To create a thick-provisioned LU of 5 GB in the specific storage pool, type the command as follows:  

```
mkbdsp -clustername newcluster -sp viossp 5G -bd THICK_LU -thick
```

The system displays the output as follows:

```
Lu Name:THICK_LU
Lu Udid:7f9ce0be4d5b5c8ddeb339fc1c71e0bf
```
5. To create and map a thick-provisioned LU to the specified VSCSI server adapter, type the command as follows:  

```
mkbdsp -clustername newcluster -sp viossp 2G -bd THICK_LU -vadapter vhost0 -thick
```

The system displays the output as follows:

```
Lu Name:THICK_LU
```

Lu Udid:510004e3d0e90c1d10e13be130b3cd34  
Assigning file "THICK\_LU" as a backing device.  
VTD:vtscsi0

---

## IVM mkgencfg command

### Purpose

Performs the initial logical partition configuration for the managed system. This command is operable only in an Integrated Virtualization Manager environment.

### Syntax

```
mkgencfg -o init [-i "ConfigurationData"] [-m ManagedSystem]
```

### Description

The **mkgencfg** command performs the initial logical partition configuration for the managed system. As part of the initial configuration, virtual Ethernet adapters are created in the management partition. The virtual Ethernet MAC address prefix can optionally be configured through this command.

### Flags

- o** *Operation*                      The type of operations:
- i** *ConfigurationData*              *init* - Perform initial logical partition configuration actions for the managed system  
The configuration data consists of attribute name/value pairs, which are in a comma separated value (CSV) format. The format of a configuration record is as follows:  
"attribute-name=value,attribute-name=value,..."
- Note that certain attributes accept a comma separated list of values, as follows:  
"attribute-name=value,value,...",..."
- When you specify a list of values, the attribute name/value pair must be enclosed in double quotation marks. Depending on the shell being used, nested double quotation marks might need to be preceded by an escape character.
- Valid Attributes for Configuration Data:**
- mac\_prefix**  
The prefix needs to be a 3 byte hexadecimal value. It specifies the first 2.5 bytes of the MAC address to be assigned to all virtual Ethernet adapters for this managed system. The value cannot be a multicast address (010000 bit must be off), and needs to be a private address (020000 bit must be on). For example, a valid MAC address prefix is 0642A0.
- pend\_configured\_max\_lpars**  
The maximum number of partitions supported by the management partition after the next restart.

### **virtual\_eth\_mac\_base\_value**

The virtual Ethernet MAC address base value is a per partition value. The base value is used to build the MAC addresses for each virtual Ethernet adapter in that partition. The base value is the first 5 bytes of the MAC address. The virtual slot number of the virtual Ethernet adapter makes up the last byte. If you do not assign a base value, one is generated automatically using the following format: Base value = 0xSSSSSBBBBB

Where SSSSS is the system wide MAC address prefix andBBBBB is a randomly generated sequence of bits (guaranteed to be unique within this physical system).

#### **Note:**

1. The system wide MAC prefix is randomly generated unless you override it using `mkgencfg -o init -i mac_prefix`.
2. If the base value is automatically generated, then the MAC address format for a virtual Ethernet adapter is 0xSSSSSBBBBBNN, where NN is the slot number. If you specify the base value using **mkgencfg** for partition 1 or **mksyscfg/chsyscfg** for any other partition, then the format is 0xBBBBBBBBBBNN, whereBBBBBBBBB is the base value that you specified.
3. If a virtual Ethernet adapter is in a slot greater than or equal to 256, then the slot number overflows into the base value, since it does not fit in 1 byte. For example, Integrated Virtualization Manager handles it as 0xBBBBBBBBBB00 + 0x00000000NNNN.
4. The **mkgencfg** command is the only way to specify the base value for partition 1. After this value is set, you cannot change it without losing all partition configuration. If you need to change the value, use the **lpcfgop** command. The value is set implicitly when **mkgencfg** is running. The value will be generated automatically, unless you specify the value. The **mkgencfg** command runs implicitly when you run your first **change** command. You can change the base value for any other partition, when the partition is powered off.

A side effect of setting the base value for partition 1 with **mkgencfg**, is that the system wide prefix is set to the first 2.5 bytes of the partition 1 base value. As a result, there are restrictions on setting both `mac_prefix` and `virtual_eth_mac_base_value` with **mkgencfg** at the same time. If you specify both, the `mac_prefix` value needs to be equal to the first 2.5 bytes of the `virtual_eth_mac_base_value` value.

### **-m ManagedSystem**

The name of the managed system. This attribute is optional because there is only one system to manage. The name might either be the user-defined name for the managed system, or be in the form `ttt-mmm*sssssss`, where `ttt` is the machine type, `mmm` is the model, and `sssssss` is the serial number of the managed system.

## **Exit Status**

This command has a return code of zero on success.

## **Security**

This command is not accessible by users with the ViewOnly role.

## **Examples**

1. To initialize the logical partition configuration for the managed system using defaults, type:  
`mkgencfg -o init`
2. To initialize the logical partition configuration for the managed system with support for 17 partitions and a MAC prefix of 0x06ABC0, type:

```
mkgencfg -o init -i "pend_lpm_max_lpars=17,mac_prefix=06ABC0"
```

---

## mkkrb5clnt command

### Purpose

Configures a Kerberos client.

### Syntax

```
mkkrb5clnt -h | -r <realm> { -c <KDC> -s <server> | -l {ldapsrvr | ldapsrvr:port} [-c <KDC> -s<server>] } [-a<admin>] -d <domain> [-A] [-i <database>] [-K] [-T] | -i <database> | -U [-a <admin>]
```

### Description

The `mkkrb5clnt` command configures a Kerberos client.

### Flags

|                                           |                                                                                                                                                                                                                                                                                                                                            |
|-------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>-h</b>                                 | Specifies that the command will only display the valid command syntax                                                                                                                                                                                                                                                                      |
| <b>-c</b> <i>KDC</i>                      | Specifies the KDC server                                                                                                                                                                                                                                                                                                                   |
| <b>-r</b> <i>realm</i>                    | Specifies the full realm name for which the Kerberos client is to be configured.                                                                                                                                                                                                                                                           |
| <b>-s</b> <i>server</i>                   | Specifies the fully qualified host name for Kerberos admin server                                                                                                                                                                                                                                                                          |
| <b>-U</b>                                 | Undo the setup from the previous configuration command                                                                                                                                                                                                                                                                                     |
| <b>-a</b> <i>admin</i>                    | Specifies the principal name of the Kerberos server admin                                                                                                                                                                                                                                                                                  |
| <b>-d</b> <i>domain</i>                   | Specifies the complete domain name for the Kerberos client                                                                                                                                                                                                                                                                                 |
| <b>-A</b>                                 | Specifies root to be added as a Kerberos administrative user                                                                                                                                                                                                                                                                               |
| <b>-i</b> <i>database</i>                 | Configures integrated Kerberos authentication                                                                                                                                                                                                                                                                                              |
| <b>-K</b>                                 | Specifies Kerberos to be configured as the default authentication scheme                                                                                                                                                                                                                                                                   |
| <b>-T</b>                                 | Specifies the flag to acquire the server administration TGT based administration ticket.                                                                                                                                                                                                                                                   |
| <b>-l</b> <i>ldapsrvr   ldapsrvr:port</i> | For clients, specifies the LDAP directory server to use for Administration server and KDC discovery using LDAP. If the <b>-l</b> flag is used, then the <b>-c</b> and <b>-s</b> flags are optional. If the <b>-l</b> option is not used, the <b>-c</b> and <b>-s</b> flags must be specified. The port number can optionally be specified. |

### Exit Status

| Return code | Description                                |
|-------------|--------------------------------------------|
| 0           | Success                                    |
| 1           | Invalid flag, argument, or command failure |

### Examples

- To display the command syntax, type the following command:  

```
mkkrb5clnt -h
```
- To configure `testbox.com` as the client, make root as the server admin, configure integrated login, configure Kerberos as default authentication scheme, type the following command:  

```
mkkrb5clnt -c bob.kerberso.com -r KERBER.COM -s bob.kerberso.com -d testbox.com -A -i files -K -T
```

---

## mkldap command

### Purpose

Sets up Virtual I/O Server as a Lightweight Direct Access Protocol (LDAP) client.

## Syntax

```
mkldap -host serverlist -bind bindDN -passwd bindpwd [-base baseDN] [-port serverport] [-timeout cacheTimeout] [-csize cacheSize] [-threads NumberOfThreads] [-hbeatint heartBeatInt] [-keypath SSL_database_path] [-keypasswd SSL_password] [-auth authType] [-users userlist | ALL]
```

**mkldap -deconfig**

## Description

The **mkldap** command is used to set up the Virtual I/O Server as an LDAP client. The server bind distinguished name (DN) and password for client access to the LDAP server. The **mkldap** command saves server bind DN, password, server name, the SSL key path and password, and other configuration attributes to the `/etc/security/ldap/ldap.cfg` file. The **mkldap** command saves the bind password and SSL key password (if SSL is configured) to the `/etc/security/ldap/ldap.cfg` file in encrypted format.

**Note:** These encrypted passwords are system specific, and can only be used by the `secldapclntd` daemon on the system where they are generated.

You can supply multiple LDAP servers to the **mkldap** command during client setup. In this case, the client contacts the servers in the supplied order and establishes connection to the first server to which the client can successfully bind.

The LDAP client communicates to the LDAP server through a client side daemon, **secldapclntd**.

The **secldapclntd** command is enabled or disabled by using the **startnetsvc** and **stopnetsvc** commands.

## Flags

|                                          |                                                                                                                                                                                   |
|------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>-host</b> <i>serverlist</i>           | Specifies comma-separated list of host names.                                                                                                                                     |
| <b>-bind</b> <i>bindDN</i>               | Specifies the DN (distinguished name) to bind to the LDAP server.                                                                                                                 |
| <b>-passwd</b> <i>bindpwd</i>            | Specifies the clear text password for the bindDN used to bind to the LDAP server.                                                                                                 |
| <b>-base</b> <i>baseDN</i>               | Specifies the base DN for the <b>mkldap</b> command, in which to search for the user base DN and group base DN. If you do not specify this flag, the entire database is searched. |
| <b>-port</b> <i>serverport</i>           | Specifies the port number to which the LDAP server is listening.                                                                                                                  |
| <b>-timeout</b> <i>cachetimeout</i>      | Specifies the maximum time length that a cache entry expires. Set this value to 0 to disable caching.                                                                             |
| <b>-csize</b> <i>cacheSize</i>           | Specifies the maximum number of user entries used in the client side daemon cache.                                                                                                |
| <b>-threads</b> <i>NumberOfThreads</i>   | Specifies the numbers of threads that the client side daemon uses.                                                                                                                |
| <b>-hbeatint</b> <i>heartBeatInt</i>     | Specifies the time interval of heartbeats between the client and the LDAP server.                                                                                                 |
| <b>-keypath</b> <i>SSL_database_path</i> | Specifies the full path to the SSL database.<br><b>Note:</b> This requires the <code>ldap.max_crypto_client</code> fileset to be installed.                                       |
| <b>-keypasswd</b> <i>SSL_password</i>    | Specifies the password for the SSL key.<br><b>Note:</b> This requires the <code>ldap.max_crypto_client</code> fileset to be installed.                                            |
| <b>-auth</b> <i>authType</i>             | Specifies the authentication mechanism that is used to authenticated users. Valid values are <code>unix_auth</code> and <code>ldap_auth</code> .                                  |
| <b>-users</b> <i>userlist</i>            | Specifies the comma-separated list of user names to enable for LDAP authentication. Specify ALL to enable all users on the client.                                                |
| <b>-deconfig</b> <i>bindpwd</i>          | Specifies that the previous client setup to the LDAP client configuration file should be undone.                                                                                  |

## Exit Status

0 Success

## Examples

Run one of the following **mkldap** commands with the **-u** flag for a user ID to become an LDAP user ID at client setup time:

- ```
mkldap -host ldapserv1 -bind cn=admin -passwd adminpwd -users user1,user2
mkldap -host ldapserv1 -bind cn=admin -passwd adminpwd
```
- To setup the client to talk to the server3.your_company.com LDAP server using SSL, enter the following command:


```
mkldap -bind cn=admin -passwd adminpwd -host server3.your_company.com
-base o=mycompany,c=us -keypath /usr/ldap/clientkey.kdb
-keypasswd keypwd -users user1,user2
```

Either of these commands set up the local host as the client of the LDAP server running on host **ldapserv1** and **cn=admin** and **-passwd adminpwd** are the LDAP server administrator DN and password.

Files

`/etc/security/ldap/ldap.cfg` Contains the **mkldap** command, and server bind DN, password, server name, the SSL key path and password, and other configuration attributes.

Related Information

The **ldapadd** command and the **ldapsearch** command.

mklv command

Purpose

Creates a logical volume.

Syntax

```
mklv [-mirror] [-lv LogicalVolume | -prefix Prefix] [-type Type] VolumeGroup Size [PhysicalVolume ... ]
```

Description

The **mklv** command creates a new logical volume within the *VolumeGroup*. If you specify one or more physical volumes with the *PhysicalVolume* parameter, only those physical volumes are available for allocating physical partitions; otherwise, all the physical volumes within the volume group are available.

The allocation policy is to use a minimum number of physical volumes.

The *type* parameter specifies the logical volume type. The standard types are **jfs** (journaled file systems), **jfslog** (journaled file system logs), **jfs2** (enhanced journaled file system), **jfs2log** (enhanced journaled file system logs), and **paging** (paging spaces). You can define other logical volume types with this flag. You cannot create a striped logical volume of type **boot**. The default is **jfs**.

The *Size* parameter specifies the minimum size the logical volume should be. When specifying *Size* the following conventions must be used:

Size	Minimum logical volume size
###M/m	### MB
###G/g	### GB

Flags

- lv** Specifies the logical volume name to use instead of using a system-generated name. Logical volume names must be unique system wide name, and can range from 1 to 15 characters.
- mirror** Activates mirroring for this logical volume.
- prefix** Specifies the Prefix to use instead of the prefix in a system-generated name for the new logical volume. The prefix must be less than or equal to 13 characters. The name cannot begin with a prefix already defined in the **PdDv** class in the Device Configuration Database for other devices, nor be a name already used by another device.
- type** Sets the logical volume type.

Exit Status

See “Exit status for Virtual I/O Server commands” on page 3.

Examples

1. To make a logical volume in volume group **vg02** with a minimum size of 1 Mb , type the following command:

```
mklv vg02 1M
```
2. To make a logical volume in volume group **vg03** with 1GB chosen from physical volumes **hdisk5**, **hdisk6**, and **hdisk9**, type the following command:

```
mklv vg03 1G hdisk5 hdisk6 hdisk9
```
3. To request a logical volume with a minimum size of 10MB, type the following command:

```
mklv VGNAME 10m
```

where *VGNAME* is the name of your logical volume.

4. To make a logical volume in volume group **vg04** with a minimum size of 10 Mb whose type is paging, type the following command:

```
mklv -lv lv01 -type paging vg04 10M
```

The system displays output as follows:

```
lv01
```

Related Information

The **lslv** command, the **extendlv** command, and the **rmlv** command.

mklvcopy command

Purpose

Creates a mirror of a logical volume.

Syntax

```
mklvcopy LogicalVolume [ PhysicalVolume ... ]
```


Description

The **mklvcopy** command creates a mirror (an additional copy) of a *LogicalVolume*. The *LogicalVolume* parameter can be a logical volume name or logical volume ID. You can request that the new copy of the logical volume be allocated on specific physical volumes (within the volume group) with the *PhysicalVolume* parameter; otherwise, all the physical volumes within the volume group are available for allocation. The new copy of the logical volume will be placed on a separate physical volume.

Note: Only one additional copy of a logical volume can be created.

Exit Status

See “Exit status for Virtual I/O Server commands” on page 3.

Examples

1. To create a copy of the logical volume **lv01**, so that a total of two copies exist, type:

```
mklvcopy lv01
```

Related Information

The **extendlv** command, the **lslv** command, the **mklv** command, the **rmlv** command, and the **rmlvcopy** command.

mkpath command

Purpose

Adds to the system another path to an MPIO capable device.

Syntax

```
mkpath { [ -dev Name ] [ -pdev Parent ] [ -conn Connection ] } [ -def ]
```

Description

The **mkpath** command defines, and possibly configures, one or more paths to the target device (**-dev *Name***). The paths are identified by a combination of the **-dev *Name***, **-pdev *Parent***, and **-conn *Connection*** flags. Both the target device and parent must be previously defined in the system to define a path. They both must be **AVAILABLE** to configure a path.

If the **-def *flag*** is specified, the **mkpath** command only defines the new path definition to the system. If the **-def** flag is not specified, the **mkpath** command attempts to define the path, if it does not already exist, before it attempts to configure the path. Configuring a path requires the path to already be defined and both the device and the parent device to already be configured.

The **mkpath** command displays a status message upon completion. It is possible for some paths to configure and others to fail.

Note that not all devices will be able to have paths manually defined by using the **mkpath** command. These limitations are due to the way that path information is stored for these devices. Fiber channel devices fall into this category.

The **mkpath** command provides status messages about the results of operation. Messages in one of the following formats will be generated:

path [available | defined]

This message is displayed when **mkpath** is run on a single path. If the path is successfully configured the message **path available** is displayed. If the path is not successfully configured and there is no explicit error code returned by the method, the message **path defined** is displayed.

paths available

This message is displayed if multiple paths were identified and all paths were successfully configured.

some paths available

This message is displayed if multiple paths were identified, but only some of them were successfully configured.

no paths processed

This message is generated if no paths were found matching the selection criteria.

Flags

-conn <i>Connection</i>	Indicates the connection information associated with the path to be added. This flag is required if the -def flag is specified.
-def	Defines a new path to a device by adding a path definition to the system. The new path will not automatically be configured when the -def flag is specified. Note that only one path may be defined at a time. The -conn and the -pdev flags are required when the -def flag is used.
-dev <i>Name</i>	Specifies the logical device name of the target device to which the path(s) are being added. The path(s) to be added are qualified by the -pdev and -conn flags.
-pdev <i>Parent</i>	Indicates the logical device name of the parent device associated with the path(s) to be added. This flag is required if the -def flag is specified.

Exit Status

See "Exit status for Virtual I/O Server commands" on page 3.

Examples

1. To define and configure an already defined path between **scsi0** and the **hdisk1** device at **SCSI ID 5** and **LUN 0** (connection 5,0), enter:

```
mkpath -dev hdisk1 -pdev scsi0 -conn 5,0
```

The system displays a message similar to the following:

```
path available
```

2. To configure an already defined path from **fcscsi0** to fiber channel disk **hdisk1**, type:

```
mkpath -dev hdisk1 -pdev fcscsi0
```

The system displays a message similar to the following:

```
path available
```

3. To only add to the Customized Paths object class a path definition between **scsi0** and the **hdisk1** disk device at **SCSI ID 5** and **LUN 0**, enter:

```
mkpath -def -dev hdisk1 -pdev scsi0 -conn 5,0
```

The system displays a message similar to the following:

```
path defined
```

Related Information

The **lspath** command and the **rmpath** command.

mkrole command

Purpose

Creates new roles.

Syntax

```
mkrole [Attribute=Value...] [newrole]
```

Description

The **mkrole** command creates a role. The *newrole* parameter must be a unique role name. You cannot use the **ALL** or **default** keywords as the role name.

Every role must have a unique role ID that is used for security decisions. If the *id* attribute is not specified when a role is created, the **mkrole** command automatically assigns a unique ID to the role.

For the current VIOS commands and their associated authorizations, see command authorization chart, Authorizations corresponding to Virtual I/O Server commands.

Restrictions

The *role* parameter cannot contain any space, tab, or newline characters.

To prevent inconsistencies, restrict role names to characters with the POSIX portable file name character set. You cannot use the keywords **ALL** or **default** as a role name. Additionally, do not use any of the following characters within a role-name string:

- : (colon)
- " (quotation mark)
- # (pound sign)
- , (comma)
- = (equal sign)
- \ (backslash)
- / (forward slash)
- ? (question mark)
- ' (single quotation mark)

Note: After creating roles, run the **setkst** command to update the kernel database with the new roles.

Security

The **mkrole** command is a privileged command. To run the command successfully, assume a role that has the authorization **vios.security.role.create**.

Exit Status

0 for success.

Nonzero for command error.

Examples

1. To create the *ManageRoles* role and have the command automatically generate a role ID, type the following command:
`mkrole authorizations=vios.security.role ManageRoles`
2. To create the *ManageRoles* role and set the role ID to 127, type the following command:
`mkrole id=127 authorizations=vios.security.role ManageRoles`

mkrep command

Purpose

Create the Virtual Media Repository.

Syntax

`mkrep -sp ParentStoragePool -sizeSize`

Description

The **mkrep** command creates the virtual media repository in the specified parent storage pool. The virtual media repository is used to store virtual optical media which can be conceptually inserted into file-backed virtual optical devices. See the **mkvdev** command for detail on how to create file-backed virtual optical devices.

The **-size** flag specifies the minimum size the repository should be. When specifying the *Size* the following conventions must be used:

Size	Minimum file storage pool size
###M/m	###MB
###G/g	###GB

Flags

- size** *Size* Specifies the minimum size the repository should be.
- sp** *ParentStoragePool* Specifies the parent storage pool the repository should be created within. The parent storage pool must be a logical volume pool.

Examples

To create the virtual media repository within logical volume storage pool *client_data* with a size of at least 100 megabytes , type the following command:

```
mkrep -sp client_data -size 100m
```

mksp command

Purpose

Create a storage pool.

Syntax

Create a logical volume storage pool:

```
mksp [ -f ] StoragePool PhysicalVolume ...
```

Create a file storage pool:

```
mksp -fb StoragePool -sp ParentStoragePool-size Size [-mirror]
```

Description

The **mksp** command creates a new logical volume or file storage pool. Logical volume pools are used to store logical volume backing devices, file storage pools and the Virtual Media Repository. These pools are created using the physical volumes represented by the *PhysicalVolume* parameter.

If the system detects a description area from a volume group that is not varied on, it prompts you for confirmation in continuing with the command. The previous contents of the physical volume are lost, so you must be cautious when using the override function. By specifying the **-f** flag, you force the volume group to be created without sending a confirmation message.

File pools are used to store backing device files. A file pool is created within a logical volume pool, specified by the **-sp** ParentStorage Pool parameter.

The **-size** *Size* flag specifies the minimum size the pool should be. When specifying *Size* the following conventions must be used:

Size	Minimum file storage pool size
###M/m	###MB
###G/g	###GB

Flags

- f** Forces the storage pool to be created on the specified physical volume unless the physical volume is part of another storage pool or volume group in the Device Configuration Database or is an active volume group.

In addition, the physical volume that you specify cannot be assigned to a shared memory pool (to be used as a paging space device by a shared memory partition).
- fb** *StoragePool* Specifies the name of the file storage pool to be created. The name must be a unique system wide name, and can range from 1 to 15 characters.
- mirror** Activates mirroring for this file storage pool.
- size** *Size* Specifies the minimum size the file storage pool should be.
- sp** *ParentStoragePool* Specifies the parent storage pool the file pool should be created within. The parent storage pool must be a logical volume pool

Examples

1. To create a new logical volume storage pool from physical volumes `hdisk3` and `hdisk4` and with the name `client_data`, type:

```
mksp -f client_data hdisk3 hdisk4
```

The new storage pool is created with the name `client_data`.

2. To create a new file storage pool within logical volume storage pool `client_data` with a size of at least 100 MB and with the name `client2_data`, type:

```
mksp -fb client2_data -sp client_data -size 100m
```

The new storage pool is created with the name `client2_data`.

IVM `mksvcevent` command

Purpose

Creates a new serviceable event. This command is operable only in an Integrated Virtualization Manager environment.

Syntax

```
mksvcevent -d Description --reporting_mtms ReportingMTMS
```

Description

The `mksvcevent` command creates a serviceable event with the specified description. This event will show up in the list of serviceable events obtained by the `lssvcevents` command.

Flags

<code>-d</code> <i>Description</i>	The description or text of the event.
<code>-reporting_mtms</code> <i>ReportingMTMS</i>	The type-model*serial of the reporting system. This should be in the form <code>ttt-mmm*sssssss</code> , where <code>ttt</code> is the machine type, <code>mmm</code> is the model, and <code>sssssss</code> is the serial number of the managed system.

Exit Status

This command has a return code of zero on success.

Security

This command is not accessible by users with the `ViewOnly` role.

Examples

1. To create a serviceable event, type:

```
mksvcevent -d This is a test event entry -reporting_mtms 9111-520*XXXXXXX
```

Related Information

The `lssvcevents` command, and the `chsvcevent` command.

IVM `mksyscfg` command

Purpose

Creates a logical partition on the managed system. This command is operable only in an Integrated Virtualization Manager environment.

Syntax

To create a logical partition on the managed system

```
mksyscfg -r lpar { -f ConfigurationFile | -i ConfigurationData } [ -m ManagedSystem ]
```

Description

The `mksyscfg` command creates a logical partition on the managed system.

Flags

-r <i>ResourceType</i>	The type of resources to create:
-m <i>ManagedSystem</i>	lpar - Logical partition resources The name of the managed system. This attribute is optional because there is only one system to manage. The name might either be the user-defined name for the managed system, or be in the form <i>ttt-mmm*sssssss</i> , where <i>ttt</i> is the machine type, <i>mmm</i> is the model, and <i>sssssss</i> is the serial number of the managed system.
-f <i>ConfigurationFile</i>	The name of the file containing the configuration data needed to change the resources. The configuration data consists of attribute name/value pairs, which are in comma separated value (CSV) format. These attribute name/value pairs form a configuration record. A line feed marks the end of a configuration record. The file must contain one configuration record for each resource to be changed, and each configuration record must be for the same resource type. If the resource type is the managed system or the managed frame, then the file must contain only one configuration record.

The format of a configuration record is as follows:

```
attribute-name=value,attribute-name=value,...<LF>
```

Note that certain attributes accept a comma separated list of values, as follows:

```
"attribute-name=value,value,...",...<LF>
```

When a list of values is specified, the attribute name/value pair must be enclosed in double quotation marks. Depending on the shell being used, nested double quotation marks might need to be preceded by an escape character.

Required attributes for partitions

desired_mem

The assigned megabytes of memory for this partition

name Name of the partition to create.

max_mem

The maximum megabytes of memory for this partition

min_mem

The minimum megabytes of memory for this partition.

Optional attributes for partitions

allow_perf_collection

Permission for the partition to retrieve shared processor pool utilization information Valid values are:

- 0: do not allow authority
- 1: allow authority

auto_start

Valid values are:

- 0 - do not automatically start with system power on
- 1 - automatically start with system power on

boot_mode

Partition power on mode. Valid values are:

- norm - normal
- dd - diagnostic with default boot list
- ds - diagnostic with stored boot list
- of - Open Firmware OK prompt
- sms - System Management Services

desired_io_entitled_mem

The amount of I/O entitled memory for a shared memory partition. This is the portion of memory that is reserved for I/O mappings.

- auto (automatically manage)
- *Number of megabytes*

If the value is auto, the entitlement is calculated based on the virtual I/O configuration of the logical partition. If the virtual I/O configuration is changed, the entitlement is updated automatically. If auto is not used, no automatic adjustments are made. The default value is auto.

desired_proc_units

The assigned number of processing units for this partition

desired_procs

The assigned number of processors for this partition. In shared processing mode, this refers to virtual processors.

io_slots

Comma separated list of I/O slots for the partition. Each item in this list has the format: *drc_index/slot_io_pool_id/is_required*

The attribute names are not present in the list, just their values are present. For example, 21010003/none/1, specifies an I/O slot with a DRC index of 0x21010003 which is not assigned to an I/O pool, and it is a required slot.

Valid values for *is_required*:

- 0 - no
- 1 - yes

lhea_capabilities

Comma-separated list of logical Host Ethernet Adapter capabilities, with each capability having one of the following formats: *adapter-ID/capability* or *adapter-ID/5/ieq/nieq/qp/cq/mr* where *ieq* (interruptible event queues), *niesq* (non-interruptible event queues), *qp* (queue pairs), *cq* (completion queues), and *mr* (memory regions) each specify the resource amount in addition to the base minimum. Valid values are:

- 0 - minimum
- 1 - low
- 2 - medium
- 3 - high
- 4 - dedicated
- 5 - custom

lhea_logical_ports

Comma-separated list of Logical Host Ethernet Adapter (LHEA) logical ports, with each logical port having the following format:

*adapter-ID/port-group/physical-port-ID/
logical-port-ID/allowed -VLAN-IDs*

All four '/' characters must be present, but optional values may be omitted. Optional values are *allowed-VLAN-IDs*.

lpar_avail_priority

The priority of the partition to maintain its entitled processors. If a processor failure occurs, processing resources will be removed first from the lowest priority partition.

lpar_id Unique integer ID for the partition. If this attribute is not specified, the lowest available partition will be assigned.

lpar_proc_compat_mode

The requested compatibility mode. Use `lssyscfg -r sys -F lpar_proc_compat_modes` to retrieve a list of valid values.

max_procs

The maximum number of processors for this partition. In shared processing mode, this refers to virtual processors.

max_proc_units

The maximum number of processing units for this partition

max_virtual_slots

Maximum number of virtual I/O adapter slots.

Note: The system determines this value.

mem_mode

The partition memory mode.

- ded - dedicated memory
- shared - shared memory

If the memory mode is shared, then the logical partition cannot be assigned any physical I/O slots or host Ethernet adapter resources, the `proc_mode` attribute on the logical partition must be shared, and a memory pool must exist. If the memory mode is omitted, then it is assumed to be a dedicated memory mode.

Note: The Virtual I/O Server logical partition only supports the dedicated memory mode.

mem_weight

The shared memory weight of the logical partition that uses shared memory.

This attribute is used for determining the priority of logical partitions in a memory pool for distribution of memory. The weight must be a value from 0 - 255. If a value is not specified, the default is 128.

min_procs

The minimum number of processors for this partition. In shared processing mode, this refers to virtual processors.

min_proc_units

The minimum number of processing units for this partition

paging_device

The paging space device to use if a memory pool is used. A paging device is a block storage device that has been added to the memory pool and is not designated as a paging device for any other logical partition. This attribute is optional. If omitted, an appropriate paging device is selected automatically. If the `paging_device` value is a blank string, no device is assigned.

proc_mode

Valid values are:

- ded- dedicated processor mode
- shared- shared processor mode

profile_name

Name of the profile to create. This attribute is not required, but if specified, must be the same as the **name** attribute.

sharing_mode

Value values are:

- keep_idle_procs- never share processors
- share_idle_procs- share processors only when the partition is inactive
- share_idle_procs_always- share processors only when the partition is active
- share_idle_procs_active- always share processors
- cap- capped mode
- uncap- uncapped mode

uncap_weight

A weighted average of processing priority when in uncapped sharing mode. The smaller the value, the lower the weight. Valid values are: 0 - 255

virtual_eth_adapters

Comma-separated list of virtual Ethernet adapters, with each adapter having the following format:

```
slot_number/is_ieee/port_vlan_id/additional_vlan_ids/  
is_trunk/is_required
```

All 5 '/' characters must be present, but optional values might be omitted. Optional values are `additional_vlan_ids`, and `is_trunk`. Valid values for `is_ieee`, `is_trunk`, and `is_required`:

0 - no
1 - yes

For example, `4/0/2//0/0` specifies a virtual Ethernet adapter with a virtual slot number of 4, is not IEEE 802.1Q enabled, has a port virtual LAN ID of 2, no additional virtual LAN IDs, it is not a trunk adapter, and is not required.

virtual_fc_adapters

Comma-separated list of virtual fibre channel adapters. Each item in this list has the following format:

```
virtual_slot_num/adapter_type/remote_lpar_id/  
remote_lpar_name/remote_slot_num/wwpn_list/is_required
```

Required values: `remote_lpar_id`, `remote_lpar_name`, `adapter_type`, `virtual_slot_num`

Note: You can specify either `remote_lpar_id`, `remote_lpar_name`, or use them both, but at least one of the values is required.

Valid values for `adapter_type`:

- client
- server

Note: If you specify a value for adapter type, the Integrated Virtualization Manager (IVM) requires the adapter type to be a client.

Optional values: `wwpn_list`, `is_required`, `remote_slot_num`

When you add a virtual fibre channel adapter, `wwpn_list` can be left blank to allow IVM to automatically assign worldwide port names to the client adapter. If you leave `wwpn_list` blank and the virtual slot number specified for this adapter already contains a virtual fibre channel adapter, IVM will use the worldwide port names that are already assigned. New worldwide port names will only be generated if it is a new adapter. If you specify a `wwpn_list` value, there must be exactly two values. Each worldwide port name must be a 16-character hexadecimal value. These values are not case sensitive.

Valid values for `is_required`:

0 - no
1 - yes

A value of none or an empty string indicates that no virtual fibre channel adapters should be assigned.

Note: If the logical partition being changed is the Virtual I/O Server, you cannot change the current configuration. The IVM handles the client and server adapter as a pair; therefore, IVM handles changes automatically.

virtual_scsi_adapters

Comma-separated list of virtual SCSI adapters. Each item in this list has the format:

```
slot_num/adapter_type/remote_lpar_id/remote_lpar_name/  
remote_slot_num/is_required
```

The attribute names are not present in the list, just their values are present. If an attribute is optional and is not to be included, then no value is specified for that attribute. For example, `2/client//lpar2/3/0` specifies a virtual client SCSI adapter with a virtual slot number of 2, a server partition name of `lpar2`, a server slot number of 3, and is not required. The server partition ID was omitted.

Required values: `slot_num`, `adapter_type`, `remote_lpar_id`, `remote_lpar_name`

Note: You can specify `remote_lpar_id`, `remote_lpar_name`, or both, but at least one of the values is required.

Optional values: `is_required`, `remote_slot_num`

Note: IVM requires that virtual slot number 2 always contain a virtual SCSI adapter, so if you specify adapters in any other slots, the default adapter is still created in slot number 2. If you use the `chsyscfg` command with an empty list for the `virtual_scsi_adapters` attribute, all virtual SCSI adapters are removed, except the default adapter.

Valid values for `adapter_type`:

- `client`: client adapter
- `server`: server adapter, valid for Virtual I/O Server logical partitions only

Valid values for `is_required`:

- 0 - no
- 1 - yes

work_group_id

Valid values:

- `none`: do not participate in the workload management group
- `1`: participate in the workload management group

-i ConfigurationData

This option allows you to enter configuration data on the command line, instead of using a file. Data entered on the command line must follow the same format as data in a file, and must be enclosed in double quotation marks. When this option is used, only a single resource can be changed. The `-i` and the `-f` options are mutually exclusive.

Exit Status

This command has a return code of 0 on success.

Security

This command is not accessible by users with the `ViewOnly` role.

Examples

1. To create a partition named `lp3` with 128 megabytes, type:

```
mksyscfg -r lpar -i "name=lp3,lpar_env=aixlinux,min_mem=128,desired_mem=128,\  
max_mem=128"
```
2. To create a partition with 128 megabytes and a dedicated processor, type:

```
mksyscfg -r lpar -i "name=lp4,lpar_env=aixlinux,min_mem=128,desired_mem=128,\
max_mem=128,proc_mode=ded, sharing_mode=share_idle_procs,min_procs=1,\
desired_procs=1,max_procs=2"
```

3. To create a partition with 128 megabytes and 0.2 shared processing units, type:

```
mksyscfg -r lpar -i "name=lp2,lpar_env=aixlinux,min_mem=128,desired_mem=128,\
max_mem=128,proc_mode=shared, sharing_mode=uncap,min_procs=1,\
desired_procs=1,max_procs=2,min_proc_units=0.1,desired_proc_units=0.2,\
max_proc_units=2"
```

4. To create a partition with 128 megabytes and 0.2 shared processing units, and a virtual Ethernet adapter on VLAN 1, type:

```
mksyscfg -r lpar -i "name=lp2,lpar_env=aixlinux,min_mem=128,desired_mem=128,\
max_mem=128,proc_mode=shared, sharing_mode=uncap,min_procs=1,desired_procs=1,\
max_procs=2,min_proc_units=0.1,desired_proc_units=0.2, max_proc_units=2,\
virtual_eth_adapters=4/0/1//0/0"
```

5. To create a partition with 128 megabytes and 0.2 shared processing units, a virtual Ethernet adapter on VLAN 1, and HEA ports 3 and 4, type:

```
mksyscfg -r lpar -i "name=lp2,lpar_env=aixlinux,min_mem=128,desired_mem=128,\
max_mem=128,proc_mode=shared, sharing_mode=uncap,min_procs=1,desired_procs=1,\
max_procs=2,min_proc_units=0.1,desired_proc_units=0.2, max_proc_units=2,\
virtual_eth_adapters=4/0/1//0/0,\ \
"lhea_logical_ports=23000000/1/0/3/all,23000000/1/1/4/all",\
lhea_capabilities=23000000/0////"
```

6. To create a logical partition with 1 gigabyte, 2 virtual processors, and 2 virtual Ethernet adapters, type:

```
mksyscfg -r lpar -i 'name=lp2,lpar_env=aixlinux,min_mem=256,desired_mem=1024,\
max_mem=2048,proc_mode=shared,sharing_mode=uncap,min_procs=1,desired_procs=2,\
max_procs=2,min_proc_units=0.1,desired_proc_units=0.2,max_proc_units=2,\
"virtual_eth_adapters=4/0/1//0/0,""5/1/2/212,313/0/0""'
```

Related Information

The `lssyscfg` command, the `chsyscfg` command, and the `rmsyscfg` command.

mktcpip command

Purpose

Sets the required values for starting TCP/IP on a host.

Syntax

```
mktcpip -hostname HostName -inetaddr Address -interface Interface [ -start ] [ -netmask SubnetMask ] [ -cabletype CableType ] [ -gateway Gateway ] [ -nsrvaddr NameServerAddress -nsrvdomain Domain ]
```

To run stateless autoconfiguration of link-local addresses:

```
mktcpip [-interface Interface] -auto
```

To add a static IPv6 address:

```
mktcpip -interface Interface -inetaddr IPv6 address [-plen PrefixLen] [-gateway gateway]
```

Description

The `mktcpip` command sets the minimal values required for using TCP/IP on a host machine. The basic functions of the `mktcpip` command include:

- Setting the host name

- Setting the IP address of the interface
- Setting the domain name and IP address of the name server, if applicable
- Setting the subnetwork mask, if applicable
- Starting the specified TCP/IP daemons

If the shared storage pool cluster is present on the server, you cannot perform any action by using this command.

Flags

-auto	Enables automatic stateless configuration of link-local addresses. Link-local addresses can coexist with other preexisting IPv6 addresses. Therefore, if an address is assigned by using the -auto flag, it does not replace existing IPv6 addresses.
-cabletype <i>CableType</i>	Specifies the cable size for Standard Ethernet or IEEE 802.3 Ethernet networks. Valid values for the <i>CableType</i> variable are <i>dix</i> for thick cable, <i>bnc</i> for thin cable, or <i>N/A</i> for not applicable. The -cabletype <i>CableType</i> flag should be used only for Standard Ethernet (<i>en</i>) and IEEE 802.3 Ethernet (<i>et</i>) interfaces. The default is <i>N/A</i> .
-gateway <i>Gateway</i>	Sets the gateway address for a static route. Specify the address in dotted decimal notation.
-hostname <i>Hostname</i>	Sets the name of the host. If using a domain naming system, the domain and any subdomains must be specified. The following is the standard format for setting the host name: hostname The standard format for setting the host name in a domain naming system is as follows: : hostname.subdomain.subdomain.rootdomain
-inetaddr <i>Address</i>	Sets the IP address of the host. Specify the address in dotted decimal notation. Each network interface on the host should have a unique IP address. The standard format for setting the IP address as follows: : 127.10.31.2
-interface <i>Interface</i>	Specifies a particular network interface, for example: en1
-netmask <i>SubnetMask</i>	Specifies the mask that the gateway should use in determining the appropriate subnetwork for routing. The subnet mask is a set of 4 bytes, as in the IP address. The subnet mask consists of high bits (1 s) corresponding to the bit positions of the network and subnetwork address, and low bits (0s) corresponding to the bit positions of the host address.
-nsrvaddr <i>NameserverAddress</i>	Specifies the IP address of the name server that the host uses for name resolution, if applicable. The address should be entered in dotted decimal notation, as follows: 127.1.0.1
-nsrvdomain <i>Domain</i>	Specifies the domain name of the name server that the host should use for name resolution, if any. The domain name should be in the format as follows: : subdomain.subdomain.rootdomain
-plen <i>prefixLen</i>	Specifies the prefix length of the IPv6 interface.
-start	Starts the TCP/IP daemons.

Exit Status

See “Exit status for Virtual I/O Server commands” on page 3.

Examples

1. To set the required values for starting TCP/IP, type:

```
mktcpip -hostname fred.austin.century.com -inetaddr 192.9.200.4 -interface en0 \  
-nsrvaddr 192.9.200.1 -nsrvdomain austin.century.com -start
```

2. To set the required values for starting TCP/IP by using an IPv6 address, type:

```
mktcpip -interface en0 -auto
```

Note: If an IPv6 address exists on the specified interface, this command format does not replace it. You can use the `-auto` flag to allow multiple IPv6 addresses to coexist on the interface.

3. To set the required values for starting TCP/IP, by using an IPv6 address to replace any pre-existing IPv6 addresses, type:

```
mktcpip -interface en0 -hostname host -inetaddr ipv6_address
```

Note: This method replaces any IPv6 addresses that were assigned to the interface.

Related Information

The `hostname` command, the `startnetsh` command, the `stopnetsh` command, the `cfglnagg` command, the `netstat` command, the `entstat` command, the `cfgnamesrv` command, the `hostmap` command, the `traceroute` command, the `ping` command, the `optimizenet` command.

mkuser command

Purpose

Creates a new user account.

Syntax

```
mkuser [-ldap] [-de | -sr ] [-attr Attribute=Value [ Attribute=Value... ] ] Name
```

Description

The `mkuser` command creates a new user account. Upon completion of creating a the new account you are prompted to set the new account password. User accounts are created with the `-attr pgrp=view` are designated as read only. These users do not have the authority to change the system configuration and do not have write permission to their home directories.

Flags

<code>-attr Attribute=Value</code>	Identifies the attribute to set, as well as the new value for the attribute. The <code>Attribute=Value</code> parameter can use one attribute value pair or multiple attribute value pairs for one <code>-attr</code> flag.
<code>-de</code>	For a complete list of supported attributes, see “chuser command” on page 65. Creates a Development Engineer (DE) user account. This type of account allows developers to log into the Virtual I/O Server and debug problems.
<code>-ldap</code>	Identifies the user as an LDAP user account. LDAP user accounts are authenticated through the LDAP load module.
<code>-sr</code>	Creates a service representative (SR) user account. This type of account enables a service representative to run commands required to service the system without being logged in as root. This includes the following command types: <ul style="list-style-type: none">• Run diagnostics, including service aids (for example, hot plug tasks, certify, format, and so forth).• Run all commands that can be run by a group system.• Configure and unconfigure devices that are not busy.• Use the service aid to update system microcode.• Perform the shut down and reboot operations. The recommended SR login user name is <code>qserv</code> .

fsize Defines the soft limit for the largest file a user's process can create or extend. The *Value* parameter is an integer representing the number of 512-byte blocks. To make files greater than 2G, specify -1 or unlimited. The minimum value for this attribute is 8192.

Exit Status

See "Exit status for Virtual I/O Server commands" on page 3.

Security

This command can only be run by the prime administrator (padmin) user.

Examples

1. To create the **davis** user account with default values, type the following command:
`mkuser davis`
2. To create the **davis** user account and set the **maxage** attribute to a value of 52, type the following command:
`mkuser -attr maxage=52 davis`
3. To create a user with read only authority, type the following command:
`mkuser -attr pgrp=view View1`

Related Information

The **chuser** command, the **lsuser** command, the **rmuser** command, and the **passwd** command.

mkvdev command

Purpose

Adds a virtual device to the system.

Syntax

To create a virtual target device:

```
mkvdev [ -f ] { -vdev TargetDevice | -dplc TDPhysicalLocatonCode } { -vadapter VirtualServerAdapter | -aplc VSAPhysicalLocationCode } [ -dev DeviceName ]
```

To create a virtual target device for a file backed virtual optical device:

```
mkvdev -fbo { -vadapter VirtualServerAdapter | -aplc VSAPhysicalLocationCode } [ -dev DeviceName ]
```

To create a virtual target device when the backing device is part of a Peer-to-Peer Remote Copy (PPRC) pair:

```
mkvdev [ -f ] { -vdev TargetDevice | -dplc TDPhysicalLocatonCode } { -vadapter VirtualServerAdapter | -aplc VSAPhysicalLocationCode } { -attr mirrored=true } [ -dev DeviceName ]
```

To create a shared Ethernet Adapter:

```
mkvdev -sea TargetDevice -vadapter VirtualEthernetAdapter... -default DefaultVirtualEthernetAdapter -defaultid SEADefaultPVID [ -attr Attribute=Value [ Attribute=Value... ] ]
```


Note: The shared Ethernet adapter attributes, *ha_mode* and *ctl_chan*, must be specified together. The command fails if only one attribute is specified. These attributes are used to create a shared Ethernet adapter in a failover configuration.

To create an Link Aggregation adapter:

```
mkvdev -lnagg TargetAdapter... [ -attr Attribute=Value [ Attribute=Value... ] ]
```

To create a VLAN Ethernet adapter:

```
mkvdev -vlan TargetAdapter -tagid TagID
```

Description

The **mkvdev** command creates a virtual device. The name of the virtual device is automatically generated and assigned unless the **-dev** *DeviceName* flag is specified, in which case *DeviceName* becomes the device name.

If the backing device is part of a PPRC pair, set the mirrored attribute to TRUE to create the virtual target device. This allows a PPRC secondary device to be exported to the client partition.

If the **-lnagg** flag is specified, a Link Aggregation or IEEE 802.3 Link Aggregation (automatic Link Aggregation) device is created. To create an IEEE 802.3 Link Aggregation set the mode attribute to 8023ad. If the **-sea** flag is specified, a Shared Ethernet Adapter is created. The *TargetDevice* may be a Link Aggregation adapter (note, however, that the *VirtualEthernetAdapter* may not be Link Aggregation adapters). The default virtual Ethernet adapter, *DefaultVirtualEthernetAapter*, must also be included as one of the virtual Ethernet adapters, *VirtualEthernetAdapter*.

The **-fbo** flag is used to create a virtual target device which maps the *VirtualServerAdapter* to a file backed virtual optical device. Virtual optical devices cannot be used until virtual media is loaded into the device. See the **loadopt** command for details.

The **-vlan** flag is used to create a VLAN device and the **-vdev** flag creates a virtual target device which maps the *VirtualServerAdapter* to the *TargetDevice*.

If the backing device that is specified by the **-vdev** or **-dplc** flags is already in use, an error is returned unless the **-f** flag is also specified. Also, the backing device that is specified cannot be a physical or logical volume that is assigned to a shared memory pool (to be used as a paging space device by a shared memory partition).

The **mkvdev** command also configures virtual optical and tape devices, where the **-vdev** or **-dplc** flag specifies the physical optical or tape device and the **-vadapter** or **-aplc** flag specifies the virtual SCSI adapter. If the specified optical or tape device is already assigned to a virtual SCSI adapter, an error is returned unless the **-f** flag is also specified. If the **-f** flag is specified, the optical or tape device will be removed from the virtual SCSI adapter it is currently assigned to before reassigning it to the new virtual SCSI adapter.

When an additional disk drive that has a lower *max_transfer_size* than the current setting is added, and it is set as a virtual target device (processing **mkvdev**), the client cannot recognize this device until the VIOS is rebooted and the *max_transfer_size* is reestablished to the new setting. Because the *max_transfer_size* cannot be dynamically changed, the *mkvdev* command checks the current *max_transfer_size* of the device being added. If it is lower, it posts a message stating that the VIOS must be rebooted before the client can see this device.

Attention: To protect the Configuration Database, the **mkvdev** command is not interruptible. Stopping this command before execution is complete could result in a corrupted database.

Flags

-aplc <i>VSAPhysicalLocationCode</i>	Specifies the virtual SCSI adapter using the physical location code
-attr <i>Attribute=Value</i>	Specifies the device attribute value pairs to be used instead of the defaults. The <i>Attribute=Value</i> variable can be used to specify one attribute value pair or multiple attribute value pairs for one -attr flag. If you use an -attr flag with multiple attribute value pairs, the list of pairs must be enclosed in quotation marks with a blank space between the pairs. For example, entering -attr Attribute=Value lists one attribute value pair per flag, while entering -attr 'Attribute1=Value1 Attribute2=Value2' lists more than one attribute value pair.
-default <i>DefaultVirtualEthernetAdapter</i>	Default virtual adapter to use for non-VLAN-tagged packets. This flag maps to the SEA device attribute <i>pvid_adapter</i> .
-defaultid <i>SEADefaultPVID</i>	The SEADefaultPVID is the VID used for untagged frames. All untagged packets are assigned the SEADefaultPVID value. When a tagged frame is received by a port, the tag is used. Otherwise if the frame is untagged, the value contained in the PVID is considered as a tag. This flag maps to the SEA device attribute <i>pvid</i> .
-dev <i>DeviceName</i>	By using the -dev flag, you can specify the name you want the device to be known by. If you do not use the -dev flag, a name will be automatically generated and assigned. Not all devices support user-supplied names.
-dplc <i>TDPhysicalLocatonCode</i>	Specifies the physical device using the physical location code
-f	Force the physical volume specified to be used as a backing device even if its already associated with a virtual SCSI adapter. If the specified backing device is an optical device, -f forces the optical device will be removed from the virtual SCSI adapter it is currently assigned to before reassigning it to the new virtual SCSI adapter.
-fbo	Create a virtual optical device.
-lnagg <i>TargetAdapter...</i>	Creates a Link Aggregation device.
-sea <i>TargetDevice</i>	Creates a Shared Ethernet Adapter which maps <i>VirtualEthernetAdapter</i> to the adapter <i>TargetDevice</i> . <i>TargetDevice</i> can be a physical adapter or a Link Aggregation adapter.
-tagid <i>TagID</i>	Specifies the VLAN tag ID.
-vadapter <i>VirtualEthernetAdapter</i> or <i>VirtualServerAdapter</i>	Specifies the virtual server adapter or virtual Ethernet adapter the new device will be mapped to. If multiple virtual Ethernet adapters are specified, separate the adapter names with commas and no spaces.
-vdev <i>TargetDevice</i>	Creates a virtual device mapped to the physical/logical device <i>TargetDevice</i> and the virtual server adapter <i>VirtualServerAdapter</i> . The <i>TargetDevice</i> can be either a physical volume, logical volume, tape, or optical device. Physical volumes assigned to volume groups cannot be used as target devices.
-vlan <i>TargetAdapter</i>	Creates a Virtual Local Area Network device.

Exit Status

13	Specified physical or logical volume is not a valid.
21	Device is already in use. Use the -f flag to force assignment.
22	Logical volumes cannot be assigned as backing devices more than once.

Examples

1. To create a virtual target device that maps the logical volume **lv20** as a virtual disk for a client partition hosted by the **vhost0** virtual server adapter, type the following command:

```
mkvdev -vdev lv20 -vadapter vhost0
```

The system displays a message similar to the following:

```
vtscsi0 available
```

2. To create a virtual target device that maps the physical volume **hdisk6** as a virtual disk for a client partition served by the **vhost2** virtual server adapter, type the following command:

```
mkvdev -vdev hdisk6 -vadapter vhost2
```

The system displays a message similar to the following:

```
vtscsi1 available
```

3. To create a virtual target device that maps physical tape device **rmt0** as a virtual tape device for a client partition served by the **vhost2** virtual server adapter, type the following command:

```
mkvdev -vdev rmt0 -vadapter vhost2
```

The system displays a message similar to the following:

```
vttape0 available
```

4. To create a shared Ethernet adapter that maps the physical Ethernet adapter **ent4** as a virtual Ethernet adapter for the client partitions served by the virtual Ethernet adapters **ent6**, **ent7**, and **ent9**, using **ent6** as the default adapter and **8** as the default ID, type the following command:

```
mkvdev -sea ent4 -vadapter ent6,ent7,ent9 -default ent6 -defaultid 8
```

The system displays a message similar to the following:

```
ent10 available
```

5. To create a shared Ethernet adapter in failover configuration with the control channel adapter **ent5**, create the shared Ethernet adapter as shown in example 4, but specify additional attributes *ha_mode* and *ctl_chan* with *-attr*. For example:

```
mkvdev -sea ent4 -vadapter ent6 -default ent6 -defaultid=1 -attr ha_mode=auto ctl_chan=ent5
```

6. To create an automatic Link Aggregation with primary adapters **ent4** and **ent5** and backup adapter **ent6**, type the following command:

```
mkvdev -lnagg ent4,ent5 -attr backup_adapter=ent6 mode=8023ad
```

The system displays a message similar to the following:

```
ent10 available
```

7. To create a virtual target device that maps the physical volume **hdisk6**, which is part of a PPRC pair as a virtual disk for a client partition served by the **vhost2** virtual server adapter, type the following command:

```
mkvdev -vdev hdisk6 -vadapter vhost2 -attr mirrored=true
```

Related Information

The **cfgdev** command, the **chdev** command, the **chpath** command, the **lsdev** command, the **lsmapi** command, and the **rmdev** command.

mkvfilt command

Purpose

Activates the virtual LAN-crossing filter rules defined by the **genvfilt** command.

Syntax

```
mkvfilt -u
```

Description

The **mkvfilt** command activates the virtual LAN-crossing filter rules defined by the **genvfilt** command.

Flags

-u Activates the filter rules in the filter rule table.

Exit Status

This command returns the following exit values:

0 Successful completion.
>0 An error occurred.

Examples

1. To activate the filter rules in the kernel, type the command as follows:

```
mkvfilt -u
```

Related Information

The **chvfilt** command, the **genvfilt** command, the **lsvfilt** command, the **rmvfilt** command, and the **vlantfw** command.

mkvg command

Purpose

Creates a volume group.

Syntax

```
mkvg [ -f ][ -vg VolumeGroup ] PhysicalVolume ...
```

Description

The **mkvg** command creates a new volume group using the physical volumes represented by the *PhysicalVolume* parameter. After creating the volume group, the **mkvg** command automatically activates the new volume group using the **activatevg** command.

Note:

1. The physical volume is checked to verify that it is not already in another volume group. If the system believes the physical volume belongs to a volume group that is active, it exits. But if the system detects a description area from a volume group that is not active, it prompts the user for confirmation in continuing with the command. The previous contents of the physical volume are lost, so the user must be cautious when using the override function.
2. This command will fail to add a disk to the volume group if the disk indicates that it is managed by a third party volume manager.
3. The physical volume that you specify cannot be assigned to a shared memory pool (to be used as a paging space device by a shared memory partition).

Flags

-f Forces the volume group to be created on the specified physical volume unless the physical volume is part of another volume group in the Device Configuration Database or a volume group that is active.

-vg *VolumeGroup*

Specifies the volume group name rather than having the name generated automatically. Volume group names must be unique system wide and can range from 1 to 15 characters. The name cannot begin with a prefix already defined in the PdDv class in the Device Configuration database for other devices. The volume group name created is sent to standard output.

The volume group name can only contain the following characters: "A" through "Z," "a" through "z," "0" through "9," or "_" (the underscore), "-" (the minus sign), or "." (the period). All other characters are considered invalid.

Exit Status

See "Exit status for Virtual I/O Server commands" on page 3.

Examples

1. To create a volume group that contains physical disks **hdisk3**, **hdisk5**, and **hdisk6**, type:

```
mkvg hdisk3 hdisk5 hdisk6
```

The volume group is created with an automatically generated name, which is displayed.

2. To create the volume group **newvg** with one physical partition, type:

```
mkvg -vg newvg hdisk1
```

Related Information

The **lsvg** command, the **chvg** command, the **extendvg** command, the **reducevg** command, the **mirrorios** command, the **unmirrorios** command, the **activatevg** command, the **deactivatevg** command, the **importvg** command, the **exportvg** command, and the **syncvg** command.

mkvlog command

Purpose

Creates a virtual log or virtual log device.

Syntax

To create a virtual log without attaching it to a virtual SCSI (VSCSI) host adapter:

```
mkvlog -name LogName -client ClientName [-sp StoragePool] [-f UUID] [-state VirtualLogState] [-lf FileCount] [-lfs FileSize] [-sf FileCount] [-sfs FileSize]
```

To attach an existing virtual log to a VSCSI host adapter:

```
mkvlog -uuid UUID {-vadapter Adapter | -aplc PhysicalLocationCode }-dev DeviceName]
```

To create a virtual log and attach it to a VSCSI host adapter:

```
mkvlog -name LogName [-client ClientName] [-sp StoragePool] {-vadapter Adapter | -aplc PhysicalLocationCode }-f UUID] [-dev DeviceName] [-lf FileCount] [-lfs FileSize] [-sf FileCount] [-sfs FileSize]
```

Description

By using the **mkvlog** command, you can create virtual logs and also attach them to specified VSCSI host adapters that are ready to use. You can specify the properties of the new virtual log by using the command parameters. If you do not specify the command parameters, the repository default values are used.

If you do not specify the client name by using the **-client** flag when you create a virtual log and attach it to the VSCSI host adapter, an attempt is made to communicate with the attached logical partition to produce a value for the client name. However, this attempt cannot be made if the client logical partition is not running or the operating system on the client logical partition does not support this feature. In such cases, you must specify the **-client** flag when you create the virtual log.

When you create a virtual log and specify the connecting VSCSI host adapter by using the **-vadapter** or **-apl** flags, the **-state** flag of the new virtual log is enabled and it is not valid to force a value by using the **-state** flag.

Flags

-apl	Specifies the adapter to which the virtual log is attached.
-client	Specifies the client name of the new virtual log and can be 1 - 96 characters.
-dev	Specifies the name of the new virtual log device.
-f	Forces the UUID of the new virtual log to the specified value.
-lf	Specifies the maximum number of log files and can be 1 - 1000.
-lfs	Specifies the maximum size of each log file and can be specified in bytes, or suffixed with K, M, or G.
-name	Specifies the log name of the new virtual log and can be 1 - 12 characters.
-sf	Specifies the maximum number of state files and can be 1 - 1000.
-sfs	Specifies the maximum size of each state file and can be specified in bytes, or suffixed with K, M, or G.
-sp	Specifies the name of the shared storage pool to use.
-state	Sets the virtual log to one of the specified states as follows: <ul style="list-style-type: none"> • Enabled • Disabled • Migrated
-uuid	Specifies the UUID of an existing virtual log to create a device for.
-vadapter	Specifies the adapter to which the virtual log is attached.

Exit Status

Table 6. Command specific return codes

Return code	Description
0	All files were written successfully.
>0	An error occurred.

Examples

1. To create a new virtual log with the log name `syslog` and the client name `lpar-01`, type the command as follows:

```
mkvlog -name syslog -client lpar-01
```

The system displays the output as follows:

```
Virtual log 00000000000000005b3f6b7cfcec4c67 created
```

2. To attach the virtual log with a UUID `00000000000000005b3f6b7cfcec4c67` to the VSCSI host adapter `vhost0`, type the command as follows:

```
mkvlog -uuid 00000000000000005b3f6b7cfcec4c67 -vadapter vhost0
```

The system displays the output as follows:

```
vtlog0 Available
```

- To create a new virtual log with the log name `audit` and attach it to the VSCSI host adapter `vhost1`, type the command as follows:

```
mkvlog -name audit -vadapter vhost1
```

The system displays the output as follows:

```
Virtual log 0000000000000000d96e956aa842d5f4 created
vtlog0 Available
```

Related Information

The `chvlog` command, the `chvlrepo` command, the `lsvlog` command, the `lsvlrepo` command, and the `rmvlog` command.

mkvopt command

Purpose

Create a virtual optical media disk in the Virtual Media Repository

Syntax

```
mkvopt -name FileName {-size Size | -dev OptDevice | -file SourceFile} [-ro ]
```

Description

The `mkvopt` command creates a new virtual optical disk in the Virtual Media Repository. If the `-size` flag is specified, the new optical disk is initialized to all zeros of the given `Size`. If the `-dev` flag is specified, the given device, `OptDevice`, must be an optical device with media loaded. The contents of the loaded media within the device, will be used to create the optical media. If the `-file` flag is specified, `SourceFile` is copied into the repository. By default, the virtual optical disk is created as DVD-RAM media. If the `-ro` flag is specified, the disk will be create as DVD-ROM media.

When specifying `Size` the following conventions must be used:

Size	Minimum file size
<i>n</i> M/m	<i>n</i> MB
<i>n</i> G/g	<i>n</i> GB

Flags

<code>-dev</code> <i>OptDevice</i>	Specifies a physical optical device with media loaded.
<code>-file</code> <i>SourceFile</i>	Specifies the name of an existing file to by copied into the repository.
<code>-name</code> <i>FileName</i>	Specifies the name of the new virtual optical media file.
<code>-ro</code>	Create the new disk as read-only (DVD-ROM).
<code>-size</code> <i>Size</i>	Specifies how big to make the new virtual optical media.

Examples

- To create a read-only virtual optical disk with the name `pressData` from existing optical media in device `cd0`, type the following command:

```
mkvopt -name pressData -dev cd0 -ro
```
- To create a new virtual optical disk with the name `blankDVD` with 1 Gigabyte storage capacity, type the following command:

```
mkvopt -name blankDVD -size 1g
```

mkvt command

Purpose

Create a virtual terminal connection to a partition.

Syntax

```
mkvt { -id lparID }
```

Description

The **mkvt** command opens a virtual terminal connection to the target partition. You can terminate the virtual terminal connection in one of the following ways:

- The virtual terminal contains an escape sequence that allows you to break out of the command. The escape sequence is `<cr>~.`, or more explicitly: the Enter key, the tilde (~), and the period (.).
- You can use the **rmvt** command to force the session to be closed.

A partition can only have one open virtual terminal session.

This command requires additional HMC configuration if used in an HMC environment.

Flags

-id *lparID* The ID of the partition for which to open the virtual terminal session.

Exit Status

27	Unexpected error
28	Virtual terminal is already connected
29	Virtual terminal device is not found
30	Permission denied
31	Specified device does not exist

Security

This command is not accessible by users with the ViewOnly role.

Examples

1. Create a virtual terminal connection to the partition with ID 3:

```
mkvt -id 3
```

Related Information

The **rmvt** command.

motd command

Purpose

Displays or modifies the partition's message of the day file.

Syntax

motd

```
motd { -append | -overwrite } { -file Filename | "Message of the day string" }
```

Description

The **motd** command writes or appends to the partitions message of the day file. The new message can be specified on the command line or in a file with the **-file** flag. If no flags are specified the current message of the day is displayed.

Flags

-append	Appends the specified message to the current message of the day
-file <i>FileName</i>	Replaces the current message of the day with the contents of <i>FileName</i>
-overwrite	Replaces the current message of the day with the specified message

mount command

Purpose

Makes a file system available for use.

Syntax

```
mount [ [ Node:Directory ] Directory ]
```

```
mount -cd DeviceDirectory
```

Description

The **mount** command instructs the operating system to make a file system available for use at a specified location (the mount point). The **mount** command mounts a file system expressed as a directory using the *Node:Directory* parameter on the directory specified by the *Directory* parameter. After the **mount** command has finished, the directory specified becomes the root directory of the newly mounted file system.

If you enter the mount command without flags, the command displays the following information for the mounted file systems:

- the node (if the mount is remote)
- the object mounted
- the mount point
- the virtual-file-system type
- the time mounted
- any mount options

The **/mnt** directory can be used as a local mount point, or you can create a directory using the **mkdir** command. Any directories created with the **mkdir** command must be a sub-directory of your home directory.

Flags

-cd	Specifies the cd device name on which to mount.
------------	---

Exit Status

See “Exit status for Virtual I/O Server commands” on page 3.

Examples

1. To list the mounted file systems, type:

```
mount
```

This command produces output similar to the following:

node	mounted	mounted	vfs	date	options	over
----	-----	-----	---	-----	-----	-----
	/dev/hd0	/	jfs	Dec 17 08:04	rw, log	=/dev/hd8
	/dev/hd3	/tmp	jfs	Dec 17 08:04	rw, log	=/dev/hd8
	/dev/hd1	/home	jfs	Dec 17 08:06	rw, log	=/dev/hd8
	/dev/hd2	/usr	jfs	Dec 17 08:06	rw, log	=/dev/hd8
sue	/home/local/src	/usr/code	nfs	Dec 17 08:06	ro, log	=/dev/hd8

For each file system, the **mount** command lists the node name, the device name, the name under which it is mounted, the virtual-file-system type, the date and time it was mounted, and its options.

2. To mount the remote directory on to a local directory, enter:

```
mount testsys3:/test /mnt
```

This command mounts the `/test` directory located on **testsys3** onto the local `/mnt` directory.

Related Information

The **backupios** command.

netstat command

Purpose

Shows network status.

Syntax

To display active sockets for each protocol or routing table information:

```
netstat [ -num ] [ -routtable ] [ -routinfo ] [ -state ] [ -protocol Protocol ] [ Interval ]
```

To display the contents of a network data structure:

```
netstat [ -stats | -cdlistats ] [ -protocol protocol ] [ Interval ]
```

To display the address resolution protocol:

```
netstat -arp
```

To clear all statistics:

```
netstat -clear
```

To display network sockets:

```
netstat -socket
```

Description

The **netstat** command symbolically displays the contents of various network-related data structures for active connections.

Flags

-arp	Displays address resolution interfaces.
-cdlistats	Shows statistics for CDLI-based communications adapters.
-clear	Clears all statistics.
-num	Shows network addresses as numbers. When this flag is not specified, the netstat command interprets addresses where possible and displays them symbolically. This flag can be used with any of the display formats.
-protocol <i>protocol</i>	Shows statistics about the value specified for the <i>protocol</i> variable, which is either a well-known name for a protocol or an alias for it. A null response means that there are no numbers to report. The program report of the value specified for this variable is unknown if there is no statistics routine for it.
-routinfo	Shows the routing tables, including the user-configured and current costs of each route.
-routtable	Shows the routing tables. When used with the -stats flag, the -routtable flag shows routing statistics. See Routing Table Display.
-socket	Specifies that network sockets are to be displayed.
-state	Shows the state of all configured interfaces.

The interface display format provides a table of cumulative statistics for the following items:

- Errors
- Collisions
Note: The collision count for Ethernet interfaces is not shown.
- Packets transferred

The interface information that is displayed also provides the interface name, number, and address, as well as the maximum transmission units (MTUs).

-stats	Shows statistics for each protocol.
<i>Interval</i>	Continuously displays information, in seconds, regarding packet traffic on the configured network interfaces.

Default Display

The default display for active sockets shows the following items:

- Local and remote addresses
- Send and receive queue sizes (in bytes)
- Protocol
- Internal state of the protocol

Internet address formats are of the form *host.port* or *network.port* if a socket's address specifies a network but no specific host address. If the address can be resolved to a symbolic host name, the host address, as well as network addresses, are displayed symbolically.

NS addresses are 12-bytes in length, consisting of a 4-byte network number, a 6-byte host number and a 2-byte port number, all stored in network standard format. For VAX architecture, the word and byte are reversed.

If a symbolic name for a host is not known or if you specified the **-num** flag, the address is printed numerically, according to the address family. Unspecified addresses and ports appear as an * (asterisk).

Interface Display

The interface display format provides a table of cumulative statistics for the following items:

- Errors
- Collisions

Note: The collision count for Ethernet interfaces is not applicable.

- Packets transferred

The interface display also provides the interface name, number, and address as well as the maximum transmission units (MTUs).

Routing table display

The routing table display indicates the available routes and their statuses. Each route consists of a destination host or network and a gateway to use in forwarding packets.

A route is given in the format A.B.C.D/XX, which presents two pieces of information. A.B.C.D indicates the destination address and XX indicates the netmask associated with the route. The netmask is represented by the number of bits that are set. For example, the route 9.3.252.192/26 has a netmask of 255.255.255.192, which has 26 bits set.

The routing table contains the following fields:

Flags

The flags field in the routing table shows the state of the route:

A	An Active Dead Gateway Detection is enabled on the route
U	Up
H	The route is to a host rather than to a network
G	The route is to a gateway
D	The route was created dynamically by a redirect
M	The route has been modified by a redirect
L	The link-level address is present in the route entry
c	Access to this route creates a cloned route
W	The route is a cloned route
1	Protocol-specific routing flag #1
2	Protocol-specific routing flag #2
3	Protocol-specific routing flag #3
b	The route represents a broadcast address
e	Has a binding cache entry
l	The route represents a local address
m	The route represents a multicast address
P	Pinned route
R	Host or net unreachable
S	Manually added
u	Route usable
s	The group routing stopsearch option is enabled on the route

Gateway Refs

Direct routes are created for each interface attached to the local host.

Shows the address of the outgoing interface.

Use

Shows the current number of active uses for the route. Connection-oriented protocols hold on to a single route for the duration of a connection, while connectionless protocols obtain a route while sending to the same destination.

PMTU

Provides a count of the number of packets sent using that route.

Interface

Lists the Path Maximum Transfer Unit (PMTU).

Exp

Indicates the network interfaces utilized for the route.

Groups

Displays the time (in minutes) remaining before the route expires.

Netmasks

Provides a list of group IDs associated with that route.

Route Tree for Protocol Family

Lists the netmasks applied on the system.

Specifies the active address families for existing routes. Values for this field are as follows:

1	Specifies the UNIX address family
2	Specifies the Internet address family (for example, TCP and UDP)
3	Specifies the Xerox Network System (XNS) address family

When a value is specified for the *Interval* variable, the **netstat** command displays a running count of statistics related to network interfaces. This display contains two columns: a column for the primary interface (the first interface found during autoconfiguration) and a column summarizing information for all interfaces. The first line contains a summary of statistics accumulated since the system was last restarted. The subsequent lines of output show values accumulated over intervals of the specified length.

Exit Status

See “Exit status for Virtual I/O Server commands” on page 3.

Examples

1. To display routing table information for an Internet interface, type:

```
netstat -routtable
```

This produces the output similar to the following:

```
Routing tables
Destination      Gateway          Flags  Refs    Use  If  PMTU Exp Groups

Route tree for Protocol Family 2 (Internet):
default          129.3.141.1     UGc    0        0  en0  -  -
129.33.140/23    127.0.0.1       U       6        53  en0  -  -
129.33.41.2      localhost       UGHS    6       115  lo0  -  -
129.45.41.2      129.3.41.1     UGHW    1       602  en0  1500 -
dcefs100         129.31.41.1    UGHW    1         2  en0  -  -
192.100.61       localhost       U       7     14446  lo0  -  -
```

```
Route tree for Protocol Family 24 (Internet v6):
::1              ::1             UH      0         0  lo0  16896 -
```

2. To display interface information for an Internet interface, type:

```
netstat -state
```

This produces the output similar to the following:

```
Name  Mtu  Network  Address          Ipkts  Ierrs  Opkts  Oerrs  Coll
en0   1500  link#2   0.5.20.4.0.4e   874986  0      22494  0      0
en0   1500  90.34.14 hostname         874986  0      22494  0      0
lo0   16896 link#1                    14581  0      14590  0      0
lo0   16896 129      localhost        14581  0      14590  0      0
lo0   16896 ::1                    14581  0      14590  0      0
```

3. To display network sockets:

```
netstat -socket
```

Related Information

The **mktcpip** command, the **hostname** command, the **startnetsvc** command, the **stopnetsvc** command, the **cfglnagg** command, the **entstat** command, the **cfgnamesrv** command, the **hostmap** command, the **traceroute** command, the **ping** command, and the **optimizenet** command.

oem_platform_level command

Purpose

Returns the operating system level of the OEM installation and setup environment.

Syntax

```
oem_platform_level
```

Description

The **oem_platform_level** command displays the name and version of the underlying Virtual I/O Server operating system.

This command can only be executed by the prime administrator.

Exit Status

See “Exit status for Virtual I/O Server commands” on page 3

Examples

To get the operating system level of the OEM installation and setup environment, run the following command:

```
oem_platform_level
```

Related Information

The `lssw` command, the `ioslevel` command, the `remote_management` command, the `oem_setup_env` command, and the `updateios` command.

oem_setup_env command

Purpose

Initiates the OEM installation and setup environment.

Syntax

```
oem_setup_env
```

Description

The `oem_setup_env` command places the user into the OEM software installation and setup environment. In this environment, the user can install and setup OEM software by following the installation instructions provided with each software package. After the software is installed, the user will need to create a link in the `/usr/ios/oem/` directory to any new commands that will run from the Virtual I/O Server command line. After these links have been created, the commands will be accessible by all Virtual I/O Server users. Note however that these commands will not run with root authority.

After the software has been installed,, typing `exit` will return the user to the Virtual I/O Server prompt.

Only the prime administrator can execute this command.

Note: The `oem_setup_env` command will place the `padmin` user in a non-restricted UNIX root shell with a home directory in the `/home/padmin` directory. The user can then run any command available to the root user. This is not a supported Virtual I/O Server administration method. The purpose of this command is to allow installation of vendor software, such as device drivers.

Exit Status

See “Exit status for Virtual I/O Server commands” on page 3

Examples

To initiate the OEM setup and installation environment, type the following:

```
oem_setup_env
```

Related Information

The `lssw` command, the `ioslevel` command, the `remote_management` command, the `updateios` command, and the `oem_platform_level` command.

optimizenet command

Purpose

Manages network tuning parameters.

Syntax

```
optimizenet [ -reboot | -perm ] { -set Tunable[=NewValue] | -default Tunable }
```

```
optimizenet -list [ Tunable ]
```

```
optimizenet -h [ Tunable ]
```

Description

The **optimizenet** command is used to configure network tuning parameters. The **optimizenet** command sets or displays current or next boot values for network 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 **-set** flag performs both actions. It can either display the value of a parameter or set a new value for a parameter.

If the **-list** flag is specified without any *Tunables*, only *Tunables* modifiable by this command will be displayed.

Flags

- default** *Tunable* Resets *Tunable* to its default value. If the *Tunable* needs to be changed (not currently set to its default value) and it is of type Reboot, it will not be changed; a warning displays instead.
- h** *Tunable* Displays help about *Tunable* parameter if one is specified.
- list** [**Tunable**] Lists the characteristics of one or all *Tunables*, one per line, using the following format:

NAME	CUR	DEF	BOOT	MIN	MAX	UNIT	TYPE
DEPENDENCIES							

General Network Parameters							

sockthresh	85	85	85	0	100	%_of_thewall	D

fasttimo	200	200	200	50	200	millisecond	D

inet_stack_size	16	16	16	1		kbyte	R

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 (Dynamic), S (Static), R (Reboot), B (Bosboot), M (Mount), I (Incremental), and C (Connect)

DEPENDENCIES

List of dependent tunable parameters, one per line

-perm Makes changes apply to both current and reboot values when used in combination with **-set** or **-default**. These combinations cannot be used on Reboot type parameters because their current value can't be changed.

When used with **-set** without specifying a new value, values displays only if the current and next boot values for a parameter are the same. Otherwise NONE displays as the value.

-reboot Makes changes apply to reboot values when used in combination with **-set** or **-default**. When used with **-set** without specifying a new value, next boot values for *tunables* display instead of the current values.

-set Tunable
[=*NewValue*] Displays the value or sets the Tunable to *NewValue*. If a tunable needs to be changed (the specified value is different than current value), and is of type **Reboot** it will not be changed but a warning displays instead.

When **-reboot** is used in combination without a new value, the **nextboot** value for Tunable is displayed. When **-perm** is used in combination without a new value, a value displays only if the current and next boot values for tunable are the same Otherwise NONE displays as the value.

Network tunable parameters

arptab_bsiz	<p>Purpose: Specifies Address Resolution Protocol (ARP) table bucket size.</p> <p>Values: Default: 7 Range: 1 to MAXSHORT Type: Reboot</p> <p>Diagnosis netstat -protocol arp will show the number of ARP packets sent and the number of ARP entries purged from the ARP table. If large number of entries are being purged, the ARP table size should be increased.</p>
arptab_nb	<p>Purpose: Specifies the number of ARP table buckets.</p> <p>Values: Default: 73 Range: 1 to MAXSHORT Type: Reboot</p> <p>Diagnosis: netstat -protocol arp will show the number of ARP packets sent and the number of ARP entries purged from the ARP table. If large number of entries are being purged, the ARP table size should be increased.</p> <p>Increase this value for systems that have a large number of clients or servers. The default provides for 73 x 7 = 511 ARP entries, but assumes an even hash distribution.</p>

clean_partial_conns	<p>Purpose: Specifies whether or not SYN (synchronizes the sequence number) attacks are being avoided.</p> <p>Values: Default: 0 (off) Range: 0 or 1 Type: Dynamic</p> <p>Tuning: This option should be turned on for servers that need to protect against network attacks. If on, randomly removes partial connections to make room for new non-attack connections.</p>
net_malloc_police	<p>Purpose: Specifies the size of the net_malloc and net_free trace buffers.</p> <p>Values: Default: 0 Range: 0 to MAXINT Type: Dynamic</p> <p>Tuning: If the value of this variable is non-zero, all net_malloc and net_free buffers will be traced in a kernel buffer and by system trace hook HKWD_NET_MALLOC. Additional error-checking will also be enabled. This includes checks for freeing a free buffer, alignment, and buffer overwrite. Enable this parameter only when investigating some network problem, because performance is affected negatively when turned on. The default value is zero (policing off). Values of net_malloc_police larger than 1024 will allocate that many items in the kernel buffer for tracing.</p>
rfc1323	<p>Purpose: Enables window scaling and time stamps as specified by RFC 1323 (TCP Extensions for High Performance). Window scaling allows the TCP window sizes (tcp_recvspace and tcp_sendspace) to be larger than 64KB (65536) and is typically used for large MTU networks.</p> <p>Values: Default: 0 (off) Range: 0 or 1 Type: Connect</p> <p>Tuning: The default value of 0 disables the RFC enhancements on a systemwide scale. A value of 1 specifies that all TCP connections will attempt to negotiate the RFC enhancements. Make changes before attempting to set tcp_sendspace and tcp_recvspace to more than 64 KB.</p>

route_expire	<p>Purpose: Specifies whether unused routes created by cloning, or created and modified by redirects expire.</p> <p>Values: Default: 1 (on) Range: 0 or 1 Type: Dynamic</p> <p>Tuning: A value of 1 allows route expiration, which is the default. Negative values are not allowed for this option.</p>
tcp_pmtu_discover	<p>Purpose: Enables or disables path MTU discovery for TCP applications.</p> <p>Values: Default: 1 Range: 0 or 1 Type: Dynamic</p> <p>Tuning: A value of 0 disables path MTU discovery for TCP applications, while a value of 1 enables it.</p>
tcp_recvspace	<p>Purpose: Specifies the system default socket buffer size for receiving data. This affects the window size used by TCP.</p> <p>Values: Default: 16384 bytes Range: 4096 to 1048576 Type: Connect</p> <p>Diagnosis: Setting the socket buffer size to 16 KB (16,384) improves performance over standard Ethernet and Token-Ring networks. Lower bandwidth networks, such as Serial Line Internet Protocol (SLIP), or higher bandwidth networks, such as Serial Optical Link, should have different optimum buffer sizes. The optimum buffer size is the product of the media bandwidth and the average round-trip time of a packet. For high-speed networks, like gigabit Ethernet or ATM 622, a value of 65536 should be used for the minimum size for best performance.</p> <p>For values larger than 65536, you must enable rfc1323 (rfc1323=1) to enable TCP window scaling.</p>

tcp_sendspace	<p>Purpose: Specifies the system default socket buffer size for sending data.</p> <p>Values: Default: 16384 bytes Range: 4096 to 1048576 Type: Connect</p> <p>Tuning: This affects the window size used by TCP. Setting the socket buffer size to 16 KB (16,384) improves performance over standard Ethernet networks. Lower bandwidth networks, such as Serial Line Internet Protocol (SLIP), or higher bandwidth networks, such as Serial Optical Link, should have different optimum buffer sizes. The optimum buffer size is the product of the media bandwidth and the average round-trip time of a packet: (optimum_window=bandwidth * average_round_trip_time) For high-speed networks, like gigabit Ethernet or ATM 622, a value of 65536 should be used for the minimum size for best performance. For values larger than 65536, you must enable rfc1323 (rfc1323=1) to enable TCP window scaling.</p>
udp_recvspace	<p>Purpose: Specifies the system default socket-buffer size for receiving UDP data.</p> <p>Values: Default: 42080 bytes Range: 4096 to 1048576 Type: Connect</p> <p>Diagnosis: Nonzero n in netstat -stats report of udp: n socket buffer overflows</p> <p>Tuning: Increase size, preferably to multiple of 4096.</p>
udp_sendspace	<p>Purpose: Specifies the system default socket-buffer size for sending UDP data.</p> <p>Values: Default: 9216 bytes Range: 4096 to 1048576 Type: Connect</p> <p>Diagnosis: Increase size, preferably to multiple of 4096.</p>

Exit Status

See “Exit status for Virtual I/O Server commands” on page 3.

Examples

1. To display the maximum size of the mbuf pool, type:
optimizenet -set thewall
2. To change the default socket buffer sizes on your system, type:
optimizenet -reboot -set tcp_sendspace=32768
3. To use a machine as an internet work router over TCP/IP networks, type:
optimizenet -set ipforwarding=1

4. To list the current and reboot value, range, unit, type and dependencies of the **arptab_bsize** parameter, type:

```
optimizenet -list arptab_bsize
```
5. To display help information on **arptab_bsize**, type:

```
optimizenet -h arptab_bsize
```

Related Information

The **entstat** command, the **lsnetsh** command, the **mktcpip** command, the **netstat** command, and the **tracert** command.

IVM os_install Command

Purpose

Performs network installation operations on **OS_install** objects. This command is operable only in an Integrated Virtualization Manager environment.

Syntax

Traditional usage:

```
OS_install {-o Operation} [ -a attr=value... ] {ObjectName}
```

For listing OS_install objects (list mode):

```
OS_install -l [-v] [-t object_type | object_name]
```

Description

The **OS_install** command performs a network installation operation on an **OS_install** object. The type of operation is dependent on the type of object specified by the ObjectName parameter. The object pointed to by the ObjectName parameter can be one of three types: client, **OS_Resource**, or **ControlHost**. Command operations involve the creation and management of **OS_install** objects in order to perform network installation of operating systems on client systems.

The list mode of the **OS_install** command is used to list the current configuration of objects in the **OS_install** environment.

Flags

-a attr=value	Assigns the specified value to the specified attribute. "Operations" on page 286 lists the required and optional attributes for a specific operation.
-l	Lists all OS_install objects in the environment by default.
-o Operation	Specifies an operation to perform on an OS_install object.
-t object_type object_name	Narrows the list returned by the -l flag to only objects of type <i>object_type</i> or to the single OS_install object specified by <i>object_name</i> .
-v	Displays the list returned by the -l flag.

Operations

Operation	Description	Required Attributes	Optional Attributes
define_client [-a <i>attr=value...</i>] { <i>ClientObjectName</i> }	Defines a new client object.	ip_addr Client's IP address. mac_addr MAC address of client's network interface. gateway Client's IP gateway. subnet_mask Client's IP subnet mask.	adapter_speed Speed of client's network adapter. adapter_duplex Duplex setting of client's network adapter. lpar LPAR name to install client. profile LPAR profile to use for client. managed_system Name of managed system that contains LPAR. disk_location Location of disk to install client. ctrl_host Name of Hardware Control Host object for this client.
define_resource [-a <i>attr=value...</i>] { <i>ResourceObjectName</i> }	Defines a new OS_Resource object.	type AIX or Linux version Operating system version	location Absolute path where OS_Resource will reside. source Source of installation images. configfile Installation configuration file.
define_ctrl_host [-a <i>attr=value...</i>] { <i>ControlHostObjectName</i> }	Defines a new Hardware Control Host object.	communication_method ssh, rsh, or local hostname Host name of control host. type hmc or ivm	None.

Operation	Description	Required Attributes	Optional Attributes
allocate [-a <i>attr=value...</i>] { <i>ClientObjectName</i> }	Allocates an OS_Resource object to a client object. Both objects must already exist in the OS_install environment. An error will occur if the client object has an OS_Resource object already allocated to it.	os_resource Existing OS_Resource object to allocate to the client object.	None.
netboot { <i>ClientObjectName</i> }	Instructs the hardware control host of the client object to initiate a network boot.	None.	None.
monitor_installation { <i>ClientObjectName</i> }	Monitors the installation status of the client object.	None.	None.
deallocate { <i>ClientObjectName</i> }	Deallocates the OS_Resource object that was allocated to the client object by an allocate operation.	None.	None.
remove { <i>ObjectName</i> }	Removes the object from the OS_install environment.	None.	None.

Exit Status

The following exit values are returned:

0 The command completed successfully.
>0 An error occurred.

Examples

1. To define a client object, type a command similar to the following:

```
OS_install -o define_client -a ip_addr=128.0.64.117
-a mac_addr=ab:cc:de:10:23:45 -a gateway=128.0.64.1
-a subnet_mask=255.255.255.0 -a ctrl_host=myhmc -a lpar=AIX1
-a profile=AIX1 -a managed_system=myMngSys myclient01
```

2. To define an **OS_Resource** object, type a command similar to the following:

```
OS_install -o define_resource -a location=/images/AIX/53ML3 -a type=AIX
-a version=53ML3 my53resource
```

3. To allocate the **OS_Resource** object defined in the preceding example to a client object, type a command similar to the following:

```
OS_install -o allocate -a os_resource=my53resource myclient01
```

4. To deallocate the my53resource client object that was allocated in the preceding example, type:

```
OS_install -o deallocate myclient01
```

- To define a **ControlHost** object to be specified for the **ctrl_host** attribute, type a command similar to the following:

```
OS_install -o define_ctrl_host -a type=ivm -a hostname=ivm_hostname
-a communication_method=ssh myivm
```

- To view a myclient01 installation, type:
OS_install -o monitor_installation myclient01
- To remove the definition of the my53resource object, type:
OS_install -o remove my53resource
- To remove the definition of the myclient01 object, type:
OS_install -o remove myclient01

Note: If an **OS_Resource** object is specified, the remove operation also removes any OS images that exist in the file system directory that is specified by the object's location attribute.

Related Information

The **lpar_netboot** command.

passwd command

Purpose

Changes a user's password.

Syntax

```
passwd [ User ]
```

Description

The **passwd** command sets and changes passwords for users. Use this command to change your own password (all users) or another user's password (padmin only). To change your own password, enter the **passwd** command. The **passwd** command prompts the non-padmin user for the old password and then prompts for the new password twice. The password is never displayed on the screen. If the two entries of the new password do not match, the **passwd** command prompts for the new password again.

Password policy is checked during a password change. Construct locally defined passwords according to the following password restrictions:

minother	Specifies the minimum number of other characters.
minlen	Specifies the minimum number of characters.
maxrepeats	Specifies the maximum number of times a single character can be used in a password.
maxage	Specifies the maximum age of a password. A password must be changed after a specified amount of time measured in weeks.
maxexpired	Specifies the maximum number of weeks beyond the maxage value that a password can be changed by the user.
histexpire	Specifies the number of weeks that a user cannot reuse a password.
histsize	Specifies the number of previous passwords that the user cannot reuse.

Exit Status

See “Exit status for Virtual I/O Server commands” on page 3.

Security

Changing a password other than your own requires prime administrator authority.

Examples

1. To change the password for user account **heerlen**, type:
passwd heerlen

The user will then be prompted to enter the new password.

Related Information

The **chuser** command, the **lsuser** command, the **mkuser** command, and the **rmuser** command.

pdump command

Purpose

Perform platform (hardware and firmware) dump-related actions.

Syntax

```
pdump -reserve fstype | -enable | -disable | -ls | -size
```

Description

The **pdump** command helps the operating system save firmware and hardware related dumps. This command also provides an estimate of the disk space required for storing these dumps. Note that platform and scan dumps are saved in order to capture the state of the firmware and the hardware for analysis.

Flags

-disable	Disables platform dumps.
-enable	Enables platform dumps.
-ls	Lists the current configuration of platform dump.
-reserve <i>fstype</i>	Reserves enough disk space on the system for platform dumps. The -enable option will create a file system (if one does not exist) exclusively for platform dumps. If a file system already exists and the size is not enough, the file system size will be increased. The <i>fstype</i> must be a valid file system type. If the file system already exists, any may be specified.
-size	Provides an estimate of disk space required to save the platform dumps when they occur. This option will interact with the firmware to provide this estimate. It is expected that, based on this space information, the user will have enough disk space allocated for platform dumps to be saved. The value output will be the required size in bytes.

ping command

Purpose

Sends an echo request to a network host.

Syntax

```
ping [ -n ] [ -r ] [ -s PacketSize ] [ -src hostname/IP_addr ] Host [ Count ]
```

Description

The **ping** command sends an Internet Control Message Protocol (ICMP) ECHO_REQUEST to obtain an ICMP ECHO_RESPONSE from a host or gateway. The **ping** command is useful for:

- Determining the status of the network and various foreign hosts.
- Tracking and isolating hardware and software problems.
- Testing, measuring, and managing networks.

If the host is operational and on the network, it responds to the echo. Each echo request contains an Internet Protocol (IP) and ICMP header, followed by a timeval structure, and enough bytes to fill out the packet. The default is to continuously send echo requests until an Interrupt is received (Ctrl-C).

The **ping** command sends one datagram per second and prints one line of output for every response received. The **ping** command calculates round-trip times and packet loss statistics, and displays a brief summary on completion. The ping command completes when the program times out or on receipt of a SIGINT signal. The Host parameter is either a valid host name or Internet address.

By default, the **ping** command will continue to send echo requests to the display until an Interrupt is received (Ctrl-C). Because of the load that continuous echo requests can place on the system, repeated requests should be used primarily for problem isolation.

Flags

-n	Specifies numeric output only. No attempt is made to look up symbolic names for host addresses.
-r	Bypasses the routing tables and sends directly to a host on an attached network. If the Host is not on a directly connected network, the ping command generates an error message. This option can be used to ping a local host through an interface that no longer has a route through it.
-s PacketSize	Specifies the number of data bytes to be sent. The default is 56, which translates into 64 ICMP data bytes when combined with the 8 bytes of ICMP header data.
-src hostname/IP_addr	Uses the IP address as the source address in outgoing ping packets. On hosts with more than one IP address, the -src flag can be used to force the source address to be something other than the IP address of the interface on which the packet is sent. If the IP address is not one of the machine's interface addresses, an error is returned and nothing is sent.

Parameters

<i>Count</i>	Specifies the number of echo requests to be sent (and received). This parameter is included for compatibility with previous versions of the ping command.
--------------	--

Exit Status

See "Exit status for Virtual I/O Server commands" on page 3.

Examples

1. To check the network connection to host **canopus** and specify the number of echo requests to send, enter:

```
ping canopus 5
```

Information similar to the following is displayed:

```
PING canopus.austin.century.com: (128.116.1.5): 56 data bytes
64 bytes from 128.116.1.5: icmp_seq=0 ttl=255 time=2 ms
64 bytes from 128.116.1.5: icmp_seq=1 ttl=255 time=2 ms
64 bytes from 128.116.1.5: icmp_seq=2 ttl=255 time=3 ms
64 bytes from 128.116.1.5: icmp_seq=3 ttl=255 time=2 ms
64 bytes from 128.116.1.5: icmp_seq=4 ttl=255 time=2 ms
```

```
----canopus.austin.century.com PING Statistics----
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max = 2/2/3 ms
```

2. To obtain information about host **opus** and specify the number of data bytes to be sent, enter:

```
ping -s 2000 opus
```

Information similar to the following is displayed:

```
PING opus.austin.century.com: (129.35.34.234): 2000 data bytes
2008 bytes from 129.35.34.234: icmp_seq=0 ttl=255 time=20 ms
2008 bytes from 129.35.34.234: icmp_seq=1 ttl=255 time=19 ms
2008 bytes from 129.35.34.234: icmp_seq=2 ttl=255 time=20 ms
2008 bytes from 129.35.34.234: icmp_seq=3 ttl=255 time=20 ms
2008 bytes from 129.35.34.234: icmp_seq=4 ttl=255 time=20 ms
2008 bytes from 129.35.34.234: icmp_seq=5 ttl=255 time=19 ms
2008 bytes from 129.35.34.234: icmp_seq=6 ttl=255 time=19 ms
^C
```

```
----opus.austin.century.com PING Statistics----
7 packets transmitted, 7 packets received, 0% packet loss
round-trip min/avg/max = 19/19/20 ms
```

Note: The output is repeated until an Interrupt (Ctrl-C) is received.

Related Information

The **optimizenet** command and the **traceroute** command.

postprocesssvc command

Purpose

Performs actions on agents after the agent is started by the **startsvc** command.

Syntax

```
postprocesssvc AgentName { -attr Attribute[=Value] ... | -ls }
```

Description

The **postprocesssvc** command performs actions for an agent based on the set of attributes specified for the command.

Flags

-attr	Lists actionable agent options, which can be any of the following attributes: D Generates the <i>stats.send</i> file from the specified date. A value of 0 (zero) specifies to generate <i>data.send</i> from all the available data files. T Enables or disables the data transmission. S Prints the absolute path to the <i>stats.send</i> file. A Lists all the stats files in data directory, except the <i>stats.send</i> file. V Prints the version.
-ls	Lists the attributes that can be processed.

perfmgr

The performance manager agent (**perfmgr**) collects data about system performance and sends it to support with the Electronic Service Agent (ESA) agent for processing.

When the agent is started, it runs a set of utilities to collect metrics to measure performance. This command generates a file from the list of available individual performance data files that the ESA agent can understand.

The list of available attributes for this agent is **D**, **T**, **S**, **A**, and **V**.

IBM License Metric Tool

Performs additional operations on IBM License Metric Tool (ILMT) after the initial configuration with **cfgsvc** command. The additional operations, which are performed are changing ILMT server, setting port number, and security level.

Attributes

capture	Copies all the agent trace, message, and configuration files to the <tivoli_common_dir>\ffdc\agent subdirectory under the Tivoli Common Directory. These files are available for remote support analysis.
clone	Prepares the agent to be cloned together with the operating system image for serial deployment. The command performs various actions as follows: <ul style="list-style-type: none">• It stops the agent if it is running.• It removes the agent cache and all the data prepared to be uploaded to the server.• It removes the agent traces.• It replaces the agent ID in the <i>tlmagent.ini</i> file with <i>%AGENT_ID</i> string.

cmds	Displays the most important information about the schedule of the commands in the standard output. The command displays the information as follows: <ul style="list-style-type: none"> • The commands that are currently active. • The current period of each command. • The time of the last command execution for each command. • The planned time of the next command execution for each command. • The end of the schedule for a given command.
i	Installs the agent as a service.
impcert=<filename.kdb>:<cms_password>	Imports the agent CMS certificate and the associated private key from the location provided. The required <i>cms_keystore</i> file is the file <code>\outputdirectory\organization_name\cms\agent_id.kdb</code> created by the <i>convertcertificate</i> command during the agent authentication process. The required <i>cms_password</i> is the corresponding CMS password.
p	Plugs into the Asset Discovery for Distributed server. The agent must be active and the task is performed immediately.
r	Removes the agent service from the Services list.
reload	Reloads the agent configuration file and restarts the trace by using the new configuration. The agent must be active.
s	Runs an inventory scan. The result of the scan is sent to the server in a scheduled agent-server communication time window. The agent must be active. This option requests that the scan task be run as soon as possible by the agent.
set=<property_name>:<property_value>	Sets custom data or modifies other properties in the <i>tlmagent.ini</i> file
sslreload	Reloads the security settings from the agent configuration file and attempts to import the private certificate file from the default location by setting the security level to 2 (MAX) and agent ID to the one specified in the private certificate file. This option is valid only if the 0 (MIN) is not used.
v	Displays the agent version.

Exit Status

0	The command completed successfully
>0	An error occurred.

Examples

1. To list the available attributes for the **perfmgr** agent, type the command as follows:

```
postprocesssvc perfmgr -ls
```
2. To generate the stats.send file from the data available on 06/14/2006, type the command as follows:

```
postprocesssvc perfmgr -attr D=20090614
```
3. To generate the stats.send file from the data available, type the command as follows:

```
postprocesssvc perfmgr -attr D=0
```
4. To list the all the stats files in the data directory, type the command as follows:

```
postprocesssvc perfmgr -attr A
```
5. To print the absolute path to the stats.send file, type the command as follows:

```
postprocesssvc perfmgr -attr S
```
6. To print the version, type the command as follows:

```
postprocesssvc perfmgr -attr V
```

ILMT

1. To reload the agent configuration file, type the command as follows:
postprocesssvc ILMT -attr reload
2. To check whether the agent plugs in to the server, type the command as follows:
postprocesssvc ILMT -attr p
3. To perform an inventory scan, type the command as follows:
postprocesssvc ILMT -attr s

Related Information

The **startsvc** command and the **stopsvc** command.

prepdev command

Purpose

Prepares a device to be assigned to a shared memory pool (to be used as a paging space device by a shared memory partition).

Syntax

```
prepdev { -dev devicename }
```

Description

The **prepdev** command prepares a device to be assigned to a shared memory pool (to be used as a paging space device by a shared memory partition). This command determines whether the device specified by the **-dev** flag is already assigned or used in the following manner:

- Assigned to a shared memory pool (being used as a paging device)
- Used as a backup device for virtual SCSI
- Used by a file system, logical volume, or volume group.

If the device is already being used or assigned to a shared memory pool, you receive system output that tells you how to prepare the device to be used as a paging space device.

Flags

-dev <i>devicename</i>	Specifies the logical device to be used as a paging space device.
--help	Displays the help text for this command and exit.

Exit Status

See “Exit status for Virtual I/O Server commands” on page 3.

Example

Related Information

The **lspv** command and the **rmvdev** command.

redefvg command

Purpose

Redefines the set of physical volumes of the given volume group in the device configuration database.

Syntax

```
redefvg { -dev Device | -vgid Vgid } VolumeGroup
```

Description

During normal operations, the device configuration database remains consistent with the Logical Volume Manager (LVM) information in the reserved area on the physical volumes. If inconsistencies occur between the device configuration database and the LVM, the **redefvg** command determines which physical volumes belong to the specified volume group and enters this information in the device configuration database. The **redefvg** command checks for inconsistencies by reading the reserved areas of all the configured physical volumes attached to the system.

Flags

-dev <i>Device</i>	The volume group ID, <i>Vgid</i> , is read from the specified physical volume device. You can specify the <i>Vgid</i> of any physical volume belonging to the volume group that you are redefining.
-vgid <i>Vgid</i>	The volume group identification number of the volume group to be redefined.

Exit Status

See “Exit status for Virtual I/O Server commands” on page 3.

Examples

1. To synchronize the copies on physical volumes `hdisk04`, type:

```
redefvg -dev hdisk04
```
2. To synchronize the copies on volume groups `vg04` and `vg05`, type:

```
redefvg -vgid vg04 vg05
```

Related Information

The **mkvg** command, the **syncvg** command, the **chvg** command, the **extendvg** command, the **lsvg** command, the **mirrorios** command, the **unmirrorios** command, the **activatevg** command, the **deactivatevg** command, the **importvg** command, the **exportvg** command, and the **reducevg** command.

reducevg command

Purpose

Removes physical volumes from a volume group. When all physical volumes are removed from the volume group, the volume group is deleted.

Syntax

```
reducevg [ -f ] [ -rmlv ] VolumeGroup PhysicalVolume ...
```

Description

The **reducevg** command removes one or more physical volumes represented by the *PhysicalVolume* parameter from the *VolumeGroup*. When you remove all physical volumes in a volume group, the volume group is also removed.

Note:

- Sometimes a disk is removed from the system without first running the **reducevg** command. The VGDA still has this removed disk in its memory, but the *PhysicalVolume* name no longer exists or has been reassigned. To remove references to this missing disk you can still use **reducevg**, but with the Physical Volume ID (PVID) instead of the disk name: `reducevg VolumeGroup PVID`.
- The specified physical volume cannot contain a logical volume that is assigned to a shared memory pool (to be used as a paging space device by a shared memory partition).

Flags

-f	Removes the requirement for user confirmation when the -rmlv flag is used.
-rmlv	Deallocates the existing logical volume partitions and then deletes resultant empty logical volumes from the specified physical volumes. User confirmation is required unless the -f flag is added.

Attention: The **reducevg** command with the **-rmlv** flag automatically deletes all logical volume data on the physical volume before removing the physical volume from the volume group. If a logical volume spans multiple physical volumes, the removal of any of those physical volumes may jeopardize the integrity of the entire logical volume.

Exit Status

See “Exit status for Virtual I/O Server commands” on page 3.

Examples

1. To remove physical volume **hdisk1** from volume group **vg01**, type:
`reducevg vg01 hdisk1`
2. To remove physical volume **hdisk1** and all residing logical volumes from volume group **vg01** without user confirmation, type:

Attention: The **reducevg** command with the **-rmlv** flag automatically deletes all logical volume data before removing the physical volume.

```
reducevg -rmlv -f vg01 hdisk1
```

The physical volume **hdisk1** and all residing logical volumes are removed.

Related Information

The **mkvg** command, the **chvg** command, the **extendvg** command, the **lsvg** command, the **mirrorios** command, the **unmirrorios** command, the **activatevg** command, the **deactivatevg** command, the **importvg** command, the **exportvg** command, and the **syncvg** command.

remote_management command

Purpose

Enables the Virtual I/O Server to be remotely managed by an AIX NIM (Network Installation Management) master.

Syntax

To enable the Virtual I/O Server to be remotely managed by an AIX NIM master:

```
remote_management [ -interface Interface ] Master
```

To disable remote management:

```
remote_management -disable
```

Description

The **remote_management** command will setup the Virtual I/O Server to allow remote management from a NIM master. The *Master* parameter specifies the NIM master *hostname*. The *Interface* parameter specifies the network interface to be used to connect to the NIM master. If *Interface* is not specified, the default network interface used will be **en0**.

The **remote_management** command utilizes the NIM service handler for client communication (nimsh), so the NIM master must be nimsh capable.

Once remote management has been enabled on the Virtual I/O Server, typical NIM functions, such as update, backup, and reinstall, can be initiated from the NIM master.

Flags

-disable	Disables the Virtual I/O Server NIM client daemon.
-interface	Specifies which network interface to use. If no network interface is specified, interface en0 will be used.

Exit Status

See “Exit status for Virtual I/O Server commands” on page 3

Examples

1. To enable remote_management using NIM master **nimsys01**, type:
remote_management nimsys01
2. To disable remote_management, type:
remote_management -disable

Related Information

The **lssw** command, the **ioslevel** command, the **updateios** command, the **oem_setup_env** command, and the **oem_platform_level** command.

replphyvol command

Purpose

Replaces a physical volume in a volume group with another physical volume.

Syntax

```
replphyvol [-force ] {SourcePhysicalVolume | SourcePhysicalVolumeID} {DestinationPhysicalVolume | DestinationPhysicalVolumeID}
```

replphyvol [-retry] *dir_name* [DestinationPhysicalVolume]

Description

The **replphyvol** command replaces allocated physical partitions and the data they contain from the *SourcePhysicalVolume* to *DestinationPhysicalVolume*. The specified source physical volume cannot be the same as *DestinationPhysicalVolume*.

Note:

1. The *DestinationPhysicalVolume* size must be at least the size of the *SourcePhysicalVolume*.
2. The **replphyvol** command cannot replace a *SourcePhysicalVolume* with stale logical volume unless this logical volume has a non-stale mirror.
3. The **replphyvol** command is not allowed on a snapshot volume group or a volume group that has a snapshot volume group. The allocation of the new physical partitions follows the policies defined for the logical volumes that contain the physical partitions being replaced.
4. The *DestinationPhysicalVolume* cannot be a physical volume that is assigned to a shared memory pool (to be used as a paging space device by a shared memory partition).

The allocation of the new physical partitions follows the policies defined for the logical volumes that contain the physical partitions being replaced.

Flags

-force	Runs the command without prompting the user to continue.
-retry	Displays specified boot list, after any specified alterations
<i>SourcePhysicalVolume</i>	Specifies the source physical volume name.
<i>SourcePhysicalVolumeID</i>	Specifies the source physical volume ID.
<i>DestinationPhysicalVolume</i>	Specifies the destination physical volume name.
<i>DestinationPhysicalVolumeID</i>	Specifies the destination physical volume ID.
<i>dir_name</i>	Recovers replphyvol if it is interrupted by CTL+C, a system lock up, or a loss of quorum. Specify the directory name given during the initial run of replphyvol if you are using the -retry flag. Using this flag, you can also change the <i>DestinationPhysicalVolume</i>

Exit Status

See “Exit status for Virtual I/O Server commands” on page 3.

Examples

To create a boot list of logical devices to be used in the next normal boot, type the following command:
`replphyvol -force hdisk0 hdisk4`

To recover a boot list of logical devices if the normal boot was interrupted, type the following command:
`replphyvol -retry <directory_name> hdisk4`

Related Information

The **migratepv** command.

restore command

Purpose

Extracts files from archives created with the **backup** command.

Syntax

To restore files archived by file name:

```
restore -x [ d M n Q v q e ] [ -b Number ] [ -f Device ] [ -s SeekBackup ] [ -E { force | ignore | warn } ] [ File ... ]
```

To list files archived by file name:

```
restore -T | -t [ a l n q v Q ] [ -b Number ] [ -f Device ] [ -s SeekBackup ]
```

To restore files archived by file system:

- **restore -r** [B n q v y] [-b Number] [-f Device] [-s SeekBackup]
- **restore -R** [B n v y] [-b Number] [-f Device] [-s SeekBackup]
- **restore -i** [h m n q v y] [-b Number] [-f Device] [-s SeekBackup]
- **restore -x** [B h n m q v y] [-b Number] [-f Device] [-s SeekBackup] [File ...]

To restore files beginning at a specified volume number:

```
restore -X Number [ -Md n q v e Q ] [ -b Number ] [ -f Device ] [ -s Number ] [ -E { force | ignore | warn } ] [ File ... ]
```

To list files archived by file system:

```
restore -t | -T [ B a l n h q v y ] [ -b Number ] [ -f Device ] [ -s SeekBackup ] [ File ... ]
```

To restore file attributes archived by file name:

```
restore -Pstring [ B d q v Q ] [ b Number ] [ s SeekNumber ] [ -f Device ] [ File ... ]
```

To restore file attributes archived by file system:

```
restore -Pstring [ h q v ] [ b Number ] [ s SeekNumber ] [ -f Device ] [ File ... ]
```

Description

The **restore** command reads archives created by the backup operation and extracts the files stored on them. These archives can be in either file name or file system format. An archive can be stored on disk, diskette, or tape. Files must be restored using the same method by which they were archived. This requires that you know the format of the archive. The archive format can be determined by examining the archive volume header information that is displayed when using the **-T** flag. When using the **-x**, **-r**, **-T**, or **-t** flags, the **restore** command automatically determines the archive format.

Note: The **restore** command determines the actively sparses files that are being restored. If a file has block aligned and sized areas that are NULL populated, then the **restore** command does not cause physical space for those file system blocks to be allocated. The size in bytes of the file remain the same, but the actual space taken within the file system is only for the non-NULL areas.

Individual files can be restored from either file name or file system archives by using the **-x** flag and specifying the file name. The file name must be specified as it exists on the archive. Files can be restored interactively from file system archives using the **-i** flag. The names of the files on an archive can be written to standard output using the **-T** flag.

Users must have write access to the file system device or have restore authorization in order to extract the contents of the archive.

The diskette device, `/dev/rfd0`, is the default media for the **restore** command. To restore from standard input, specify a dash (-) with the **-f** flag. You can also specify a range of devices, such as `/dev/rmt0-2`.

Note:

1. If you are restoring from a multiple-volume archive, the **restore** command reads the volume mounted, prompts you for the next volume, and waits for your response. After inserting the next volume, press the Enter key to continue restoring files.
2. If an archive was created using the **backup** command and was made to a tape device with the device block size set to 0, it might be necessary for you to have explicit knowledge of the block size that was used when the tape was created to restore from the tape.
3. Multiple archives can exist on a single tape. When you are restoring multiple archives from tape, the **restore** command expects the input device to be a no-retension-on-open, no-rewind-on-close tape device. Do not use a no-rewind tape device for restoring unless the **-B**, **-s**, or **-X** flag is specified.

File system archives

File system archives are also known as i-node archives due to the method used to archive the files. A file system name is specified with the **backup** command, and the files within that file system are archived based on their structure and layout within the file system. The **restore** command restores the files on a file system archive without any special understanding of the underlying structure of the file system.

When restoring file system archives, the **restore** command creates and uses a file named `restoresymtable`. This file is created in the current directory. The file is necessary for the **restore** command to do incremental file system restore operation.

Note: Do not remove the `restoresymtable` file if you perform incremental file system backups and restores.

The *File* parameter is ignored when using either the **-r** or the **-R** flag.

File-name archives

File-name archives are created by specifying a list of file names to archive by using the **backup** command. The restore operation restores the files from a file name archive without any special understanding of the underlying structure of the file system. The restore operation allows for metacharacters to be used when specifying files for archive extraction. This provides the capability to extract files from an archive based on pattern matching. A pattern file name should be enclosed in single quotations, and patterns should be enclosed in parentheses (...).

About sparse files

File system files for the operating system that contain long strings of NULLs can be stored much more efficiently than other files. To be specific, if a string of NULLs spans an entire allocation block, that whole block is not stored on disk at all. Files where one or more blocks are omitted in this way are called sparse files. The missing blocks are also known as holes.

Note: Sparse files are not the same as compressed files. Sparse files are exactly the same as their nonsparse equivalents when they are read.

Sparse files are generally created by database applications. Whenever a database file is created, it is formatted with NULLs. These fragments might be either filled with allocated or unallocated NULLs.

Flags

- a** Displays the **-a** option displays the list of files in the archive, along with their permissions when specified with the **-t** or **-T** flag.
- B** Specifies that the archive should be read from standard input. Normally, the **restore** operation examines the actual medium to determine the backup format. When using a | (pipe), this examination cannot occur. As a result, the archive is assumed to be in file system format, and the device is assumed to be standard input (**-f -**).
- b *Number***
- For backups done by name, specifies the number of 512-byte blocks. For backups done by i-node, specifies the number of 1024-byte blocks to read in a single output operation. When the **restore** operation reads from tape devices, the default is 100 for backups by name and 32 for backups by i-node.
- The read size is the number of blocks multiplied by the block size. The default read size for the **restore** operation reading from tape devices is 51200 (100 × 512) for backups by name and 32768 (32 × 1024) for backups by i-node. The read size must be an even multiple of the tape's physical block size. If the read size is not an even multiple of the tape's physical block size and it is in fixed block mode (nonzero), the **restore** command tries to determine a valid value for *Number*. If successful, the **restore** operation changes *Number* to the new value, writes a message about the change to standard output, and continues. If unsuccessful in finding a valid value for *Number*, the **restore** operation writes an error message to standard error and exits with a nonzero return code. Larger values for the *Number* parameter result in larger physical transfers from the tape device.
- The value of the **-b** flag is always ignored when the **restore** command reads from diskette. In this case, the command always reads in clusters that occupy a complete track.
- d** Indicates that, if the *File* parameter is a directory, all files in that directory should be restored. This flag can only be used when the archive is in file name format.
- e**
- Restores the nonsparse files because they were archived by the name format of the backup operation for both packed and unpacked files. It is necessary to know the sparseness or nonsparseness of the files before archiving the files, since enabling this flag restores the sparse files as nonsparse.
- This flag should be enabled only if the files being restored are nonsparsed, consisting of more than 4K NULLs. If the **-e** flag is specified during the restore operation, it successfully restores all normal files normally and nonsparse database files as nonsparse.
- E**
- Extracts beginning at a specified volume number and requires one of the following arguments. If you omit the **-E** flag, **warn** is the default behavior.
- force** Fails the restore operation on a file if the fixed extent size or space reservation of the file cannot be preserved.
- ignore** Ignores any errors in preserving extent attributes.
- warn** Issues a warning if the space reservation or the fixed size of the file cannot be preserved.

-f *Device*

Specifies the input device. To receive input from a named device, specify the *Device* variable as a path name (such as `/dev/rmt0`). To receive input from the standard output device, specify a dash (-)

You can also specify a range of archive devices. The range specification must be in the following format:

`/dev/deviceXXX-YYY`

where XXX and YYY are whole numbers, and XXX must always be less than YYY; for example:

`/dev/rfd0-3`

All devices in the specified range must be of the same type. For example, you can use a set of 8 mm, 2.3 GB tapes or a set of 1.44 MB diskettes. All tape devices must be set to the same physical tape block size.

If the *Device* variable specifies a range, the restore operation automatically goes from one device in the range to the next. After exhausting all of the specified devices, the restore operation halts and requests that new volumes be mounted on the range of devices.

-h

Restores only the actual directory, not the files contained in it. This flag can only be used when the archive is in file system format. This flag is ignored when used with the **-r** or **-R** flag.

-i

Allows you to interactively restore selected files from a file system archive. The subcommands for the **-i** flag are:

add [*File*]: Specifies that the *File* parameter is added to the list of files to extract. If *File* is a directory, that directory and all the files contained in it are added to the extraction list (unless the **-h** flag is used). If *File* is not specified, the current directory is added to the extraction list.

cd *Directory*: Changes the current directory to the specified directory.

delete [*File*]: Specifies that the *File* parameter is to be removed from the list of files to be extracted. If *File* is a directory, that directory and all the files contained in it are removed from the extraction list (unless the **-h** flag is used).

extract: Restores all the directories and files on the extraction list.

help: Displays a summary of the subcommands.

ls [*Directory*]: Displays the directories and files contained within the *Directory* parameter. Directory names are displayed with a slash (/) after the name. Files and directories, within the specified directory, that are on the extraction list are displayed with an asterick (*) before the name. If verbose mode is on, the i-node number of the files and directories is also displayed. If the *Directory* parameter is not specified, the current directory is used .

pwd: Displays the full path name of the current directory.

quit: Causes the restore operation to exit immediately. Any files on the extraction list are not restored.

setmodes: Sets the owner, mode, and time for all directories added to the extraction list.

verbose: Causes the **ls** subcommand to display the i-node number of files and directories. Additional information about each file is also displayed as it is extracted from the archive.

-l

Displays a detailed list of files, which includes the time stamp, file permissions, file size, owner, and group when specified with the **-t** or **-T** flag.. The **-l** option overrides the **-a** option.

- M** Sets the access and modification times of restored files to the time of restoration. You can specify the **-M** flag only when you are restoring individually named files and only if the **-x** or **-X** flags are also specified. When the **-M** flag is not specified, the restore operation maintains the access and modification times as they appear on the backup medium.
- m** Renames restored files to the file's i-node number as it exists on the archive. This is useful if a few files are being restored and you want these files restored under a different file name. Since any restored archive members are renamed to their i-node numbers, directory hierarchies and links are not preserved. Directories and hard links are restored as regular files. The **-m** flag is used when the archive is in file system format.
- n** By default the restore operation restores any Access Control List (ACL)s, Printing Color Files (PCL)s, or named extended attributes in the archive. The **-n** flag causes the restore operation to skip over any ACLs, PCLs, or named extended attributes in the archive and not restore them.
- Pstring** Restore only the file attributes. This flag does not restore the file contents. If the file specified does not exist in the target directory path, the file is not created. This flag restores file attributes selectively depending on the flags specified in the string parameter. The *String* parameter can be a combination of the following characters:
- A** Restore all attributes.
 - a** Restore only the permissions of the files.
 - o** Restore only the ownership of the files.
 - t** Restore only the time stamp of the files.
 - c** Restore only the ACL attributes of the files
- Note:** Among the existing flags for the **restore** command, flags **v**, **h**, **b**, **s**, **f**, **B**, **d**, and **q** are valid with the **P** flag. The **P** flag can be used with both file name and file system archives. If the *File* parameter is a symbolic link, then the metadata of the target file is modified and not that of the symbolic link.
- Attention:** Usage of **-P** flag overwrites the attributes of files owned by another user when run by the superuser.
- Q** For backups done by name, specifies that the command should exit upon encountering any type of error rather than attempt to recover and continue processing the archive.
- q** Specifies that the first volume is ready to use and that the restore operation should not prompt you to mount the volume and press Enter. If the archive spans multiple volumes, the **restore** command prompts you for the subsequent volumes.
- r** Restores all files in a file system archive. The **-r** flag is only used to restore complete level 0 backups or to restore incremental backups after a level 0 backup is restored. The *restoresymtable* file is used by the restore operation to pass information between incremental restore operations. This file should be removed after the last incremental backup is restored. The *File* parameter is ignored when using the **-r** flag.
- R** Requests a specific volume of a multiple-volume, file system archive. The **-R** flag allows a previously interrupted restore operation to be restarted. The *File* parameter is ignored when using the **-R** flag. Once restarted, the **restore** command behavior is the same as with the **-r** flag.
- s *SeekBackup*** Specifies the backup to seek and restore on a multiple-backup tape archive. The **-s** flag is only applicable when the archive is written to a tape device. To use the **-s** flag properly, a no-rewind-on-close and no-retension-on-open tape device, such as */dev/rmt0.1* or */dev/rmt0.5*, must be specified. If the **-s** flag is specified with a rewind tape device, the **restore** command displays an error message and exits with a nonzero return code. If a no-rewind tape device is used and the **-s** flag is not specified, a default value of **-s 1** is used. The value of the *SeekBackup* parameter must be in the range of 1 to 100 inclusive. It is necessary to use a no-rewind-on-close, no-retension-on-open tape device because of the behavior of the **-s** flag. The value specified with **-s** is relative to the position of the tape's read/write head and not to an archive's position on the tape. For example, to restore the first, second, and fourth backups from a multiple-backup tape archive, the respective values for the **-s** flag would be **-s 1**, **-s 1**, and **-s 2**.

- t Displays information about the backup archive. If the archive is in file system format, a list of files found on the archive is written to standard output. The name of each file is preceded by the i-node number of the file as it exists on the archive. The file names displayed are relative to the root (/) directory of the file system that was backed up. If the *File* parameter is not specified, all the files on the archive are listed. If the *File* parameter is used, then just that file is listed. If the *File* parameter refers to a directory, all the files contained in that directory are listed. If the archive is in file name format, information contained in the volume header is written to standard error. This flag can be used to determine if the archive is in file name or file system format.
- T Displays information about the backup archive. If the archive is in file name format, the information contained in the volume header is written to standard error, and a list of files found on the archive is written to standard output. The *File* parameter is ignored for file name archives. If the archive is in file system format, the behavior is identical to the **-t** flag.
- v Displays additional information when restoring. If the archive is in file name format and either the **-x** or **-T** flag is specified, the size of the file as it exists on the archive is displayed in bytes. Directory, block, or character device files are archived with a size of 0. Symbolic links are listed with the size of the symbolic link. Hard links are listed with the size of the file, which is how they are archived. Once the archive is read, a total of these sizes is displayed. If the archive is in file system format, directory and nondirectory archive members are distinguished.
- x Restores individually named files specified by the *File* parameter. If the *File* parameter is not specified, all the archive members are restored. If the *File* parameter is a directory and the archive is in file name format, only the directory is restored. If the *File* parameter is a directory and the archive is in file system format, all the files contained in the directory are restored. The file names specified by the *File* parameter must be the same as the names shown by the restore **-T** flag. Files are restored with the same name they were archived with. If the file name was archived using a relative path name (./filename), the file is restored relative to the current directory. If the archive is in file system format, files are restored relative to the current directory.

The **restore** command automatically creates any needed directories. When using this flag to restore file system backups, you are prompted to enter the beginning volume number.

The **restore** command allows for shell-style pattern matching metacharacters to be used when specifying files for archive extraction. The rules for matching metacharacters are the same as those used in shell pathname "globbing," namely:

asterick* (*)
Matches zero or more characters, but not a . (period) or / (slash).

question mark (?)
Matches any single character, but not a . (period) or / (slash).

brackets ([])
Matches any one of the characters enclosed within the brackets. If a pair of characters separated by a dash are contained within the brackets, the pattern matches any character that lexically falls between the two characters in the current local. Additionally, a . (period) or a / (slash) within the brackets will not match a . (period) or a / (slash) in a file name.

backslash (\)
Matches the immediately following character, preventing its possible interpretation as a metacharacter.
- X *VolumeNumber* Begins restoring from the specified volume of a multiple-volume, file name backup. Once started, the **restore** command behavior is the same as with the **-x** flag. The **-X** flag applies only to file name archives.
- y Continues restoring when tape errors are encountered. Normally, the **restore** command asks you whether or not to continue. In either case, all data in the read buffer is replaced with zeros. The **-y** flag applies only when the archive is in file system format.

? Displays a usage message.

Exit Status

This command returns the following exit values:

- 0 Successful completion.
- >0 An error occurred.

Examples

1. To list the names of files in either a file name or file system archive on diskette device `/dev/rfd0`, type the following command:

```
restore -Tq
```

The archive is read from the `/dev/rfd0` default restore device. The names of all the files and directories contained in the archive are displayed. For file system archives, the file name is preceded by the i-node number of the file as it exists on the archive. The `-q` flag tells the `restore` command that the first volume is available and is ready to be read. As a result, you are not prompted to mount the first volume.

2. To restore a specific file, type the following command:

```
restore -xvqf myhome.bkup system.data
```

This command extracts the file `system.data` into the current directory from the archive `myhome.bkup`. The archive in this example is in the current directory. File and directory names must be specified as they are displayed when using the `-T` flag. The `-v` flag displays additional information during the extraction. This example applies to both file name and file system archives.

3. To restore a specific file from tape into the virtual media repository, type the following command:

```
restore -xvqf /dev/rmt0 /var/vio/VMLibrary/media_file
```

4. To restore a specific directory and the contents of that directory from a file name archive, type the following command:

```
restore -xdvqf /dev/rmt0 /home/mike/tools
```

The `-x` flag extracts files by their file name. The `-d` flag tells the `restore` command to extract all the files and subdirectories in the `/home/mike/tools` directory. File and directory names must be specified as they are displayed when using the `-T` flag. If the directories do not exist, they are created.

5. To restore a specific directory and the contents of that directory from a file system archive, type the following command:

```
restore -xvqf /dev/rmt0 /home/mike/tools
```

This command extracts files by file name. File and directory names must be specified as they are displayed when using the `-T` flag. If the directories do not exist, they are created.

6. To restore an entire file system archive, type the following command:

```
restore -rvqf /dev/rmt0
```

This command restores the entire file system archived on tape device `/dev/rmt0` the current directory. This example assumes you are in the root directory of the file system to be restored. If the archive is part of a set of incremental file system archives, the archives should be restored in increasing backup-level order beginning with level 0 (for example, 0, 1, 2).

7. To restore the fifth and ninth backups from a single-volume, multiple-backup tape, type the following command:

```
restore -xvqs 5 -f/dev/rmt0.1  
restore -xvqs 4 -f/dev/rmt0.1
```

The first command extracts all files from the fifth archive on the multiple-backup tape specified by `/dev/rmt0.1`. The `.1` designator specifies that the tape device will not be retensioned when it is opened and that it will not be rewound when it is closed. It is necessary to use a no-rewind-on-close, no-retension-on-open tape device because of the behavior of the `-s` flag. The second command extracts all the files from the fourth archive (relative to the current location of the tape head on the tape). After the fifth archive has been restored, the tape read/write head is in a position to read the archive. Since you want to extract the ninth archive on the tape, you must specify a value of 4 with the `-s` flag. This is because the `-s` flag is relative to your position on the tape and not to an archive's position on the tape. The ninth archive is the fourth archive from your current position on the tape.

8. To restore the fourth backup, which begins on the sixth tape on a 10-tape multiple-backup archive, put the sixth tape into the tape drive and type the following command:

```
restore -xcs 2 -f /dev/rmt0.1 /home/mike/manual/chap3
```

Assuming the fourth backup is the second backup on the sixth tape, specifying `-s 2` advances the tape head to the beginning of the second backup on this tape. The `restore` command then restores the specified file from the archive. If the backup continues onto subsequent volumes and the file has not been restored, the `restore` command instructs you to insert the next volume until the end of the backup is reached. The `-f` flag specifies the no-rewind, no-retension tape device name.

Note: The `-s` flag specifies the backup number relative to the tape inserted in the tape drive, not to the overall 10-tape archive.

9. To improve the performance on streaming tape devices, pipe the `dd` command to the `restore` command and type the following command:

```
dd if=/dev/rmt0 bs=64b | restore -xf- -b64
```

The `dd` command reads the archive from the tape using a block size of 64 512-byte blocks and writes the archive to standard output. The `restore` command reads the standard input using a block size of 64 512-byte blocks. The value of the block size used by the `dd` command to read the archive from the tape must be an even multiple of the block size that was used to create the tape with the `backup` command. For example, the following `backup` command could be used to create the archive that this example extracts:

```
find /home -print | backup -ivqf/dev/rmt0 -b64
```

This example applies to archives in file name format only. If the archive was in file system format, the `restore` command should include the `-B` flag.

10. To improve the performance of the `restore` command on 9348 Magnetic Tape Unit Model 12, you can change the block size by typing the following command:

```
chdev -l DeviceName -a BlockSize=32k
```

11. To restore nonsparse database files, type the following command:

```
restore -xef /dev/rmt0
```

12. To restore files that were archived as sparse, type the following command:

```
restore -xf /dev/rmt0
```

13. To restore only the permissions of the files from the archive, type the following command:

```
restore -Pa -vf /dev/rmt0
```

14. To restore only the ACL attributes of the files from the archive, type the following command:

```
restore -Pc -vf /dev/rmt0
```

15. To view the table of contents along with the file permissions, type the following command:

```
restore -Ta -vf /dev/rmt0
```

16. To view the table of contents of a file name archive along with the time stamps and file permissions, type the following command:

```
restore -Tl -vf /dev/rmt0
```

17. To view the table of contents of a file system archive along with the time stamps and file permissions, type the following command:

```
restore -t1 -vf /dev/rmt0
```

Related Information

The **backup** command.

restorevgstruct command

Purpose

Restores the user volume group.

Syntax

```
restorevgstruct { -ls | -vg VolumeGroupLabel [ DiskName ... ] }
```

Description

The **restorevgstruct** command restores the structure of a previously saved user volume group. If the **-ls** flag is specified, a list of previously-saved volume groups and the date each volume group was saved is displayed. This command does not work on rootvg.

Flags

DiskName...	Specifies the names of disk devices to be used instead of the disk devices saved in the volume group structure. Target disk devices must be defined as empty physical volumes; that is, they must contain a physical volume identifier and must not belong to a volume group. If the target disk devices belong to a volume group, they must be removed from the volume group using the reducevg command. Also, the target disk device cannot be a physical volume that is assigned to a shared memory pool (to be used as a paging space device by a shared memory partition).
-ls	Displays a list of previously saved volume groups.
-vg	Specifies the name of the VolumeGroup to restore.

Exit Status

See “Exit status for Virtual I/O Server commands” on page 3.

Examples

1. To restore the volume group **myvg**, onto the **hdisk2** and **hdisk3** disks, enter:

```
restorevgstruct myvg hdisk2 hdisk3
```

2. To list all previously saved volume groups, enter:

```
restorvgstruct -ls
```

The message generated would be similar to:

```
-rw-r--r-- 1 root    system    51200 Jun 18 10:53 myvg.data
-rw-r--r-- 1 root    system    51200 Jun 18 10:53 myvg2.data
```

Related Information

The **activatevg** command, the **savevgstruct** command, the **chvg** command, the **deactivatevg** command, the **exportvg** command, the **extendvg** command, the **importvg** command, the **reducevg** command, the **lsvg** command, the **mkvg** command, and the **syncvg** command.

rmauth command

Purpose

Removes one or more user-defined authorizations.

Syntax

```
rmauth [-h]
```

Description

The **rmauth** command removes the user-defined authorization identified by the *auth* parameter. The command only removes existing user-defined authorizations in the authorization database. You cannot remove system-defined authorizations with this command. If an authorization is being referenced in the privileged command database, it cannot be removed until the authorization is no longer referenced by the database.

By default, the **rmauth** command only attempts to remove the specified authorization from the authorization database. You must remove authorizations from the lowest level of a hierarchy before the higher level can be removed. If you specify a higher level authorization and lower-level authorizations still exist, the command fails. To remove a hierarchy of authorizations, specify the **-h** flag. With the **-h** flag, any lower-level authorization beneath the specified authorization is also removed. If any of the lower level authorizations is being referenced in the privileged command database, no authorizations are removed and the entire operation fails.

Important: You cannot remove any authorizations that start with “**vios.**” or “**aix.**”

Note: After removing authorizations, you must run the **setkst** command to update the kernel database with the new authorizations.

Security

The **rmauth** command is a privileged command. To run the command successfully, assume a role that has the authorization **vios.security.auth.remove**.

Flags

-h Allows removal of a hierarchy of authorizations.

Exit Status

0 for success.

Non-zero for command error.

Examples

1. To remove the *custom.test* authorization, type the following command:

```
rmauth custom.test
```
2. To remove the custom authorization and all of its children authorizations, type the following command:

```
rmauth -h custom
```

rmbdsp command

Purpose

Removes a backing device or VIOS objects and returns the storage back to the shared storage pool.

Syntax

To remove backing device:

```
rmbdsp [-sp StoragePool] -bd BackingDevice | -vtd VirtualTargetDevice [-savebd][-f]
```

To remove logical units of the shared storage pool:

```
rmbdsp -clustername ClusterName -sp StoragePool {-bd LogicalUnit | -luudid LUUDID | -all }
```

Description

The **rmbdsp** command removes a backing device from a virtual SCSI server adapter by removing its associated virtual target device. By default the backing device is also removed and its storage returned to the storage pool. If the **-savebd** flag is included then the backing device is not removed. The backing device can be identified by either specifying the name, **-bd**, or the virtual target device, **-vtd**. If the backing device is identified by the name, **-bd**, then the storage pool the device is contained within must also be specified, **-sp**. If the storage pool is not given the default storage pool is assumed.

This command also removes logical units within a specified shared storage pool. In this case, the cluster name and the storage pool must be specified. If **-bd** option is specified, then it will remove the specified logical unit. If the **-all** option is specified, it will remove all the logical units within the specified shared storage pool.

Note:

- If the specified logical unit (LU) is not unique, then the system lists all the LUs with the same name and the user has to pass the **LU UDID**, which needs to be removed.
- A storage pool must be specified if the backup device name is a file backing device or a VIOS object. The default storage pool is not assumed.
- The logical volume specified for the **-bd** option cannot be assigned to a shared memory pool because the shared memory pool must be used as a paging space device by a shared memory partition.

If **-vtd** flag is given only the specified virtual target device is removed. If **-bd** flag is given all virtual target devices associated with the backing device is removed.

Flags

-all	Specifies to remove all the logical units belonging to a storage pool.
-bd	Specifies the backing device to be removed.
-clustername	Specifies the cluster name.
-f	Forces the file-backed disk meta data file to be removed. If -savebd is also specified then the backing device will not be removed.

CAUTION:

Use this flag only when the metadata file associated with the file-backed disk is corrupted.

-luudid	Specifies the LU UDID in case the specified logical unit name is not unique.
-savebd	Indicates the backing device should not be deleted.

-sp	Specifies the storage pool rather than utilizing the default storage pool.
-vtd	Specifies the virtual target device that maps a backing device to a server virtual SCSI adapter.

Exit Status

25 Specified logical volume is not a backing device.

Examples

1. To remove the virtual target device `vtscsi4` and not remove the backing device associated with it, type the following:

```
rmbdsp -vtd vtscsi4 -savebd
```

2. To remove a logical object within a shared storage pool, type the command as follows:

```
rmbdsp -clustername newcluster -sp viossp -bd LU
```

The system displays the following output:

Specified LU is not unique. Please select the LU UDID from the below list.

```
rmbdsp -clustername newcluster -sp SP -luudid 5fc34cf4172d7cb4143fdea2a9477811
```

Type the command as follows:

```
rmbdsp -clustername newcluster -sp viossp -luudid 5fc34cf4172d7cb4143fdea2a9477811
```

The system displays the following output:

Logical unit ""with udid "5fc34cf4172d7cb4143fdea2a9477811" is removed.

3. To remove all the logical units within a shared storage pool, type the command as follows:

```
rmbdsp -clustername newcluster -sp SP -all
```

rmdev command

Purpose

Removes a device from the system.

Syntax

```
rmdev { -dev | -pdev } Name [ -recursive ] [ -ucfg ]
```

Description

The **rmdev** command unconfigures and undefines the device specified with the device logical name. If you specify the **-recursive** flag, the **rmdev** command acts on any children of the device as well. By specifying the **-ucfg** flag the device will be unconfigured but not undefined.

Use the **-pdev** flag along with the parent device's logical name to delete all of the children devices. The children are deleted in the same recursive fashion as described above for the **-recursive** flag. The only difference is that the specified device itself is not deleted. Thus, the **-recursive** flag is redundant and need not be specified with the **-pdev** flag.

Attention: To protect the Configuration database, the **rmdev** command cannot be interrupted. Stopping this command before execution is complete could result in a corrupted database.

Note: The device specified or the children of the device specified cannot be a physical or logical volume that is assigned to a shared memory pool (to be used as a paging space device by a shared memory partition).

Flags

-dev <i>Name</i>	Specifies the logical device, indicated by the <i>Name</i> parameter. This flag may not be used with the -pdev flag.
-pdev <i>Name</i>	Specifies the parent logical device (indicated by the <i>Name</i> parameter whose children need to be removed. This flag may not be used with the -dev flag.
-recursive	Unconfigures the device and its children.
-ucfg	Unconfigures, but does not undefine, the specified device. The device's state will be moved from Available to Defined. To move the device back to Available state run <code>cfgdev -dev <i>Name</i></code>

Exit Status

See “Exit status for Virtual I/O Server commands” on page 3.

Examples

1. To unconfigure the **cd0** CD-ROM device, type the command as follows:

```
rmdev -dev cd0
```

2. To unconfigure the SCSI adapter **scsi1** and all of its children, type the command as follows:

```
rmdev -recursive -dev scsi1
```

3. To unconfigure just the children of the SCSI adapter **scsi1**, but not the adapter itself, type the command as follows:

```
rmdev -pdev scsi1
```

4. To unconfigure the children of PCI bus **pci1** and all other devices under them, type the command as follows:

```
rmdev -pdev pci1
```

5. To remove the suspended virtual adapter and all of its children, type the command as follows:

```
rmdev -dev susadpt0 -recursive
```

The system displays the following output:

```
vtscsi0          deleted
client_rootvg    deleted
susadpt0         deleted
```

Related Information

The **cfgdev** command, the **chdev** command, the **lsdev** command, the **mkvdev** command, and the **rmdev** command.

rmlv command

Purpose

Removes logical volumes from a volume group.

Syntax

```
rmlv [ -f ] LogicalVolume ...
```

Description

The `rmlv` command removes a logical volume. The *LogicalVolume* parameter can be a logical volume name or logical volume ID.

Attention: This command destroys all data in the specified logical volumes. The specified logical volume cannot be assigned to a shared memory pool (to be used as a paging space device by a shared memory partition).

Flags

`-f` Removes the logical volumes without requesting confirmation.

Exit Status

See “Exit status for Virtual I/O Server commands” on page 3.

Examples

1. To remove logical volume `lv05` without requiring user confirmation, enter the following command:

```
rmlv -f lv05
```

The logical volume is removed from the volume group.

Related Information

The `mklv` command, the `extendlv` command, and the `lslv` command.

rmlvcopy command

Purpose

Removes a copy of a logical volume.

Syntax

```
rmlvcopy LogicalVolume [ PhysicalVolume ... ]
```

Description

The `rmlvcopy` command removes one of the copies (disabling mirroring) of the logical volume. The *LogicalVolume* parameter can be a logical volume name or logical volume ID. The *PhysicalVolume* parameter can be the physical volume name or the physical volume ID. If the *PhysicalVolume* parameter is used, then only the copy from that physical volume will be removed.

Note: If the LVM (Logical Volume Manager) has not recognized that a disk has failed, it is possible that the LVM will remove a different mirror. Therefore, if you know that a disk has failed and the LVM does not show those disks as missing, you should specify the failed disks on the command line.

Exit Status

See “Exit status for Virtual I/O Server commands” on page 3.

Examples

To remove mirroring from the logical volume **lv0112**, type:

```
rm1vcopy lv0112
```

Related Information

The **mk1v** command, the **extend1v** command, the **rm1v** command, and the **ls1v** command.

rm1v command

Purpose

Removes from the system a path to an MPIO-capable device.

Syntax

```
rm1v { [ -dev Name ] [ -pdev Parent ] [ -conn Connection ] } [ -rm ]
```

Description

The **rm1v** command unconfigures, and possibly undefines, one or more paths associated with the specified target device (**-dev** *Name*). The set of paths that is removed is determined by the combination of the **-dev** *Name*, **-pdev** *Parent* and **-conn** *Connection* flags. If the command will result in all paths associated with the device being unconfigured or undefined, the command will exit with an error and without unconfiguring or undefining any path. In this situation, **rmdev** command must be used instead to unconfigure or undefine the target device itself.

The default action unconfigures each specified path, but does not completely remove it from the system. If the **-rm** flag is specified, the **rm1v** command unconfigures (if necessary) and removes, or deletes, the path definition(s) from the system.

When the **rm1v** command finishes, it displays a status message. When unconfiguring paths, it is possible for this command to be able to unconfigure some paths and not others (for example, paths that are in the process of doing I/O cannot be unconfigured).

The **rm1v** command provides status messages about the results of operation. Messages in one of the following formats will be generated:

path [defined | deleted]

This message is displayed when a single path was successfully unconfigured or undefined. If the path is successfully configured the message path available displays. If the path is not successfully configured and there is no explicit error code returned by the method, the message path defined displays.

paths [defined | deleted]

This message is displayed if multiple paths were identified and all paths were successfully unconfigured or undefined. If the **-rm** flag is not specified, the message would be paths defined. If the **-rm** flag is specified, the message would be paths deleted.

some paths [defined | deleted]

This message is display if multiple paths were identified, but only some of them were successfully unconfigured or undefined. If the **-rm** flag is not specified, the message would be some paths defined. If the **-rm** flag is specified, the message would be some paths deleted.

no paths processed

This message is generated if no paths were found matching the selection criteria.

Flags

-rm	Indicates that the specified paths are to be deleted from the system.
-dev <i>Name</i>	Specifies the logical device name of the target device whose path is to be removed. The paths to be removed are qualified via the -pdev and -conn flags.
-pdev <i>Parent</i>	Indicates the logical device name of the parent device to use in qualifying the paths to be removed. Since all paths to a device cannot be removed by this command, either this flag, the -conn flag, or both must be specified.
-conn <i>Connection</i>	Indicates the connection information to use in qualifying the paths to be removed. Since all paths to a device cannot be removed by this command, either this flag, the -pdev flag, or both must be specified.

Exit Status

See “Exit status for Virtual I/O Server commands” on page 3.

Examples

1. To unconfigure the path from **scsi0** to **hdisk1** at connection 5,0, type:

```
rmpath -dev hdisk1 -pdev scsi0 -conn "5,0"
```

The message generated would be similar to:

```
path defined
```

2. To unconfigure all paths from **scsi0** to **hdisk1**, type:

```
rmpath -dev hdisk1 -pdev scsi0
```

If all paths were successfully unconfigured, the message generated would be similar to:

```
paths defined
```

However, if only some of the paths were successfully unconfigured, the message would be similar to:

```
some paths defined
```

3. To undefine the path definition between **scsi0** and **hdisk1** at connection 5,0, type:

```
rmpath -rm -dev hdisk1 -pdev scsi0 -conn "5,0"
```

The message generated would be similar to the following:

```
path deleted
```

4. To unconfigure all paths from **scsi0** to **hdisk1**, type:

```
rmpath -rm -dev hdisk1 -pdev scsi0
```

The message generated would be similar to:

```
paths deleted
```

Related Information

The **chpath** command, the **lspath** command, and the **rmpath** command.

rmrep command

Purpose

Remove the Virtual Media Repository

Syntax

```
rmrep [-f ]
```

Description

The **rmrep** command removes the Virtual Media Repository from the Virtual I/O Server. If there are any Virtual Target Devices associated with any virtual optical media in the repository the command will fail. If virtual optical media exist in the repository, but there are no Virtual Target Devices associated with them, the command will fail unless the **-f** flag is specified.

Flags

-f Forces the repository to be removed even if it contains virtual optical media.

Examples

To remove the Virtual Media Repository, type the following command:

```
rmrep
```

rmrole command

Purpose

Removes a role.

Syntax

```
rmrole role
```

Description

The **rmrole** command removes the role identified by the *role* parameter from the roles database. The role name must already exist.

Security

The **rmrole** command is a privileged command. To run the command successfully, assume a role that has the authorization **vios.security.role.remove**.

Exit Status

0 for success.

Non-zero for command error.

Examples

1. To remove the *ManageObjects* role, type the following command:

```
rmrole ManageObjects
```
2. To remove the *ManageRoles* role, type the following command:

```
rmrole ManageRoles
```

rmsecattr command

Purpose

Removes the definition of the security attributes for a command, a device, or a privileged file in the database.

Syntax

```
rmsecattr {-c | -d | -f } Name
```

Description

The **rmsecattr** command removes the security attributes for a command, a device, or a file entry identified by the *Name* parameter from the appropriate database. The command interprets the *Name* parameter as either a command or a device based on whether the **-c** (command), **-d** (device), or **-f** (privileged file) flag is specified.

If you specify the **-c** flag, the *Name* parameter must include the full path to the command and the command must at that time have an entry in the **/etc/security/privcmds** privileged command database. If you specify the **-d** flag, the *Name* parameter must include the full path to the device and the device must at that time have an entry in the **/etc/security/privdevs** privileged device database. If you specify the **-f** flag, the *Name* parameter must include the full path to the file and the file must have an entry in the **/etc/security/privfiles** privileged file database.

Important: You can remove the security attributes for your own commands and scripts, but you cannot remove security attributes for VIOS or AIX commands.

Security

The **rmsecattr** command is a privileged command. To run the command successfully, assume a role that has the following authorizations:

vios.security.cmd.remove

Required to list the attributes of a command with the **-c** flag.

vios.security.device.remove

Required to list the attributes of a device with the **-d** flag.

vios.security.file.remove

Required to list the attributes of a device with the **-f** flag.

Flags

-c	When used with the <i>Name</i> parameter, it specifies the full paths to one or more commands on the system that have entries in the privileged command database.
-d	When used with the <i>Name</i> parameter, it specifies the full paths to one or more devices on the system that have entries in the privileged device database.
-f	When used with the <i>Name</i> parameter, it specifies the full path to a privileged file on the system.
-R	-R load_module specifies the loadable module to use for the deletion of the <i>Name</i> entry.

Exit Status

0 for success.

Non-zero for command error.

Examples

1. To remove the `/usr/sbin/mytest` command from the privileged command database, type the following command:

```
rmsecattr -c /usr/sbin/mytest
```
2. To remove the `/dev/mydev` device from the privileged device database, type the following command:

```
rmsecattr -d /dev/mydev
```
3. To remove the `/etc/testconf` file from the privileged file database, type the following command:

```
rmsecattr -f /etc/testconf
```

rmfsp command

Purpose

Removes a file storage pool.

Syntax

```
rmfsp [ -f StoragePool ]
```

Description

The **rmfsp** command removes the specified storage pool from the Virtual I/O Server. This command only works for file storage pools. To remove a logical volume storage pool, use the **chsp** command to remove all physical volumes from the pool. If the pool contains any files, the **-f** flag must be specified or the command will fail. Also, any virtual target devices associated with files in the pool must be removed prior to running the **rmfsp** command.

Flags

-f Forces the pool to be removed even if it contains files.

Example

To remove the storage pool `Client_Data`, type the following command:

```
rmfsp Client_Data
```

IVM rmsyscfg command

Purpose

Removes a logical partition from the managed system. This command is operable only in an Integrated Virtualization Manager environment.

Syntax

```
rmsyscfg -r lpar { -n PartitionName | --id PartitionID }
```

Description

The **rmsyscfg** command removes a logical partition from the managed system.

Flags

-r <i>ResourceType</i>	The type of resources to remove:
-m <i>ManagedSystem</i>	lpar - Logical partition resources The name of the managed system. This attribute is optional because there is only one system to manage. The name may either be the user-defined name for the managed system, or be in the form <code>ttt-mmm*sssssss</code> , where <code>ttt</code> is the machine type, <code>mmm</code> is the model, and <code>sssssss</code> is the serial number of the managed system.
-n <i>PartitionName</i>	The name of the partition which you want to remove.
--id <i>PartitionID</i>	The partition ID of the partition to remove.

Exit Status

This command has a return code of zero on success.

Security

This command is not accessible by users with the ViewOnly role.

Examples

1. To delete a partition with an lpar ID of 3, type:
`rmsyscfg -r lpar --id 3`
2. To delete a partition a name of lp3, type:
`rmsyscfg -r lpar -n lp3`

Related Information

The `lssyscfg` command, the `mksyscfg` command, and the `chsyscfg` command.

rmtcpip command

Purpose

Removes the Virtual I/O Server TCP/IP configuration.

Syntax

```
rmtcpip [-f] [-nextboot] {-all | [-hostname] [-routing] [-interface ifnameList]}
```

```
rmtcpip [-f] {-all | [-namesrv] [-hostname] [-routing] [-interface ifnameList]}
```

Description

The `rmtcpip` command removes Virtual I/O Server TCP/IP settings, such as name server information, network interface configuration, routing information, and host name.

Flags

-all	Removes all TCP/IP settings, effectively resetting it to a newly installed system.
-f	Executes the command with out prompting for user confirmation.
-interface	Removes TCP/IP configuration from listed interfaces.
-hostname	Resets the hostname to <code>ioserver</code> .
-namesrv	Removes DNS information and clears the hosts file.

-nextboot	Removes the specified information from the configuration files, leaving the current network parameters intact (all except DNS information and hosts file).
-routing	Removes static routing tables.

Examples

1. To remove all Virtual I/O Server TCP/IP configuration, type:

```
rmtcpip -all
```

Answer yes when prompted
2. To unconfigure a network interface en0 without confirmation, type:

```
rmtcpip -f -interface en0
```
3. To cleanup the static routing table, type:

```
rmtcpip -f -routing
```
4. To remove IP information from a network interface on the next boot, keeping the current configuration running execute:

```
rmtcpip -f -interface en0 -nextboot
```

rmuser command

Purpose

Removes a user account.

Syntax

```
rmuser [-ldap] [-rmdir] Name
```

Description

The **rmuser** command removes the user account that is identified by the *Name* parameter. This command removes a user account attribute. If you specify the **rmdir** flag, this command also removes the user's home directory and files.

Flags

-rmdir	Removes the specified user's home directory and files.
	Attention: This will delete all data stored in this user account's home directory.
-ldap	Identifies the user as an LDAP user account.

Exit Status

See "Exit status for Virtual I/O Server commands" on page 3.

Security

This command can only be run by the prime administrator (padmin) user.

Examples

1. To remove user account **haarlem**, type the following command:

```
rmuser haarlem
```

2. To remove the user account and home directory of user account **emmen**, type the following command:

```
rmuser -rmdir emmen
```

Attention: This will delete all data stored in this user account's home directory.

Related Information

The **chuser** command, the **lsuser** command, the **mkuser** command, and the **passwd** command.

rmvdev command

Purpose

To remove the connection between a physical device and its associated virtual SCSI adapter.

Syntax

```
rmvdev [ -f ] { -vdev TargetDevice | -vtd VirtualTargetDevice } [-rmlv]
```

Description

The **rmvdev** command removes the connection between a physical device and its associated virtual SCSI adapter. The connection can be identified by specifying the backing (physical) device or the virtual target device. If the connection is specified by the device name and there are multiple connections between the physical device and virtual SCSI adapters and error is returned unless the **-f** flag is also specified. If **-f** is included then all connections associated with the physical device are removed.

If the backing (physical) device is a logical volume and the **-rmlv** flag is specified, then logical volume will be removed as well.

Flags

-f	Forces the removal of all virtual target devices associated with the give backing device.
-vdev <i>TargetDevice</i>	Specifies the physical backing device
-rmlv	Deletes the backing device. This flag is valid only for logical volume backing devices.
-vtd <i>VirtualTargetDevice</i>	Specifies the virtual target device to remove.

Exit Status

See “Exit status for Virtual I/O Server commands” on page 3.

Examples

1. To remove the connection between a physical volume and its associated virtual SCSI adapter by specifying the physical volume, type the following command:

```
rmvdev -vdev hdisk4
```
2. To remove the connection between a logical volume and its associated virtual SCSI adapter by specifying the logical volume and then to remove the logical volume after the connection is removed, type the following command:

```
rmvdev -vdev lv1 -rmlv
```
3. To remove the connection between a backing device and its associated virtual SCSI adapter by specifying the virtual target device, type the following command:

```
rmvdev -vtd vtscsi0
```


4. To remove multiple connections between a physical volume and virtual SCSI adapters by specifying the physical volume, type the following command:

```
rmvdev -vdev hdisk0 -f
```

rmvfilt command

Purpose

Removes the virtual LAN-crossing filter rules from the filter table.

Syntax

```
rmvfilt -n [fid | all> ]
```

Description

The **rmvfilt** command is used to remove the virtual LAN-crossing filter rules from the filter table.

Flags

-n Specifies the ID of the filter rule that will be removed. The **all** option is used to remove all the filter rules.

Exit Status

This command returns the following exit values:

0	Successful completion.
>0	An error occurred.

Examples

1. To remove all the filter rules or to deactivate all the filter rules in the kernel, type the command as follows:

```
rmvfilt -n all
```

Related Information

The **chvfilt** command, the **genvfilt** command, the **lsvfilt** command, the **mkvfilt** command, and the **vlanfw** command.

rmvlog command

Purpose

Removes the virtual log and virtual log device.

Syntax

To unconfigure the virtual log device that has the specified device name or the virtual log UUID:

```
rmvlog {-dev DeviceName | -uuid UUID}
```

To remove the virtual log device that has the specified device name or virtual log UUID:

rmvlog -d {-dev *DeviceName* | -uuid *UUID*}

To remove the virtual log and any virtual log device that has the specified device name or the virtual log UUID:

rmvlog -db {-dev *DeviceName* | -uuid *UUID*}

To remove the virtual log, log data, and any virtual log device that has the specified device name or virtual log UUID:

rmvlog -dbdata {-dev *DeviceName* | -uuid *UUID*}

Description

You can use the **rmvlog** command to detach the virtual logs from virtual SCSI (VSCSI) host adapters, leaving the virtual log intact. You can also delete the virtual log, optionally including all user data in it, by using this command.

Flags

-d	Removes the device definition.
-db	Removes the virtual log.
-dbdata	Removes the virtual log and associated data in the repository.
-dev	Specifies the name of the virtual log device to be removed.
-uuid	Specifies the UUID of an existing virtual log to be removed.

Exit Status

Table 7. Command-specific return codes

Return code	Description
0	All files were written successfully.
>0	An error occurred.

Examples

1. To remove virtual log device *vtlog0*, type the command as follows:

```
rmvlog -dev vtlog0 -d
```

The system displays the output as follows:

```
vtlog0 deleted
```

2. To remove the virtual log with UUID *9705340b31a7883573a1cd04b2254efd*, leaving the data intact, type the command as follows:

```
rmvlog -uuid 9705340b31a7883573a1cd04b2254efd -db
```

The system displays the output as follows:

```
Virtual log 9705340b31a7883573a1cd04b2254efd deleted.
```

3. To remove the virtual log device *vtlog1*, the virtual log, and the virtual log data, type the command as follows:

```
rmvlog -dev vtlog0 -dbdata
```

The system displays the output as follows:

```
vtlog0 deleted
Virtual log 02392437473b6c552680a9ddd2fd8d06 deleted.
Log files deleted.
```

Related Information

The **chvlog** command, the **chvrepo** command, the **lsvlog** command, the **lsvlrepo** command, and the **mkvlog** command.

rmvopt command

Purpose

Remove a virtual optical media disk from the Virtual Media Repository.

Syntax

```
rmvopt [-f ] -name FileName
```

Description

The **rmvopt** command removes the specified media from the Virtual Media Repository. If the media is currently loaded into one or more virtual optical devices the command will fail, unless the **-f** flag is specified.

Flags

-f	Force the media to be removed even if it is loaded into one or more virtual optical devices.
-name <i>FileName</i>	The name of the virtual optical media to remove.

Examples

To remove the file `clientData` from the Virtual Media Repository, type the following command:

```
rmvopt -name clientData
```

rmvt command

Purpose

Closes a virtual terminal connection to a partition. This command is valid only in an Integrated Virtualization Manager environment.

Syntax

```
rmvt { -id lparID }
```

Description

The **rmvt** command closes a virtual terminal connection to the target logical partition. To close the virtual terminal session normally, enter the `~.` character sequence in the terminal window.

This command requires additional HMC configuration if used in an HMC environment.

Flags

-id *lparID* The ID of the partition for which to close the virtual terminal session.

Exit Status

29 Virtual terminal device is not found

Security

This command is not accessible by users with the ViewOnly role.

Examples

1. To close a virtual terminal connection to the partition with ID 3, type:

```
rmvt -id 3
```

Related Information

The **mkvt** command.

rolelist command

Purpose

Displays role information for a user or process.

Syntax

```
Rolelist [-a][-e | -u username | -p PID]
```

Description

The **rolelist** command provides role and authorization information to invokers about their current roles or the roles assigned to them. If no flags or arguments are specified, the **rolelist** command displays the list of roles assigned to the invoker with the text description of each role if one is provided in the roles database. Specifying the **-e** flag provides information about the current effective active role set for the session. If the invoker is not currently in a role session and specifies the **-e** flag, no output is displayed. Specifying the **-a** flag displays the authorizations associated with the roles instead of the text description.

The **rolelist** command also permits a privileged user to list the role information for another user or for a process. Specifying a user name with the **-u** flag allows a privileged user to list the roles assigned to another user. The active role set of a given user cannot be determined because the user can have multiple active role sessions. Therefore, if the **-u** flag is specified, the **-e** flag is not allowed. Specifying a process ID with the **-p** flag allows a privileged user to display the roles associated with a process. The command fails immediately if called by a nonprivileged user when the **-u** or **-p** flag is specified.

The authorization information displayed by the **rolelist** command is retrieved from the kernel security tables. The information can differ with the current state of the roles database if it is modified after the kernel security tables are updated.

Security

The **rolelist** command is a privileged command. To run the command successfully, assume a role that has the authorization **vios.security.role.list** or **vios.security.proc.role.list**.

Flags

-a	Displays the authorizations assigned to each role instead of the role description.
-e	Displays information about the effective active role set of the session.
-u	Displays role information for the specified user.
-p	Displays role information of the specified process.

Exit Status

0 for success.

Nonzero for command error.

Examples

1. To display the list of roles that is assigned to you and their text descriptions, type the following command:
`rolelist`
2. To display the authorizations associated with the assigned roles, type the following command:
`rolelist -a`
3. As a privileged user, type the following command to display the roles assigned to a specific user:
`rolelist -u user1`

IVM `rsthwres` command

Purpose

Restores hardware resources. This command is operable only in an Integrated Virtualization Manager environment.

Syntax

```
rsthwres -r io | mem | proc | virtualio | mempool [-m <managed system>] [-p <partition name> | --id <partition ID>] [-l <DRC index>] [-s <virtual slot number>] [-a "<attributes>"] [--help]
```

```
rsthwres -rhea | -m <managed system>] [-p <partition name> | --id <partition ID>] [-l<HEA DRC index>] [-g <port group ID>] --logport <logical port ID>]--help]
```

To restore memory pool resources

```
rsthwres -r mempool [-a "<attributes>"] [--help]
```

Description

The `rsthwres` command restores the hardware resource configuration in the managed system. This operation may need to be performed after a dynamic LPAR operation fails.

You also need to perform this operation to restore memory pool resources in the following situations:

- When you reinstall the VIOS base code and applied fix packs from original VIOS media.
- When you restore the VIOS from backup media prior to the time you created the memory pool, but the firmware still has a memory pool configured.

Flags

-r	Shows the type of hardware resources to restore: io I/O slot (physical) hea Host Ethernet Adapter mem Memory mempool Memory pools proc Processing virtualio Virtual I/O
-m <managed system>	Indicates the name of the managed system which has the partitions for which to restore the hardware resources. The name may either be the user-defined name for the managed system, or be in the form <i>ttt-mmm*sssssss</i> , where <i>ttt</i> is the machine type, <i>mmm</i> is the model, and <i>sssssss</i> is the serial number of the managed system.
-p <partition name>	Indicates the name of the partition for which to restore the hardware resources. To restore hardware resources for a single partition, you must either use this option to specify the name of the partition, or use the --id option to specify the partition's ID. Otherwise, hardware resources for all partitions in the managed-system will be restored. The -p and the --id options are mutually exclusive.
--id <partition ID>	Indicates the ID of the partition for which to restore the hardware resources. To restore hardware resources for a single partition, you must either use this option to specify the ID of the partition, or use the -p option to specify the partition's name. Otherwise, hardware resources for all partitions in the managed-system will be restored. The --id and the -p options are mutually exclusive.
-l <DRC index>	Indicates the DRC index of the physical I/O slot to restore. This option is only valid when restoring physical I/O slots.
-s <virtual slot>	This option is also used to specify the physical Host Ethernet Adapter to restore.
-a "<attributes>"	Indicates the slot number of the virtual I/O. The configuration data needed to set hardware resource related attributes. The configuration data consists of attribute name/value pairs, which are in comma separated value (CSV) format. The configuration data must be enclosed in quotation marks. Possible values are: paging_storage_pool Note: <ul style="list-style-type: none">• If you intend to use a paging storage pool, you must specify the value here. After you run the rsthwres command, you cannot set or change the paging storage pool value without deleting your memory pool and recreating it.• The -a parameter is valid only with -r mempool, and the only supported attribute is paging_storage_pool.
-g port group ID	Indicates the logical ports that are part of the specified port group that will be restored.
--logport <virtual slot>	Indicates the ID of the logical port to recover.
--help	Displays the help text for this command and exit.

Exit Status

This command has a return code of 0 on success.

Examples

To restore the physical I/O slots for all partitions, type the following command:

```
rsthwres -r io -m mySystem
```

To restore the physical I/O slot with DRC index **21010003**, type the following command:

```
rsthwres -r io -m 9406-570*12345678 -l 21010003
```

To restore memory resources for partition **p1**, type the following command:

```
rsthwres -r mem -m 9406-570*12345678 -p p1
```

To restore processing resources for the partition with ID **1**, type the following command:

```
rsthwres -r proc -m mySystem --id 1
```

To clean up all recoverable Host Ethernet Adapter resources, type the following command:

```
rsthwres -r hea
```

To clean up a specific logical port with a given partition, type the following command:

```
rsthwres -r hea -l <HEA DRC> -g <PORT GROUP>  
--logport <LP ID> {-p <LPAR NAME | --id <LPAR ID>}
```

To recover a memory pool using rootvg as the paging storage pool, type the following command:

```
rsthwres -r mempool -a paging_storage_pool=rootvg
```

Note: If you intend to use a paging storage pool, you must specify the value here. After you run the **rsthwres** command, you cannot set or change the paging storage pool value without deleting your memory pool and recreating it.

IVM rstprofdata command

Purpose

Restores profile data. This command is operable only in an Integrated Virtualization Manager environment.

Syntax

To restore logical partition configuration data from a file:

```
rstprofdata -l RestoreType [-f RestoreFile] [ --ignoremtms ] [ --ignoremac ] [ --ignorehea ] [ -m  
ManagedSystem ]
```

Description

The **rstprofdata** command performs a restore of logical partition configuration information from a file which was created with the **bkprofdata** command. Logical partition configuration must be cleared using the **lpcfgop** command before performing a restore operation. The restore operation may require a restart of the managed system, after which the **rstprofdata** command should be run again with the same parameters as were used the first time (before the system restart).

Flags

-l <i>RestoreType</i>	The type of restore to perform. Valid options are: 1 - full restore from the backup file
-f <i>RestoreFile</i>	The name of the file to read from in the current working directory. If not specified, the default file is <code>"/var/adm/lpm/profile.bak"</code>
--ignoremtms	Do not fail the restore operation if the type, model, or serial number of the managed system does not match the values in the backup file. This flag is used for the following primary reasons: <ul style="list-style-type: none">• To clone the configuration from one system to another.• To force new worldwide port names to be allocated for any logical partitions that use virtual fibre channel. Note: If you use the -ignoremtms flag, new worldwide port names are generated.
--ignoremac	Do not try to restore the virtual Ethernet MAC addresses from the backup file. Default to the automatically generated MAC addresses.
--ignorehea	Do not try to restore the Host Ethernet Adapter resources.
-m <i>ManagedSystem</i>	The name of the managed system. This attribute is optional because there is only one system to manage. The name may either be the user-defined name for the managed system, or be in the form <code>ttt-mmm*sssssss</code> , where <code>ttt</code> is the machine type, <code>mmm</code> is the model, and <code>sssssss</code> is the serial number of the managed system.
--help	Display the help text for this command and exit.

Exit Status

This command has a return code of 0 on success.

Security

This command is not accessible by users with the ViewOnly role.

Examples

1. To restore the partition configuration data from `/var/adm/lpm/profile.bak`, type:

```
rstprofdata -l 1
```
2. To restore the partition configuration data from `lparData.bak` without validating that the type, model, and serial number match, type:

```
rstprofdata -l 1 -f lparData.bak --ignoremtms
```

Related Information

The `bkprofdata` command.

save_base command

Purpose

Saves information about base-customized devices in the Device Configuration database for the boot device.

Syntax

```
save_base [-path Directory][-file File][-verbose]
```


Description

The **save_base** command stores customized information for base devices to use during phase 1 of system boot. By default, the **save_base** command retrieves this information from the `/etc/objrepos` directory. However, you can override this action by using the **-o** flag to specify an ODM directory. The **save_base** command is typically run without any parameters. It uses the `/dev/ipl_blv` special file link to identify the output destination.

Alternatively, use the **-d** flag to specify a destination file or a device, such as the `/dev/hdisk0` device file. To identify a specific output destination, the **-d** flag identifies the file to which **save_base** writes the base customized device data. This file can be either a regular file or a device special file. The device special file identifies either a disk device special file or a boot logical volume device special file.

A disk device special file can be used where there is only one boot logical volume on the disk. The **save_base** command ensures that the given disk has only one boot logical volume present. If neither of these conditions is true, **save_base** does not save the base customized device data to the disk and exits with an error.

When a second boot logical volume is on a disk, the boot logical volume device special file must be used as the destination device to identify which boot image the base customized device data will be stored in. A boot logical volume device special file can be used even if there is only one boot logical volume on the disk. The **save_base** command ensures that the given device special file is a boot logical volume and it is bootable before saving any data to it. If either of these checks fails, **save_base** exits with an error.

The **save_base** command determines what device information to save using the `PdDv.base` field corresponding to each entry in the `CuDv` object class. Specifically, the `PdDv.base` field is a bit mask which represents the type of boot for which this device is a base device. The **save_base** command determines the current type of boot by accessing the `boot_mask` attribute in the `CuAt` object class. The value of this attribute is the bit mask to apply to the `PdDv.base` field to determine which devices are base.

Note:

- Base devices are those devices that get configured during phase 1 boot; they may vary depending on the type of boot (mask). For example, if the mask is `NETWORK_BOOT`, network devices are considered base; for `DISK_BOOT`, disk devices are considered base. The type of boot masks are defined in the `/usr/include/sys/cfgdb.h` file.
- The **-m** flag is no longer used by the **save_base** command. For compatibility reasons, the flag can be specified, but is not used.

Flags

-path <i>Directory</i>	Specifies a directory containing the Device Configuration database. Note: By default, the <code>save_base</code> command retrieves this information from the <code>/etc/objrepos</code> directory.
-file <i>file</i>	Specifies the destination file or device to which the base information will be written.
-verbose	Specifies that verbose output is to be written to standard output (STDIN).

Exit Status

See “Exit status for Virtual I/O Server commands” on page 3.

Examples

1. To save the base customized information and see verbose output, type the following command:

```
save_base -verbose
```

2. To specify an ODM directory other than the `/usr/lib/objrepos` directory, type the following command:
`save_base -o /tmp/objrepos`
3. To save the base customized information to the `/dev/hdisk0` device file instead of to the boot disk, type the following command:
`save_base -d /dev/hdisk0`

savevgstruct command

Purpose

Backs up a volume group.

Syntax

```
savevgstruct VolumeGroupLabel
```

Description

The `savevgstruct` command will make a backup of a volume group structure.

Exit Status

See “Exit status for Virtual I/O Server commands” on page 3.

Examples

1. To save the structure of the user defined volume group `myvg`, enter:
`savevgstruct myvg`

Related Information

The `activatevg` command, the `restorevgstruct` command, the `chvg` command, the `deactivatevg` command, the `exportvg` command, the `extendvg` command, the `importvg` command, the `lsvg` command, the `mkvg` command, and the `syncvg` command.

seastat command

Purpose

Generates a report to view, per client, Shared Ethernet Adapter statistics.

Syntax

To display Shared Ethernet Adapter statistics, per client.

```
seastat -d Shared Ethernet Adapter device name [-n]
```

To clear all Shared Ethernet Adapter statistics that have been gathered per client.

```
seastat -d Shared Ethernet Adapter device name -c
```

Description

The `seastat` command generates a report to view, per client, shared ethernet adapter statistics. To gather network statistics at a per-client level, advanced accounting can be enabled on the Shared Ethernet

Adapter to provide more information about its network traffic. To enable per-client statistics, the VIOS administrator can set the Shared Ethernet Adapter accounting attribute to enabled. The default value is disabled. When advanced accounting is enabled, the Shared Ethernet Adapter keeps track of the hardware (MAC) addresses of all of the packets it receives from the LPAR clients, and increments packet and byte counts for each client independently. After advanced accounting is enabled on the Shared Ethernet Adapter, the VIOS administrator can generate a report to view per-client statistics by running the **seastat** command.

Note: Advanced accounting must be enabled on the Shared Ethernet Adapter before the **seastat** command can print any statistics.

To enable advanced accounting on the SEA, enter the following command:

```
chdev -dev <SEA device name> -attr accounting=enabled
```

Flags

-d <i>Shared Ethernet Adapter device name</i>	Specifies the device name of the Shared Ethernet Adapter.
-n	Disables name resolution on the IP addresses.
-c	Clears all of the per client Shared Ethernet Adapter statistics that have been gathered.

Exit Status

The following exit values are returned:

0	Successful completion.
>0	Invalid flag, argument, or command failure

Examples

1. To display Shared Ethernet Adapter statistics for **sea0**, type:

```
seastat -d sea0
```
2. To clear the Shared Ethernet Adapter statistics for **sea0**, type:

```
seastat -d sea0 -c
```

setkst command

Purpose

Sets the entries in the kernel security tables (KST).

Syntax

```
setkst [-q][-b | -t table1, table2, ...]
```

Description

The **setkst** command reads the security databases and loads the information from the databases into the kernel security tables. By default, all of the security databases are sent to the KST. Alternatively, you can specify a specific database using the **-t** flag. If the authorization database is the only one you specified, the role and privileged command databases are updated in the KST because they are dependent on the authorization database.

The **setkst** command checks the tables before updating the KST. If any severe error in the database is found, the **setkst** command warns the user by sending a message to the stderr stream, and exits without

resetting the KST. If a minor error is found in the database, a warning message is displayed, and the entry is skipped.

Security

The **setkst** command is a privileged command. To run the command successfully, assume a role that has the authorization **vios.security.kst.set**.

Flags

-b	Loads the KST with the information that is stored in the backup binary file on the system. If information in the binary file cannot be loaded, the tables are regenerated from the security databases.
-q	Specifies quiet mode. Warning messages that occur are not displayed when the security databases are parsed.
-t <i>table1, table2</i>	Sends the specified security databases to the KST. The parameter for the -t flag is a comma-separated list of security databases. Valid values follow: auth Authorizations database role Role database cmd Privileged command database dev Privileged device database

Exit Status

0 for success.

Non-zero for command error.

Examples

1. To send all of the security databases to the KST, type the following command:

```
setkst
```
2. To send the role and privileged command databases to the KST, type the following command:

```
setkst -t role,cmd
```

setsecattr command

Purpose

Sets the security attributes of a command, a device, a privileged file, or a process.

Syntax

```
setsecattr { [-c | -d | -p | -f ] Attribute=Value [ Attribute =Value ...] Name
```

Description

The **setsecattr** command sets the security attributes of the command, device, or process that is specified by the *Name* parameter. The command interprets the *Name* parameter as either a command, a device, a privileged file, or a process based on whether the **-c** (command), **-d** (device), **-f** (privileged file), or **-p** (process) flag is specified.

To set a value for an attribute, specify the attribute name and the new value with the *Attribute=Value* parameter. To clear an attribute, specify the *Attribute=* for the *Attribute=Value* pair. To make incremental

changes to attributes, whose values are lists, specify the *Attribute=Value* pairs as *Attribute+=Value*, or *Attribute=-Value*. If you specify the *Attribute+=Value*, the value is added onto the existing value for the attribute. If you specify the *Attribute=-Value*, the value is removed from the existing value for the attribute.

Important: You can set security attributes for your own commands and scripts, but you cannot set security attributes for VIOS or AIX commands. To identify the privileges required to run a user's command or script, use the **tracepriv** command.

Note: When security attributes are set for a command or device, run the **setkst** command to update the kernel database.

Security

The **setsecattr** command is a privileged command. To run the command successfully, assume a role that has the following authorizations:

vios.security.cmd.set

Required to modify the attributes of a command with the **-c** flag.

vios.security.device.set

Required to modify the attributes of a device with the **-d** flag.

vios.security.file.set

Required to modify the attributes of a device with the **-f** flag.

vios.security.proc.set

Required to modify the attributes of a process with the **-p** flag.

Flags

- | | |
|-----------|---|
| -c | Specifies that the security attributes of a command on the system are to be set. If the command name that you specified by using the <i>Name</i> parameter is not in the privileged command database, a command entry is created in the <code>/etc/security/privcmds</code> privileged command database. If an attribute is being cleared and is the only attribute set for the command, the command is removed from the privileged command database. |
| -d | Specifies that the security attributes of a device on the system are to be set. If the device name that you specified by using the <i>Name</i> parameter is not in the privileged device database, a device entry is created in the <code>/etc/security/privdevs</code> privileged device database. If an attribute is being cleared and is the only attribute set for the device, the device is removed from the privileged device database. |
| -f | Specifies that the security attributes of a privileged file on the system are to be set. Changes requested through the <i>Attribute=Value</i> pairs are made in the <code>/etc/security/privfiles</code> privileged file database. If the specified file is not in the privileged file database, a file entry is created in the database. If an attribute is being cleared and is the only attribute set for the command, the command is removed from the privileged file database. |
| -p | Specifies that the numeric process identifier (PID) of an active process on the system is to be set. Changes that you specify with the <i>Attribute=Value</i> pairs immediately affects the state of the specified active process. Modifications are not saved in a database. |

Attributes

Attribute=Value

Sets the value of a security attribute for the object. The list of valid attribute names is dependent on the object type as specified by using the **-c**, **-d**, **-p**, and **-f** flags.

Use the following attributes for the privileged command database (**-c**) flag:

accessauths

Specifies access authorizations, which includes a comma-separated list of authorization names. You can specify a total of sixteen authorizations. A user with any of the authorizations that you specified can run the command. This attribute has these special additional values: **ALLOW_OWNER**, **ALLOW_GROUP**, and **ALLOW_ALL** that allows a command owner, a group, or all users to run the command without checking for access authorizations.

authprivs

Specifies authorized privileges, which includes a list of authorizations and privilege pairs that grant additional privileges to the process. The authorization and its corresponding privileges are separated by an equal sign (=), individual privileges are separated by a plus sign (+), and authorization or privilege pairs are separated by a comma (,), as shown in the following examples:

```
auth=priv+priv+...,auth=priv+priv+.....
```

You can specify a maximum of sixteen pairs of authorizations or privileges. Specifies roles, the users of which need to be authenticated before a command can be executed successfully. Specifies a comma separated list of roles. Each role should be authenticated by different users, such as no user can perform the authentication for more than one role at a time.

authroles

Specifies the user roles that need to be authenticated before the command can run successfully. If listing multiple roles, separate each role with a comma. For example:

```
authroles=so,isso
```

Each role must be authenticated by different users. For example, no one user can perform the authentication for more than one role.

innateprivs

Specifies the innate privileges, which includes a comma-separated list of privileges that are assigned to the process when the command is run.

inheritprivs

Specifies inheritable privileges, which includes a comma-separated list of privileges that are passed to child processes.

eid Specifies the effective user ID to assume when the command is run.

egid Specifies the effective group ID to assume when the command is run.

ruid Specifies the real user ID to assume when the command is run. The only valid value is 0. This attribute value is ignored if the command provides access to all users by specifying the special value **ALLOW_ALL** in its **accessauths** attribute.

secflags

Specifies the file security flags, which includes a comma-separated list of security flags. The possible value follows:

FSF_EPS

Causes the maximum privilege set to be loaded into the effective privilege set when the command is run.

Use the following attributes for the privileged device database (**-d**) flag:

readprivs

Specifies a comma-separated list of privileges that a user or a process must have for read access to the device. You can specify a maximum of eight privileges. The user or process must have one of the listed privileges to read from the device.

writeprivs

Specifies a comma-separated list of privileges that a user or a process must have for write access to the device. You can specify a maximum of eight privileges. The user or process must have one of the listed privileges to write to the device.

Use the following attributes for the privileged file (**-f**) flag:

readauths

Specify the read access authorizations, which includes a comma-separated list of authorization names. A user with any of the authorizations can read the file.

writeauths

Specify the write access authorizations, which includes a comma-separated list of authorization names. A user with any of the authorizations can read or write the file.

Use the following attributes for the privileged process (**-p**) flag:

eprivs Specify the effective privilege set, which includes a comma-separated list of privileges that are to be active for the process. The process might remove the privileges from this set and add the privileges from the maximum privilege set to its effective privilege set.

iprivs Specifies the inheritable privilege set, which includes a comma-separated list of privileges that are passed to child processes' effective and maximum privilege sets. The inheritable privilege set is a subset of the limiting privilege set.

mprivs Specify a maximum privilege set, which includes a comma-separated list of privileges that the process can add to its effective privilege set. The maximum privilege set is a superset of the effective privilege set.

lprivs Specify the limiting privilege set, which includes a comma-separated list of privileges that make up the maximum possible privilege set for a process. The limiting privilege set is a superset of the maximum privilege set.

uprivs Specify the used privilege set, which includes a comma-separated list of privileges that are used during the life of the process. This set is mainly used by the **tracepriv** command.

Name

Specify the object to modify. The *Name* parameter is interpreted according to the flags that you specify.

Exit Status

0 for success.

Non-zero for command error.

Examples

1. To set an authorized privilege pair for the user's binary *mybinary*, type the following command:

```
setsecattr -c authprivs=custom.user.mybin=PV_FS_MOUNT /home/padmin/mybinary
```
2. To incrementally add the PV_AU_WRITE and PV_DAC_W privileges to the existing set of writing privileges for the /dev/mydev device, type the following command:

```
setsecattr -d writeprivs=+PV_AU_WRITE,PV_DAC_W /dev/mydev
```

3. To set a read authorization for the `/etc/security/user` file, type the following command:
`setsecatr -f readauths=vios.security.user.change /etc/security/user`
4. To incrementally remove the `PV_DAC_R` privilege from the effective privilege set of an active process, type the following command:
`setsecatr -p eprivs=-PV_DAC_R 35875`

showmount command

Purpose

Displays a list of exported directories.

Syntax

`showmount Host`

Description

The `showmount` command displays a list of all exported directories from a specified machine in the *Host* parameter.

Parameters

Host Host name of the systems to display exported directories.

Exit Status

See “Exit status for Virtual I/O Server commands” on page 3.

Examples

1. To display all exported directories on the host **middelburg**, type:
`showmount middelburg`

Related Information

The `mount` command and the `unmount` command.

shutdown command

Purpose

Ends system operation.

Syntax

`shutdown [-force] [-restart]`

Description

The `shutdown` command halts the operating system. When the shut down is complete, you receive a shutdown completion message.

Attention: Do not attempt to restart the system or turn off the system before the shutdown completion message is displayed; otherwise, file system damage can result.

The **-force** flag is used to bypass the following user prompt:
"Shutting down the VIOS could affect Client Partitions. Continue [y|n]?"

Flags

-force	Begins a system shut down without a user prompt.
-restart	Restarts the system after being shut down.

Exit Status

See "Exit status for Virtual I/O Server commands" on page 3.

Security

This command can only be run by the prime administrator (padmin).

Examples

1. To shut down the system, type the following command:
shutdown
2. To restart the system, type the following command:
shutdown -restart
3. To shut down the system and restart, without the user prompt, type the following command:
shutdown -force -restart

snap command

Purpose

Gathers system configuration information.

Syntax

```
snap [-general] [-dev DeviceName]
```

```
snap script1 "script2 arg1 arg2" ...
```

Description

The **snap** command gathers system configuration information and compresses the information into a pax file (snap.pax.Z). The file can then be transmitted to a remote system. The information gathered with the **snap** command may be required to identify and resolve system problems.

At least 8 MB of temporary disk space is required to collect all system information, including contents of the error log.

Flags

-dev <i>OutputDevice</i>	Copies the compressed image onto the specified device.
-general	Gathers general system information which is a subset of the system configuration information

Exit Status

See "Exit status for Virtual I/O Server commands" on page 3.

Examples

1. To gather all system configuration information, enter the following command:

```
snap
```

The output of this command is written to the users home directory.

2. To gather general system configuration information, including the output of the **lspp -hBc** command, enter the following command :

```
snap -general -dev /dev/rfd0
```

3. To run the scripts foo1, foo2 and foo3. where foo1 takes no argument, foo2 takes three arguments and foo3 takes one argument, type the following"

```
snap foo1 "foo2 -x -y 3" "foo3 6" foo4
```

Files

snap.pax.Z

snapshot command

Purpose

Creates, deletes, or rolls back a snapshot image file of a single logical unit (LU) or multiple LUs.

Syntax

```
snapshot -clustername ClusterName -create <filename | ""> -spname StoragePool {-lu LU_Names | -luudid LUUIDIDs}
```

```
snapshot -clustername ClusterName -delete <filename>-spname StoragePool {-lu LU_Names | -luudid LUUIDIDs}
```

```
snapshot -clustername ClusterName -rollback <filename> -spname StoragePool {-lu LU_Names | -luudid LUUIDIDs}
```

```
snapshot -clustername ClusterName -list -spname StoragePool [-lu LuName | -luudid LUUIDID]
```

Description

The **snapshot** command can create, delete, or roll back a snapshot image file of a single LU or multiple LUs. The **create** flag generates a snapshot image that can either have a user-supplied name or with a system-generated name. The **delete** flag is used to delete a previously generated snapshot image. The **rollback** option is used to roll back the snapshot file to a previously generated version of a snapshot image.

Note: You must synchronize the virtual disk of the client before taking a snapshot. If the rolled back LU is a rootvg, the client must be shut down. If the rolled back LU is a datavg, all the volume groups present in the virtual disk must be varied off.

Flags

-clustername Specifies the cluster name.

- create** *SnapshotName* Specifies the name of the snapshot file to be created. This argument can be left blank, in which case a system-generated name is used.
- delete** *SnapshotName* Specifies the name of the snapshot file to be deleted. This argument is required with the **delete** flag.
- list** Lists the snapshot within the LU.
- lu** *LUnames* Specifies the LU names that are contained in the snapshot file.
- luudid** *luudid* Specifies the LU-UDID for this LU.
- rollback** *SnapshotName* Specifies the name of the snapshot file to be rolled back. This argument is required with the rollback flag.
- spname** *StoragePool* Specifies the storage pool name.

Examples

1. To create a snapshot with a single LU and a user-supplied snapshot name, type the command as follows:

```
snapshot -clustername mycluster -spname mypool -create bob -lu lu_A
```

The system displays the output as follows:

```
bob
```

2. To create a snapshot image file that contains a single LU, with a system-generated name, type the command as follows:

```
snapshot -clustername mycluster -spname mypool -create -lu lu_A
```

The system displays the output as follows:

```
SS_487532_2010-05-08_01:23:23
```

3. To create a snapshot that contains multiple LUs with a system-generated name, type the command as follows:

```
snapshot -clustername mycluster -spname mypool -create -lu lu_A lu_B
```

The system displays the output as follows:

```
SS_487555_2010-05-09_01:23:45
```

4. To delete a previously generated snapshot image called *lu_image01*, type the command as follows:

```
snapshot -clustername mycluster -spname mypool -delete lu_image01 -lu lu_A
```

5. To roll back to a previously generated snapshot image called *lu_image02*, type the command as follows:

```
snapshot -clustername mycluster -spname mypool -rollback lu_image02 -lu lu_A
```

6. To list all snapshots, type the command as follows:

```
snapshot -clustername mycluster -list -spname mypool
```

The system displays the output as follows:

LuName	Size(mb)	ProvisionType	Lu Udid
LUA	1024	THIN	c49cf79726dc42ed3787c878bf4fd30c
Snapshot			
snap1			
snap2			
LUB	1024	THIN	7f6c2fed80c77bad1f309de16f68abff
Snapshot			
snap1			
snap2			
LUC	1024	THIN	dbc88fd787e41be55480f63244ffc99
Snapshot			
snap1			
LUD	1024	THIN	387ca4f5eb73be0ac74e436be3c78557
Snapshot			

<u>LuName</u>	<u>Size(mb)</u>	<u>ProvisionType</u>	<u>Lu Udid</u>
snap2			

snmp_info command

Purpose

Requests or modifies values of Management Information Base (MIB) variables managed by a Simple Network Management Protocol (SNMP) agent.

Syntax

snmp_info [-mode get | next] [-verbose] [-com *Community*] [-debug *Level*] [-host *HostName*] [-file *ObjectFile*] **-retry** *Tries*] [-wait*WaitTime*] [*Variable Instance*]

The following syntax is for the set option:

snmp_info -mode set [-verbose] [-com *Community*] [-debug *Level*] [-host *HostName*] [-file *ObjectFile*] **-retry** *Tries*] [-wait*WaitTime*] *Variable Instance*=*Value*

The following syntax is for the dump option:

snmp_info -mode dump[-verbose] [-com *Community*] [-debug *Level*] [-host *HostName*] [-file *ObjectFile*] **-retry** *Tries*] [-wait*WaitTime*] [*Variable Instance*]

Description

The **snmp_info** command requests or modifies values for one or more Management Information Base (MIB) variables managed by a Simple Network Management Protocol (SNMP) agent. This command may only be issued by a user with root privileges or by a member of the system group.

If you specify the get option, the **snmp_info** command requests information about one or more MIB variables from an SNMP agent.

If you specify the next option, the **snmp_info** command requests information from an SNMP agent about the instances following the specified instances. The next option makes it possible to obtain MIB values without knowledge of the instance qualifiers.

If you specify the set option, the **snmp_info** command modifies values for one or more MIB variables for an SNMP agent. Only a few MIB variables are designated read-write. The agent that manages the MIB database may take various actions as a side effect of modifying MIB variables. For example, setting the `ifAdminStatus` MIB variable to 2 will normally shut down a network interface. The action taken is determined by the implementation of the SNMP agent that manages the database.

If you specify the dump option, the **snmp_info** command can be used to traverse the entire MIB tree of a given agent. If a group is passed in as the `Variable` parameter, the **snmp_info** command will traverse that specified path of the MIB tree.

The **snmp_info** command has a debug facility that will dump debug information for transmitted and received packets. The facility is enabled with the **-d** flag.

Flags

-com <i>Community</i>	Specifies the community name to be used to query. If you do not specify this flag, the default community name is public.
-debug <i>Level</i>	Specifies the level of I/O debug information. Use one of the following values: 0 No debug information. 1 Port bindings and the number of bytes transmitted and received. 2 Level 1 plus a hexadecimal dump of incoming and outgoing packets. 3 Level 2 plus an English version of the request and response packets. If you do not specify this flag, the default debug level is 0.
-host <i>HostName</i>	Specifies the host name of the SNMP agent to be queried. If you do not specify this flag, the default host name is the host name of the system on which the user is currently logged.
-file <i>ObjectFile</i>	Specifies the name of the objects definition file. If you do not specify this flag, the default objects definition file name is /etc/mib.defs.
-mode <i>Option</i>	Specifies the mode by which to access the MIB variables. The value can be one of the following options: get Requests information about the specified MIB variables next Requests the instances following the specified instances set Modifies the specified write access MIB variables dump Dumps the specified section of the MIB tree Note: 1. The option name can be specified by the minimum number of characters required to make it unique. 2. If you do not specify this flag, the default mode is get.
-retry <i>Tries</i>	Specifies the number of times the snmp_info command transmits the SNMP request to the SNMP agent before ending with the message no SNMP response. If you do not specify this flag, the default number of tries is 3.
-verbose	Specifies that the output from the snmp_info command be displayed in verbose mode. If you do not specify this flag, the information will not be displayed in verbose mode.
-wait	Specifies the wait time in seconds for the response from the snmpd agent. If you do not specify this flag, the default wait time is 15 seconds.

Parameters

Value	Specifies the value to which the MIB Variable parameter is to be set. A value must be specified for each variable. If a value is not specified, the request packet will be invalid.
Variable	Specifies the name in text format or numeric format of a specific MIB variable as defined in the /etc/mib.defs file. If the option to the -m flag is next or dump, the Variable parameter may be specified as a MIB group.
Instance	Specifies the instance qualifier for the MIB Variable parameter. The Instance parameter is required if the option to the -m flag is get or set. The Instance parameter is optional if the option to the -m flag is next or dump. Note: 1. There should be no blank spaces in the <i>Variable.Instance</i> parameter sequence. 2. If the Instance parameter is not specified, do not place a period after the Variable parameter.

For further information, consult RFC 1213, which defines the Management Information Base (MIB) for network management, and RFC 1157, which defines the SNMP protocol for creating requests for MIB information and formatting responses.

Exit Status

See “Exit status for Virtual I/O Server commands” on page 3.

Examples

To display the current system information and snmp configuration information, type the following command:

```
snmp_info -mode dump system
```

Related Information

The `snmp_trap` command and the `cl_snmp` command.

snmp_trap command

Purpose

The `snmp_trap` command generates a notification (trap) to report an event to the SNMP manager with the specified message.

Syntax

```
snmp_trap [-debug] [-host HostName] [-com Community] [-target TargetHost] [-msg message]
```

Description

The `snmp_trap` command generates a notification (trap) to report an event to the SNMP manager with the specified message.

Flags

<code>-host</code> <i>HostName</i>	Specifies to connect to the SNMP agent on the specified host. If you do not specify this flag, the default host is the local host. Host can be an IP address or a host name.
<code>-com</code> <i>Community</i>	Specifies the community name to use. This community must have been set in the <code>/etc/snmpdv3.conf</code> file for SNMP version 3 file. You also have to have read-access privilege to at least the SNMP agent running on the specified host or local host. If you do not specify this flag, the default community name is "public".
<code>-debug</code>	Enables the debug facility.
<code>-msg</code> <i>Message</i>	Defines the message that the snmptrap command will send. This value specifies the information the trap will hold. This information is in the text format. You must order this flag as the last one in the list when you specify this command.
<code>-target</code> <i>TargetHost</i>	Specifies the target network-manger host to which the trap message is sent. It is different from <code>-host</code> flag. If you do not specify the <code>-host</code> and <code>-target</code> flags, the trap is sent to the VIOS SNMP agent on the local host.

Exit Status

See “Exit status for Virtual I/O Server commands” on page 3.

Examples

To display the current system information and snmp configuration information, type the following command:

```
snmp_trap -mode hello world
```

Related Information

The `snmp_trap` command and the `cl_snmp` command.

snmpv3_ssw command

Purpose

Switch the symbolic links among the non-encrypted snmpdv3 agent, encrypted snmpdv3 agent, and the snmpdv1 agent.

Syntax

```
snmpv3_ssw [-e | -n | -1 ]
```

Description

The `snmpv3_ssw` command switches the symbolic links among the non-encrypted snmpdv3 agent and snmpdv1 agent. The `snmpv3_ssw` command then starts the newly chosen SNMP agent. You can choose which version of SNMP agent to run.

Flags

<code>-e</code>	Switches to the encrypted version of snmpdv3 agent.
<code>-n</code>	Switches to the non-encrypted version of snmpdv3 agent.
<code>-1</code>	Switches to the snmpdv1 agent.

Exit Status

Return code	Description
0	Success
1	Invalid flag, argument, or command failure

Examples

To switch to the encrypted version of snmpdv3 agent, enter the following command:

```
snmp3_ssw -e
```

startnetsvc command

Purpose

Starts the ndpd-host, telnet, ftp, xntpd, ssh, snmp, ldap, or cimserver daemon.

Syntax

```
startnetsvc [NetworkService][TracingSelection]
```

Description

The `startnetsvc` command starts the ndpd-host, telnet, ftp, xntpd, ssh, snmp, ldap or cimserver daemon. By enabling the network service daemon, users can use that service to connect to the Virtual I/O Server.

Parameters

NetworkService

Use the following values:

ndpd-host

Enables the **ndpd-host** daemon.

telnet Enables the **telnet** daemon.

ftp Enables the **ftp** daemon.

xntpd Enables the **xntpd** daemon.

ssh Enables the **ssh** daemon.

snmp Enables the **snmp** daemon.

ldap Enables the **ldap** daemon.

cimserver

Enables the **cimserver** daemon.

ALL Enables all service daemons.

Specifying **ALL** starts all services, but does not enable the **tracelog** or **errorlog** options.

TraceSelection

Use the following values:

tracelog

Sends CLI tracing information to the system log.

Attention: Enabling the **tracelog** or **errorlog** can make the system log grow quickly.

errorlog

Sends system error log information to the system log.

Attention: Enabling the **tracelog** or **errorlog** can make the system log grow quickly.

Note: A padmin user can use the **vi** command to view files sent to the system log.

Exit Status

See “Exit status for Virtual I/O Server commands” on page 3.

Examples

1. To start the **telnet** daemon, type the following command:

```
startnetsvc telnet
```
2. To start the **ftp** daemon, type the following command:

```
startnetsvc ftp
```
3. To start the **tracelog** option, type the following command:

```
startnetsvc tracelog
```
4. To start all service daemons, type the following command:

```
startnetsvc ALL
```
5. To start the **ndpd-host** daemon, type the following command:

```
startnetsvc ndpd-host
```

This command produces output similar to the following:

```
0513-059 The ndpd-host Subsystem has been started. Subsystem PID is 356522.
```


Related Information

The **mktcpip** command, the **hostname** command, the **stopnetsvc** command, the **cfglnagg** command, the **netstat** command, the **entstat** command, the **cfgnamesrv** command, the **hostmap** command, the **traceroute** command, the **ping** command, and the **optimizenet** command.

startsvc command

Purpose

Starts the agent or service that is specified by the given name.

Syntax

```
startsvc AgentName
```

Description

The **startsvc** command starts the specified agent or service. Use the **lssvc** command to obtain a list of valid agent or service names. The agent determines the operations that occur while the start command is running. If you attempt to restart an agent that is already started, you will receive a message indicating that the agent has already been started.

Note: For the TotalStorage Productivity Center (TPC) agent, the **lssvc** command displays only the valid agent that is used by the **cfgsvc** command during configuration.

Agent or service names

The following agents can be managed by the Virtual I/O Server.

DIRECTOR_agent	Starts the Systems Director agent.
ITM_premium	Starts the specified Tivoli Monitoring agent.
ITM_cec	The ITM_premium and ITM_cec agents are Tivoli Monitoring agents. These agents provide system information including I/O virtualization mappings and system utilization.
TSM_base	Starts the Tivoli Storage Manager agent.
ITUAM_base	Starts the Tivoli Usage and Accounting Manager agent.
TPC	Starts the TotalStorage Productivity Center agent.
	The TPC_fabric and TPC_data agents are valid TotalStorage Productivity Center agents for the startsvc command.
perfmgr	Starts the performance management data collector agent.
ipsec_tunnel	Creates secure tunnels between local and remote Virtual I/O Servers that are configured with the cfgsvc command. Note: The clirc.rte file set has to be installed before you can create tunnels.
ILMT	Starts the IBM License Metric Tool (ILMT) agent on the Virtual I/O Server.

Exit Status

See “Exit status for Virtual I/O Server commands” on page 3.

Examples

1. To start the **ITM_premium** agent, type the command as follows:
startsvc ITM_premium
2. To start the **ITUAM** base agent, type the command as follows:

```
startsvc ITUAM_base
```

3. To start the perfmgr agent, type the command as follows:

```
startsvc perfmgr
```

4. To start secure tunnels, type the command as follows:

```
startsvc ipsec_tunnel
```

IBM License Metric Tool

1. To start the IBM License Metric Tool agent, type the command as follows:

```
startsvc ILMT
```

Related Information

The **cfgsvc** command, the **lssvc** command, and the **stopsvc** command.

For more information about the various agents, see the following information:

- Tivoli software and the Virtual I/O Server
- Configuring the Tivoli agents and clients on the Virtual I/O Server
- Systems Director software
- Configuring the Systems Director agent

startsysdump command

Purpose

Starts a kernel dump to the primary dump device.

Syntax

```
startsysdump
```

Description

The **startsysdump** command provides a command line interface to start a kernel dump to the primary dump device. Any previous kernel dumps will be erased before the dump is created. During a kernel dump, the following values can be displayed on the three-digit terminal display as follows. The user will be required to run the **snap** command to obtain the system dump.

0c0	Indicates that the dump completed successfully.
0c1	Indicates that an I/O occurred during the dump.
0c2	Indicates that the dump is in progress.
0c4	Indicates that the dump is too small.
0c5	Indicates a dump internal error.
0c6	Prompts you to make the secondary dump device ready.
0c7	Indicates that the dump process is waiting for a response from the remote host.
0c8	Indicates that the dump was disabled. In this case, no dump device was designated in the system configuration object for dump devices. The startsysdump command halts, and the system continues running.
0c9	Indicates that a dump is in progress.
0cc	Indicates that the system switched to the secondary dump device after attempting a dump to the primary device.

Note: When the dump completes, the system reboots.

Exit Status

See “Exit status for Virtual I/O Server commands” on page 3.

Examples

1. To start a kernel dump, type:
`startsysdump`

starttrace command

Purpose

Records selected system events.

Syntax

```
starttrace [ -event Event[, Event ] ...]
```

Description

The **starttrace** command starts the **trace** daemon which configures a trace session and starts the collection of system events. The data collected by the trace function is recorded in the trace log. A report from the trace log can be generated with the **cattracerpt** command.

Flags

-event *Event*[,*Event*]

Specifies the user-defined events for which you want to collect trace data. The Event list items should be separated by commas.

Note: The following events are used to determine the pid, the cpuid and the exec path name in the **cattracerpt** report:

```
106 DISPATCH  
10C DISPATCH IDLE PROCESS  
134 EXEC SYSTEM CALL  
139 FORK SYSTEM CALL  
465 KTHREAD CREATE
```

If any of these events is missing, the information reported by the **cattracerpt** command will be incomplete. When using the **-event** flag, you should include all these events in the *Event* list.

Exit Status

See “Exit status for Virtual I/O Server commands” on page 3.

Examples

1. To trace hook 234 and the hooks that will allow you to see the process names, enter:
`starttrace -event 234,106,10C,134,139,465`

Related Information

The **stoptrace** command and the **cattracerpt** command.

stopnetsvc command

Purpose

Disables the `ndpd-host`, `telnet`, `ftp`, `xntpd`, `ssh`, `snmp`, `ldap`, or `cimserver` daemon.

Syntax

```
stopnetsvc [NetworkService][TracingSelection]
```

Description

The **stopnetsvc** command stops the `ndpd-host`, `telnet`, `ftp`, `xntpd`, `ssh`, `snmp`, `ldap`, or `cimserver` daemon. By disabling a service daemon, users can prevent anyone from being able to connect through the associated network service. The **tracelog** or **errorlog** options stop sending CLI tracing information or system error log information to the system log.

Parameters

NetworkService

The following values can be used:

ndpd-host

Disables the **ndpd-host** daemon.

telnet Disables the **telnet** daemon.

ftp Disables the **ftp** daemon.

xntpd Disables the **xntpd** daemon.

ssh Disables the **ssh** daemon.

snmp Disables the **snmp** daemon.

ldap Disables the **ldap** daemon.

cimserver

Disables the **cimserver** daemon.

ALL Disables all service daemons, but does not disable the **tracelog** or **errorlog** option.

TraceSelection

The following values can be used:

tracelog

Stops sending CLI tracing information to the system log.

errorlog

Stops sending system error log information to the system log.

Note: A `padmin` user can use the `vi` command to view files sent to the system log.

Exit Status

See “Exit status for Virtual I/O Server commands” on page 3.

Examples

1. To disable the **telnet** daemon, type the following command:

```
stopnetsvc telnet
```
2. To disable the **ftp** daemon, type the following command:

```
stopnetsvc ftp
```
3. To disable the **ndpd-host** daemon, type the following command:

```
stopnetsvc ndpd-host
```

This command produces output similar to the following:
0513-044 The /usr/sbin/ndpd-host Subsystem was requested to stop.

4. To stop the **tracelog** option, type the following command:
stopnetsvc tracelog
5. To disable all service daemons, type the following command:
stopnetsvc ALL

Related Information

The **mktcpip** command, the **hostname** command, the **startnetsvc** command, the **cfglnagg** command, the **netstat** command, the **entstat** command, the **cfgnamesrv** command, the **hostmap** command, the **traceroute** command, the **ping** command, and the **optimizenet** command.

stopsvc command

Purpose

Stops the agent or service that is specified by the given name.

Syntax

stopsvc *AgentName*

Description

The **stopsvc** command stops the specified agent or service. Use the **lssvc** command to obtain a list of valid agent and service names. The agent determines the operations that occur during the stop command. If you attempt to stop an agent or service that is already stopped, you will receive a message indicating that the agent or service has already been stopped.

Note: For the TotalStorage Productivity Center (TPC) agent, the **lssvc** command displays only the valid agent used by the **cfgsvc** command during configuration.

Agent or service names

The following agents can be managed by the Virtual I/O Server (VIOS).

DIRECTOR_agent	Stops the Systems Director agent.
ITM_premium	Stops the specified Tivoli Monitoring agent.
ITM_cec	The ITM_premium and ITM_cec agents are Tivoli Monitoring agents. These agents provide system information including I/O virtualization mappings and system utilization.
TSM_base	Stops the Tivoli Storage Manager agent.
ITUAM_base	Stops the Tivoli Usage and Accounting Manager agent.
TPC	Stops the TotalStorage Productivity Center agent.
	The TPC_fabric and TPC_data agents are valid TotalStorage Productivity Center agents for the stopsvc command.
perfmgr	Stops the performance management data collector agent.
ipsec_tunnel	Deactivates and deletes the secure tunnels.
ILMT	Stops the IBM License Metric Tool (ILMT) agent on the VIOS.

Exit Status

See “Exit status for Virtual I/O Server commands” on page 3.

Examples

1. To stop the TPC_data agent, type the command as follows:
`stopsvc TPC_data`
2. To stop the TPC_fabric agent, type the command as follows:
`stopsvc TPC_fabric`
3. To stop the DIRECTOR_agent agent, type the command as follows:
`stopsvc DIRECTOR_agent`
4. To stop the ITM_premium agent, type the command as follows:
`stopsvc ITM_premium`
5. To stop the ITUAM_base agent, type the command as follows:
`stopsvc ITUAM_base`
6. To stop the perfmgr agent, type the command as follows:
`stopsvc perfmgr`
7. To stop and delete secure tunnels, type the command as follows:
`stopsvc ipsec_tunnel`

IBM License Metric Tool

1. To stop the IBM License Metric Tool agent, type the command as follows:
`stopsvc ILMT`

Related Information

The `lssvc` command, the `cfgsvc` command, and the `startsvc` command.

For more information about the various agents, see the following information:

- Tivoli software and the Virtual I/O Server
- Configuring the Tivoli agents and clients on the Virtual I/O Server
- Systems Director software
- Configuring the Systems Director agent

stoptrace command

Purpose

Stops the trace function.

Syntax

```
stoptrace
```

Description

The `stoptrace` command ends a trace session.

Exit Status

See “Exit status for Virtual I/O Server commands” on page 3.

Examples

1. To end a trace session, type:
stoptrace

Related Information

The **starttrace** command and the **cattracerpt** command.

svmon command

Purpose

Captures and analyzes a snapshot of virtual memory.

Syntax

Global Report

```
svmon [-G [-i Intvl [ NumIntvl ]][-z]]
```

Process Report

```
svmon[-P [pid1...pidn] [-r] [-u | -g | -v] [-ns] [-wfc] [-q[s|m|L|S]] [-t Count] [ -i Intvl [NumIntvl] ] [-l] [-j] [-z] [-m] ]
```

Segment Report

```
svmon [-S [sid1...sidn] [-r] [-u | -P | -g | -v] [-ns] [-wfc] [-g [s|m|L|S]] [-t Count] [ -i Intvl [NumIntvl] ] [-l] [-j] [-z] [-m] ]
```

Detailed Report

```
svmon [-D sid1...sidn [-b] [-q [s|m|L|S]] [ -i Intvl [NumIntvl] ]][-z]]
```

Framed Report

```
svmon [-F [fr1...frn] [-q [s|m|L|S]] [-i Intvl [NumIntvl] ]][-z] ]
```

Command Report

```
svmon [-C cmd1...cmdn [-r] [-u | -p | -g | -v] [-ns] [-wfc] [-q [s|m|L|S]] [-t Count] [ -i Intvl [NumIntvl] ] [-d] [-l] [-j] [-z] [-m] ]
```

User Report

```
svmon [-U [lognm1...lognmn] [-r] [-u | -p | -g | -v] [-ns] [-wfc] [-t Count] [ -i Intvl [NumIntvl] ] [-d] [-l] [-j] [-z] [-m] ]
```

Workload Management Class Report

```
svmon [-W [class1...classn] [-e] [-r]] [-u | -p | -g | -v] [-ns] [-wfc] [-q [s|m|L|S]] [-t Count] [-i Intvl [NumIntvl] ] [-d] [-l] [-j] [-z] [-m] ]
```

Workload Management Tier Report

svmon [-T [*tier1...tiern*] [-a *superclass*] [-x] -e [-r] [-u | -p | -g | -v] [-ns] [-wfc] [-q [*s | m | L | S*]] [-t *Count*] [-i *Intvl* [*NumIntvl*]] [-d] [-l] [-j] [-z] [-m]]

Description

Flags

If no command line flag is given, then the **-G** flag is the default.

-a <i>SupClassName</i>	Restricts the scope to the subclasses of the <i>SupClassName</i> class parameter (in the Tier report -T). The parameter is a superclass name. No list of class is supported.
-b	Shows the status of the reference and modified bits of all the displayed frames (detailed report -D). When shown, the reference bit of the frame is reset. When used with the -i flag, it detects which frames are accessed between each interval. Note: Use this flag with caution because of its impact on performance.
-C <i>Command1...CommandN</i>	Displays memory usage statistics for the processes running the command name <i>Commandnm1...CommandnmN</i> . <i>Commandnm</i> is a string. It is the exact basename of an executable file.
-d	Displays for a given entity, the memory statistics of the processes belonging to the entity (user name or command name).
-D <i>SID1...SIDN</i>	Displays memory-usage statistics for segments <i>SID1...SIDN</i> , and a detail status of all frames of each segment. Segment ids specified must be of primary segments.
-e	Displays the memory usage statistics of the subclasses of the <i>Class</i> parameter in the Workload Class report -W and in the Tier report -T . The class parameter of -W or -a needs to be a superclass name.
-f	Indicates that only persistent segments (files) are to be included in the statistics. By default all segments are analyzed.
-F [<i>Frame1...FrameN</i>]	Displays the status of frames <i>Frame1...FrameN</i> including the segments to which they belong. If no list of frames is supplied, the percentage of memory used displays.
-g	Indicates that the information to be displayed is sorted in decreasing order by the total number of pages reserved or used on paging space. This flag, in conjunction with the segment report, shifts the non-working segment at the end of the sorted list
-G	Displays a global report.
-i <i>Interval</i> [<i>NumIntervals</i>]	Instructs the svmon command to display statistics repetitively. Statistics are collected and printed every <i>Interval</i> seconds. <i>NumIntervals</i> is the number of repetitions; if not specified, svmon runs until user interruption, Ctrl-C. Note: It may take a few seconds to collect statistics for some options. The observed interval may be larger than the specified interval.
-j	Shows, for each persistent segment, the file path referred. Note: This flag should be used with caution because of its potential performance impacts (especially with svmon -S).

- l** Shows, for each displayed segment, the list of process identifiers that use the segment and, according to the type of report, the entity name (login, command, tier, or class) to which the process belongs. For special segments, a label is displayed instead of the list of process identifiers.
- System segment**
This label is displayed for segments that are flagged *system*.
- Unused segment**
This label is displayed for segments which are not used by any existing processes. For example, persistent segments relative to files no longer in usage.
- Shared library text**
This label is displayed for segments that contain text of shared library, and that may be used by most of the processes (for example, *libc.a*). This is to prevent the display of a long list of processes.
- m** Displays information about source segment and mapping segment when a segment is mapping a source segment. The default is to display only information about the mapping segment.
- n** Indicates that only non-system segments are to be included in the statistics. By default, all segments are analyzed.
- s** Indicates that only system segments are to be included in the statistics. By default, all segments are analyzed.
- p** Indicates that the information to be displayed is sorted in decreasing order by the total number of pages pinned.
- P [PID1...PIDN]** Displays memory usage statistics for process *PID1...PIDN*. PID is a decimal value. If no list of process IDs (PIDs) is supplied, memory usage statistics are displayed for all active processes.
- q [small | Large]** Filters results regarding whether they deal with pages of the requested size or not. The requested page size is specified through the option sub-argument. 4KB pages are specified with *s* and 16MB pages with *Large*. If no sub-argument is specified, the default is 16MB pages (*Large*). Metrics per page size are also displayed.
- r** Displays the ranges within the segment pages which have been allocated. A working segment may have two ranges because pages are allocated by starting from both ends and moving towards the middle.
- S [SID1...SIDN]** Displays memory-usage statistics for segments *SID1...SIDN*. SID is a hexadecimal value. Segment IDs specified must be of primary segments. If no list of segment IDs (SIDs) is supplied, memory usage statistics are displayed for all defined segments.
- t Count** Displays memory usage statistics for the top *Count* object to be printed
- T [Tier1...TierN]** Displays memory usage statistics of all the classes of the tier numbers *Tier1...TierN*. If no list of tier is supplied, memory usage statistics displayed for all the defined tiers.
- u** Indicates that the information to be displayed is sorted in decreasing order by the total number of pages in real memory. It is the default sorting criteria if none of the following flags are present: **-p**, **-g** and **-v**.
- U [LogName1...LogNameN]** Displays memory usage statistics for the login name *LogName1...LogNameN*. *LogName* is a string, it is an exact login name. If no list of login identifier is supplied, memory usage statistics are displayed for all defined login identifiers.
- v** Indicates that the information to be displayed is sorted in decreasing order by the total number of pages in virtual space. This flag, in conjunction with the segment report, shifts the non-working segment at the end of the sorted list.
- w** Indicates that only working segments are to be included in the statistics. By default all segments are analyzed.

-W [<i>Clmn1...ClmnN</i>]	Displays memory usage statistics for the workload management class <i>Clmn1...ClmnN</i> . <i>Clmn</i> is a string. It is the exact name of a class. For a subclass, the name should have the form <i>superclassname.subclassname</i> . If no list of class name is supplied, memory usage statistics display for all defined class names.
-c	Indicates that only client segments are to be included in the statistics. By default, all segments are analyzed.
-x	Displays memory usage statistics for the segments for every class of a tier in the Tier report -T .
-z	Displays the maximum memory size dynamically allocated by svmon during its execution.

Examples

To print out global statistics, type :

```
svmon -G
      size      inuse      free      pin      virtual
memory 262144    99728    162416    52690    81741
pg space 393216      706
      work      pers      clnt
pin    52690      0        0
in use 81741      0        17987
```

Related Information

Logical volume storage and device management

swrole command

Purpose

Switches to a specified role session.

Syntax

```
swrole {ALL | Role [,Role ] ... } [ Argument ... ]
```

Description

The **swrole** command creates a new role session with the roles that are specified by the *Role* parameter. The *Role* parameter must be composed of the names of roles in the *roles* attribute of the user. Before creating a new role session, the **swrole** command performs authentication according to the *auth_mode* attribute of the **chrole** command for the specified roles. If any of the specified roles requires authentication, the user must be successfully authenticated for the action to be performed. If none of the specified roles require authentication, no authentication is requested.

The **swrole** command creates a new role session with the specified roles added to the active role set of the session. The **ALL** keyword specifies that a role session is created with all the roles that are assigned to the user. Role sessions are limited to eight roles per session. If a user has more than eight roles, only the first eight roles are assigned to the role session when the **ALL** keyword is specified. Creation of a new role session preserves the user environment for the current session.

Any argument, such as a flag or a parameter, which is specified by the *Arguments* parameter, must relate to the login shell that is defined for the user. The arguments are passed to the login shell that is created for the role session.

To restore the previous session, type `exit` or press the `Ctrl-D`. The action ends the shell that was created by the `swrole` command and returns the user to the previous shell and environment.

Each time the `swrole` command is run, an entry is made in the `/var/adm/rolelog` file. The `/var/adm/rolelog` file records the following information: date, time, system name, login name, and role name. The `/var/adm/rolelog` file also records whether or not the role initiation attempt is successful: A plus sign (+) indicates a successful role initiation and a minus sign (-) indicates an unsuccessful role initiation.

Exit Status

0 for success.

Non-zero for command error.

Examples

1. To assume the `RoleAdmin` and `FSAdmin` roles as a user who has been assigned the roles, type the following command:

```
swrole - RoleAdmin,FSAdmin
```

syncvg command

Purpose

Synchronizes logical volume copies that are not current.

Syntax

```
syncvg { -lv | -pv | -vg } Name ...
```

Description

The `syncvg` command synchronizes the physical partitions, which are copies of the original physical partition, that are not current. The `syncvg` command can be used with logical volumes, physical volumes, or volume groups, with the `Name` parameter representing the logical volume name, physical volume name, or volume group name.

Unless disabled, the copies within a volume group are synchronized automatically when the volume group is activated by the `activatevg` command.

Note: For the `syncvg` command to be successful, at least one good copy of the logical volume should be accessible, and the physical volumes that contains this copy should be in `ACTIVE` state.

The `syncvg` command will check for the `NUM_PARALLEL_LPS` environment variable. The value of `NUM_PARALLEL_LPS` will be used to set the number of logical partitions to be synchronized in parallel.

Flags

<code>-lv</code>	Specifies that the <code>Name</code> parameter represents a logical volume device name.
<code>-pv</code>	Specifies that the <code>Name</code> parameter represents a physical volume device name.
<code>-vg</code>	Specifies that the <code>Name</code> parameter represents a volume group device name.

Exit Status

See “Exit status for Virtual I/O Server commands” on page 3.

Examples

1. To synchronize the copies on physical volumes **hdisk04** and **hdisk05**, type:
`syncvg -pv hdisk04 hdisk05`
2. To synchronize the copies on volume groups **vg04** and **vg05**, type:
`syncvg -vg vg04 vg05`

Related Information

The **mkvg** command, the **chvg** command, the **extendvg** command, the **lsvg** command, the **mirrorios** command, the **unmirrorios** command, the **activatevg** command, the **deactivatevg** command, the **importvg** command, the **exportvg** command, and the **reducevg** command.

sysstat command

Purpose

Displays a summary of current system activity.

Syntax

```
sysstat [ -long | -short ] [ User]
```

Description

The **sysstat** command provides the following information: time of day, amount of time since last system startup, number of users logged in and number of processes running.

Flags

-long	Prints the summary in long form. This is the default
-short	Prints the time of day, amount of time since last system startup, number of Users logged in and number of processes running.

Parameters

<i>User</i>	Specify an existing user name.
-------------	--------------------------------

topas command

Purpose

Reports selected local system statistics.

Syntax

```
topas [ -disks number_of_monitored_hot_disks ] [ -interval monitoring_interval_in_seconds ] [ -nets  
number_of_monitored_hot_network_interfaces ] [ -procs number_of_monitored_hot_processes ] [ -wlms  
number_of_monitored_hot_WLM_classes ] [ -cpus number_of_monitored_hot_CPUs ] [ -procsdisp | -wlmdisp |  
-cecdisp ] [ -filesvs number_of_monitored_file_systems ] [ -tape ] [ -fullscreen { process | WLM | lpar disk |  
net | CEC | vg | fileys | tape } ]
```

Description

The **topas** command reports selected statistics about the activity on the local system. The command displays its output in a format suitable for viewing on an 80x25 character-based display.

If the **topas** command is invoked without flags, it runs as if invoked with the following command line:
`topas -disks 20 -interval 2 -nets 20 -procs 20 -w1ms 20 -cpus 20 -filesys 2`

When you specify a single flag, the **topas** command will take that value and the other values will be the default values. This is applicable to all the flags.

The program extracts statistics from the system with an interval specified by the *monitoring_interval_in_seconds* argument. The default output, as shown below, consists of two fixed parts and a variable section. The top two lines at the left of the display show the name of the system the **topas** command runs on, the date and time of the last observation, and the monitoring interval.

The second fixed part fills the rightmost 25 positions of the display. It contains five subsections of statistics, as follows:

EVENTS/QUEUES

Displays the per-second frequency of selected system-global events and the average size of the thread run and wait queues:

Cswitch

The number of context switches per second over the monitoring interval.

Syscalls

The total number of system calls per second executed over the monitoring interval.

Reads The number of read system calls per second executed over the monitoring interval.

Writes The number of write system calls per second executed over the monitoring interval.

Forks The number of fork system calls per second executed over the monitoring interval.

Execs The number of exec system calls per second executed over the monitoring interval.

Runqueue

The average number of threads that were ready to run but were waiting for a processor to become available.

Waitqueue

The average number of threads that were waiting for paging to complete.

FILE/TTY

Displays the per-second frequency of selected file and tty statistics.

Readch The number of bytes read per second through the **read** system call over the monitoring interval.

Writech

The number of bytes written per second through the **write** system call over the monitoring interval.

Rawin The number of raw bytes read per second from TTYs over the monitoring interval.

Ttyout The number of bytes written to TTYs per second over the monitoring interval.

Igets The number of calls per second to the inode lookup routines over the monitoring interval.

Namei The number of calls per second to the pathname lookup routines over the monitoring interval.

Dirblk The number of directory blocks scanned per second by the directory search routine over the monitoring interval.

PAGING

Displays the per-second frequency of paging statistics.

Faults Total number of page faults taken per second over the monitoring interval. This includes page faults that do not cause paging activity.

Steals Physical memory 4K frames stolen per second by the virtual memory manager over the monitoring interval.

PgspIn Number of 4K pages read from paging space per second over the monitoring interval.

PgspOut

Number of 4K pages written to paging space per second over the monitoring interval.

PageIn Number of 4K pages read per second over the monitoring interval. This includes paging activity associated with reading from file systems. Subtract **PgspIn** from this value to get the number of 4K pages read from file systems per second over the monitoring interval.

PageOut

Number of 4K pages written per second over the monitoring interval. This includes paging activity associated with writing to file systems. Subtract **PgspOut** from this value to get the number of 4K pages written to file systems per second over the monitoring interval.

Sios The number of I/O requests per second issued by the virtual memory manager over the monitoring interval.

MEMORY

Displays the real memory size and the distribution of memory in use.

Real,MB

The size of real memory in megabytes.

% Comp

The percentage of real memory currently allocated to computational page frames. Computational page frames are generally those that are backed by paging space.

% Noncomp

The percentage of real memory currently allocated to non-computational frames. Non-computational page frames are generally those that are backed by file space, either data files, executable files, or shared library files.

% Client

The percentage of real memory currently allocated to cache remotely mounted files.

PAGING SPACE

Display size and utilization of paging space.

Size,MB

The sum of all paging spaces on the system, in megabytes.

% Used

The percentage of total paging space currently in use.

% Free The percentage of total paging space currently free.

NFS

Display NFS stats in calls per second

- Server V2 calls/sec
- Client V2 calls/sec
- Server V3 calls/sec
- Client V3 calls/sec

The variable part of the **topas** display can have one, two, three, four, or five subsections. If more than one subsection displays, they are always shown in the following order:

- CPU
- Network Interfaces
- Physical Disks
- WorkLoad Management Classes
- Processes

When the **topas** command is started, it displays all subsections for which hot entities are monitored. The exception to this is the WorkLoad Management (WLM) Classes subsection, which is displayed only when WLM is active.

CPU Utilization

This subsection displays a bar chart showing cumulative CPU usage. If more than one CPU exists, a list of CPUs can be displayed by pressing the **c** key *twice*. Pressing the **c** key only once will turn this subsection off. The following fields are displayed by both formats:

User% This shows the percent of CPU used by programs executing in user mode.
(Default sorted by User%)

Kern% This shows the percent of CPU used by programs executing in kernel mode.

Wait% This shows the percent of time spent waiting for IO.

Idle% This shows the percent of time the CPU(s) is idle.

When this subsection first displays the list of hot CPUs, the list is sorted by the User% field. However, the list can be sorted by the other fields by moving the cursor to the top of the desired column.

Network Interfaces

This subsection displays a list of hot network interfaces. The maximum number of interfaces displayed is the number of hot interfaces being monitored, as specified with the **-nets** flag. A smaller number of interfaces will be displayed if other subsections are also being displayed. Pressing the **n** key turns off this subsection. Pressing the **n** key again shows a one-line report summary of the activity for all network interfaces. Both reports display the following fields:

- Interf** The name of the network interface.
- KBPS** The total throughput in megabytes per second over the monitoring interval. This field is the sum of kilobytes received and kilobytes sent per second.
- I-Pack** The number of data packets received per second over the monitoring interval.
- O-Pack** The number of data packets sent per second over the monitoring interval.
- KB-In** The number of kilobytes received per second over the monitoring interval.
- KB-Out**
The number of kilobytes sent per second over the monitoring interval.

When this subsection first displays the list of hot network interfaces, the list is sorted by the KBPS field. However, the list can be sorted by the other fields by moving the cursor to the top of the desired column. Sorting is only valid for up to 16 network adapters.

Physical Disks

This subsection displays a list of hot physical disks. The maximum number of physical disks displayed is the number of hot physical disks being monitored as specified with the **-disks** flag. A smaller number of physical disks will be displayed if other subsections are also being displayed. Pressing the **d** key turns off this subsection. Pressing the **d** key again shows a one-line report summary of the activity for all physical disks. Both reports display the following fields:

- Disk** The name of the physical disk.
- Busy%** Indicates the percentage of time the physical disk was active (bandwidth utilization for the drive).
- KBPS** The number of kilobytes read and written per second over the monitoring interval. This field is the sum of **KB-Read** and **KB-Writ**.
- TPS** The number of transfers per second that were issued to the physical disk. A transfer is an I/O request to the physical disk. Multiple logical requests can be combined into a single I/O request to the disk. A transfer is of indeterminate size.
- KB-Read**
The number of kilobytes read per second from the physical disk.
- K -Writ**
The number of kilobytes written per second to the physical disk.

When this subsection first displays the list of hot physical disks, the list is sorted by the KBPS field. However, the list can be sorted by the other fields by moving the cursor to the top of the desired column. Sorting is only valid for up to 128 physical disks.

WLM Classes

This subsection displays a list of hot WorkLoad Management (WLM) Classes. The maximum number of WLM classes displayed is the number of hot WLM classes being monitored as specified with the **-wlmdisp** flag. A smaller number of classes will be displayed if other subsections are also being displayed. Pressing the **w** key turns off this subsection. The following fields are displayed for each class:

% CPU Utilization

The average CPU utilization of the WLM class over the monitoring interval.

% Mem Utilization

The average memory utilization of the WLM class over the monitoring interval.

% Blk I/O

The average percent of Block I/O of the WLM class over the monitoring interval.

When this subsection first displays the list of hot WLM classes, the list will be sorted by the CPU% field. However, the list can be sorted by the other fields by moving the cursor to the top of the desired column.

Processes

This subsection displays a list of hot processes. The maximum number of processes displayed is the number of hot processes being monitored as specified with the **-procs** flag. A smaller number of processes will be displayed if other subsections are also being displayed. Pressing the **p** key turns off this subsection. The process are sorted by their CPU usage over the monitoring interval. The following fields are displayed for each process:

Name The name of the executable program executing in the process. The name is stripped of any pathname and argument information and truncated to 9 characters in length.

Process ID

The process ID of the process.

% CPU Utilization

The average CPU utilization of the process over the monitoring interval. The first time a process is shown, this value is the average CPU utilization over the lifetime of the process.

Paging Space Used

The size of the paging space allocated to this process. This can be considered an expression of the footprint of the process but does not include the memory used to keep the executable program and any shared libraries it may depend on.

Process Owner (if the WLM section is off)

The user name of the user who owns the process.

WorkLoad Management (WLM) Class (if the WLM section is on)

The WLM class to which the process belongs.

Implementation Specifics

Changes to WLM that are shown by **topas** (like adding new classes, or changing existing class names) will not be reflected after starting **topas**. You must stop **topas** and all clients which use **Spmi**, then restart after the WLM changes are made. This is also the case for Disks and Network Adapters added after **topas** or any other **Spmi** consumer is started.

Sample Default Output

The following is an example of the display generated by the **topas** command:

```

Topas Monitor for host:  niller          EVENTS/QUEUES  FILE/TTY
Mon Mar 13 15:56:32 2000  Interval: 2    Cswitch      113  Readch  1853576
                               Syscall      2510 Writech  49883
CPU   User%   Kern%  Wait%  Idle%      Reads       466  Rawin   0
cpu0  7.0     4.0   0.0   89.0      Writes      12  Ttyout  706
cpu1  1.0     8.0   0.0   91.0      Forks        0  Igets   0
cpu2  0.0     0.0   0.0  100.0     Execs        0  Namei   0
                               Runqueue    0.0  Dirblk  0
                               Waitqueue   0.0

Interf  KBPS   I-Pack  O-Pack  KB-In  KB-Out
lo0     100.4  45.7   45.7   50.2  50.2
tr0     2.0    4.4    3.4    1.4   0.6

Disk   Busy%   KBPS    TPS  KB-Read  KB-Writ
hdisk0 0.0     0.0    0.0   0.0    0.0
hdisk1 0.0     0.0    0.0   0.0    0.0

WLM-Class (Active) CPU%   Mem%   Disk%
System             8     41    12
Shared            1     24     9

Name  PID    CPU%  PgSP  Class
topas (35242) 3.0   0.3   System
X     (3622) 1.4  44.4   System
notes (25306) 1.3 123.3 System

EVENTS/QUEUES  FILE/TTY
PAGING         MEMORY
Faults         1  Real,MB  1024
Steals         0  % Comp  81.0
PgspIn         0  % Noncomp 19.0
PgspOut        0  % Client  3.0
PageIn         0
PageOut        0  PAGING SPACE
Sios           0  Size,MB  0
               % Used
               % Free

NFS  calls/sec
ServerV2  0  Press:
ClientV2  0  "h" for help.
ServerV3  0  "q" to quit.
ClientV3  0

```

Sample Full Screen Process Output

```

Topas Monitor for host:  mothra          Interval: 2    Wed Nov 8 12:27:34 2000
                               DATA  TEXT  PAGE
USER  PID  PPID  PRI  NI   RES  RES  SPACE  TIME CPU%  I/O  OTH  COMMAND
root  1806  0  37  41   16  3374  16  13:25 1.0  0  0  gil
root  1032  0  16  41    3  3374  3  0:00 0.0  0  0  lrud
root  1290  0  60  41    4  3374  4  0:02 0.0  0  0  xmgc
root  1548  0  36  41    4  3374  4  0:26 0.0  0  0  netm
root   1  0  60  20  197  9  180  0:24 0.0  0  0  init
root  2064  0  16  41    4  3374  4  0:04 0.0  0  0  wlmsched
root  2698  1  60  20   14    2  14  0:00 0.0  0  0  shlap
root  3144  1  60  20   40    1  36  5:19 0.0  0  0  syncd
root  3362  0  60  20    4  3374  4  0:00 0.0  0  0  lvmbb
root  3666  1  60  20  135  23  123  0:00 0.0  0  0  errdemon
root  3982  0  60  20    4  3374  4  0:01 0.0  0  0  rtcmd
root  4644  1  17  20    6  3374  6  0:00 0.0  0  0  dog
root  4912  1  60  20  106  13  85  0:00 0.0  0  0  srcmstr
root  5202  4912  60  20   94    8  84  0:01 0.0  0  0  syslogd
root  5426  4912  60  20  195  76  181  0:12 0.0  0  0  sendmail
root  5678  4912  60  20  161  11  147  0:01 0.0  0  0  portmap
root  5934  4912  60  20  103  11  88  0:00 0.0  0  0  inetd
root  6192  4912  60  20  217  61  188  0:21 0.0  0  0  snmpd
root  6450  4912  60  20  137  10  116  0:00 0.0  0  0  dpid2
root  6708  4912  60  20  157  29  139  0:06 0.0  0  0  hostmibd
root   0  0  16  41    3  3374  3  7:08 0.0  0  0
root  6990  1  60  20  106  10  86  0:06 0.0  0  0  cron

```

Sample Full-Screen WorkLoad Management Classes Output

```

Topas Monitor for host:  mothra          Interval: 2    Wed Nov 8 12:30:54 2000
WLM-Class (Active)      CPU%   Mem%   Disk-I/O%
System                  0      0      0
Shared                  0      0      0
Default                 0      0      0
Unmanaged               0      0      0
Unclassified            0      0      0

```

```

=====
USER      PID  PPID  PRI  NI   DATA  TEXT  PAGE   TIME  CPU%  I/O  OTH  COMMAND
root       1    0 108  20   197    9   180   0:24  0.0  0   0  init
root    1032    0  16  41    3  3374    3   0:00  0.0  0   0  lrud
root    1290    0  60  41    4  3374    4   0:02  0.0  0   0  xmgc
root    1548    0  36  41    4  3374    4   0:26  0.0  0   0  netm
root    1806    0  37  41   16  3374   16  13:25  0.0  0   0  gil
root    2064    0  16  41    4  3374    4   0:04  0.0  0   0  wlmsched
root    2698    1 108  20   14    2   14   0:00  0.0  0   0  shlap
root    3144    1 108  20   40    1   36   5:19  0.0  0   0  syncd
root    3362    0 108  20    4  3374    4   0:00  0.0  0   0  lvmbb
root    3666    1 108  20   135   23  123   0:00  0.0  0   0  errdemon
root    3982    0 108  20    4  3374    4   0:01  0.0  0   0  rtcmd
=====

```

Flags

- cecdisp** Displays the cross-partition panel.
- cpus** Specifies the number of hot CPUs to be monitored. This is also the maximum number of CPUs displayed when enough room is available on the screen. If this number exceeds the number of CPUs available, only the installed CPUs will be monitored and displayed. If this argument is omitted, a default of 2 is assumed. If a value of 0 (zero) is specified, no CPU information is monitored.

When you specify a single flag, the **topas** command will take that value and the other values will be the default values. This is applicable to all the flags.
- disks** Specifies the number of disks to be monitored. This is also the maximum number of disks displayed when enough room is available on the screen. When this number exceeds the number of disks installed, only the installed disks are monitored and displayed. If this argument is omitted, a default of 2 is assumed. If a value of 0 (zero) is specified, no disk information is monitored.
- interval** Sets the monitoring interval in seconds. The default is 2 seconds.
- nets** Specifies the number of hot network interfaces to be monitored. This is also the maximum number of network interfaces displayed when enough room is available on the screen. When this number exceeds the number of network interfaces installed, only the installed network interfaces will be monitored and displayed. If this argument is omitted, a default of 2 is assumed. If a value of 0 (zero) is specified, no network information is monitored.
- procsdisp** Displays the full-screen process display. This display shows a list of the busiest processes, similar to the process subsection on the default display, only with more columns showing more metrics per process. This list can be sorted by any column.
- procs** Specifies the number of hot processes to be monitored. This is also the maximum number of processes shown when enough room is available on the screen. If this argument is omitted, a default of 20 is assumed. If a value of 0 is specified, no process information will be monitored. Retrieval of process information constitutes the majority of the **topas** overhead. If process information is not required, always use this option to specify that you do not want process information.
- wlmdisp** Displays the full-screen WLM class display, which is a split display. The top part of the display shows a list of hot WLM classes, similar to the WLM classes subsection on the default display, but with enough space available to display the full class names. This list can be sorted on any column.

The bottom part of the display shows a list of busiest processes, similar to the full screen process display, but only displays processes belonging to one WLM class (selected with the **f** key).

-wlms	Specifies the number of hot WorkLoad Management (WLM) classes to be monitored. This is also the maximum number of WLM classes displayed when enough room is available on the screen. If this number exceeds the number of WLM classes installed, only the installed WLM classes will be monitored and displayed. If this argument is omitted, a default of 2 is assumed. If a value of 0 (zero) is specified, no WLM class information is monitored.
-fileSYS	Specifies the number of file systems to be monitored. This is also the maximum number of file systems displayed when enough room is available. When this number exceeds the number of file systems mounted, only the mounted file systems are monitored and displayed. If you do not specify the -fileSYS flag, the default value is 2. If you specify a value of 0, the file system information is monitored.
-tape	Switches the tape display section between on and off in the main topas display.
-fullScreen	Specifies the full-screen display of the specified parameter. The following parameters can be specified: process , WLM , lpar , disk , net , CEC , vg , fileSYS , or tape . This list can be sorted by any column.
	process Displays the full screen Process screen. This display shows a list of the busiest processes. This is similar to the process subsection on the default panel, only with more columns showing more metrics per process.
	WLM Displays the full screen WLM Classes panel. The top part of the screen shows a list of hot WLM classes. This is similar to the WLM classes subsection on the default screen, but with enough space available to display the full class names.
	lpar Displays the full screen Logical Partition panel. In shared-memory mode, this screen displays information about I/O memory entitlement of the partition.
	disk Displays the full screen Disk Metrics panel (Disk Panel View). The screen reports disk service times, disk queuing metrics, and disk throughput.
	net Displays the full screen statistics for the shared Ethernet adapter on a Virtual I/O Server.
	CEC Displays the Cross-partition Panel. The topas command collects a set of metrics from AIX partitions that are running on the same hardware platform. Dedicated and shared partitions are displayed, and a set of aggregated values provide an overview of the entire hardware systems partition set. Certain values that are only available from the HMC platform can be set through the HMC console.
	vg Displays the Volume Group Panel. The panel reports the following metrics of the volume groups in the top section of the screen, and the same metrics of the logical volumes in the bottom section of the screen.
	fileSYS Displays full file system information. The screen reports file system service times, file system queuing metrics, and file system throughput.
	tape Displays the full screen tape screen. This reports bandwidth use of the tape device, the amount of data transferred (read or written) to the tape device, and the average number of transfers per second issued to the tape device.

Subcommands

While **topas** is running, it accepts one-character subcommands. Each time the monitoring interval elapses, the program checks for one of the following subcommands and responds to the action requested.

a	The a key shows all of the variable subsections being monitored (CPU, network, disk, WLM, and process). Pressing the a key always returns the topas command to the initial main display.
c	The c key toggles the CPU subsection between the cumulative report, off, and a list of the busiest CPUs. The number of busiest CPUs displayed will depend upon the space available on the screen.

d	The d key toggles the disk subsection between a list of busiest disks, off, and the report on the total disk activity of the system. The number of busiest disks displayed will depend upon the space available on the screen.
h	Show the help screen.
n	The n key toggles the network interfaces subsection between a list of busiest interfaces, off, and the report on the total network activity of the system. The number of busiest interfaces displayed will depend upon the space available on the screen.
w	The w key toggles the WorkLoad Management (WLM) classes subsection on and off. The number of busiest WLM classes displayed will depend upon the space available on the screen.
p	The p key toggles the hot processes subsection on and off. The number of busiest processes displayed will depend upon the space available on the screen.
P	The uppercase P key replaces the default display with the full-screen process display. This display provides more detailed information about processes running on the system than the process section of the main display. When the P key is pressed again, it toggles back to the default main display.
W	The uppercase W key replaces the default display with the full-screen WLM class display. This display provides more detailed information about WLM classes and processes assigned to classes. When the W key is pressed again, it toggles back to the default main display.
L	The uppercase L key replaces the current display with the logical partition display.
f	Moving the cursor over a WLM class and pressing the f key displays the list of top processes in the class at the bottom of the WLM screen. This key is valid only when <code>topas</code> is in the full-screen WLM display (by using the W key or the <code>-wlms</code> flag).
q	Quit the program.
r	Refresh the display.
Arrow and Tab keys	Subsections from the main display such as the CPU, Network, Disk, WLM Classes, and the full-screen WLM and Process displays can be sorted by different criteria. Positioning the cursor over a column activates sorting on that column. The entries are always sorted from highest to lowest value. The cursor can be moved by using the Tab key or the arrow keys. Sorting is only valid for 128 disks and 16 network adapters.

Exit Status

See “Exit status for Virtual I/O Server commands” on page 3.

Examples

1. To display up to twenty "hot" disks every five seconds and omit network interface, WLM classes, and process information, type:

```
topas -interval 5 -nets 0 -procs 0 -wlms 0
```
2. To display the five most active processes and up to twenty most active WLM classes (which is the default when omitting the `-w` flag) but no network or disk information, type:

```
topas -procs 5 -nets 0 -disks 0
```
3. To run the program with default options, type:

```
topas
```
4. To go directly to the process display, type:

```
topas -procdisp
```
5. To go directly to the WLM classes display, type:

```
topas -wlmdisp
```
6. To display the top two file systems, type:

```
topas -fileysys 2
```
7. To go directly to the WLM classes display, type:

```
topas -wlms 2
```

8. To go directly to the tape device display, type:
topas -tape
9. To go directly to the full screen WLM classes display, type:
topas -fullscreen WLM
10. To go directly to the full screen CEC display, type:
topas -fullscreen CEC
11. To go directly to the full screen volume group display, type:
topas -fullscreen vg

topasrec command

Purpose

Generates a binary recording of local system statistics, central electronics process (CEC) statistics, and cluster statistics.

Syntax

```
topasrec { -local [ -trace trace_level ] | -cec | -cluster } [ -count sample_count ] [ -interval seconds ] [ -out filename ]
```

Description

The **topasrec** command records the local system data and the cross-partition data (CEC statistics) in binary format.

When you run the **topasrec** command to record CEC statistics, the command collects a set of metrics from the AIX logical partitions that are running on the same CEC. The command collects dedicated and shared logical partition data, as well as a set of aggregated values to provide an overview of the partition set on the same CEC.

When you run the command to record cluster statistics, the command collects a set of metrics for the list of hosts that are specified in the cluster configuration file.

Flags

-count	Specifies the number of records to generate.
-interval	Specifies the recording interval in seconds. The default is 2 seconds.
-out	Specifies the name of the output file.
-cec	Records the CEC statistics in binary format.
-local	Records local system statistics in binary format.
-cluster	Records cluster statistics in binary format.
-trace	Specifies a trace level. Trace data is stored under <i>/home/padmin/topasrec</i> . You can specify a value from 1 - 9. More trace data is generated at higher trace levels. This data is useful for recording status and for debugging purposes. Trace data is collected only for local system recording.

Exit Status

Return Code	Description
0	Success
>0	An error occurred

Examples

1. To start recording local system statistics with a sample count of 5 and a sample interval of 60 seconds, type:
`topasrec -local -count 5 -interval 60`
2. To start recording CEC statistics with a sample count of 5 and a sample interval of 60 seconds, type:
`topasrec -cec -count 5 -interval 60`
3. To start recording cluster statistics with a sample count of 5 and a sample interval of 60, type:
`topasrec -cluster -count 5 -interval 60`

Related Information

The `wkldout` command.

tracepriv command

Purpose

Traces the privileges that a command needs for a successful run.

Syntax

```
tracepriv [-d][-e][-f] [-o outputfile ] Command [args]
```

Description

The `tracepriv` command records the privileges that a command attempts to use when the command is run. The `tracepriv` command is used for command investigation when entries are added to the privileged command database. The `tracepriv` command runs the command that is specified by the *Command* parameter with the specified arguments (with the *args* parameter). Generally, run the `tracepriv` command with the PV_ROOT privilege so that any attempt to use a privilege succeeds. In this case, the `tracepriv` command can keep track of all of the privileges that the *Command* needs for a successful run without the PV_ROOT privilege. After the *Command* is run or when an `exec` subroutine within the command occurs, the list of used privileges is written to standard output (**stdout**).

Important: The `tracepriv` command cannot be used for any VIOS commands. The purpose of this command is to provide the ability to find the privileges to run your commands and scripts.

Security

The `tracepriv` command is a privileged command. To run the command successfully, assume a role that has the authorization `vios.security.priv.trace`.

Flags

- d Displays the output of the `truss` command with the privileges that are required by the command.
- e Follows the `exec` subroutine. If the command specified by the *Command* parameter runs an `exec` subroutine, the `tracepriv` command reports the privileges needed so far (and sets them if the `-a` flag is used), and then proceeds with recording (and setting) the privileges associated with the new executable file. If the file that is run by the `exec` subroutine has its setuid bit set and is not owned by root, the `tracepriv` command cannot properly trace the privilege use of the file.
- f Follows the `fork` subroutine. If the controlled process calls the fork subroutine, the `tracepriv` command also reports the privileges used by the new child process.
- o Writes the output to the specified file instead of the standard output (**stdout**).

Parameters

<i>args</i>	Specifies the arguments.
<i>command</i>	Specifies the command.
<i>outputfile</i>	Specifies the file to record the output.

Exit Status

0 for success.

Non-zero for command error.

Examples

1. To find the privileges to run the command `myexe`, type the following command:

```
tracepriv /home/padmin/myexe
```

traceroute command

Purpose

Prints the route that IP packets take to a network host.

Syntax

```
traceroute [ -hops Hops ] [ -num ] [ -port Port ] [ -src Address ] Host [ PacketSize ]
```

Description

The **traceroute** command attempts to trace the route an IP packet follows to an Internet host by launching UDP probe packets with a small maximum time-to-live (*Hops* parameter), then listening for an **ICMP TIME_EXCEEDED** response from gateways along the way. Probes are started with a Hops value of one hop, which is increased one hop at a time until an **ICMP PORT_UNREACHABLE** message is returned. The **ICMP PORT_UNREACHABLE** message indicates either that the host has been located or the command has reached the maximum number of hops allowed for the trace.

The **traceroute** command sends three probes at each Hops setting to record the following:

- Hops value
- Address of the gateway
- Round-trip time of each successful probe

If the probe answers come from different gateways, the command prints the address of each responding system. If there is no response from a probe within a 3-second time-out interval, an * (asterisk) is printed for that probe.

Note: The **traceroute** command is intended for use in network testing, measurement, and management. It should be used primarily for manual fault isolation. Because of the load it imposes on the network, the **traceroute** command should not be used during normal operations or from automated scripts.

The **traceroute** command prints an ! (exclamation mark) after the round-trip time if the Hops value is one hop or less. A maximum time-to-live value of one hop or less generally indicates an incompatibility in the way ICMP replies are handled by different network software. The incompatibility can usually be resolved by doubling the last Hops value used and trying again.

Other possible annotations after the round-trip notation are as follows:

!H	Host unreachable
!N	Network unreachable
!P	Protocol unreachable
!S	Source route failed
!F	Fragmentation needed

If the majority of probes result in an error, the **traceroute** command exits.

The only mandatory parameter for the **traceroute** command is the destination host name or IP number. The **traceroute** command will determine the length of the probe packet based on the Maximum Transmission Unit (MTU) of the outgoing interface. The UDP probe packets are set to an unlikely value so as to prevent processing by the destination host.

Flags

-hops <i>Max_ttl</i>	Sets the maximum time-to-live (maximum number of hops) used in outgoing probe packets. The default is 30 hops (the same default used for TCP connections).
-num	Prints hop addresses numerically rather than symbolically and numerically. This flag saves a name-server address-to-name lookup for each gateway found on the path.
-port <i>Port</i>	Sets the base UDP port number used in probes. The default is 33434. The traceroute command depends on an open UDP port range of base to base + nhops - 1 at the destination host. If a UDP port is not available, this option can be used to pick an unused port range.
-src <i>SRC_Addr</i>	Uses the next IP address in numerical form as the source address in outgoing probe packets. On hosts with more than one IP address, the -src flag can be used to force the source address to be something other than the IP address of the interface on which the probe packet is sent. If the next IP address is not one of the machine's interface addresses, an error is returned and nothing is sent.

Parameters

<i>Host</i>	Specifies the destination host, either by host name or IP number. This parameter is required.
<i>PacketSize</i>	Specifies the probe datagram length. The default packet size is determined by the traceroute command based on the MTU of the outgoing interface.

Exit Status

See "Exit status for Virtual I/O Server commands" on page 3.

Examples

- To print the route to host **nis.nsf.net**, type:

```
traceroute nis.nsf.net
```

The output will look similar to the following:

```
traceroute to rotterdam (35.1.1.48), 30 hops max, 56 byte packet
 1 helios.ee.lbl.gov (128.3.112.1) 19 ms 19 ms 0 ms
 2 lilac-dmc.Berkeley.EDU (128.32.216.1) 39 ms 39 ms 19 ms
 3 lilac-dmc.Berkeley.EDU (128.32.216.1) 39 ms 39 ms 19 ms
 4 ccngw-ner-cc.Berkeley.EDU (128.32.136.23) 39 ms 40 ms 39 ms
 5 ccn-nerif22.Berkeley.EDU (128.32.168.22) 39 ms 39 ms 39 ms
 6 128.32.197.4 (128.32.197.4) 40 ms 59 ms 59 ms
 7 131.119.2.5 (131.119.2.5) 59 ms 59 ms 59 ms
```

```

8 129.140.70.13 (129.140.70.13) 99 ms 99 ms 80 ms
9 129.140.71.6 (129.140.71.6) 139 ms 239 ms 319 ms
10 129.140.81.7 (129.140.81.7) 220 ms 199 ms 199 ms
11 nic.merit.edu (35.1.1.48) 239 ms 239 ms 239 ms

```

Related Information

The `ping` command and the `optimizenet` command.

uname command

Purpose

Writes to standard output the name of the operating system that you are using

Syntax

```
uname [ -a ] [-f ] [-F] [-l] [-L] [-m] [-M] [-n] [-p] [-r] [-s] [-u] [-x]
```

Description

The machine ID number contains 12 characters in the following digit format: *xyyyyyyyymmss*. The *xx* positions indicate the system and is always 00. The *yyyyyy* positions contain the unique ID number for the entire system. The *mm* position represents the model ID. The *ss* position is the submodel number and is always 00. The model ID describes the ID of the CPU Planar, not the model of the System as a whole. You can use the `uname -m` command sometimes to determine which model you are using. The following list is not complete. Refer to hardware vendor supplied documentation for values in the range E0 - FF. Also note that not all machine types have a machine ID. Many new machines share a common machine ID of 4C.

Flags

<code>-a</code>	Displays all information specified with the <code>-m</code> , <code>-n</code> , <code>-r</code> , <code>-s</code> , and <code>-v</code> flags. Cannot be used with the <code>-x</code> flag. If the <code>-x</code> flag is specified with the <code>-a</code> flag, the <code>-x</code> flag overrides it.
<code>-F</code>	Displays a system identification string comprised of hexadecimal characters. This identification string is the same for all partitions on a particular system.
<code>-f</code>	Similar to the <code>F</code> flag, except that the partition number is also used in the calculation of this string. The resulting identification string is unique for each partition on a particular system.
<code>-l</code>	Displays the LAN network number.
<code>-L</code>	Displays LPAR number and LPAR name. If LPAR does not exist, <code>-1</code> is displayed for LPAR number and <code>NULL</code> for LPAR name.
<code>-m</code>	Displays the machine ID number of the hardware running the system. Note: The <code>-m</code> flag cannot be used to generate a unique machine identifier for partitions in an LPAR environment.
<code>-M</code>	Displays the system model name. If the model name attribute does not exist, a null string is displayed.
<code>-n</code>	Displays the name of the node. This may be a name the system is known by to a UUCP communications network.
<code>-p</code>	Displays the architecture of the system processor.
<code>-r</code>	Displays the release number of the operating system.
<code>-s</code>	Displays the system name. This flag is on by default.
<code>-u</code>	Displays the system ID number. If this attribute is not defined, the output is the same as the output displayed by <code>uname -m</code> .
<code>-x</code>	Displays the information specified with the <code>-a</code> flag as well as the LAN network number, as specified by the <code>-l</code> flag.

Exit Status

The following exit values are returned:

0	Successful completion.
>0	An error occurred.

Examples

1. To display the complete system name and version banner, enter:

```
uname -a
AIX vios_bat 3 5 00CD1B0E4C00
```

2. To display the operating system name, enter:

```
uname
AIX
```

Virtual I/O Server **uname** command calls the AIX **uname** command. The flags are the same except that **-S**, **-T**, and **-v** are not allowed.

Related Information

uname subroutine

unloadopt command

Purpose

Remove a virtual optical media disk from a virtual optical device.

Syntax

```
unloadopt [ -release ] -vtd VirtualTargetDevice
```

Description

The **unloadopt** command removes a virtual optical disk from the specified virtual optical device (*VirtualTargetDevice*). After the command completes the specified optical device will contain no media.

Flags

-release	Forces the virtual optical device to be unlocked even if the client has a reserve on the device.
-vtd <i>VirtualTargetDevice</i>	The name of the virtual target device.

Examples

To unload the virtual optical disk loaded into the virtual optical device vopt1, type the following command:

```
unloadopt -vtd vopt1
```

unmirrorios command

Purpose

Removes the mirrors that exist on the rootvg volume group.

Syntax

unmirrorios [*PhysicalVolume* ...]

Description

The **unmirrorios** command unmirrors all the logical volumes detected on the rootvg volume group. By default, **unmirrorios** will pick the set of mirrors to remove from a mirrored volume group. To control which drives no longer are to contain mirrors, you must include the list of disks in the input parameters, *PhysicalVolume*.

At the completion of this command, Quorum will be disabled until the system is rebooted.

Note: If LVM has not recognized that a disk has failed, it is possible that it will remove a different mirror. If you know that a disk has failed, and LVM does not show those disks as missing, you should specify the failed disks on the command line or you should use the **reducevg** command to remove the disk.

Exit Status

See “Exit status for Virtual I/O Server commands” on page 3.

Security

This command can only be executed by the prime administrator.

Related Information

The **activatevg** command, the **chvg** command, the **deactivatevg** command, the **exportvg** command, the **importvg** command, the **lsvg** command, the **mkvg** command, the **syncvg** command, and the **mirrorios** command.

unmount command

Purpose

Unmounts a previously mounted file system, directory, or file.

Syntax

unmount { *Directory* | *File* | *FileSystem* }

Description

The **unmount** command unmounts a previously mounted directory, file, or file system. Processing on the file system, directory, or file completes and it is unmounted.

To unmount local mounts you can specify the device, directory, file, or file system on which it is mounted.

Exit Status

See “Exit status for Virtual I/O Server commands” on page 3.

Examples

1. To unmount files and directories, type the following command:

```
umount /home/user/test
```

This unmounts the filesystem mounted at `/home/user/test`.

Related Information

The `mount` command.

updateios command

Purpose

Updates the Virtual I/O Server to latest maintenance level.

Syntax

```
updateios -dev Media [-f] [ -install ] [ -accept ]
```

```
updateios -commit
```

```
updateios -cleanup
```

```
updateios -remove { -file RemoveListFile | RemoveList }
```

```
updateios -list -dev Media
```

```
updateios -fs FilesetName -dev Media
```

Description

The `updateios` command is used to install fixes, or to update the VIOS to the latest maintenance level. Before installing a fix or updating the maintenance level, the `updateios` command runs a preview installation and displays the results. You are then prompted to continue or exit. If the preview fails for any reason, do not install the update.

Note: A fix pack or service pack cannot be applied if the VIOS partition is part of a shared storage pool and the cluster node state is **UP**. To apply a fix pack or service pack, either remove the VIOS partition from the cluster or change the cluster node state to **DOWN** by running the `clstartstop` command.

Note: Do not do any cluster configuration activities until the `updateios` command has finished running.

Note: If you use the File Backed Optical Media repository, ensure to offload the media images before you apply a fix pack.

Flags

<code>-accept</code>	Specifies that you agree to required software license agreements for software to be installed.
<code>-cleanup</code>	Specifies the cleanup flag to remove all incomplete pieces of the previous installation. Perform cleanup processing whenever any software product or update is after an interrupted installation or update is in a state of either applying or committing. You can run this flag manually, as needed.
<code>-commit</code>	The <code>-commit</code> flag will commit all uncommitted updates to the VIOS.
<code>-dev <i>Media</i></code>	Specifies the device or directory containing the images to install.

-f	Forces all uncommitted updates to be committed prior to applying the new updates. When combined with the -dev flag, this flag commits all updates prior to applying any new ones.
-file file	Specifies the file containing a list of entries to uninstall.
-fs FilesetName	Specifies the name of the fileset to be installed from the VIOS installation media.
-install	Installs new and supported filesets onto the VIOS.
	Attention: The log file, install.log in the user's home directory, will be overwritten with a list of all filesets that were installed.
-list	Lists the filesets on the VIOS installation media, that are available to be installed.
-remove	Removes the listed filesets from the system. The filesets to be removed must be listed on the command line or in the RemoveListFile file.

Exit Status

The following exit values are returned:

19	All uncommitted updates must be committed
20	There are no uncommitted updates

Examples

- To update the VIOS to the latest level, where the updates are located on the mounted file system `/home/padmin/update`, type the following command:

```
updateios -dev /home/padmin/update
```
- To update the VIOS to the latest level, when previous levels are not committed, type the following command:

```
updateios -f -dev /home/padmin/update
```
- To cleanup partially installed updates, type the following command:

```
updateios -cleanup
```
- To commit the installed updates, type the following command:

```
updateios -commit
```
- To list the available filesets on the VIOSs installation media, type the following command:

```
updateios -list -dev /dev/cd0
```
- To install a fileset from the VIOS installation media, type the following command:

```
updateios -fs ILMT-TAD4D-agent -dev /dev/cd1
```

Related Information

The **lssw** command, the **ioslevel** command, the **remote_management** command, the **oem_setup_env** command, and the **oem_platform_level** command.

vasistat command

Purpose

Shows VASI device driver and device statistics.

Syntax

```
vasistat [ -all ] [ -reset ] [ -debug ] Device
```

Description

The **vasistat** command displays the statistics gathered by the specified VASI device driver. The user can optionally specify that the device-specific statistics be displayed in addition to the device generic statistics. This will display statistics for all the operations, such as migrations, in that VASI device. If no flags are specified, only the device generic statistics are displayed.

Flags

-all	Displays all the statistics, including the device-specific statistics.
-reset	Resets all the statistics back to their initial values.
-debug	Toggles internal debug tracing in device driver

Parameters

-all	The name of the VASI device, for example, vasi0.
-------------	--

Exit Status

The statistic fields displayed in the output of the **vasistat** command and their descriptions are as follows:

Device Type

Displays the description of the adapter type.

Elapsed Time

Displays the real time period which has elapsed since last time the statistics were reset.

Transmit statistics fields

Table 8. Transmit statistics fields and descriptions.

Statistic	Description
Packets	The number of packets transmitted successfully to PHY by the device.
Bytes	The number of bytes transmitted successfully to PHY by the device.
Transmit Errors	The number of output errors encountered on this device. This is a counter for unsuccessful transmissions due to errors returned by PHY.
Bad Packets	The number of outgoing packets that could not be sent because they were malformed (for example, packets exceeding the VASI MTU size).
No Buffers	The number of times a packet could not be sent to PHY because there were no transmit buffers available for sending.
Interrupts	The number of transmit interrupts (for example, attempts at transmitting a packet to PHY).

Receive Statistics Fields

Table 9. Receive statistics fields and descriptions.

Statistic	Description
Packets	The number of packets received successfully from PHY by the device.
Bytes	The number of bytes received successfully from PHY by the device.
Receive Errors	The number of receive errors encountered on this device (for example, bad packets).
Bad Packets	The number of incoming packets that could not be processed because they were malformed (for example, VASI could not figure out to which operation the data belonged to).

Table 9. Receive statistics fields and descriptions. (continued)

Statistic	Description
No Buffers	The number of times VASI tried to allocate a system buffer but there was no memory to do so. This is a soft error since the data will still be delivered in the original buffer.
Interrupts	The number of receive interrupts (for example, number of CRQ entries received by VASI).
System Buffers	The number of system buffers allocated when the receive buffer pools are low on buffers.

Miscellaneous Statistics Fields

Table 10. Miscellaneous statistics fields and descriptions.

Statistic	Description
Interrupt Processing Exceeded	The number of times this device attempted to process more packets on the interrupt context than the maximum allowed.
Offlevel Interrupt Scheduled	The number of times an offlevel interrupt was scheduled on this device, in order to handle packets that could not be handled on the interrupt context.
Maximum Operations	The maximum number of operations (for example, migrations) this device can handle simultaneously.
Maximum Receive Pools	The maximum number of different receive pool sizes this device can handle.
Active Operations	The number of operations (for example, migrations) that are currently active on this device.

Operation-specific Statistics Fields

These statistics are displayed for each operation, either active or inactive. Inactive applications will have the text `INACTIVE` in the title header, and the text `INVALID STREAM ID` will appear in the Stream ID field. None of these statistics will be zeroed out by the `-reset` flag.

Table 11. Operation-specific statistics fields and descriptions.

Statistic	Description
Operation Type	The type of this operation (for example, migration) as well as any pertinent information (for example, whether it is the source or target of the migration).
Stream ID	The unique number that identifies this operation; the text <code>INVALID STREAM ID</code> denotes that this operation is no longer active.
TOP/BOTTOM	The opaque identifiers used by the VASI device and PHYP to refer to this operation.
Elapsed Time	Displays the real time period which has elapsed since the operation was started. This time will stop when the operation is terminated, so it can be used to measure how long the operation took to complete.
Flags	Denotes values that may be used to describe this operation: <ol style="list-style-type: none"> 1. <code>RUNNABLE</code> - this operation has been initialized and is ready to receive CRQ commands 2. <code>TERMINATED</code> - this operation has been terminated internally due to some error within VASI

Table 11. Operation-specific statistics fields and descriptions. (continued)

Statistic	Description
Operation State	State the operation is in with regards to its setup: <ol style="list-style-type: none"> 1. Not Operational - no memory has been allocated for this operation 2. Work Queues Allocated - the work queues (where incoming CRQ responses are queued) have been allocated 3. PHYP Registered - VASI has established a connection with PHYP on behalf of this operation 4. Pools Partially Allocated - some receive buffer pools have been allocated 5. Pools Allocated - all receive buffer pools have been allocated 6. Buffers Partially Registered - some receive buffers have been registered with PHYP 7. Operational - all setup is complete
Stream State	State the stream (for example, migration) is in; these values will vary depending on the operation in question. For migration, these states are as follows: <ol style="list-style-type: none"> 1. Invalid - the migration has not been started 2. Enabled - the migration has been started 3. Aborted - the migration has aborted (may be due to internal VASI error, user-initiated abort, Mover-initiated abort, or PHYP-initiated abort) 4. Suspending - the suspend trigger percentage has been reached and the migrating partition may be suspended 5. Suspended - the migrating partition has been suspended at the source 6. Resumed - the migrating partition has been resumed at the target 7. Completed - the migration has completed successfully
Total Bytes to Transfer	The estimated total number of bytes to transfer for this operation
Bytes Left to Transfer	The estimated number of bytes that are left to transfer for this operation

Receive Buffer Usage Statistics Fields

These statistics are displayed only for operations that are currently active. For every receive buffer pool the following information is shown in tabular form:

Table 12. Receive buffer usage statistics fields and descriptions.

Statistic	Description
Size	The size in bytes of the packets of the buffer pool.
Reg	The number of buffers from the buffer pool that are currently registered with PHYP.
Alloc	The number of buffers that have been allocated for the buffer pool.
Max	The maximum number of buffers that could ever be allocated to this buffer pool.
LowReg	The lowest number of buffers from this pool that have ever been registered with PHYP.

Transmit Buffer Usage Statistics Fields

These statistics are displayed only for operations that are currently active.

Table 13. Transmit buffer usage statistics fields and descriptions.

Statistic	Description
Number of Buffers	The number of buffers that have been allocated for the transmit buffer pool.
Buffer Size	The size in bytes of the packets of the transmit buffer pool.
Mapped	The number of buffers in the transmit buffer pool that have been DMA-mapped.

Table 13. Transmit buffer usage statistics fields and descriptions. (continued)

Statistic	Description
Next Available Buffer	The index of the next buffer to be used in the transmit buffer pool.
In Use	The number of buffers in the transmit buffer pool that are currently being used.
Maximum Used	The maximum number of buffers in the transmit buffer pool that have ever been used concurrently.

Examples

1. To display the device generic statistics for vasi0, type:

```
vasistat vasi0
```

This produces output similar to the following:

```
VASI STATISTICS (vasi0) :
```

```
Device Type: Virtual Asynchronous Services Interface (VASI)
```

```
Elapsed Time: 0 days 0 hours 10 minutes 38 seconds
```

```
Transmit Statistics:          Receive Statistics:
```

```
-----                    -----  
Packets: 2                  Packets: 179098
```

```
Bytes: 224                  Bytes: 753605020
```

```
Transmit Errors: 0         Receive Errors: 0
```

```
Bad Packets: 0            Bad Packets: 0
```

```
No Buffers: 0            No Buffers: 0
```

```
Interrupts: 2            Interrupts: 214876
```

```
System Buffers: 0
```

```
Interrupt Processing Exceeded: 0
```

```
Offlevel Interrupt Scheduled: 0
```

```
Driver Flags: Up Running 64BitSupport
```

```
Maximum Operations: 4
```

```
Maximum Receive Pools: 3
```

```
Active Operations: 1
```

2. To display the VASI device generic statistics and the statistics for all the operations for vasi0, type:

```
vasistat -all vasi0
```

This produces results similar to the following:

```
VASI STATISTICS (vasi0) :
```

```
Device Type: Virtual Asynchronous Services Interface (VASI)
```

Elapsed Time: 0 days 0 hours 10 minutes 38 seconds

Transmit Statistics:	Receive Statistics:
-----	-----
Packets: 2	Packets: 179098
Bytes: 224	Bytes: 753605020
Transmit Errors: 0	Receive Errors: 0
Bad Packets: 0	Bad Packets: 0
No Buffers: 0	No Buffers: 0
Interrupts: 2	Interrupts: 214876
	System Buffers: 0

Interrupt Processing Exceeded: 0

Offlevel Interrupt Scheduled: 0

Driver Flags: Up Running 64BitSupport

Maximum Operations: 4

Maximum Receive Pools: 3

Active Operations: 1

Statistics for each operation:

=====

Operation #0 (ACTIVE):

Operation Type: Migration (Source)

Stream ID: 0000000000000001

TOP/BOTTOM: 00000000/00040000

Elapsed Time: 0 days 0 hours 0 minutes 11 seconds

Flags: <RUNNABLE>

Operation State: Operational

Stream State: Enabled

Total Bytes to Transfer: 546832384

Bytes Left to Transfer: 360833024

Transmit Statistics: Receive Statistics:

```
-----  
Packets: 1                                Packets: 45415  
Bytes: 112                                Bytes: 191086638  
Transmit Errors: 0                        Receive Errors: 0  
Bad Packets: 0                            Bad Packets: 0  
No Buffers: 0                             No Buffers: 0  
Interrupts: 1                             Interrupts: 0
```

System Buffers: 0

Receive Buffer Usage:

```
-----  
  
          Size      Reg      Alloc      Max      LowReg  
Pool #0  8192      198      256      256      193  
Pool #1  2048      2048     2048     2048     2044
```

Transmit Buffer Usage:

```
-----
```

Number of Buffers: 64

Buffer Size: 16384 bytes

Mapped: 64

Next Available Buffer: 0

In Use: 0

Maximum Used: 1

Operation #1 (INACTIVE):

```
-----
```

Operation Type: Unknown (Unknown)

Stream ID: ** INVALID STREAM ID **

TOP/BOTTOM: 00000000/00000000

Elapsed Time: 0 days 0 hours 0 minutes 0 seconds

Flags:

Operation State: Not Operational

Stream State: Unknown

Total Bytes to Transfer: 0

Bytes Left to Transfer: 0

Transmit Statistics:	Receive Statistics:
----------------------	---------------------

-----	-----
-------	-------

Packets: 0	Packets: 0
------------	------------

Bytes: 0	Bytes: 0
----------	----------

Transmit Errors: 0	Receive Errors: 0
--------------------	-------------------

Bad Packets: 0	Bad Packets: 0
----------------	----------------

No Buffers: 0	No Buffers: 0
---------------	---------------

Interrupts: 0	Interrupts: 0
---------------	---------------

System Buffers: 0

Operation #2 (INACTIVE):

Operation Type: Unknown (Unknown)

Stream ID: ** INVALID STREAM ID **

TOP/BOTTOM: 00000000/00000000

Elapsed Time: 0 days 0 hours 0 minutes 0 seconds

Flags:

Operation State: Not Operational

Stream State: Unknown

Total Bytes to Transfer: 0

Bytes Left to Transfer: 0

Transmit Statistics:	Receive Statistics:
----------------------	---------------------

-----	-----
-------	-------

Packets: 0	Packets: 0
------------	------------

Bytes: 0	Bytes: 0
----------	----------

Transmit Errors: 0	Receive Errors: 0
--------------------	-------------------

```
Bad Packets: 0          Bad Packets: 0
No Buffers: 0          No Buffers: 0
Interrupts: 0          Interrupts: 0
                        System Buffers: 0
```

Operation #3 (INACTIVE):

```
Operation Type: Unknown (Unknown)
Stream ID: ** INVALID STREAM ID **
TOP/BOTTOM: 00000000/00000000
Elapsed Time: 0 days 0 hours 0 minutes 0 seconds
Flags:
Operation State: Not Operational
Stream State: Unknown
Total Bytes to Transfer: 0
Bytes Left to Transfer: 0
```

```
Transmit Statistics:      Receive Statistics:
-----
Packets: 0                Packets: 0
Bytes: 0                   Bytes: 0
Transmit Errors: 0        Receive Errors: 0
Bad Packets: 0            Bad Packets: 0
No Buffers: 0             No Buffers: 0
Interrupts: 0             Interrupts: 0
                        System Buffers: 0
```

vfcmap command

Purpose

Maps the virtual fibre channel adapter to the physical fibre channel port.

Syntax

```
vfcmap -vadapter virtual fibre channel adapter -fcp fibre channel port name
```

Description

The **vfcmmap** command maps or unmaps the virtual fibre channel adapter to the physical fibre channel port.

Flags

-vadapter *virtual fibre channel adapter* Specifies the virtual server adapter.
-fcp *fibre channel port name* Specifies the physical fibre channel port.
Note: If no parameter is specified with this flag, the command unmaps the virtual fibre channel adapter from the physical fibre channel port.

Exit Status

The following exit values are returned:

0 Successful completion.
>0 Invalid flag, argument, or command failure.

Examples

1. To map the virtual fibre channel, **vfchost7**, to the physical fibre channel port, **fcs0**, type:
vfcmmap -vadapter vfchost7 -fcp fcs0

The system displays a message similar to the following:
vfchost0 changed

2. To unmap the virtual fibre channel, **vfchost7**, from any physical fibre channel port, type:
vfcmmap -vadapter vfchost7 -fcp

The system displays a message similar to the following:
vfchost0 changed

Related Information

The **lsmmap** and **lsnports** commands.

viosbr command

Purpose

Performs the operations for backing up the virtual and logical configuration, listing the configuration, and restoring the configuration of the Virtual I/O Server (VIOS).

The **viosbr** command can be run only by the *padmin* user.

Syntax

To perform a backup:

```
viosbr -backup -file FileName [-frequency daily | weekly | monthly [-numfiles fileCount]]
```

```
viosbr -backup -file FileName -clustername ClusterName [-frequency daily | weekly | monthly [-numfiles fileCount]]
```

To view a backup file:

viosbr -view -file *FileName* [-**type** *devType*] [-**detail** | -**mapping**]

viosbr -view -file *FileName* -**clustername** *ClusterName* [-**type** *devType*] [-**detail** | -**mapping**]

To view the listing of backup files:

viosbr -view -list [*UserDir*]

To restore a backup file:

viosbr -restore -file *FileName* [-**validate** | -**inter**] [-**type** *devType*]

viosbr -restore -file *FileName* [-**type** *devType*] [-**force**]

viosbr -restore -clustername *ClusterName* -**file** *FileName* -**subfile** *NodeFileName* [-**validate** | -**inter** | -**force**] [-**type** *devType*] [-**skipcluster**]

viosbr -restore -clustername *ClusterName* -**file** *FileName* -**repopvs** *list_of_disks* [-**validate** | -**inter** | -**force**] [-**type** *devType*] [-**currentdb**]

viosbr -restore -clustername *ClusterName* -**file** *FileName* -**subfile** *NodeFile* -**xmlvtds**

viosbr -restore -file *FileName* [-**skipcluster**]

To disable a scheduled backup:

viosbr -nobackup

To recover from a corrupted shared storage pool (SSP) database:

viosbr -recoverdb -clustername *ClusterName* [-**file** *FileName*]

To migrate a backup file from an older release level to a current release level:

viosbr -migrate -file *FileName*

Description

The **viosbr** command uses the parameters **-backup**, **-view**, and **-restore** to perform backup, list, and recovery tasks for the VIOS.

This **viosbr** command backs up all the relevant data to recover VIOS after a new installation. The **-backup** parameter backs up all the device properties and the virtual devices configuration on the VIOS. This includes information regarding logical devices, such as storage pools, file-backed storage pools, the virtual media repository, and PowerVM Active Memory Sharing (AMS) paging devices. It also includes the virtual devices, such as Etherchannel, shared Ethernet adapters (SEAs), virtual server adapters, the virtual log repository, and server virtual Fibre Channel (SVFC) adapters. Additionally, it includes the device attributes, such as the attributes for disks, optical devices, tape devices, Fibre Channel SCSI controllers, Ethernet adapters, Ethernet interfaces, and logical Host Ethernet adapters (HEAs). All the configuration information is saved in a compressed XML file. If a location is not specified with the **-file** option, the file is placed in the default location */home/padmin/cfgbackups*. This command can be run once or can be run in a stipulated period by using the **-frequency** parameter with the *daily*, *weekly*, or *monthly* option. Daily backups occur at 00:00, weekly backups on Sunday at 00:00, and monthly backups on the first day of the month at 00:01. The **-numfile** parameter specifies the number of successive backup files that are saved, with a maximum value of 10. After reaching the given number of files, the oldest backup

file is deleted during the next backup cycle. The format of the file name is *<givenfilename>.xx.tar.gz*, where *xx* starts from 01. For cluster backups, the format is *<givenfilename>.xx.<clustername>.tar.gz*.

The **viosbr** command does not back up the parent devices of adapters or drivers, device drivers, virtual serial adapters, virtual terminal devices, kernel extensions, the Internet Network Extension (inet0), virtual I/O bus, processor, memory, or cache.

The **-view** parameter displays the information of all the backed up entities in a formatted output. This parameter requires an input file in a compressed or noncompressed format that is generated with the **-backup** parameter. The **-view** parameter uses the option flags *type* and *detail* to display information in detail or to display minimal information for all the devices or for a subset of devices. The **-mapping** option flag provides **lsmmap**-like output for Virtual Small Computer System Interface (VSCSI) server adapters, SEA, server virtual Fibre Channel (SVFC) adapters, and PowerVM Active Memory Sharing paging devices. The entities can be controllers, disks, optical devices, tape devices, network adapters, network interfaces, storage pools, repositories, Etherchannels, virtual log repositories, SEAs, VSCSI server adapters, server virtual Fibre Channel (SVFC) adapters, and paging devices. The **-list** option displays backup files from the default location */home/padmin/cfgbackups* or from a user-specified location.

The **-restore** parameter uses an earlier backup file as input and brings the VIOS partition to the same state as when the backup was created. With the information available from the input file, the command sets the attribute values for physical devices, imports logical devices, and creates virtual devices and their corresponding mappings. The attributes can be set for controllers, adapters, disks, optical devices, tape devices, and Ethernet interfaces. Logical devices that can be imported are volume groups, storage pools, logical volumes (LVs), file systems, and repositories. Virtual devices that can be created are Etherchannel, SEA, server virtual Fibre Channel (SVFC) adapters, virtual target devices, and PowerVM Active Memory Sharing paging devices. The command creates mappings between virtual SCSI server adapters and the VTD-backing devices, between a virtual Fibre Channel (VFC) server adapter and a Fibre Channel (FC) adapter, and between PowerVM Active Memory Sharing paging devices and backing devices. The **viosbr** command with the **-restore** option must be run on the same VIOS partition as the one where the backup was performed. The command uses parameters to validate the devices on the system and restores a category of devices. The **-restore** option runs interactively so that if any devices fail to restore, you can decide how to handle the failure.

Restoration of network configuration is blocked if the cluster is configured on the VIOS node. Cluster **-restore** and **-subfile** options do not restore the network configuration for the current release. To restore network configuration, **-subfile** option with **-type net** must be used to restore network settings before restoring cluster.

The **viosbr** command recovers the data that is used to reconfigure an SSP cluster. This command does not recover any of the data, such as the contents of an LU. You must take separate action to back up that data.

The **viosbr** command recovers an entire cluster configuration by using the *-clustername* option, which includes re-creating a cluster, adding all the nodes that comprise the cluster, and re-creating all cluster entities on all the nodes. If a node is down during this operation, the node is recovered when it is started if the cluster is not deleted.

If a single node is reinstalled and you want to restore the entities of that node, you must use the *-subfile* option and specify the .xml file that corresponds with the node.

If the restore of a cluster fails, rerun the command to resolve the issue. For example, while restoring a four node cluster, if the restore fails after restoring two nodes, rerun the command to restore the other two nodes.

If one of the one of the nodes are not added when restoring a cluster, do not add that node using **cluster -addnode**. The **cluster -addnode** command adds a new node to the cluster and this invalidates the existing node information in the database.

An SSP cluster might incur a database corruption. If a database corruption occurs, you must use the **-recoverdb** option. If this option is used with the **-file** option, the **viosbr** command uses the database information from the specified backup file. If the resources of the SSP cluster change after the backup file is made, those changed resources do not appear. The SSP cluster is updated to make a copy of the SSP database on a daily basis. If you prefer this copy of the database to the database stored in the backup, you can exclude the **-file** option and the backup file from the command-line call. Use the **-view** option to get the list of xml files in the cluster, choose the correct file from the list by using the MTM and Partition Number.

Note: Recovery of the database is allowed only when all the other nodes in the cluster are down except the node where recovery is initiated.

When VIOS is reinstalled with the a newer level of software then restoring the shared storage pool cluster to newer software level is a two step process as follows:

1. Migrate the existing backup.
2. Restore the share storage pool cluster using the migrated backup.

The **-migrate** option makes a backup file from the **-file** option and migrates the file to form a new backup that can be used for a shared storage pool configuration restore copy on a current release of VIOS. This option must be called before restoring and an shared storage pool cluster must not be present.

Flags

-backup	Takes the backup of VIOS configurations.
-clustername	Specifies the Cluster name to back up, restore, or view; including all of its associated nodes.
-currentdb	Restores the cluster without restoring the database from the backup. When restoring mapping, some of the mapping can fail if the mappings are not in the current database.
-detail	Displays all the devices from the XML file with all their attribute values.
-file	Specifies the absolute path or relative path and file name of the file that has backup information. If the file name starts with a slash (/) it is considered an absolute path; otherwise, it is a relative path. For backup, compressed file is created with .tar.gz extension and for cluster backups, compressed file is created with <clustername>.tar.gz extension.
-force	If this option is specified in noninteractive mode, restoration of a device that has not been successfully validated is attempted. This option cannot be used in combination with the -inter or -validate options.
-frequency	Specifies the frequency of the backup to run automatically. Note: You can add or edit the <i>crontab</i> entry for backup frequencies other than daily, weekly, monthly. A compressed file in the form file_name.XX.tar.gz is created, where <file_name> is the argument to -file and XX is a number from 01 to numfiles provided by you. The maximum numfiles value is 10. The format of the cluster backup file is file_name.XX.clustername.tar.gz
-inter	Interactively deploys each device with your confirmation. Note: User input can be taken to set properties of all drivers, adapters, and interfaces (disks, optical devices, tape devices, fibre channel SCSI controllers controllers, Ethernet adapters, Ethernet interfaces, and logical HEAs) or each category of logical or virtual devices. This includes logical devices, such as storage pools, file-backed storage pools, and optical repositories, and virtual devices such as Etherchannel, SEA, virtual server adapters, and virtual server fibre channel adapters.
-list	This option displays backup files from either the default location <i>/home/padmin/cfgbackups</i> or a user-specified location.
-mapping	Displays mapping information for SEA, virtual SCSI adapters, VFC adapters, and PowerVM Active Memory Sharing paging devices.

-migrate	Migrates earlier cluster version of backup file to the current version. A new file is created with <code>_MIGRATED</code> string appended to the given filename.
-nobackup	This option removes any previously scheduled backups and stops any automatic backups.
-numfiles	When backup runs automatically, this number indicates the maximum number of backup files that can be saved. The oldest file is deleted during the next cycle of backup. If this flag is not given, the default value is 10.
-recoverdb	Recovers from the shared storage pool database corruption, either from the backup file or from the solid database backup.
-repopvs	Lists <i>hdisks</i> to be used as <i>repository</i> disks for restoring the cluster (space-separated list of <i>hdiskX</i>). The given disks must not contain a repository signature. Note: First release supports only one physical volume.
-restore	Takes backup file as input and brings the VIOS partition to the same state when the backup was taken.
-skipcluster	Restores all local devices, except cluster0.
-subfile	Specifies the node configuration file to be restored. This option must be used when the valid cluster repository exists on the disks. It cannot be used with the -repopvs option. This option is ignored if the backup file is not a cluster backup.
-type	Displays information corresponding to all instances of the device type specified. The devType can be <i>pv</i> , <i>optical</i> , <i>tape</i> , <i>controller</i> , <i>interface</i> , <i>sp</i> , <i>fbsp</i> , <i>repository</i> , <i>ethchannel</i> , <i>sea</i> , <i>svsa</i> , <i>svfca</i> , <i>vlogrepo</i> , <i>pool</i> , or <i>paging</i> . With the -restore option, the devType option can be <i>net</i> , <i>vscsi</i> , <i>npiv</i> , <i>cluster</i> , <i>vlogrepo</i> , or <i>ams</i> . When deploying a given type of device, all the dependent devices also are deployed. For example, when deploying <i>vscsi</i> , related disks, attributes are set, the corresponding storage pool is imported, and all file-backed storage pools are mounted.
-validate	Validates the devices on the server against the devices on the backed-up file. If the -inter option is specified, you are prompted to specify how to handle items that do not validate successfully. Without the -inter option, if items do not validate successfully, the -restore operation fails.
-view	Display the information of all the backed up entities.
-xmlvtds	Allows you to restore SSP mappings, which are not in SSP database but are in the backup .xml file. This option is valid only while restoring a node and not while restoring clusters.

A cluster cannot be restored on a system if the cluster or node from the cluster is removed by using the **cluster** command with the **-delete** or **-rmnode** option.

When the cluster backup is taken, the file name of the individual node backup .xml file is in the format as follows:

```
<cluster Name>MTM<Machine TYPE MODEL>P<partitionId>.xml
```

Exit Status

Table 14. Command specific return codes

Return code	Description
0	Success
-1	Failure

Examples

1. To back up all the device attributes and logical and virtual device mappings on the VIOS file called */tmp/myserverbackup*, type the command as follows:

```
viosbr -backup -file /tmp/myserverbackup
```
2. To back up all the device attributes and virtual device mappings daily on the VIOS and keep the last five backup files, type the command as follows:

```
viosbr -backup -file mybackup -frequency daily -numfiles 5
```

The backup files resulting from this command are located under *home/padmin/cfgbackups* with the names *mybackup.01.tar.gz*, *mybackup.02.tar.gz*, *mybackup.03.tar.gz*, *mybackup.04.tar.gz*, and *mybackup.05.tar.gz* for the five most recent files.

3. To display information about all the entities in a backup file called *myserverbackup.012909.tar.gz*, type the command as follows:

```
viosbr -view -file myserverbackup.012909.tar.gz
```

The system displays the output as follows:

Controllers:

Name	Phys Loc
scsi0	U787B.001.DNWFPMH-P1-C3-T1
scsi1	U787B.001.DNWFPMH-P1-C3-T2
fscsi0	U789D.001.DQD42T5-P1-C1-T1
iscsi0	U787B.001.DNWFPMH-P1-T10
lhea0	U789D.001.DQD42T5-P1
fcs0	U789D.001.DQD42T5-P1-C1-T1

Physical Volumes:

Name	Phys loc
hdisk1	U787B.001.DNWFPMH-P1-C3-T2-L4-L0
hdisk2	U789D.001.DQD90N4-P3-D2

Optical Devices:

Name	Phys loc
cd0	U78A0.001.DNWGLV2-P2-D2

Tape devices:

Name	Phys loc
rmt0	U78A0.001.DNWGLV2-P2-D1

Ethernet Interface(s):

Name
en0
en1

Etherchannels:

Name	Prim adapter(s)	Backup adapter
ent4	ent0	NONE
	ent1	

Shared Ethernet Adapters:

Name	Target Adapter	Virtual Adapter(s)
ent3	ent0	ent1 ent2

Storage Pools (*-default SP):

SP name	PV Name
testsp	hdisk1
	hdisk2

mysp*	hdisk3
	hdisk4

File-backed Storage Pools:

Name	Parent SP
myfbsp	mysp

Optical Repositories:

Name	Parent SP
VMLibrary_LV	mysp

```
VSCSI Server Adapters:
SVSA      VTD      Phys loc
vhost0    vtscsi0  U9133.55A.063368H-V4-C3
          vtopt1
vhost1    vtopt0   U9133.55A.063368H-V4-C4
          vttape0
```

```
SVFC Adapters:
Name      FC Adapter  Phys loc
vfchost0  fcs0        U9117.MMA.06AB272-V5-C17
vfchost1  -           U9117.MMA.06AB272-V5-C18
```

```
VBSD Pools:
Name
pool0
pool1
```

```
VRM Pages:
Name      StreamID
vrmpage0  0x2000011b7ec18369
vrmpage1  0x2000011b7dec9128
```

```
Virtual Log Repositories:
=====
Virtual Log Repository      State
-----
vlogrepo0                   AVAILABLE
```

- To display information for only physical disks, type the command as follows:
`viosbr -view -file myserverbackup.002.tar.gz -type pv`

The system displays the output as follows:

```
Physical Volumes:
=====
Name      Phys Loc
-----
hdisk0    U789D.001.DQD42T5-P1-C1-T1-W500507630513402B-L401040000000000
hdisk1    U789D.001.DQD42T5-P1-C1-T1-W500507630513402B-L4010400100000000
hdisk2    U789D.001.DQD42T5-P1-C1-T1-W500507630513402B-L4010400400000000
hdisk3    U789D.001.DQD42T5-P1-C1-T1-W500507630513402B-L4010405C00000000
```

- To restore all the possible devices and display a summary of deployed and nondeployed devices, type the command as follows:
`viosbr -restore -file /home/padmin/cfgbackups/myserverbackup.002.tar.gz`

The system displays the output as follows:

```
Deployed/changed devices:
<Name(s) of deployed devices>

Unable to deploy/change devices:
<Name(s) of non-deployed devices>
```

- To back up a cluster and all the nodes (that are UP), type the command as follows:
`viosbr -backup -clustername mycluster -file systemA`
- To view the contents of a cluster backup and associated nodes, type the command as follows:
`viosbr -view -clustername mycluster -file /home/padmin/cfgbackups/systemA.mycluster.tar.gz`

The system displays the output as follows:

```
Files in the cluster Backup
=====
myclusterDB
myclusterMTM8233-E8B02HV32001P2.xml
```

myclusterMTM8233-E8B02HV32001P3.xml

Details in: /home/ios/mycluster.9240654/myclusterMTM8233-E8B02HV32001P2.xml

Controllers:

```
=====
Name          Phys Loc
----          -
iscsi0
pager0        U8233.E8B.HV32001-V3-C32769-L0-L0
vasi0         U8233.E8B.HV32001-V3-C32769
vbsd0         U8233.E8B.HV32001-V3-C32769-L0
fcs0          U5802.001.00H1180-P1-C8-T1
fcs1          U5802.001.00H1180-P1-C8-T2
sfwcomm0      U5802.001.00H1180-P1-C8-T1-W0-L0
sfwcomm1      U5802.001.00H1180-P1-C8-T2-W0-L0
fscsi0        U5802.001.00H1180-P1-C8-T1
ent0          U5802.001.00H1180-P1-C2-T1
fscsi1        U5802.001.00H1180-P1-C8-T2
ent1          U5802.001.00H1180-P1-C2-T2
ent2          U5802.001.00H1180-P1-C2-T3
ent3          U5802.001.00H1180-P1-C2-T4
sfw0
fcnet0        U5802.001.00H1180-P1-C8-T1
fcnet1        U5802.001.00H1180-P1-C8-T2
```

Physical Volumes:

```
=====
Name          Phys loc
----          -
caa_private0  U5802.001.00H1180-P1-C8-T1-W500507630508C075-L4012400400000000
hdisk0        U5802.001.00H1180-P1-C8-T1-W500507630508C075-L4003402500000000
hdisk1        U5802.001.00H1180-P1-C8-T1-W500507630508C075-L4003402600000000
hdisk2        U5802.001.00H1180-P1-C8-T1-W5005076305088075-L4004400100000000
hdisk5        U5802.001.00H1180-P1-C8-T1-W500507630508C075-L4012400600000000
hdisk6        U5802.001.00H1180-P1-C8-T1-W500507630508C075-L4012400700000000
cldisk1       U5802.001.00H1180-P1-C8-T1-W500507630508C075-L4012400500000000
```

Optical Devices:

```
=====
Name          Phys loc
----          -
```

Tape devices:

```
=====
Name          Phys loc
----          -
```

Ethernet Interfaces:

```
=====
Name
----
en0
en1
en2
en3
```

Storage Pools:

```
=====
SP name        PV Name
-----
rootvg         hdisk2
caavg_private  caa_private0
```

Virtual Server Adapters:

```
=====
```

SVSA	Phys Loc	VTD
vhost0	U8233.E8B.HV32001-V3-C2	
vhost1	U8233.E8B.HV32001-V3-C3	
vhost2	U8233.E8B.HV32001-V3-C4	
vhost3	U8233.E8B.HV32001-V3-C5	

Cluster:

```

=====
Name      State
----      -
cluster0  UP

```

```

Cluster Name      Cluster ID
-----
mycluster         ce7dd2a0e70911dfac3bc32001017779

```

```

Attribute Name      Attribute Value
-----
node_uuid           77ec1ca0-a6bb-11df-8cb9-00145ee81e01
clvdisk             16ea129f-0c84-cdd1-56ba-3b53b3d45174

```

Virtual Log Repositories:

```

=====
Virtual Log Repository  State
-----
vlogrepo0              AVAILABLE

```

Details in: /home/ios/mycluster.9240654/myclusterMTM8233-E8B02HV32001P3.xml

Controllers:

```

=====
Name      Phys Loc
----      -
iscsi0
pager0    U8233.E8B.HV32001-V3-C32769-L0-L0
vasi0     U8233.E8B.HV32001-V3-C32769
vbsd0     U8233.E8B.HV32001-V3-C32769-L0
fcs0      U5802.001.00H1180-P1-C8-T1
fcs1      U5802.001.00H1180-P1-C8-T2
sfwcomm0  U5802.001.00H1180-P1-C8-T1-W0-L0
sfwcomm1  U5802.001.00H1180-P1-C8-T2-W0-L0
fscsi0    U5802.001.00H1180-P1-C8-T1
ent0      U5802.001.00H1180-P1-C2-T1
fscsi1    U5802.001.00H1180-P1-C8-T2
ent1      U5802.001.00H1180-P1-C2-T2
ent2      U5802.001.00H1180-P1-C2-T3
ent3      U5802.001.00H1180-P1-C2-T4
sfw0
fcnet0    U5802.001.00H1180-P1-C8-T1
fcnet1    U5802.001.00H1180-P1-C8-T2

```

Physical Volumes:

```

=====
Name      Phys Loc
----      -
caa_private0  U5802.001.00H1180-P1-C8-T1-W500507630508C075-L4012400400000000
disk0        U5802.001.00H1180-P1-C8-T1-W500507630508C075-L4003402500000000
hdisk1       U5802.001.00H1180-P1-C8-T1-W500507630508C075-L4003402600000000
hdisk2       U5802.001.00H1180-P1-C8-T1-W5005076305088075-L4004400100000000
hdisk5       U5802.001.00H1180-P1-C8-T1-W500507630508C075-L4012400600000000
hdisk6       U5802.001.00H1180-P1-C8-T1-W500507630508C075-L4012400700000000
cldisk1      U5802.001.00H1180-P1-C8-T1-W500507630508C075-L4012400500000000

```

Optical Devices:

```

=====
Name          Phys Loc
-----          -

```

Tape Devices:

```

=====
Name          Phys Loc
-----          -

```

Ethernet Interfaces:

```

=====
Name
-----
en0
en1
en2
en3

```

Storage Pools:

```

=====
SP Name      PV Name
-----      -
rootvg      hdisk2
caavg_private caa_private0

```

Virtual Server Adapters:

```

=====
SVSA      Phys Loc      VTD
-----
vhost0    U8233.E8B.HV32001-V3-C2
vhost1    U8233.E8B.HV32001-V3-C3
vhost2    U8233.E8B.HV32001-V3-C4
vhost3    U8233.E8B.HV32001-V3-C5

```

Cluster:

```

=====
Cluster    State
-----    -
cluster0   UP

```

```

Cluster Name      Cluster ID
-----
mycluster         ce7dd2a0e70911dfac3bc32001017779

```

```

Attribute Name      Attribute Value
-----
node_uuid           77ec1ca0-a6bb-11df-8cb9-00145ee81e01
clvdisk             16ea129f-0c84-cdd1-56ba-3b53b3d45174

```

8. To view the details of a cluster backup and associated nodes, type the command as follows:

```
viosbr -view -clustername mycluster -file /home/padmin/cfgbackups/systemA.mycluster.tar.gz -detail
```
9. To restore a particular node within the cluster, type the command as follows:

```
viosbr -restore -clustername mycluster -file systemA.mycluster.tar.gz -subfile myclusterMTM8233-E8B02HV32001P3.xml
```
10. To restore a cluster and its nodes, type the command as follows:

```
viosbr -restore -clustername mycluster -file systemA.mycluster.tar.gz -repopvs hdisk5
```
11. To restore shared storage pool virtual target devices that are in the backup file but not in the shared storage pool database, type the command as follows:

```
viosbr -restore -clustername mycluster -file systemA.mycluster.tar.gz -subfile myclusterMTM8233-E8B02HV32001P3.xml -xmlvtds
```
12. To restore only the shared storage pool database from the backup file, type the command as follows:

```
viosbr -recoverdb -clustername mycluster -file systemA.mycluster.tar.gz
```


13. To restore only the shared storage pool database from the automated database backups, type the command as follows:

```
viosbr -recoverdb --clustername mycluster
```
14. To migrate the older cluster backup file, type the command as follows:

```
viosbr -migrate -file systemA.mycluster.tar.gz
```

A new file `systemA_MIGRATED.mycluster.tar.gz` is created.
15. To restore legacy device mappings on a node, which is in cluster using cluster backup file, type the command as follows:

```
viosbr -restore -clustername mycluster -file systemA.mycluster.tar.gz -subfile myclusterMTM8233-E8B02HV32001P3.xml -skipcluster
```
16. To restore cluster from backup file but use the database, which exists on the system, type the command as follows:

```
viosbr -restore -clustername mycluster -file systemA.mycluster.tar.gz -repopvs hdisk5 -currentdb
```

Related Information

The `alt_root_vg` command, the `backup` command, the `backupios` command, the `lsdev` command, the `lsmmap` command, the `restore` command, the `restorevgstruct` command, the `savevgstruct` command, and the `save_base` command.

viosecure command

Purpose

Activates, deactivates, and displays security hardening rules. Configures, unconfigures, or displays firewall settings.

Syntax

```
viosecure -level LEVEL [-apply] [ -rule ruleName] [-outfile filename]
```

```
viosecure -view [ -actual | -latest] [-rule ruleName | -nonint]
```

```
viosecure -file rulesFile
```

```
viosecure -changedRules
```

```
viosecure -undo
```

```
viosecure -firewall on [[ -force] -reload]
```

```
viosecure -firewall { allow | deny } -port number [-interface ifname] [-address IPaddress] [-timeout Timeout] [-remote]
```

```
viosecure -firewall view [-fmt delimiter]
```

Description

The `viosecure` activates, deactivates, and displays security hardening rules. By default, none of the security strengthening features are activated after installation. Upon running the `viosecure` command, the command guides the user through the proper security settings, which can be high, medium, or low. After this initial selection, a menu is displayed itemizing the security configuration options that are associated

with the selected security level in sets of 10. These options can be accepted in whole, individually toggled off or on, or ignored. After any changes, **viosecure** continues to apply the security settings to the computer system.

The **viosecure** command also configures, unconfigures, and displays network firewall settings. Using the **viosecure** command, you can activate and deactivate specific ports and specify the interface and IP address from which connections are allowed.

Note: For a complete listing of rules that apply to each security level, see AIX Security Expert.

Flags

-level <i>LEVEL</i>	Specifies the security <i>LEVEL</i> settings to choose, where <i>LEVEL</i> is <i>low</i> , <i>medium</i> , <i>high</i> , or <i>default</i> . The default <i>LEVEL</i> deactivates any previous security <i>LEVEL</i> system settings. Except for the default <i>LEVEL</i> , ten security <i>LEVEL</i> settings are displayed at a time. The user can then choose the desired security settings by entering comma-separated numbers, the word ALL to choose all of the settings, A to apply the selected settings, NONE to choose none of the settings, q to exit, or h for help. The security settings chosen are then applied to the system.
-view	Displays the current security level settings. All of the security setting names start with 3 characters <i>Xls</i> , where <i>X</i> means <i>l</i> (<i>low</i>), <i>m</i> (<i>medium</i>), <i>h</i> (<i>high</i>), or <i>d</i> (<i>default</i>). For example, the security level name <i>lls_minlen</i> is the low-level security setting for the minimum length of a password.
-apply	Applies all of the <i>LEVEL</i> security settings to the system. There is no user-selectable option.
-nonint	Specifies non-interactive mode.
-outfile	Specifies that security rules be sent to a specific file.
-file	Specifies the security rules file to be applied.
-rule	Specifies the name of the rule, for example, <i>lls_maxexpired</i> , <i>hls_telnet</i> .
-changedRules	Displays new values, if they are changed by any other commands.
-latest	Displays last applied rules.
-actual	Displays actual values for the rules that are set.
-undo	Undoes the latest security settings that have been applied. Use -latest to view the latest security settings.
-firewall on [[-force] -reload]	Configures the default firewall settings from the filter rules in ODM. If you use the reload option, then the ODM rules are deleted and the default values are loaded from the <i>/home/ios/security/viosecure.ctl</i> file. If the <i>viosecure.ctl</i> file does not exist, the force option specifies to use the hard-coded, default firewall settings.
-firewall off	Unconfigures the firewall settings and saves all the firewall filter rules to the <i>/home/padmin/viosfirewall.rules</i> file.
-firewall allow -port <i>Port</i> [-interface <i>ifname</i>] [-address <i>IPaddress</i>] [-timeout <i>Timeout</i>] [-source]	Permits IP activity per port with optional parameters according to the interface, IP address, and time that the activity is effective. The <i>Port</i> argument can be a number or a service name from the <i>/etc/services</i> file. The remote option specifies that the port is a remote port. All IP activity to and from that remote port is allowed. The default is all IP activity to and from a local port is allowed. The timeout period can be specified as a number (in seconds), or with a number followed by <i>m</i> (minutes), <i>h</i> (hours), or <i>d</i> (days). The maximum timeout period is 30 days.
-firewall deny -port <i>Port</i> [-interface <i>Ifname</i>] [-address <i>IPaddress</i>] [-timeout <i>Timeout</i>] [-source]	Removes a previous firewall allow setting. The <i>Port</i> argument can be a number or a service name from the <i>/etc/services</i> file. If -port 0 is specified, then all allow settings are removed. The remote option specifies that the port is the remote port. The default is local port. The timeout period can be specified as a number (in seconds), or with a number followed by <i>m</i> (minutes), <i>h</i> (hours) or <i>d</i> (days). The maximum timeout period is 30 days.
-firewall view [-fmt <i>delimiter</i>]	Displays the current allowable ports. If the -fmt option is specified, then it divides output by a user-specified delimiter.

Examples

1. To display the high system security settings, and to select which of the high security settings to apply to the system, type:
`viosecure -level high`
2. To apply all of the 'high' system security settings to the system, type:
`viosecure -level high -apply`
3. To display the current system security settings, type:
`viosecure -view`
4. To unconfigure the previous system security settings, type:
`viosecure -level default`
5. To allow IP activity on the ftp-data, ftp, ssh, www, https, rmc, and cimon ports, and to deny other IP activity, type:
`viosecure -firewall on`
6. To allow IP activity on all ports, type:
`viosecure -firewall off`
7. To allow users from IP address 10.10.10.10 to rlogin, type:
`viosecure -firewall allow -port login -address 10.10.10.10`
8. To allow users to rlogin for seven days, type:
`viosecure -firewall allow -port login -timeout 7d`
9. To allow rsh client activity through interface en0, type:
`viosecure -firewall allow -port 514 -interface en0 -remote`
10. To remove the rule that allows users from IP address 10.10.10.10 to rlogin, type:
`viosecure -firewall deny -port login -address 10.10.10.10`
11. To display the list of allowed ports, type:
`viosecure -firewall view`
12. To undo the security settings that have been applied, type:
`viosecure -undo /etc/security/aixpert/core/undo.xml`

Note: This command removes all the security settings specified in the `undo.xml` file.
13. To write low-level security rules to *myfile*, type:
`viosecure -level low -outfile myfile`
14. To apply security rules from *myfile*, type:
`viosecure -file myfile`
15. To display recently applied rules, type:
`viosecure -view -latest`
16. To display rules that are changed after they are applied with the **viosecure** command, type:
`viosecure -changedRules`
17. To apply the single rule *lls_maxage*, type:
`viosecure -level low -rule lls_maxage -apply`
18. To view the applied rule *ll_maxage*, type:
`viosecure -view -rule lls_maxage`
19. To view the rule *lls_maxage* if it exists during last applied rules, type:
`viosecure -view -rule lls_maxage -latest`
20. To display the actual values of rules, even if they are changed by another command, type:
`viosecure -view -actual`

viostat command

Purpose

Reports Central Processing Unit (CPU) statistics, asynchronous input/output (AIO) and input/output statistics for the entire system, adapters, tty devices, disks and CD-ROMs.

Syntax

viostat

viostat [**-sys**] [**-adapter**] [**-tty** | **-disk**] [**-path**] [**-time**] [*PhysicalVolume ...*] [*Interval* [*Count*]]

Description

The **viostat** command is used for monitoring system input/output device loading by observing the time the physical disks are active in relation to their average transfer rates. The **viostat** command generates reports that can be used to change system configuration to better balance the input/output load between physical disks and adapters.

The first report generated by the **viostat** command provides statistics concerning the time since the system was booted. Each subsequent report covers the time since the previous report. All statistics are reported each time the **viostat** command is run. The report consists of a tty and CPU header row followed by a row of tty and CPU statistics. On multiprocessor systems, CPU statistics are calculated system-wide as averages among all processors.

If the **-sys** flag is specified, a system-header row is displayed followed by a line of statistics for the entire system. The hostname of the system is printed in the system-header row.

If the **-adapter** flag is specified, an adapter-header row is displayed followed by a line of statistics for the adapter. This will be followed by a disk-header row and the statistics of all the disks/CD-ROMs connected to the adapter. Such reports are generated for all the disk adapters connected to the system.

A disks header row is displayed followed by a line of statistics for each disk that is configured. If the *PhysicalVolume* parameter is specified, only those names specified are displayed.

If the *PhysicalVolume* parameter is specified, one or more alphabetic or alphanumeric physical volumes can be specified. If the *PhysicalVolume* parameter is specified, the tty and CPU reports are displayed and the disk report contains statistics for the specified drives. If a specified drive name is not found, the report lists the specified name and displays the message Drive Not Found. If no Drive Names are specified, the report contains statistics for all configured disks and CD-ROMs. If no drives are configured on the system, no disk report is generated. The first character in the *PhysicalVolume* parameter cannot be numeric.

The *Interval* parameter specifies the amount of time in seconds between each report. The first report contains statistics for the time since system startup (boot). Each subsequent report contains statistics collected during the interval since the previous report. The *Count* parameter can be specified in conjunction with the *Interval* parameter. If the *Count* parameter is specified, the value of count determines the number of reports generated at *Interval* seconds apart. If the *Interval* parameter is specified without the *Count* parameter, the **viostat** command generates reports continuously.

The **viostat** command reports number of physical processors consumed (*phyc*) and the percentage of entitlement consumed (*% entc*) in shared processor environments. These metrics will only be displayed in shared processor environments.

The **viostat** command is useful in determining whether a physical volume is becoming a performance bottleneck and if there is potential to improve the situation. The % utilization field for the physical volumes indicates how evenly the file activity is spread across the drives. A high % utilization on a physical volume is a good indication that there may be contention for this resource. Since the CPU utilization statistics are also available with the viostat report, the percentage of time the CPU is in I/O wait can be determined at the same time. Consider distributing data across drives if the I/O wait time is significant and the disk utilization is not evenly distributed across volumes.

Reports

The viostat command generates four types of reports, the tty and CPU Utilization report, the Disk Utilization report, the System throughput report and the Adapter throughput report.

tty and CPU Utilization Report

The first report generated by the viostat command is the tty and CPU Utilization Report. For multiprocessor systems, the CPU values are global averages among all processors. Also, the I/O wait state is defined system-wide and not per processor. The report has the following format:

Statistic	Description
tin	Shows the total number of characters read by the system for all ttys.
tout	Shows the total number of characters written by the system to all ttys.
% user	Shows the percentage of CPU utilization that occurred while executing at the user level (application).
% sys	Shows the percentage of CPU utilization that occurred while executing at the system level (kernel).
% idle	Shows the percentage of time that the CPU or CPUs were idle and the system did not have an outstanding disk I/O request.
% iowait	Shows the percentage of time that the CPU or CPUs were idle during which the system had an outstanding disk I/O request.

This information is updated at regular intervals by the kernel (typically sixty times per second). The tty report provides a collective account of characters per second received from all terminals on the system as well as the collective count of characters output per second to all terminals on the system.

Disk Utilization Report

The second report generated by the viostat command is the Disk Utilization Report. The disk report provides statistics on a per physical disk basis. The report has a format similar to the following:

Statistic	Description
% tm_act	Indicates the percentage of time the physical disk was active (bandwidth utilization for the drive).
Kbps	Indicates the amount of data transferred (read or written) to the drive in KB per second.
tps	Indicates the number of transfers per second that were issued to the physical disk. A transfer is an I/O request to the physical disk. Multiple logical requests can be combined into a single I/O request to the disk. A transfer is of indeterminate size.
Kb_read	The total number of KB read.
Kb_wrtn	The total number of KB written.

Statistics for CD-ROM devices are also reported.

For large system configurations where a large number of disks are configured, the system can be configured to avoid collecting physical disk input/output statistics when the **viostat** command is not executing. If the system is configured in the above manner, the first Disk report displays the message `Disk History Since Boot Not Available` instead of the disk statistics. Subsequent interval reports generated by the **viostat** command contain disk statistics collected during the report interval. Any tty and CPU statistics after boot are unaffected.

System Throughput Report

This report is generated if the `-sys` flag is specified. This report provides statistics for the entire system. This report has the following format:

Statistic	Description
Kbps	Indicates the amount of data transferred (read or written) in the entire system in KB per second.
tps	Indicates the number of transfers per second issued to the entire system.
Kb_read	The total number of KB read from the entire system.
Kb_wrtn	The total number of KB written to the entire system.

Adapter Throughput Report

This report is generated if the `-adapter` flag is specified. This report provides statistics on an adapter-by-adapter basis. This report has the following format:

Statistic	Description
Kbps	Indicates the amount of data transferred (read or written) in the adapter in KB per second.
tps	Indicates the number of transfers per second issued to the adapter.
Kb_read	The total number of KB read from the adapter.
Kb_wrtn	The total number of KB written to the adapter.

Disk Input/Output History

To improve performance, the collection of disk input/output statistics has been disabled. To enable the collection of this data, type:

```
chdev -dev sys0 -attr iostat=true
```

To display the current settings, type:

```
lsdev -dev sys0 -attr iostat
```

If the collection of disk input/output history is disabled, the first disk report of **viostat** output displays the message `Disk History Since Boot Not Available` instead of disk statistics. As before, subsequent interval reports generated by the **viostat** command contain disk statistics collected during the report interval.

Flags

-adapter	Displays the adapter throughput report. If the -adapter flag is specified with the -tty flag, the tty and CPU report is displayed, followed by the adapter throughput report. Disk Utilization reports of the disks connected to the adapters, will not be displayed after the Adapter throughput report. If the -adapter flag is specified with the -disk flag, tty and CPU report will not be displayed. If Physical Volume parameter is specified, the Disk Utilization Report of the specified Physical volume will be printed under the corresponding adapter to which it belongs.
-disk -path	The -disk flag is exclusive of the -tty flag and displays only the disk utilization report. The -m flag will print the path statistics for the following <ul style="list-style-type: none">• Paths to MPIO (Multi-Path I/O) enabled devices.• Paths in the ESS machines. The throughput is per device. The throughout for all the paths to that device follow the throughput of that device. For ESS machines, the vpaths will be treated as disks and hdisks will be treated as Paths. Internally the vpaths are actually disks and hdisks are the paths to them. For MPIO enabled devices, the path name will be represented as Path0, Path1, Path2 and so on. The numbers 0, 1, 2, and so on are the path IDs provided by the lspath command. Since paths to a device can be attached to any adapter, the adapter report will report the path statistics under each adapter. The disk name will be a prefix to all the paths. For all MPIO enabled devices, the adapter report will print the path names as hdisk10_Path0, hdisk0_Path1, and so on. For all ESS Machines, the adapter report will print the path names as vpath0_hdisk3, vpath10_hdisk25, and so on.
-sys	Displays the system throughput report.
-time	Prints the time stamp next to each line of output of viostat. The time stamp displays in the HH:MM:SS format.
-tty	The -tty flag is exclusive of the -disk flag and displays only the tty and cpu usage reports.

Examples

1. To display a single history since boot report for all tty, CPU, and Disks, type:
viostat
2. To display a continuous disk report at two second intervals for the disk with the logical name disk1, type:
viostat -disk disk1 2
3. To display six reports at two second intervals for the disk with the logical name disk1, type:
viostat disk1 2 6
4. To display six reports at two second intervals for all disks, type:
viostat -disk 2 6
5. To display six reports at two second intervals for three disks named disk1, disk2, disk3, type:
viostat disk1 disk2 disk3 2 6
6. To print the System throughput report, type:
viostat -sys
7. To print the Adapter throughput report, type:
viostat -adapter
8. To print the System and Adapter throughput reports, with only the tty and CPU report (no disk reports), type:
viostat -sys -adapter -tty

- To print the System and Adapter throughput reports with the Disk Utilization reports of hdisk0 and hdisk7, type
`viostat -sys -adapter -disk hdisk0 hdisk7`
- To display time stamp next to each line of output of viostat, type:
`viostat -time`

vlantfw command

Purpose

Displays or clears the IP and MAC mappings

Syntax

```
vlantfw -h | -s | -t | -d | -f | -q
```

Description

The **vlantfw** command displays or clears the IP and Media Access Control (MAC) mappings.

Flags

<code>-d</code>	Displays all the IP mappings.
<code>-f</code>	Removes or flushes all the IP mappings.
<code>-s</code>	Starts the trusted firewall.
<code>-t</code>	Stops the trusted firewall.
<code>-q</code>	Queries the secure virtual machine status.

Exit Status

This command returns the following exit values:

0	Successful completion.
>0	An error occurred.

Examples

- To display all the IP mappings, type the command as follows:
`vlantfw -d`
- To remove all the IP mappings, type the command as follows:
`vlantfw -f`
- To start trusted firewall, type the command as follows:
`vlantfw -s`
- To stop trusted firewall, type the command as follows:
`vlantfw -t`
- To check the status of a secure virtual machine, type the command as follows:
`vlantfw -q`

Related Information

The **chvfilt** command, the **genvfilt** command, the **lsvfilt** command, the **mkvfilt** command, and the **rmvfilt** command.

vmstat command

Purpose

Reports statistics about kernel threads, virtual memory, disks, traps, and processor activity.

Syntax

To report virtual memory statistics:

```
vmstat [ -f ] [-i ][-s ][-I ][-t ][-v ][-w ][-l ][{-p | -P | } pagesize | ALL ] ALL] [PhysicalVolume...] [Interval  
[Count ] ]
```

Description

The **vmstat** command reports statistics about kernel threads, virtual memory, disks, traps and CPU activity. Reports 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 expressed as percentages, and as sums otherwise.

Flags

- f Reports the number of forks since system startup.
- i Displays the number of interrupts taken by each device since system startup.
Note: The **-I**, **-t**, **-w**, and **-l** flags are ignored when they are specified with the **-i** flag.
- s Writes to standard output the contents of the sum structure, which contains an absolute count of paging events since system initialization. The **-s** flag can only be used with the **-v** flag. These events are described as follows:
 - address translation faults**
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.
 - 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 represents the total amount of real I/O initiated by the virtual memory manager.
 - page outs**
Incremented for each page 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 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.

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

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.

executable-filled page faults

Incremented for each instruction page fault.

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.

revolutions of the clock hand

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

pages freed by the clock

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

backtracks

Incremented for each page fault that occurs while resolving a previous page fault.

free frame waits

Incremented each time a process is waited by VMM while free frames are gathered.

extend XPT waits

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

pending I/O waits

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

start I/Os

Incremented for each read or write I/O request initiated by VMM. This count should equal the sum of page-ins and page-outs.

iodones

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

CPU context switches

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

device interrupts

Incremented on each hardware interrupt

software interrupts

Incremented on each software interrupt. A software interrupt is a machine instruction similar to a hardware interrupt that saves some state and branches to a service routine. System calls are implemented with software interrupt instructions that branch to the system call handler routine.

-s (continued)

decrementer interrupts

Incremented on each decrementer interrupt.

mpc send interrupts

Incremented on each mpc send interrupt

mpc receive interrupts

Incremented on each mpc receive interrupt

phantom interrupts

Incremented on each phantom interrupt

traps Not maintained by the operating system.

syscalls

Incremented for each system call.

-I Displays I/O oriented view with the new columns of output, **p** under heading **kthr**, and columns **fi** and **fo** under heading **page** instead of the columns **re** and **cy** in the **page** heading.

-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 will not be printed if **-f**, **-s**, or **-i** flags are specified.

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

memory pages

Size of real memory in number of 4 KB pages.

lrutable pages

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

free pages

Number of free 4 KB pages.

memory pools

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

pinned pages

Number of pinned 4 KB pages

maxpin percentage

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

minperm percentage

Tuning parameter (managed using **vmo**) in percentage of real memory. This specifies the point below which file pages are protected from the re-page algorithm.

maxperm percentage

Tuning parameter (managed using **vmo**) in percentage of real memory. This specifies the point above which the page stealing algorithm steals only file pages.

numperm percentage

Percentage of memory currently used by the file cache.

file pages

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

compressed percentage

Percentage of memory used by compressed pages.

	compressed pages	Number of compressed memory pages.
	numclient percentage	Percentage of memory occupied by client pages.
	maxclient percentage	Tuning parameter (managed using vmo) specifying the maximum percentage of memory which can be used for client pages.
-v (continued)	client pages	Number of client pages.
	remote pageouts scheduled	Number of pageouts scheduled for the client file systems.
	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.
	paging space I/Os blocked with no psbuf	Number of paging space I/O requests blocked because no psbuf was available. Psbufs are pinned memory buffers used to hold I/O requests at the virtual memory manager layer.
	filesystem I/Os blocked with no fsbuf	Number of file system I/O requests blocked because no fsbuf was available. Fsbuf are pinned memory buffers used to hold I/O requests in the file system layer.
	client filesystem I/Os blocked with no fsbuf	Number of client file system I/O requests blocked because no fsbuf was available. NFS (Network File System) and VxFS (Veritas) are client file systems. Fsbuf are pinned memory buffers used to hold I/O requests in the file system layer.
	external pager filesystem I/Os blocked with no fsbuf	Number of external pager client file system I/O requests blocked because no fsbuf was available. JFS2 is an external pager client file system. Fsbuf are pinned memory buffers used to hold I/O requests in the file system layer.
-w		Display the report in wide mode
-l		Displays an additional "large-page" section with the <i>alp</i> and <i>flp</i> columns.
-p <i>pagesize</i>		Appends the VMM statistics for the specified page size to the regular vmstat output.
-P <i>pagesize</i>		Displays only the VMM statistics which are relevant for the specified page size.
Interval		Specifies the amount of time in seconds between each report.
Count		Determines the number of reports generated and the number of seconds apart.

Exit Status

The following exit values are returned:

0	Successful completion.
>0	Invalid flag, argument, or command failure

Examples

1. To display a summary of the statistics since boot, type: :
vmstat
2. To display five summaries at 2-second intervals, type:
vmstat 2 5

3. To display a summary of the statistics since boot including statistics for logical disks `hdisk1` and `hdisk2`, type
`vmstat hdisk1 hdisk2`
4. To display fork statistics, type:
`vmstat -f`
5. To display the count of various events, type:
`vmstat -s`
6. To display time-stamp next to each column of output of **vmstat**, type:
`vmstat -t`
7. To display the I/O oriented view with an alternative set of columns, type:
`vmstat -I`
8. To display all the VMM statistics available, type:
`vmstat -v -s`
9. To display the large-page section with the *alp* and *flp* columns at 8-second intervals, type:
`vmstat -l 8`
10. To display the VMM statistics specific to a particular page size (in the example, 4K), type:
`vmstat -p 4K`
11. To display the VMM statistics for all page sizes that are supported on the system, type:
`vmstat -p ALL`

or
`vmstat -p all`
12. To display only the VMM statistics for a particular page size (in this example, 4K), type:
`vmstat -P 4K`
13. To display only the per-page breakdown of VMM statistics for all supported page sizes, type:
`vmstat -P ALL`

or
`vmstat -P all`

Related Information

The **cfgdev** command, the **chdev** command, the **chpath** command, the **lsdev** command, the **lsmap** command, and the **rmdev** command.

wkldagent command

Purpose

Starts, stops, or queries the state of the Workload Manager Agent.

Syntax

wkldagent -start | -status | -stop

Description

The **wkldagent** command starts, stops, and queries the state of the Workload Manager Agent. The Workload Manager 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.

The Workload Manager must be started using the **wkldmgr** command before the **wkldagent** command is run. Daily recordings are stored in the **/home/ios/perf/wlm** directory with filenames **xmwlm.YYMMDD**, where **YY** is the year, **MM** is the month, and **DD** is the day. The **wkldout** command can be used to process Workload Manager-related recordings. All recordings cover 24-hour periods and are retained for only two days.

Flags

-start	Starts the Workload Manager Agent.
-status	Displays the state of the Workload Manager Agent, either running or stopped.
-stop	Stops the Workload Manager Agent.

Exit Status

0	The command completed successfully
>0	An error occurred.

Examples

1. To start the Workload Manager Agent, type:
`wkldmgr -start`
2. To check whether the Workload Manager Agent is currently active, type:
`wkldmgr -status`
3. To stop the Workload Manager Agent, type:
`wkldmgr -stop`

Related Information

The **topas** command, the **wkldmgr** command, and the **wkldout** command.

wkldmgr command

Purpose

Starts or stops Workload Manager.

Syntax

```
wkldmgr -start | -status | -stop
```

Description

The **wkldmgr** command starts, stops, and queries the state of the Workload Manager. Starting the Workload Manager is necessary for the **-cecdisp** option of the **topas** command to work properly.

Flags

-start	Starts the Workload Manager.
-status	Displays the state of the Workload Manager, either running or stopped.
-stop	Stops the Workload Manager.

Exit Status

0	The command completed successfully
---	------------------------------------

>0 An error occurred.

Examples

1. To start the Workload Manager, type:
`wkldmgr -start`
2. To check whether the Workload Manager is currently active, type:
`wkldmgr -status`
3. To stop the Workload Manager, type:
`wkldmgr -stop`

Related Information

The **topas** command, the **wkldagent** command, and the **wkldout** command.

wkldout command

Purpose

Provides post-processing of the recordings made by the Workload Manager Agent (**wkldagent**).

Syntax

```
wkldout [-report reportType] [-interval MM] [-beg HHMM] [-end HHMM] [-fmt [-mode modeType]] [-graph] [-BinToAscii] -filename <xmwlm_recording_file>
```

Description

The **wkldout** command provides post-processing of recordings made by **wkldagent**. It can generate reports in ASCII or spreadsheet format. The output can be formatted by using the flags detailed, summary, disk or LAN. The time interval can be specified to process the *xmwlm* file for a particular interval.

Flags

-report <i>report</i> < <i>type</i> >	Generates the report based on the type specified. The report type can be detailed, summary, LAN, or disk. The detailed report type produces the comprehensive output of the <i>xmwlm</i> file and displays the report on the screen. The summary report type produces the abstract output of the <i>xmwlm</i> file and displays on the report on the screen. The LAN report type generates the summary output related to system LAN from the given <i>xmwlm</i> file and displays on the report on the screen. The disk report type generates I/O summary information and generates the <i>xmwlm</i> file and displays on the report on the screen.
-interval <i>MM</i>	Specifies how the recording reports are split into equal size time periods. Allowed values (in minutes) are 5, 10, 15, 30, 60. If <i>-interval</i> flag is not specified, the default value is 5 minutes.
-beg <i>HHMM</i>	Indicates the beginning time in hours (HH) and minutes (MM). The range is between 0000 and 2400.
-end <i>HHMM</i>	Indicates the ending time in hours (HH) and minutes (MM). The range is between 0000 and 2400 and is greater than the begin time.
-fmt	Generates a spreadsheet format output in a file under <i>/home/ios/perf/wlm</i> in the format <i>xmwlm.YYMMDD_01</i> .
-mode < <i>type</i> >	Specifies the post-processor only output mean values by default. Other values and the full set are available via other options (min, max, mean, stdev, set).

-graph	Generates the .csv file under <i>/home/ios/perf/wlm</i> in the format <i>xmwlm.YYMMDD.csv</i> , which can be sent to the Nmon analyzer to produce graphs to help analyze and write reports. The Nmon analyzer requires Excel 2002 or later.
-BinToAscii	Converts the recording data from binary to ASCII text format and stores the output in the <i>/var/perf/pm/daily/</i> directory.
-filename <name>	Name of the wkldagent recording file. These files will be generated by the wkldagent command and resides in <i>/home/ios/perf/wlm</i> in the format <i>xmwlm.YYMMDD</i>

Exit Status

0	The command completed successfully
>0	An error occurred.

Examples

1. To display the system summary from *xmwlm.070731*, type the following command :
`wkldout -report summary -filename /home/ios/perf/wlm/xmwlm.070731`
2. To display the detailed report of the system from *xmwlm.070731*, type the following command:
`wkldout -report detailed -filename /home/ios/perf/wlm/xmwlm.070731`
3. To display the disk I/O summary from *xmwlm.070731*, type the following command:
`wkldout -report disk -filename /home/ios/perf/wlm/xmwlm.070731`
4. To display the system lan summary from *xmwlm.070731*, type the following command:
`wkldout -report lan -filename /home/ios/perf/wlm/xmwlm.070731`
5. To produce the csv file from *xmwlm.070731*, type the following command:
`wkldout -graph -filename /home/ios/perf/wlm/xmwlm.070731`
6. To display the system summary from *xmwlm.070731* with time interval 10 min, type the following command:
`wkldout -report summary -interval 10 -filename /home/ios/perf/wlm/xmwlm.070731`
7. To display the detailed report from *xmwlm.070731*, starting at 01:00 Hr and ending 02:00 Hr, type
`wkldout -report detailed -beg 0100 -end 0200
-filename /home/ios/perf/wlm/xmwlm.070731`
8. To convert the file *xmwlm.090610* from binary to ASCII format, type:
`wkldout -BinToAscii -filename /home/ios/perf/wlm/xmwlm.090610`

Related Information

The **topas** command, the **wkldmgr** command, and the **wkldagent** command.

Notices

This information was developed for products and services offered in the U.S.A.

The manufacturer may not offer the products, services, or features discussed in this document in other countries. Consult the manufacturer's representative for information on the products and services currently available in your area. Any reference to the manufacturer's product, program, or service is not intended to state or imply that only that product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any intellectual property right of the manufacturer may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any product, program, or service.

The manufacturer 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 the manufacturer.

For license inquiries regarding double-byte character set (DBCS) information, contact the Intellectual Property Department in your country or send inquiries, in writing, to the manufacturer.

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law: THIS PUBLICATION IS PROVIDED "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 states 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. The manufacturer 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 websites not owned by the manufacturer 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 product and use of those websites is at your own risk.

The manufacturer may use or distribute any of the information you supply 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 the manufacturer.

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, IBM License Agreement for Machine Code, or any equivalent agreement between us.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the

same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

Information concerning products not produced by this manufacturer was obtained from the suppliers of those products, their published announcements or other publicly available sources. This manufacturer has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to products not produced by this manufacturer. Questions on the capabilities of products not produced by this manufacturer should be addressed to the suppliers of those products.

All statements regarding the manufacturer's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

The manufacturer's prices shown are the manufacturer'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 the names and addresses used by an actual business enterprise 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 the manufacturer, 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. The manufacturer, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. The sample programs are provided "AS IS", without warranty of any kind. The manufacturer 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_.

If you are viewing this information in softcopy, the photographs and color illustrations may not appear.

Programming interface information

This Virtual I/O Server and Integrated Virtualization Manager commands publication documents intended Programming Interfaces that allow the customer to write programs to obtain the services of IBM Virtual I/O Server.

Trademarks

IBM, the IBM logo, and 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.

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Other product and service names might be trademarks of IBM or other companies.

Terms and conditions

Permissions for the use of these publications are granted subject to the following terms and conditions.

Applicability: These terms and conditions are in addition to any terms of use for the the manufacturer website.

Personal Use: You may reproduce these publications for your personal, noncommercial use provided that all proprietary notices are preserved. You may not distribute, display or make derivative works of these publications, or any portion thereof, without the express consent of the manufacturer.

Commercial Use: You may reproduce, distribute and display these publications solely within your enterprise provided that all proprietary notices are preserved. You may not make derivative works of these publications, or reproduce, distribute or display these publications or any portion thereof outside your enterprise, without the express consent of the manufacturer.

Rights: Except as expressly granted in this permission, no other permissions, licenses or rights are granted, either express or implied, to the Publications or any information, data, software or other intellectual property contained therein.

The manufacturer reserves the right to withdraw the permissions granted herein whenever, in its discretion, the use of the publications is detrimental to its interest or, as determined by the manufacturer, the above instructions are not being properly followed.

You may not download, export or re-export this information except in full compliance with all applicable laws and regulations, including all United States export laws and regulations.

THE MANUFACTURER MAKES NO GUARANTEE ABOUT THE CONTENT OF THESE PUBLICATIONS. THE PUBLICATIONS ARE PROVIDED "AS-IS" AND WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING BUT NOT LIMITED TO IMPLIED WARRANTIES OF MERCHANTABILITY, NON-INFRINGEMENT, AND FITNESS FOR A PARTICULAR PURPOSE.

