




Virtual I/O Server and Partition Load Manager Commands Reference





Virtual I/O Server and Partition Load Manager Commands Reference

Note

Before using this information and the product it supports, be sure to read the information in "Notices," on page 137 and the manual *IBM eServer™ Safety Information*, G229-9054.

First Edition (August 2004)

© Copyright International Business Machines Corporation 2004. All rights reserved.

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

Contents

Chapter 1. Virtual I/O Server commands 1

Command categories	1
Installation commands	1
Volume group commands	1
Logical volume commands	1
Physical volume commands	1
Network commands	2
Device commands	2
User ID commands	2
Security commands	2
Maintenance commands	2
Alphabetical listing of commands	3
Virtual I/O Server sub-command exit status	6
activatevg Command	7
Purpose	7
Syntax	7
Description	7
Flags	7
Exit Status	7
Examples	7
Related Information	7
backupios Command	8
Purpose	8
Syntax	8
Description	8
Flags	8
Exit Status	8
Examples	8
Related Information	9
bootlist Command	10
Purpose	10
Syntax	10
Description	10
Flags	11
Exit Status	12
Examples	12
Related Information	12
cattracerpt Command	13
Purpose	13
Syntax	13
Description	13
Flags	13
Exit Status	13
Examples	13
Related Information	13
cfgdev Command	14
Purpose	14
Syntax	14
Description	14
Flags	14
Exit Status	14
Examples	14
Related Information	14
cfglnagg Command	15
Purpose	15
Syntax	15

Description	15
Flags	15
Parameters	15
Exit Status	15
Examples	15
Related Information	15
cfgnamesrv Command	17
Purpose	17
Syntax	17
Description	17
Flags	17
Exit Status	18
Examples	18
Related Information	18
chdev Command	19
Purpose	19
Syntax	19
Description	19
Flags	19
Exit Status	19
Examples	19
Related Information	20
chlang Command	21
Purpose	21
Syntax	21
Description	21
Flags	21
Exit Status	21
Examples	21
Related Information	21
chpath Command	22
Purpose	22
Syntax	22
Description	22
Flags	22
Exit Status	23
Examples	23
Related Information	23
chuser Command	24
Purpose	24
Syntax	24
Description	24
Exit Status	25
Security	25
Examples	25
Related Information	25
chvg Command	26
Purpose	26
Syntax	26
Description	26
Flags	26
Exit Status	26
Examples	26
Related Information	26
cleargcl Command	27
Purpose	27

Syntax	27	Description	38
Description	27	Exit Status.	39
Exit Status.	27	Examples	39
Security	27	hostmap Command.	40
Examples	27	Purpose	40
Related Information	27	Syntax	40
deactivatevg Command	28	Description	40
Purpose	28	Flags	40
Syntax	28	Exit Status.	40
Description	28	Examples	41
Exit Status.	28	Related Information	41
Examples	28	hostname Command	42
Related Information	28	Purpose	42
diagmenu Command	29	Syntax	42
Purpose	29	Description	42
Syntax	29	Parameters	42
Description	29	Exit Status.	42
Related Information	29	Examples	42
entstat Command	30	Related Information	42
Purpose	30	importvg Command	43
Syntax	30	Purpose	43
Description	30	Syntax	43
Flags	30	Description	43
Parameters	30	Flags	43
Exit Status.	30	Exit Status.	43
Examples	31	Examples	43
Related Information	33	Restrictions	43
errlog Command	34	Related Information	43
Purpose	34	installios Command	44
Syntax	34	Purpose	44
Description	34	Syntax	44
Flags	34	Description	44
Exit Status.	34	Flags	44
Examples	34	ioslevel Command	45
Related Information	34	Purpose	45
exportvg Command	35	Syntax	45
Purpose	35	Description	45
Syntax	35	Exit Status.	45
Description	35	Related Information	45
Exit Status.	35	license Command	46
Examples	35	Purpose	46
Restrictions	35	Syntax	46
Related Information	35	Description	46
extendlv Command.	36	Flags	46
Purpose	36	Exit Status.	46
Syntax	36	Examples	46
Description	36	Related Information	46
Exit Status.	36	lsdev Command.	47
Examples	36	Syntax	47
Related Information	36	Description	47
extendvg Command	37	Flags	47
Purpose	37	Exit Status.	48
Syntax	37	Examples	48
Description	37	Related Information	49
Flags	37	lsfailedlogin Command	50
Exit Status.	37	Purpose	50
Examples	37	Syntax	50
Related Information	37	Description	50
fsck Command	38	Exit Status.	50
Purpose	38	Security	50
Syntax	38	Examples	50

Related Information	50	Description	65
lsgcl Command	51	Exit Status.	65
Purpose	51	Security	65
Syntax	51	Examples	65
Description	51	Related Information	65
Exit Status.	51	lsvg Command	66
Security	51	Purpose	66
Examples	51	Syntax	66
Related Information	51	Description	66
lslv Command	52	Flags	67
Purpose	52	Exit Status.	68
Syntax	52	Examples	68
Description	52	Related Information	68
Flags	53	migratepv Command	69
Exit Status.	53	Purpose	69
Examples	53	Syntax	69
Related Information	54	Description	69
lsmapi Command	55	Flags	69
Purpose	55	Exit Status.	69
Syntax	55	Examples	69
Description	55	Related Information	70
Flags	55	mirrorios Command	71
Output Field Definitions	55	Purpose	71
Exit Status.	56	Syntax	71
Examples	56	Description	71
Related Information	56	Flags	71
lsnetvc Command	57	Exit Status.	71
Purpose	57	Examples	71
Syntax	57	Related Information	71
Description	57	mklv Command	72
Parameters	57	Purpose	72
Exit Status.	57	Syntax	72
Examples	57	Description	72
Related Information	57	Flags	72
lspath Command	58	Exit Status.	72
Purpose	58	Examples	72
Syntax	58	Related Information	72
Description	58	mklvcopy Command	73
Flags	59	Purpose	73
Exit Status.	60	Syntax	73
Examples	60	Description	73
Related Information	60	Exit Status.	73
lspv Command	61	Examples	73
Purpose	61	Related Information	73
Syntax	61	mkpath Command	74
Description	61	Purpose	74
Flags	61	Syntax	74
Exit Status.	62	Description	74
Examples	62	Flags	74
Related Information	62	Exit Status.	75
lssw Command	63	Examples	75
Purpose	63	Related Information	75
Syntax	63	mktcpip Command	76
Description	63	Purpose	76
Flags	64	Syntax	76
Exit Status.	64	Description	76
Examples	64	Flags	76
Related Information	64	Exit Status.	77
lsuser Command	65	Examples	77
Purpose	65	Related Information	77
Syntax	65	mkuser Command	78

Purpose	78	Examples	96
Syntax	78	Related Information	96
Description	78	passwd Command	97
Flags	78	Purpose	97
Exit Status.	78	Syntax	97
Security	78	Description	97
Examples	78	Exit Status.	97
Related Information	78	Security	97
mkvdev Command	79	Examples	97
Purpose	79	Related Information	97
Syntax	79	ping Command	98
Description	79	Purpose	98
Flags	79	Syntax	98
Exit Status.	80	Description	98
Examples	80	Flags	98
Related Information	81	Parameters	98
mkvg Command	82	Exit Status.	98
Purpose	82	Examples	99
Syntax	82	Related Information	99
Description	82	reducevg Command	100
Flags	82	Purpose	100
Exit Status.	82	Syntax.	100
Examples	82	Description	100
Related Information	82	Flags	100
mount Command	83	Exit Status	100
Purpose	83	Examples.	100
Syntax	83	Related Information	100
Description	83	remote_management Command	101
Exit Status.	83	Purpose	101
Examples	83	Syntax.	101
Related Information	84	Description	101
netstat Command	85	Flags	101
Purpose	85	Exit Status	101
Syntax	85	Examples.	101
Description	85	Related Information	101
Flags	85	restorevgstruct Command	102
Default Display	86	Purpose	102
Exit Status.	88	Syntax.	102
Examples	88	Description	102
Related Information	88	Flags	102
oem_platform_level Command	89	Exit Status	102
Purpose	89	Examples.	102
Syntax	89	Related Information	102
Description	89	rmdev Command	103
Exit Status.	89	Purpose	103
Examples	89	Syntax.	103
Related Information	89	Description	103
oem_setup_env Command	90	Flags	103
Purpose	90	Exit Status	103
Syntax	90	Examples.	103
Description	90	Related Information	103
Exit Status.	90	rmlv Command	104
Examples	90	Purpose	104
Related Information	90	Syntax.	104
optimizenet Command	91	Description	104
Purpose	91	Flags	104
Syntax	91	Exit Status	104
Description	91	Examples.	104
Flags	91	Related Information	104
Network tunable parameters	92	rmlvcopy Command	105
Exit Status.	95	Purpose	105

Syntax	105	Purpose	114
Description	105	Syntax	114
Exit Status	105	Description	114
Examples	105	Exit Status	114
Related Information	105	Examples	114
rmrpath Command	106	starttrace Command	115
Purpose	106	Purpose	115
Syntax	106	Syntax	115
Description	106	Description	115
Flags	106	Flags	115
Exit Status	107	Exit Status	115
Examples	107	Examples	115
Related Information	107	Related Information	115
rmuser Command	108	stopnetsvc Command	116
Purpose	108	Purpose	116
Syntax	108	Syntax	116
Description	108	Description	116
Flags	108	Parameters	116
Exit Status	108	Exit Status	116
Security	108	Examples	116
Examples	108	Related Information	116
Related Information	108	stoptrace Command	117
savevgstruct Command	109	Purpose	117
Purpose	109	Syntax	117
Syntax	109	Description	117
Description	109	Exit Status	117
Exit Status	109	Examples	117
Examples	109	Related Information	117
Related Information	109	syncvg Command	118
showmount Command	110	Purpose	118
Purpose	110	Syntax	118
Syntax	110	Description	118
Description	110	Flags	118
Parameters	110	Exit Status	118
Exit Status	110	Examples	118
Examples	110	Related Information	118
Related Information	110	topas Command	119
shutdown Command	111	Purpose	119
Purpose	111	Syntax	119
Syntax	111	Description	119
Description	111	Implementation Specifics	123
Flags	111	Sample Default Output	123
Exit Status	111	Sample Full Screen Process Output	124
Security	111	Sample Full-Screen WorkLoad Management	
Examples	111	Classes Output	124
snap Command	112	Flags	125
Purpose	112	Subcommands	126
Syntax	112	Exit Status	126
Description	112	Examples	127
Exit Status	112	traceroute Command	128
Examples	112	Purpose	128
Files	112	Syntax	128
startnetsvc Command	113	Description	128
Purpose	113	Flags	129
Syntax	113	Parameters	129
Description	113	Exit Status	129
Parameters	113	Examples	129
Exit Status	113	Related Information	129
Examples	113	unmirrorios Command	130
Related Information	113	Purpose	130
startsysdump Command	114	Syntax	130

Description	130
Exit Status	130
Security	130
Related Information	130
unmount Command	131
Purpose	131
Syntax.	131
Description	131
Exit Status	131
Examples.	131
Related Information	131
updateios Command	132
Purpose	132
Syntax.	132
Description	132
Flags	132
Exit Status	132
Examples.	132
Related Information	132

**Chapter 2. Partition Load Manager
commands 133**

xlplm Command	133
Purpose	133
Syntax.	133
Description	133
Flags	133
Parameters	134
Exit Status	134
Examples.	135
xlpstat Command	135
Purpose	135
Syntax.	135
Description	135
Flags	136
Exit Status	136

Appendix. Notices 137
Trademarks 138
Terms and conditions for downloading and printing information 139

Index 141

Chapter 1. Virtual I/O Server commands

This section lists all Virtual I/O Server commands. The first list has them categorized into high-level administrative categories. The next list has every command listed in alphabetical order.

Command categories

Installation commands

The following commands are the Virtual I/O Server installation sub-commands:

- “ioslevel Command” on page 45
- “license Command” on page 46
- “lssw Command” on page 63
- “oem_platform_level Command” on page 89
- “oem_setup_env Command” on page 90
- “remote_management Command” on page 101
- “updateios Command” on page 132

Volume group commands

The following commands are the Virtual I/O Server volume group sub-commands:

- “activatevg Command” on page 7
- “chvg Command” on page 26
- “deactivatevg Command” on page 28
- “exportvg Command” on page 35
- “extendvg Command” on page 37
- “importvg Command” on page 43
- “lsvg Command” on page 66
- “mirrorios Command” on page 71
- “mkvg Command” on page 82
- “syncvg Command” on page 118
- “unmirrorios Command” on page 130

Logical volume commands

- “extendlv Command” on page 36
- “lslv Command” on page 52
- “mklv Command” on page 72
- “mklvcopy Command” on page 73
- “rmlv Command” on page 104
- “rmlvcopy Command” on page 105

Physical volume commands

- “lspv Command” on page 61
- “migratepv Command” on page 69

Network commands

- “cfdnagg Command” on page 15
- “cfdnamesrv Command” on page 17
- “entstat Command” on page 30
- “hostmap Command” on page 40
- “hostname Command” on page 42
- “lsnetnvc Command” on page 57
- “mktcpip Command” on page 76
- “netstat Command” on page 85
- “optimizenet Command” on page 91
- “ping Command” on page 98
- “startnetnvc Command” on page 113
- “stopnetnvc Command” on page 116
- “traceroute Command” on page 128

Device commands

- “cfdgdev Command” on page 14
- “chdev Command” on page 19
- “chpath Command” on page 22
- “lsdev Command” on page 47
- “lsmap Command” on page 55
- “lspath Command” on page 58
- “mkpath Command” on page 74
- “mkvdev Command” on page 79
- “rmdev Command” on page 103
- “rmpath Command” on page 106

User ID commands

- “chuser Command” on page 24
- “lsuser Command” on page 65
- “mkuser Command” on page 78
- “rmuser Command” on page 108
- “passwd Command” on page 97

Security commands

- “cleargcl Command” on page 27
- “lsfailedlogin Command” on page 50
- “lsgcl Command” on page 51

Maintenance commands

- “backupios Command” on page 8
- “bootlist Command” on page 10
- “cattracerpt Command” on page 13
- “chlang Command” on page 21
- “diagmenu Command” on page 29

- “errlog Command” on page 34
- “fsck Command” on page 38
- “mount Command” on page 83
- “restorevgstruct Command” on page 102
- “savevgstruct Command” on page 109
- “shutdown Command” on page 111
- “showmount Command” on page 110
- “snap Command” on page 112
- “startsysdump Command” on page 114
- “starttrace Command” on page 115
- “stoptrace Command” on page 117
- “topas Command” on page 119
- “unmount Command” on page 131

Alphabetical listing of commands

The following commands are described in this section:

- “activatevg Command” on page 7
- “backupios Command” on page 8
- “bootlist Command” on page 10
- “cattracerpt Command” on page 13
- “cfgdev Command” on page 14
- “cfglnagg Command” on page 15
- “cfgnamesrv Command” on page 17
- “chdev Command” on page 19
- “chlang Command” on page 21
- “chpath Command” on page 22
- “chuser Command” on page 24
- “chvg Command” on page 26
- “cleargcl Command” on page 27
- “deactivatevg Command” on page 28
- “diagmenu Command” on page 29
- “entstat Command” on page 30
- “errlog Command” on page 34
- “exportvg Command” on page 35
- “extendlv Command” on page 36
- “extendvg Command” on page 37
- “fsck Command” on page 38
- “hostmap Command” on page 40
- “hostname Command” on page 42
- “importvg Command” on page 43
- “ioslevel Command” on page 45
- “license Command” on page 46
- “lsdev Command” on page 47
- “lsfailedlogin Command” on page 50
- “lsgcl Command” on page 51

- “lslv Command” on page 52
- “lsmmap Command” on page 55
- “lsnetsh Command” on page 57
- “lspv Command” on page 61
- “lssw Command” on page 63
- “lsuser Command” on page 65
- “lsvg Command” on page 66
- “migratepv Command” on page 69
- “mirrorios Command” on page 71
- “mklv Command” on page 72
- “mklvcopy Command” on page 73
- “mkpath Command” on page 74
- “mktcpip Command” on page 76
- “mkuser Command” on page 78
- “mkvdev Command” on page 79
- “mkvg Command” on page 82
- “mount Command” on page 83
- “netstat Command” on page 85
- “oem_platform_level Command” on page 89
- “oem_setup_env Command” on page 90
- “optimizenet Command” on page 91
- “passwd Command” on page 97
- “ping Command” on page 98
- “reducevg Command” on page 100
- “remote_management Command” on page 101
- “restorevgstruct Command” on page 102
- “rmdev Command” on page 103
- “rmlv Command” on page 104
- “rmlvcopy Command” on page 105
- “rmpath Command” on page 106
- “rmuser Command” on page 108
- “savevgstruct Command” on page 109
- “showmount Command” on page 110
- “shutdown Command” on page 111
- “snap Command” on page 112
- “startnetsh Command” on page 113
- “startsysdump Command” on page 114
- “starttrace Command” on page 115
- “stopnetsh Command” on page 116
- “stoptrace Command” on page 117
- “syncvg Command” on page 118
- “topas Command” on page 119
- “traceroute Command” on page 128
- “unmirrorios Command” on page 130
- “unmount Command” on page 131

- “updateios Command” on page 132

Virtual I/O Server sub-command exit status

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

Return Code	Description
0	Success
1	Syntax Error
2	Not Found (this is a place holder and is not returned at this time)
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 Executes the command with out prompting the user to continue.

Exit Status

See “Virtual I/O Server sub-command exit status” on page 6.

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.

backupios Command

Purpose

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

Syntax

backupios -file *DirectoryName*

backupios -tape *Device* [**-nopak**] [**-verify**]

backupios -cd *Device* [**-udf** | **-cdformat**] [**-accept**]

Description

The **backupios** command creates a backup of the Virtual I/O Server and places it onto a filesystem, bootable tape, a CD, or a 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 CD-Recordable (CD-R) or DVD-Recordable (DVD-R, DVD-RAM). If you need to create multi-volume CDs because the image does not fit on one CD, the **backupios** command gives instructions for CD replacement and removal until all the volumes have been created.

If the **-file** flag is specified, the **backupios** command creates a system backup image to the path specified. The filesystem must be mounted prior to running the **backupios** command (See the **mount** command for details). This backup can be reinstalled from the HMC using the **installios** command.

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.

Flags

-accept	Accepts licenses.
-cd	The Virtual I/O Server backup is placed onto a CD or DVD.
-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.
-nopak	Disables software packing of the files as they are backed up. Some tape drives use their own packing or compression algorithms.
-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 “Virtual I/O Server sub-command exit status” on page 6.

Examples

1. To generate a bootable system backup to the CD-R device named **/dev/cd1**, type:
`backupios -cd /dev/cd1`
2. To generate a backup to the DVD-R or DVD-RAM device named **/dev/cd1**, type:

```
backupios -cd /dev/cd1 -cdformat
```

3. 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.

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 “Virtual I/O Server sub-command exit status” on page 6.

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
```
3. 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 “Virtual I/O Server sub-command exit status” on page 6.

Examples

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

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

```
cattracerpt -lshid
```

Related Information

The **starttrace** command, and the **stoptrace** 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 “Virtual I/O Server sub-command exit status” on page 6.

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 a Link Aggregation attributes.

Syntax

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

```
cfglnagg [ -f ] -attr Attribute=NewValue 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.

Note: To create a Link Aggregation see the **mkvdev** command.

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	Specifies an attribute of the specified Link Aggregation.
-backup	Specifies that the <i>Adapter</i> is being added as a backup adapter.
-f	Instructs the cfglnagg command to not prompt the user.
-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.

Exit Status

See "Virtual I/O Server sub-command exit status" on page 6.

Examples

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

```
cfglnagg -add ent3 ent8
```

Related Information

The **mktcip** command, the **mkvdev** command, the **hostname** command, the **startnetsvc** command, the **stopnetsvc** command, the **netstat** command, the **entstat** command, the **cfgnamesrv** command, the

hostmap command, the **traceroute** command, the **ping** command, the **optimizenet** 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.
-chslist	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.
-slist	Specifies that the operation is on the search list. Use this flag with the -rm and -ls flag.
-slist <i>SearchList</i>	Changes the search list in the system configuration database.

Exit Status

See “Virtual I/O Server sub-command exit status” on page 6.

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 -chslist abc.aus.century.com xyz.aus.century.com
```

Related Information

The **cfglnagg** command, the **entstat** command, the **lsnetsh** command, the **mktcpip** command, the **netstat** command, and the **optimizenet** command.

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.

Flags

-attr <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 -attr flag. If you use an -attr flag with multiple attribute value pairs, the list of pairs must be enclosed in quotes with spaces between the pairs. For example, entering -attr <i>Attribute=Value</i> lists one attribute value pair per flag, while entering -attr '<i>Attribute1=Value1 Attribute2=Value2</i>' lists more than one attribute value pair.
-dev <i>Name</i>	Specifies the device logical name, indicated by the <i>Name</i> parameter, whose characteristics are to be changed.
-perm	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 -perm flag.

Exit Status

See "Virtual I/O Server sub-command exit status" on page 6.

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.

chlang Command

Purpose

Changes the language settings for the system.

Syntax

chlang {-lang *Name* | -ls}

Description

The **chlang** command is a high-level command that changes the language and keyboard settings for the entire system. The user needs to log out for language changes to take effect. The Virtual I/O Server must be rebooted for keyboard changes to take effect.

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

Flags

-lang <i>Name</i>	Specifies the language-territory (locale name) that will become the locale setting for the LANG environment variable.
-ls	Lists available languages.

Exit Status

See “Virtual I/O Server sub-command exit status” on page 6.

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.

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 quotes 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.
-dev Name	Specifies the logical device name of the target device for the path(s) affected by the change. This flag is required in all cases.

-pdev <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.
-perm	Changes the path's characteristics without actually changing the path. The change takes affect on the path the next time the path is unconfigured and then configured (possibly on the next boot).
-conn <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.
-op <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: <p>enable 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.</p> <p>disable 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 -attr Attribute=Value flag, a usage error is generated.</p>

Exit Status

See "Virtual I/O Server sub-command exit status" on page 6.

Examples

1. 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 **lsmap** command, the **lspath** command, the **mkpath** command, the **mkvdev** command, the **rmdev** command, and the **rmpath** command.

chuser Command

Purpose

Changes user attributes.

Syntax

chuser **-attr** *Attribute=Value ... Name*

Description

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

Attributes

If you have the proper authority, you can set the following user attributes:

account_locked	Indicates if the user account is locked. Possible values include: true The user's account is locked. The values yes , true , and always are equivalent. The user is denied access to the system. false The user's account is not locked. The values no , false , and never are equivalent. The user is allowed access to the system. This is the default value.
expires	Identifies the expiration date of the account. The Value parameter 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 Value parameter is 0, the account does not expire. The default is 0.
histexpire	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.
histsize	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.
loginretries	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 will not be able to log in until the prime administrator resets the user's account_locked attribute.
maxage	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. Range: 0 to 52
maxexpired	Defines the maximum time (in weeks) beyond the maxage 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 maxexpired attribute is 0, the password expires when the maxage value is met. If the maxage attribute is 0, the maxexpired attribute is ignored. Range: 0 to 52
maxrepeats	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. Range: 0 to 8
minlen	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 minlen and/or ' minalpha + minother ', whichever is greater. ' minalpha + minother ' should never be greater than 8. If ' minalpha + minother ' is greater than 8, then the effective value for minother is reduced to ' 8 - minalpha '.
minother	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. Range: 0 to 8

pwdwarntime

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 **maxage** and **minage** attributes. Values greater than this difference are ignored and a message is issued when the **minage** value is reached.

Exit Status

See “Virtual I/O Server sub-command exit status” on page 6.

Security

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

Examples

1. To change the expiration date for the **davis** user account to 8 a.m., 1 May, 1995, type:
`chuser -attr expires=0501080095 davis`

Related Information

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

chvg Command

Purpose

Sets the characteristics of a volume group.

Syntax

```
chvg [ -unlock ] [ -suspend | -resume] VolumeGroup
```

Description

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

Flags

-resume	Resumes normal I/O operations for a volume group.
-suspend	Drains I/O's for this volume group and suspends future I/O's.
-unlock	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). Note: Before using the -unlock flag, make sure that the volume group is not being used by another LVM command.

Exit Status

See “Virtual I/O Server sub-command exit status” on page 6.

Examples

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

```
chvg -suspend vg03
```
2. To resume the volume group **vg03**, type:

```
chvg -resume vg03
```

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.

cleargcl Command

Purpose

Clears the contents of the global command log.

Syntax

`cleargcl`

Description

The `cleargcl` command removes all entries from the global command log.

Exit Status

See “Virtual I/O Server sub-command exit status” on page 6.

Security

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

Examples

1. To clear all entries from the global command log, type:
`cleargcl`

Related Information

The `lsgcl` command, the `cleargcl` command, and the `lsfailedlogin` 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 “Virtual I/O Server sub-command exit status” on page 6.

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.

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

<i>Device_name</i>	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

Device Type	Displays the description of the adapter type.
Hardware Address	Displays the Ethernet network address currently used by the device.
Elapsed Time	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

Packets	The number of packets transmitted successfully by the device.
Bytes	The number of bytes transmitted successfully by the device.
Interrupts	The number of transmit interrupts received by the driver from the adapter.
Transmit Errors	The number of output errors encountered on this device. This is a counter for unsuccessful transmissions due to hardware or network errors.
Packets Dropped	The number of packets accepted by the device driver for transmission which were not (for any reason) given to the device.
Max Packets on S/W Transmit Queue	The maximum number of outgoing packets ever queued to the software transmit queue.

S/W Transmit Queue Overflow	The number of outgoing packets which have overflowed the software transmit queue.
Current S/W+H/W Transmit Queue Length	The number of pending outgoing packets on either the software transmit queue or the hardware transmit queue.
Broadcast Packets	The number of broadcast packets transmitted without any error.
Multicast Packets	The number of multicast packets transmitted without any error.
No Carrier Sense	The number of unsuccessful transmissions due to the no carrier sense error.
DMA Underrun	The number of unsuccessful transmissions due to the DMA underrun error.
Lost CTS Errors	The number of unsuccessful transmissions due to the loss of the Clear-to-Send signal error.
Max Collision Errors	The number of unsuccessful transmissions due to too many collisions. The number of collisions encountered exceeded the number of retries on the adapter.
Late Collision Errors	The number of unsuccessful transmissions due to the late collision error.
Deferred	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.
SQE Test	Contains the number of "Signal Quality Error" Tests (i.e. Heartbeat) performed successfully during transmission.
Timeout Errors	The number of unsuccessful transmissions due to adapter reported timeout errors.
Single Collision Count	The number of outgoing packets with single (only one) collision encountered during transmission.
Multiple Collision Count	The number of outgoing packets with multiple (2 - 15) collisions encountered during transmission
Current HW Transmit Queue Length	The number of outgoing packets which currently exist on the hardware transmit queue.
CRC Errors	The number of incoming packets with the Checksum (FCS) error.
DMA Overrun	The number of incoming packets with the DMA overrun error.
Alignment Errors	The number of incoming packets with the alignment error.
No Resource Errors	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.
Receive Collision Errors	The number of incoming packets with the collision errors during the reception.
Packet Too Short Errors	The number of incoming packets with the length error indicating that the packet size is less than the Ethernet minimum packet size.
Packet Too Long Errors	The number of incoming packets with the length error indicating that the packet size is bigger than the Ethernet maximum packet size.
Packets Discarded by Adapter	The number of incoming packets dropped by the hardware for any other reasons.
Receiver Start Count	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** | **-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

-ls	Displays detailed information about errors in the error log file.
-rm <i>Days</i>	Deletes error-log entries older than the number of days specified by the <i>Days</i> parameter.

Exit Status

See “Virtual I/O Server sub-command exit status” on page 6.

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 “Virtual I/O Server sub-command exit status” on page 6.

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 “Virtual I/O Server sub-command exit status” on page 6

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 **mklv** command, the **lslv** 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 *VolumeGroup* by adding one or more *PhysicalVolumes*.

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 activated, it exits. But if the system detects a description area from a volume group that is not activated, 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.

Flags

-f	Forces the physical volume to be added to the specified volume group unless it is a member of another volume group that has been activated by the activatevg command.
-----------	--

Exit Status

See “Virtual I/O Server sub-command exit status” on page 6.

Examples

1. To add physical volumes **hdisk3** and **hdisk8** to volume group **vg3**, type:

```
extendvg vg3 hdisk3 hdisk8
```

Note: The volume group must be activated before extending.

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 **exportvg** command, and the **syncvg** command.

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 fsck command was interrupted before it could complete checks or repairs.
4	The fsck command changed the file system; the user must restart the system immediately.
8	The file system contains unrepaired damage.

Exit Status

See “Virtual I/O Server sub-command exit status” on page 6.

Examples

1. To check a file system, enter:

```
fsck /dev/hd1
```

This command checks the unmounted file system located on the **/dev/hd1** device.

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.

Notes:

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

-addr *IPAddress*

Adds an IP address-to-host name mapping entry for the given Internet Protocol address in the database. Specify the host names with the **-host** flag.

-host *HostName...*

Specifies a list of host names. Entries in the list should be separated by blanks.

-ls

Shows all entries in the database.

-rm *IPAddress*

Deletes the IP address-to-host name mapping entry in the database that corresponds to the given address specified by the *IPAddress* variable.

Exit Status

See “Virtual I/O Server sub-command exit status” on page 6.

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.

Parameters

<i>HostName</i>	Sets the primary name of the host.
-----------------	------------------------------------

Exit Status

See “Virtual I/O Server sub-command exit status” on page 6.

Examples

1. To set the hostname to **rotterdam**, type:
 hostname rotterdam

Related Information

The **mktcpip** 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.

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 "Virtual I/O Server sub-command exit status" on page 6.

Examples

1. To import the volume group **bkvg** from physical volume **hdisk07**, type:

```
importvg -vg bkvg hdisk07
```

The volume group **bkvg** 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] ] | -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 operating system. 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 /dev/cdrom or the path to a system backup of the Virtual I/O Server created by the backupios command. The path may also specify a remote location mountable by NFS such as hostname:/path_to_backup
-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 installios 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.
-u	Unconfigure installios (optional). Will manually unconfigure the installios installation resources. This flag is only needed if a problem occurs during the installation and installios does not automatically unconfigure itself.

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 “Virtual I/O Server sub-command exit status” on page 6

Related Information

The **lssw** command, the **updateios** command, the **remote_management** command, the **oem_setup_env** command, and the **oem_platform_level** command.

license Command

Purpose

View and accept the license agreement.

Syntax

license [**-view**] [**-accept**]

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

-accept	Accepts the license agreement
-view	Displays the Virtual I/O Server license agreement.

Exit Status

See “Virtual I/O Server sub-command exit status” on page 6.

Examples

1. To view the license, type:
`license -view`
2. To accept the license, type:
`license -accept`
3. To view if the license has been accepted, type:
`license`

Related Information

The **chlang** command.

Isdev Command

Purpose

Displays Virtual I/O Server devices and their characteristics.

Syntax

To list devices

```
lsdev [ -type DeviceType... ] [ -virtual ] [ -field FieldName... ]
```

To display information about a specific device:

```
lsdev { -dev DeviceName | -plc PhysicalLocationCode } [ -child | -parent | -attr [ Attribute ] | -range Attribute | -slot | -vpd | -field FieldName... ]
```

To display platform-specific device information:

```
lsdev -vpd
```

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.

Table 1. 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

Flags

-attr [*Attribute*]

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:

attribute

Attribute name

value Current value of the attribute

description

Description of the attribute

user Whether attribute can be set by the user (TRUE/FALSE)

-child	Displays the name, status, physical location code, and description for each child of the specified device (-dev <i>DeviceName</i> or -plc <i>PhysicalLocationCode</i>).
-dev <i>DeviceName</i>	Specifies the device logical name for the device for which information is listed. This flag cannot be used with the -plc flag.
-field <i>FieldName</i>	Specifies the list of fields to be displayed. The following fields are supported: name Device name status Device status physloc Physical location code description Description of the device Note: The -field flag cannot be combined with the -parent , -attr , -range , -slot , or -vpd flags.
-parent	Displays the name, status, physical location code, and description of the parent device of the given device (-dev <i>DeviceName</i> or -plc <i>PhysicalLocationCode</i>).
-plc <i>PhysicalLocationCode</i>	Specifies the device physical location code for the device for which information is listed. This flag cannot be used with the -dev flag.
-range <i>Attribute</i>	Displays the allowed values for the specified attribute.
-slot	Displays the slot, description, and device name of the specified device (-dev <i>DeviceName</i> or -plc <i>PhysicalLocationCode</i>). The device must be in a PCI hot swappable slot.
-type <i>DeviceType</i>	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 tty Lists tty devices
-virtual	Restricts output to virtual devices only.
-vpd	Displays platform-specific information for all devices or for a single device when -dev <i>DeviceName</i> or -plc <i>PhysicalLocationCode</i> are specified.

Exit Status

See “Virtual I/O Server sub-command exit status” on page 6.

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
```

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 “Virtual I/O Server sub-command exit status” on page 6.

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 **cleargcl** command and the **lsgcl** command.

lsgcl Command

Purpose

Lists the contents of the global command log to the screen.

Syntax

lsgcl

Description

The **lsgcl** 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	cccccccc	Command	Command options span to 80 characters then wraps to the next row

Exit Status

See “Virtual I/O Server sub-command exit status” on page 6.

Security

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

Examples

1. To list the contents of the global command log, type:
`lsgcl`

Related Information

The **cleargcl** command and **lsfailedlogin** command.

Islv Command

Purpose

Displays information about a logical volume.

Syntax

Islv [**-map** | **-pv**] *LogicalVolume*

Description

The **Islv** 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:

Logical volume	Name of the logical volume. Logical volume names must be unique systemwide and can range from 1 to 15 characters.
Volume group	Name of the volume group. Volume group names must be unique systemwide and can range from 1 to 15 characters.
Logical volume identifier	Identifier of the logical volume.
Permission	Access permission; read-only or read-write.
Volume group state	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 activatevg command, the state is inactive.
Logical volume state	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.
Write verify	Write verify state of On or Off.
Mirror write consistency	Mirror write consistency state of Yes or No.
Max LPs	Maximum number of logical partitions the logical volume can hold.
PP size	Size of each physical partition.
Copies	Number of physical partitions created for each logical partition when allocating.
Schedule policy	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 partitions	Number of physical partitions in the logical volume that are not current.
Bad blocks	Bad block relocation policy.
Inter-policy	Inter-physical allocation policy.
Strictness	Current state of allocation. Possible values are 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, 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.
Intra-policy	Intra-physical allocation policy.
Upper bound	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 point	File system mount point for the logical volume, if applicable.
Label	Specifies the label field for the logical volume.

PV distribution	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.
striping width	The number of physical volumes being striped across.
strip size	The number of bytes per stripe.

Flags

-map	Lists the following fields for each logical partition: <ul style="list-style-type: none"> 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: <ul style="list-style-type: none"> PV Physical volume name. Copies The following three fields: <ul style="list-style-type: none"> • 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 "Virtual I/O Server sub-command exit status" on page 6.

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.

Related Information

The **mklv** command, the **extendlv** command, and the **rm_lv** command.

Ismap Command

Purpose

Displays the mapping between physical, logical, and virtual devices.

Syntax

```
Ismap { -vadapter ServerVirtualAdapter | -plc PhysicalLocationCode | -all } [ -net ] [ -fmt delimiter ]
```

Description

The **Ismap** command displays the mapping between virtual host adapters and the physical devices they are backed to. Given a device 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.

Return code	Description
15	Specified device is not a SVSA
16	Specified device is not a SVEA
17	Specified device is not in the AVAILABLE state

Flags

-all	Specifies Ismap output should be displayed for all virtual SCSI devices. If used with the -net flag, virtual Ethernet adapters will be displayed.
-fmt <i>delimiter</i>	Specifies a delimiter character to separate output fields.
-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.
-plc <i>PhysicalLocatoinCode</i>	Specifies the device physical location code of a server virtual adapter. This flag cannot be used with the -vadapter flag.
-vadapter <i>VirtualServerAdapter</i>	Specifies the device name of a server virtual adapter. This flag cannot be used with the -plc flag.

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

Examples

1. To list all virtual target devices and backing devices mapped to the server virtual SCSI adapter **vnode2**, type:

```
lsmmap -vadapter vhost2
```

The system displays a message similar to the following:

SVSA	Physloc	Client Partition ID

vhost0	U9111.520.10004BA-V1-C2	0x00000004
VTD	vtscsi0	
LUN	0x8100000000000000	
Backing device	vtd0-1	
Physloc		
VTD	vtscsi1	
LUN	0x8200000000000000	
Backing device	vtd0-2	
Physloc		
VTD	vtscsi2	
LUN	0x8300000000000000	
Backing device	hdisk2	
Physloc	U787A.001.0397658-P1-T16-L5-L0	

2. To list the shared Ethernet adapter and backing device mapped to the virtual server Ethernet adapter **ent4**, type:

```
lsmmap -vadapter ent4 -net
```

The system displays a message similar to the following:

SVSA	Physloc

ent4	P2-I1/E1
SEA	ent5
Backing device	ent1
Physloc	P2-I4/E1

3. 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:

```
lsmmap -vadapter ent5 -fmt ":"
```

The system displays a message similar to the following:

```
ent5:ent8:ent2
```

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, the **rmpath** command.

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

The following values may be used:

- | | |
|---------------|--|
| 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, executing the startnetsvc command with any of the supported network services will reactivate the inetd subsystem. |
| telnet | Returns the status of the telnet daemon |
| ftp | Returns the status of the ftp 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.

Ispath Command

Purpose

Displays information about paths to a MultiPath I/O (MPIO) capable device.

Syntax

```
lspath [ -dev DeviceName ] [ -pdev Parent ] [ -status Status ] [ -conn Connection ]
```

```
lspath -dev DeviceName -pdev Parent [ -conn Connection ] -lsattr [ -attr Attribute... ]
```

```
lspath -dev DeviceName -pdev Parent [ -conn Connection ] -range 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 Ispath 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 **-lsattr**, **-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 <i>Attribute</i>	Identifies the specific attribute to list. The ' <i>Attribute</i> ' 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 <i>Name</i>	Specifies the logical device name of the target device whose path information is to be displayed.
-pdev <i>Parent</i>	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 MPIO path selection.
 - disabled**
Display paths that are **disabled** from MPIO 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

- 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
```

- 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
```

- 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 **lsmapi** command, the **mkpath** command, the **chpath** command, and the **rmpath** command.

lspv Command

Purpose

Displays information about a physical volume within a volume group.

Syntax

lspv [**-lv** | **-map** | **-pv** | **-size**] *PhysicalVolume*

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

-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 will be 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	Returns the size of the physical volume in megabytes.		

Exit Status

See “Virtual I/O Server sub-command exit status” on page 6.

Examples

1. To display the status and characteristics of physical volume **hdisk3**, type:
lspv hdisk3
2. To display all physical volumes in the system, type:
lspv

You should see output similar to the following:

```
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.

Related Information

The **migratepv** command.

Issw Command

Purpose

Lists installed software products.

Syntax

Issw [**-hist**]

Description

The **Issw** command displays information about installed file sets or file set updates. If the **-hist** parameter is not specified, the name, most recent level, state, and description of all file sets is displayed. Part information (usr, root, and share) is consolidated into the same listing. For formatted file sets, it displays the most recent maintenance level.

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
APPLIED	The specified fileset is installed on the system. The APPLIED state means that the fileset can be removed with the updateios command and the previous level of the fileset restored.
APPLYING	An attempt was made to apply the specified fileset, but it did not complete successfully, and cleanup was not performed.
BROKEN	The specified fileset or fileset update is broken and should be reinstalled before being used.
COMMITTED	The specified fileset is installed on the system. The COMMITTED 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 updateios command.
EFIX LOCKED	The specified fileset was installed successfully and locked.
OBSOLETE	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.
COMMITTING	An attempt was made to commit the specified fileset, but it did not complete successfully, and cleanup was not performed.
REJECTING	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 **Issw** 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 “Virtual I/O Server sub-command exit status” on page 6

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.

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 users or all the attributes of specific users. If more than one user is specified, each user must be separated by a comma. If no users are specified the attributes of all users 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.

Exit Status

See "Virtual I/O Server sub-command exit status" on page 6.

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:

```
padmin account_locked=false expires=0 histexpire=0 histsize=0 loginretries=0 maxage=0
maxexpired=-1 maxrepeats=8 minlen=0 minother=0 pwdwarntime=0
help account_locked=false expires=0 histexpire=0 histsize=0 loginretries=0 maxage=0
maxexpired=-1 maxrepeats=8 minlen=0 minother=0 pwdwarntime=330
user_de account_locked=false expires=0 histexpire=0 histsize=0 loginretries=0 maxage=0
maxexpired=-1 maxrepeats=8 minlen=0 minother=0 pwdwarntime=330
user_sr account_locked=false expires=0 histexpire=0 histsize=0 loginretries=0 maxage=0
maxexpired=-1 maxrepeats=8 minlen=0 minother=0 pwdwarntime=330
```

2. To display the attributes of user delft, type:

```
lsuser delft
```

3. To display the attributes of user delft and user gouda, type:

```
lsuser delft, gouda
```

Related Information

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

lsvg Command

Purpose

Displays information about volume groups.

Syntax

lsvg [**-map** *VolumeGroup* | **-lv** *VolumeGroup* ... | **-pv** *VolumeGroup* ... | *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.

If you do not specify any flags, the following information will be displayed:

Volume group	Name of the volume group. Volume group names must be unique and can range from 1 to 15 characters.
Volume group state	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.
Permission	Access permission: read-only or read-write.
Max LVs	Maximum number of logical volumes allowed in the volume group.
LVs	Number of logical volumes currently in the volume group.
Open LVs	Number of logical volumes within the volume group that are currently open.
Total PVs	Total number of physical volumes within the volume group.
Active PVs	Number of physical volumes that are currently active.
VG identifier	The volume group identifier.
PP size	Size of each physical partition.
Total PPs	Total number of physical partitions within the volume group.
Free PPs	Number of physical partitions not allocated.
Alloc PPs	Number of physical partitions currently allocated to logical volumes.
Quorum	Number of physical volumes needed for a majority.
VGDS	Number of volume group descriptor areas within the volume group.
Auto-on	Automatic activation at IPL (yes or no).
Concurrent	States whether the volume group is Concurrent Capable or Non-Concurrent Capable.
Auto-Concurrent	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.
VG Mode	The mode of the volume group: Concurrent or Non-Concurrent.
Node ID	Node id of this node if volume group is in concurrent node.
Active Nodes	Node ids of other concurrent nodes that have this volume group active.
Max PPs Per PV	Maximum number of physical partitions per physical volume allowed for this volume group.
Max PVs	Maximum number of physical volumes allowed in this volume group.
LTG size	Logical track group size, in number of kilobytes, of the volume group.
BB POLICY	Bad block relocation policy of the volume group.
SNAPSHOT VG	Snapshot volume group name if the snapshot volume group is active else snapshot volume group identifier.

PRIMARY VG

Original volume group name of a snapshot volume group if the original volume group is active else original volume group identifier.

Flags

-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 “Virtual I/O Server sub-command exit status” on page 6.

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.

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.

mirrorios Command

Purpose

Mirrors all the logical volumes on rootvg.

Syntax

mirrorios [**-f**] [*PhysicalVolume* ...]

Description

The **mirrorios** subcommand 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.

By default, the **mirrorios** subcommand attempts to mirror the logical volumes onto any of the disks in a volume group. If you wish to control which drives are used for mirroring, you must include the list of disks in the input parameters, *PhysicalVolume*. Mirror strictness is enforced. Additionally, the **mirrorios** subcommand mirrors the logical volumes, using the default settings of the logical volume being mirrored.

The user is warned that the Virtual I/O Server will reboot upon the completion of this command and is prompted to continue. If the **-f** option is specified, the command will run without prompting the user.

Note: It is recommended that the rootvg volume group be mirrored on all Virtual I/O Server partitions.

This command can only be executed by the prime administrator.

Flags

-f Executes the command with out prompting the user to continue.

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 **hdisk8**, type:

```
mirrorios hdisk8
```

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 **unmirrorios** command.

mklv Command

Purpose

Creates a logical volume.

Syntax

```
mklv [ -lv NewLogicalVolume | -prefix Prefix ] 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 *Size* parameter specifies the minimum size the logical volume should be. When specifying *Size* the following conventions must be used:

Size	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. |
| -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. |

Exit Status

See “Virtual I/O Server sub-command exit status” on page 6.

Examples

1. To make a logical volume in volume group **vg02** with a minimum size of 1 Mb , type:

```
mklv vg02 1MB
```
2. To make a logical volume in **vg03** with 1GB chosen from physical volumes **hdisk5**, **hdisk6**, and **hdisk9**, type:

```
mklv vg03 1GB hdisk5 hdisk6 hdisk9
```
3. To request a logical volume with a minimum size of 10MB, type:

```
mklv VGNAME 10M
```

where *VGNAME* is the name of your logical volume.

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 “Virtual I/O Server sub-command exit status” on page 6.

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 *Parent* 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 “Virtual I/O Server sub-command exit status” on page 6.

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 **fscsi0** to fiber channel disk **hdisk1**, type:

```
mkpath -dev hdisk1 -pdev fscsi0
```

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.

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 ] ]
```

Description

The **mktcpip** command sets the required minimal values required for using TCP/IP on a host machine. The basic functions of the **mktcpip** command include the following:

- Setting the host name
- Setting the IP address of the interface
- Setting the domain name and IP address of the nameserver, if applicable
- Setting the subnetwork mask, if applicable
- Starting the specified TCP/IP daemons

Flags

-cabletype <i>CableType</i>	Specifies 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 following is the standard format for setting the host name in a domain naming system: hostname.subdomain.subdomain.rootdomain
-inetaddr <i>Address</i>	Sets the Internet address of the host. Specify the address in dotted decimal notation. Each network interface on the host should have a unique Internet address. The following is the standard format for setting the Internet address: 127.10.31.2
-interface <i>Interface</i>	Specifies a particular network interface, for example: tr0
-netmask <i>SubnetMask</i>	Specifies the mask 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 (1s) 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 Internet address of the name server 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 the host should use for name resolution, if any. The domain name should be in the following format: subdomain.subdomain.rootdomain
-start	Starts the TCP/IP daemons.

Exit Status

See “Virtual I/O Server sub-command exit status” on page 6.

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
```

Related Information

The **hostname** command, the **startnetsvc** command, the **stopnetsvc** command, the **cfgnagg** 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 [**-de** | **-sr**] [**-attr** *Attributes=Value* [*Attribute=Value...*]] *Name*

Description

The **mkuser** command creates a new user account. Upon completion of creating a the new account you will be prompted for set the new account's password.

For a list of supported attributes, see "chuser Command" on page 24.

Flags

-attr <i>Attribute=Value</i>	Identifies the attribute to set as well as the new value for the attribute. The <i>Attribute=Value</i> parameter can use one attribute value pair or multiple attribute value pairs for one -attr flag.
-de	Creates a Development Engineer or DE user account. This type of account allows IBM developers to log into the Virtual I/O Server and debug problems.
-sr	<p>Creates a service representative or SR user account. This type of account enables a service representatives to run commands required to service the system without being logged in as root. This includes the following:</p> <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 shutdown and reboot operations <p>The recommended SR login user name is qserv.</p>

Exit Status

See "Virtual I/O Server sub-command exit status" on page 6.

Security

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

Examples

1. To create the **davis** user account with default values, type:
`mkuser davis`
2. To create the **davis** user account and set the **maxage** attribute to a value of 52, type:
`mkuser -attr maxage=52 davis`

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 -vdev TargetDevice -vadapter VirtualServerAdapter [ -dev DeviceName ]
```

To create a Shared Ethernet Adapter:

```
mkvdev -sea TargetDevice -vadapter VirtualEthernetAdapter... -default DefaultVirtualEthernetAdapter  
-defaultid SEADefaultPVID [ -attr Attribute=Value [ Attribute=Value... ] ]
```

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 will be automatically generated and assigned unless the **-dev** *DeviceName* flag is specified, in which case *DeviceName* will become the device name. If the **-lnagg** flag is specified, an EtherChannel (manual 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, *DefaultVirtualEthernetAdapter*, must also be included as one of the virtual Ethernet adapters, *VirtualEthernetAdapter*. 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*.

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.

Table 2. Command specific return codes

Return code	Description
13	Specified physical or logical volume is not a valid.

Flags

-attr *Attribute=Value*

Specifies the device attribute value pairs to be used instead of the defaults. The *Attribute=Value* 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

DefaultVirtualEthernetAdapter Default virtual adapter to use for non-VLAN-tagged packets. This flag maps to the SEA device attribute pvid_adapter.

-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 pvid.
-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.
-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 the new device will be mapped to.
-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 or logical volume. 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

See “Virtual I/O Server sub-command exit status” on page 6.

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:

```
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:

```
mkvdev -vdev hdisk6 -vadapter vhost2
```

The system displays a message similar to the following:

```
vtscsi1 available
```

3. 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:

```
mkvdev -sea ent4 -vadapter ent6,ent7,ent9 -default ent6 -defaultid 8
```

The system displays a message similar to the following:

```
ent10 available
```

4. To create an automatic Link Aggregation with primary adapters **ent4** and **ent5** and backup adapter **ent6**, type:

```
mkvdev -lnagg ent4,ent5 -attr backup_adapter=ent6 mode=6023ad
```

The system displays a message similar to the following:

```
ent10 available
```

Related Information

The **cfgdev** command, the **chdev** command, the **chpath** command, the **lsdev** command, the **lsmap** command, and the **rmdev** 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.

Notes:

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.

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 <i>VolumeGroup</i>	<p>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.</p> <p>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.</p>

Exit Status

See "Virtual I/O Server sub-command exit status" on page 6.

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.

mount Command

Purpose

Makes a file system available for use.

Syntax

mount [[*Node:Directory*] *Directory*]

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.

Exit Status

See “Virtual I/O Server sub-command exit status” on page 6.

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 ] [ -protocol protocol ] [ -routtable ] [ -routinfo ] [ -state ] [ Interval ]
```

To display the contents of a network data structure

```
netstat [ -stats | -cdlistats ] [ -protocol protocol ] [ Interval ]
```

To display the address resolution protocol (arp)

```
netstat -arp
```

To clear all statistics

```
netstat -clear
```

Description

The **netstat** command symbolically displays the contents of various network-related data structures for active connections. The *Interval* parameter, specified in seconds, continuously displays information regarding packet traffic on the configured network interfaces.

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 the <i>Protocol</i> 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."
-state	Shows the state of all configured interfaces.
	<p>The interface display format provides a table of cumulative statistics for the following items:</p> <ul style="list-style-type: none">• Errors• Collisions <p>Note: The collision count for Ethernet interfaces is not supported.</p> <ul style="list-style-type: none">• Packets transferred <p>The interface display also provides the interface name, number, and address as well as the maximum transmission units (MTUs).</p>
-stats	Shows statistics for each protocol.

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. The host address is displayed symbolically if the address can be resolved to a symbolic host name, while network addresses are displayed symbolically.

NS addresses are 12-byte quantities, 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, these are word and byte reversed; for the Sun systems, they are not reversed.

If a symbolic name for a host is not known or if the **-num** flag is used, 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 supported.

- 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 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 ten fields:

Flags	<p>The flags field of the routing table shows the state of the route:</p> <ul style="list-style-type: none"> 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	Direct routes are created for each interface attached to the local host.
Refs	<p>The gateway field for these entries shows the address of the outgoing interface.</p> <p>Gives 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.</p>
Use	Provides a count of the number of packets sent using that route.
PMTU	Gives the Path Maximum Transfer Unit (PMTU).
Interface	Indicates the network interfaces utilized for the route.
Exp	Displays the time (in minutes) remaining before the route expires.
Groups	Provides a list of group IDs associated with that route.
Netmasks	Lists the netmasks applied on the system.
Route Tree for Protocol Family	<p>Specifies the active address families for existing routes. Supported values for this field are as follows:</p> <ul style="list-style-type: none"> 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* parameter, 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 of each screen of information 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 “Virtual I/O Server sub-command exit status” on page 6.

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

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, the **optimizenet** command.

oem_platform_level Command

Purpose

Returns the operating system level of the OEM install 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 “Virtual I/O Server sub-command exit status” on page 6

Examples

1. To get the operating system level of the OEM install 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 install and setup environment.

Syntax

oem_setup_env

Description

The **oem_setup_env** command places the user into the OEM software install and setup environment. In this environment, the user will be able to 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.

Upon completion of installing all desired software, typing **exit** will return the user to the Virtual I/O Server prompt.

Only the prime administrator can execute this command.

Exit Status

See “Virtual I/O Server sub-command exit status” on page 6

Examples

1. To initiate the OEM setup and install 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	<p>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.</p> <p>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.</p>
-reboot	<p>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 <i>tunables</i> display instead of the current values.</p>
-set <i>Tunable</i> [=NewValue]	<p>Displays the value or sets the Tunable to <i>NewValue</i>. 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.</p> <p>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.</p>

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 timestamps 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>
thewall	<p>Purpose: Specifies the maximum amount of memory, in kilobytes, that is allocated to the memory pool.</p> <p>Values: Default: 1/2 of RAM or 64 GB (whichever is smaller) Range: N/A Type: Static</p> <p>Tuning: Not settable</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 “Virtual I/O Server sub-command exit status” on page 6.

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_bsiz** parameter, type:
`optimizenet -list arptab_bsiz`
5. To display help information on **arptab_bsiz**, type:
`optimizenet -h arptab_bsiz`

Related Information

The **entstat** command, the **lsnetsh** command, the **mktcpip** command, the **netstat** command, and the **traceroute** 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.

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 "Virtual I/O Server sub-command exit status" on page 6.

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.

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 <i>PacketSize</i>	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 <i>hostname/IP_addr</i>	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 "Virtual I/O Server sub-command exit status" on page 6.

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.

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**.

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 “Virtual I/O Server sub-command exit status” on page 6.

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 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**.

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 “Virtual I/O Server sub-command exit status” on page 6

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.

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.
-ls	Displays a list of previously saved volume groups.
-vg	Specifies the name of the VolumeGroup to restore.

Exit Status

See “Virtual I/O Server sub-command exit status” on page 6.

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 **lsvg** command, the **mkvg** command, and the **syncvg** command.

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 is not interruptible. Stopping this command before execution is complete could result in a corrupted database.

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 "Virtual I/O Server sub-command exit status" on page 6.

Examples

1. To unconfigure the **cd0** CD-ROM device, type:

```
rmdev -dev cd0
```
2. To unconfigure the SCSI adapter **scsi1** and all of its children, type:

```
rmdev -recursive -dev scsi1
```
3. To unconfigure just the children of the SCSI adapter **scsi1**, but not the adapter itself, type:

```
rmdev -pdev scsi1
```
4. To unconfigure the children of PCI bus **pci1** and all other devices under them, type:

```
rmdev -pdev pci1
```

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.

Flags

-f Removes the logical volumes without requesting confirmation.

Exit Status

See “Virtual I/O Server sub-command exit status” on page 6.

Examples

Attention: The following command destroys all data in the logical volumes.

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 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 “Virtual I/O Server sub-command exit status” on page 6.

Examples

1. To remove mirroring from the logical volume **lv0112**, type:
`rmlvcopy lv0112`

Related Information

The **mklv** command, the **extendlv** command, the **rmlv** command, and the **lslv** command.

rmpath Command

Purpose

Removes from the system a path to an MPIIO-capable device.

Syntax

```
rmpath { [ -dev Name ] [ -pdev Parent ] [ -conn Connection ] } [ -rm ]
```

Description

The **rmpath** 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 **rmpath** command unconfigures (if necessary) and removes, or deletes, the path definition(s) from the system.

When the **rmpath** 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 (e.g., paths that are in the process of doing I/O cannot be unconfigured).

The **rmpath** 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 *Connection*

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 “Virtual I/O Server sub-command exit status” on page 6.

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.

rmuser Command

Purpose

Removes a user account.

Syntax

rmuser [**-rmdir**] *Name*

Description

The **rmuser** command removes the user account identified by the *Name* parameter. This command removes a user's attributes without removing the user's home directory and files unless the **rmdir** flag is specified.

Flags

-rmdir Removes the specified user's home directory.

Exit Status

See "Virtual I/O Server sub-command exit status" on page 6.

Security

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

Examples

1. To remove user account **haarlem**, type:
`rmuser haarlem`
2. To remove the user account and home directory of user account **emmen**, type:
`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.

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 “Virtual I/O Server sub-command exit status” on page 6.

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.

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 “Virtual I/O Server sub-command exit status” on page 6.

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 [-restart]`

Description

The **shutdown** command halts the operating system. When the shutdown is complete the user receives a shutdown completion message. 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.

Flags

<code>-restart</code>	Restarts the system after being shutdown.
-----------------------	---

Exit Status

See “Virtual I/O Server sub-command exit status” on page 6.

Security

This command can only be executed by the prime administrator (padmin).

Examples

1. To shut down the system, type:
`shutdown`
2. To restart the system, type:
`shutdown -restart`

snap Command

Purpose

Gathers system configuration information.

Syntax

snap

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.

Approximately 8MB of temporary disk space is required to collect all system information, including contents of the error log.

Exit Status

See “Virtual I/O Server sub-command exit status” on page 6.

Examples

1. Enter the following command to gather all system configuration information:

```
snap
```

The output of this command is written to the users home directory.

Files

snap.pax.Z

startnetsvc Command

Purpose

Starts the **telnet** and **ftp** daemons.

Syntax

startnetsvc *NetworkService*

Description

The **startnetsvc** command can start the **telnet** and **ftp** daemons. Enabling the **telnet** daemon allows users to **telnet** into the Virtual I/O Server. Enabling the **ftp** daemon allows users to **ftp** into the Virtual I/O Server.

Parameters

NetworkService

The following values may be used:

telnet Enables the **telnet** daemon

ftp Enables to **ftp** daemon

ALL Enables both the **telnet** and **ftp** daemons.

Exit Status

See “Virtual I/O Server sub-command exit status” on page 6.

Examples

1. To start the **telnet** daemon, type:
`startnetsvc telnet`
2. To start the **ftp** daemon, type:
`startnetsvc ftp`
3. To start both the **ftp** daemon and the **telnet** daemon, type:
`startnetsvc ALL`

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, the **optimizenet** command.

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 “Virtual I/O Server sub-command exit status” on page 6.

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 can be separated by commas, or enclosed in double quotation marks and separated by commas or blanks.

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 “Virtual I/O Server sub-command exit status” on page 6.

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 **telnet** and **ftp** daemons.

Syntax

stopnetsvc *NetworkService*

Description

The **stopnetsvc** command can stop the **telnet** and **ftp** daemons. Disabling the **telnet** daemon prevents anyone from being able to **telnet** into the Virtual I/O Server. Disabling the **ftp** daemon prevents anyone from being able to **ftp** into the Virtual I/O Server.

Parameters

NetworkService

The following values may be used:

telnet	Disables the telnet daemon
ftp	Disables the ftp daemon
ALL	Disables both the telnet daemon and the ftp daemon

Exit Status

See “Virtual I/O Server sub-command exit status” on page 6.

Examples

1. To disable the **telnet** daemon, type:
`stopnetsvc telnet`
2. To disable the **ftp** daemon, type:
`stopnetsvc ftp`
3. To disable both the **telnet** daemon and the **ftp** daemon, type:
`stopnetsvc ALL`

Related Information

The **mktcpip** command, the **hostname** command, the **startnetsvc** command, the **cfglagg** command, the **netstat** command, the **entstat** command, the **cfgnamesrv** command, the **hostmap** command, the **traceroute** command, the **ping** command, the **optimizenet** command.

stoptrace Command

Purpose

Stops the trace function.

Syntax

stoptrace

Description

The **stoptrace** command ends a trace session.

Exit Status

See “Virtual I/O Server sub-command exit status” on page 6.

Examples

1. To end a trace session, type:
`stoptrace`

Related Information

The **starttrace** command and the **cattracert** command.

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

-lv	Specifies that the <i>Name</i> parameter represents a logical volume device name.
-pv	Specifies that the <i>Name</i> parameter represents a physical volume device name.
-vg	Specifies that the <i>Name</i> parameter represents a volume group device name.

Exit Status

See “Virtual I/O Server sub-command exit status” on page 6.

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.

topas Command

Purpose

Reports selected local system statistics.

Syntax

```
topas [ -cpus number_of_monitored_hot_CPUs ] [ -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 ] [ -procsdisp | -wlmdisp ]
```

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 -wlms 20 -cpus 20
```

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-Write**.
- 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 -Write**
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
Mon Mar 13 15:56:32 2000  Interval: 2

CPU  User%  Kern%  Wait%  Idle%
cpu0  7.0    4.0    0.0    89.0
cpu1  1.0    8.0    0.0    91.0
cpu2  0.0    0.0    0.0    100.0

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
Cswitch 113 Readch 1853576
Syscall 2510 Writech 49883
Reads 466 Rawin 0
Writes 12 Ttyout 706
Forks 0 Igets 0
Execs 0 Namei 0
Runqueue 0.0 Dirblk 0
Waitqueue 0.0

PAGING  MEMORY
Faults 1 Real,MB 255
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

USER  PID  PPID  PRI  NI  DATA  TEXT  PAGE  TIME  CPU%  PGFAULTS
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% PGFAULTS 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

-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.
-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 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 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	<p>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.</p> <p>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).</p>

-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.
--------------	--

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; Micro-Partitioning and SMT metrics similar to what lparstat and mpstat provide are displayed.
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 topas is in the full-screen WLM display (by using the W key or the -wlms 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 “Virtual I/O Server sub-command exit status” on page 6.

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`

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 tracert 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 tracert command based on the MTU of the outgoing interface.

Exit Status

See “Virtual I/O Server sub-command exit status” on page 6.

Examples

1. To print the route to host **nis.nsf.net**, type:

```
tracert nis.nsf.net
```

The output will look similar to the following:

```
tracert 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.

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, Quarum 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 “Virtual I/O Server sub-command exit status” on page 6.

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 “Virtual I/O Server sub-command exit status” on page 6.

Examples

1. To unmount files and directories, type the following command:

```
unmount /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* [**-accept**]

updateios -remove { **-file** *RemoveListFile* | *RemoveList* }

Description

The **updateios** command is used to install fixes, or updates the Virtual I/O Server to the latest maintenance level. Before installing a fix or maintenance level, the **updateios** command will first run a preview installation and display the results. Upon completion of the preview, the user will then be prompted to continue or exit. If the preview fails for any reason, the updates should not be installed.

If the **-remove** flag is specified, previously installed updates will be removed. The updates to be removed must be listed on the command line or in the **RemoveListFile** file.

The log file, **install.log** in the user's home directory, will be appended with a list of all filesets that were installed.

Flags

-accept	Agrees to required software license agreements for software to be installed.
-dev <i>Device</i> or <i>Directory</i>	Specifies the device or directory containing the images to install.
-file <i>file</i>	Specifies the file containing a list of entries to uninstall.
-remove	Performs an uninstall of the specified software.

Exit Status

See "Virtual I/O Server sub-command exit status" on page 6

Examples

1. To update the Virtual I/O Server to the latest level, where the updates are located on the mounted filesystem **/home/padmin/update**, type:

```
updateios -dev /home/padmin/update
```

Related Information

The **lssw** command, the **ioslevel** command, the **remote_management** command, the **oem_setup_env** command, and the **oem_platform_level** command.

Chapter 2. Partition Load Manager commands

Please enter your text here. Thank-you.

xlplm Command

Purpose

Start, stop, modify, reserve, and query a cross-logical partition load manager server.

This command is for use with the Partition Load Manager on AIX only.

Syntax

```
xlplm -S -p policy_file -l log_file [ -o operation_mode ] [ configuration ]
```

```
xlplm -K [ configuration ]
```

```
xlplm -M [ -p policy_file ] [ -l log_file ] [ -o operation_mode ] [ configuration ]
```

```
xlplm -R -g group_name [ -c cpu_resource_size ] [ -m memory_resource_size ] configuration
```

```
xlplm -Q [ configuration ] [ -r ] [ -a ]
```

```
xlplm -C -p policy_file
```

Description

The **xlplm** command is used to start the **xlplmd** daemon, which does all of the resource management work according to the definitions set in the policy file.

Flags

-a

Print active attribute values. This flag is only valid if it is used with the **-Q** flag. The **xlplm** command assigns default values to some optional attributes when they are missing from the policy file. Some of those attributes provided in the policy file can be overridden by values from the logical partitioning configuration information, which is retrieved at run-time.

-c *cpu_resource_size*

This flag displays currently-active attribute values. If this flag is not specified, the original values from the policy file, with default values for non-specified attributes, are displayed.

-C

Specifies the amount of processor resources to reserve. Verifies the validity of the policy file.

-g *group_name*

Specifies the name of a group in the policy file. Use this flag when you are reserving or releasing resources. The resources that you want to reserve are removed from the specified group in the policy file. When you are releasing resources, they are placed in the free pool of the specified group in the policy file.

-K	Stop the Partition Load Manager instance. You must either have root authority or be logged in as the authorized plmuser user ID to use this flag.
-l <i>log_file</i>	Specifies the name of the file you want to contain the Partition Load Manager activity log.
-M	Modify a Partition Load Manager server. You must either have root authority or be logged in as the authorized plmuser user ID to use this flag.
-m <i>memory_resource_size</i>	Specifies the amount of memory resource to reserve.
-o <i>operation_mode</i>	Specifies whether the Partition Load Manager server is to operate in management mode, with a value of <i>M</i> , or monitoring mode, with a value of <i>N</i> . When the xlplm daemon starts, the default value is management mode, or <i>M</i> .
-p <i>policy_file</i>	Specifies the name of the Partition Load Manager policy file.
-Q	Query the Partition Load Manager server status.
-R	Reserve or release resources from a partition managed by a Partition Load Manager server. You must either have root authority or be logged in as the authorized plmuser user ID to use this flag.
-r	Use the Partition Load Manager server in raw data mode.
-S	Start a Partition Load Manager instance. You must either have root authority or be logged in as the authorized plmuser user ID to use this flag.

Parameters

<i>configuration</i>	Identifies an instance of the Partition Load Manager management. This parameter must be specified if there are multiple instances of the Partition Load Manager server on your system. If the parameter is not specified, a 'default' value is used.
----------------------	--

Exit Status

This command returns the following exit values:

0	Command completed successfully.
1	The program encountered a non-recoverable internal error, such as a memory allocation or system call failure.
2	The specified log file could not be opened or created.
3	The specified policy file could not be opened or created.
4	A required temporary file could not be created in the /tmp directory.
5	The specified policy is invalid.
6	The daemon failed to start. This could be the result of an internal error or an inability to communicate with the Hardware Management Console (HMC).
7	Command line usage error.
8	An invalid number was specified for the reservation amount.

9	The current user does not match the authorized user in the <code>/etc/plm/auth/plmuser</code> file, or the file could not be read.
10	An instance with the requested name already exists.
11	An instance with the requested name does not exist.
12	The requested mode is the same as the current mode.
13	A remote command to the HMC failed.
14	A reservation request failed due to one of the following reasons: <ul style="list-style-type: none"> • unknown group • reservation amount is already set to requested amount • could not reserve the requested amount • the requested resource is not managed

Examples

1. Start a Partition Load Manager server in management mode by typing one of the following:
 - `cd /etc/xlplm/cec1`
`xlplm -S -p policy -l log cec1`
 - `xlplm -S -p /etc/xlplm/cec1/policy -l /etc/xlplm/cec1/log -o M cec1`
2. Start a Partition Load Manager server in monitoring mode by typing the following:
`xlplm -S -p policy -l log -o N cec1`
3. Stop a Partition Load Manager server by typing the following:
`xlplm -K cec1`
4. Load a new policy into a Partition Load Manager server by typing the following:
`xlplm -M -p evening_policy cec1`
5. Start using a new log file, called 'newlog', for a Partition Load Manager server by typing the following:
`xlplm -M -l newlog cec1`
6. Display configuration names for the active Partition Load Manager server:
`xlplm -Q`

xlplstat Command

Purpose

Display logical partition load statistics for a list of host names.

This command is for use with the Partition Load Manager on AIX only.

Syntax

```
xlplm [-r] {-p | -f} [ interval ] [ count ]
```

Description

Display load statistics for one or more remote LPARs. The command will contact the remote systems every number of seconds specified by the *interval* parameter for each number of intervals specified by the *count* parameter. If the *interval* parameter and the *count* parameter are omitted, the remote systems are queried once. If only the *count* parameter is omitted, the remote systems are queried every number of seconds specified by the *interval* parameter until the command is terminated by the user.

The caller of this command must be the root user or the Partition Load Manager authorized user.

For the formatted output, the output is displayed as follows, for each host listed in the input file:

STAT	TYP	CPU			MEM			HOST
		CUR	PCT	LOAD	CUR	PCT	PGSTL	
up	D	4.00	50.15	0.65	1024	43.10	0	testlp1
up	D	2.00	95.72	0.90	2048	97.48	250	testlp2
up	D	10.00	98.31	1.03	5120	72.25	0	testlp3

STAT Partition status. May be "up" or "down".

TYP Partition type. May be "D" (dedicated) "S" (shared) or "U" (unknown).

If the type is "U", the command was unable to query the partition type and there may be a connectivity or authentication problem.

CUR The current amount of resource allocated to the partition.

PCT Percent utilization for the resource

LOAD CPU load average

PGSTL Page steals per second

HOST Managed host name

The raw output is displayed as a header containing column descriptions followed by one line of data for each host:

```
#host_name:group_name:status:cpu_type:cpu_ent:cpu_util:cpu_load:mem_ent:mem_util:mem_pgstl
testlp1.mydomain.com:group1:up:dedicated:4.00:45.05:0.38:1024:75.00:0
testlp2.mydomain.com:group1:up:dedicated:2.00:87.23:0.92:2048:92.21:123
testlp3.mydomain.com:group1:up:dedicated:10.00:95.17:1.01:5120:70.30:0
```

Flags

-r	Raw output mode. Data is printed in colon separated format, with one line per host.
-p	Retrieves the host list from the given policy file.
-f	Retrieves the host list from the given plain text file. This file has one hostname per line.

Exit Status

This command returns the following exit values:

1	Internal error.
3	Could not open input file.
5	Invalid policy file.
7	Usage error.
9	Not authorized.

Appendix. 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 give you any license to these patents. You can send license inquiries, in writing, to the manufacturer.

For license inquiries regarding double-byte (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 INFORMATION 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 Web sites not owned by the manufacturer are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this product and use of those Web sites 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 information 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.

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.

SUBJECT TO ANY STATUTORY WARRANTIES WHICH CAN NOT BE EXCLUDED, THE MANUFACTURER, ITS PROGRAM DEVELOPERS AND SUPPLIERS MAKE NO WARRANTIES OR CONDITIONS EITHER EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OR CONDITIONS OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, AND NON-INFRINGEMENT, REGARDING THE PROGRAM OR TECHNICAL SUPPORT, IF ANY.

UNDER NO CIRCUMSTANCES IS THE MANUFACTURER, ITS PROGRAM DEVELOPERS OR SUPPLIERS LIABLE FOR ANY OF THE FOLLOWING, EVEN IF INFORMED OF THEIR POSSIBILITY:

1. LOSS OF, OR DAMAGE TO, DATA;
2. SPECIAL, INCIDENTAL, OR INDIRECT DAMAGES, OR FOR ANY ECONOMIC CONSEQUENTIAL DAMAGES; OR
3. LOST PROFITS, BUSINESS, REVENUE, GOODWILL, OR ANTICIPATED SAVINGS. SOME JURISDICTIONS DO NOT ALLOW THE EXCLUSION OR LIMITATION OF INCIDENTAL OR CONSEQUENTIAL DAMAGES, SO THE ABOVE LIMITATION OR EXCLUSION MAY NOT APPLY TO YOU.

If you are viewing this information in softcopy, the photographs and color illustrations may not appear.

Trademarks

The following terms are trademarks of International Business Machines Corporation in the United States, other countries, or both:

AIX
AIX 5L
e(logo)server

eServer
i5/OS
IBM
iSeries
pSeries

Linux™ is a 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 company, product or service names may be trademarks or service marks of others.

Terms and conditions for downloading and printing information

Permissions for the use of the information you have selected for download are granted subject to the following terms and conditions and your indication of acceptance thereof.

Personal Use: You may reproduce this information for your personal, noncommercial use provided that all proprietary notices are preserved. You may not distribute, display or make derivative works of this information, or any portion thereof, without the express consent of the manufacturer.

Commercial Use: You may reproduce, distribute and display this information solely within your enterprise provided that all proprietary notices are preserved. You may not make derivative works of this information, or reproduce, distribute or display this information or any portion thereof outside your enterprise, without the express consent of the manufacturer.

Except as expressly granted in this permission, no other permissions, licenses or rights are granted, either express or implied, to the information or any 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 information 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 THIS INFORMATION. THE INFORMATION IS 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.

All material copyrighted by IBM Corporation.

By downloading or printing information from this site, you have indicated your agreement with these terms and conditions.

Index

A

activatevg command 7

B

backupios command 8
bootlist Command 10

C

cattracerpt command 13
cfgdev command 14
cfglnagg command 15
cfgnamesrv command 17
chdev command 19
chlang command 21
chpath command 22
chuser command 24
chvg command 26
cleargcl command 27

Command

mount 83
rmlv 104

commands

xlplm 133
xlpstat 135

Commands

activatevg 7
backupios 8
bootlist 10
cattracerpt 13
cfgdev 14
cfglnagg 15
cfgnamesv 17
chdev 19
chlang 21
chpath 22
chuser 24
chvg 26
cleargcl 27
deactivatevg 28
diagmenu 29
entstat 30
errlog 34
exportvg 35
extendlv 36
extendvg 37
fsck 38
hostmap 40
hostname 42
importvg 43
installios 44
ioslevel 45
license 46
lsdev 47
lsfailedlogin 50
lsgcl 51
lslv 52
lsmmap 55
lsnetvc 57

Commands (continued)

lspath 58
lspv 61
lssw 63
lsuser 65
lsvg 66
migratepv 69
mirrorios 71
mklv 72
mklvcopy 73
mkpath 74
mktcpip 76
mkuser 78
mkvdev 79
mkvg 82
netstat 85
oem_platform_level 89
oem_setup_env 90
optimizenet 91
passwd 97
ping 98
reducevg 100
remote_management 101
restorevgstruct 102
rmdev 103
rmlvcopy 105
rmpath 106
rmuser 108
savevgstruct 109
showmount 110
shutdown 111
snap 112
startnetvc 113
startsysdump 114
starttrace 115
stopnetvc 116
stoptrace 117
syncvg 118
topas 119
traceroute 128
unmirrorios 130
unmount 131
updateios 132

D

deactivatevg command 28
diagmenu command 29

E

entstat command 30
errlog command 34
exportvg command 35
extendlv command 36
extendvg command 37

F

fsck command 38

H

hostmap command 40
hostname command 42

I

importvg command 43
installios command 44
ioslevel command 45

L

license command 46
lsdev command 47
lsfailedlogin command 50
lsgcl command 51
lslv command 52
lsmmap command 55
lsnetvc command 57
lspath command 58
lspv command 61
lssw command 63
lsuser command 65
lsvg command 66

M

migratepv command 69
mirrorios command 71
mklv command 72
mklvcopy command 73
mkpath command 74
mktcpip command 76
mkuser command 78
mkvdev command 79
mkvg command 82
mount command 83

N

netstat command 85

O

oem_platform_level command 89
oem_setup_env command 90
optimizenet command 91

P

passwd command 97
ping command 98

R

reducevg command 100
remote_management command 101

restorevgstruct command 102
rmdev command 103
rmlv command 104
rmlvcopy command 105
rmpath command 106
rmuser command 108

S

savevgstruct command 109
showmount command 110
shutdown command 111
snap command 112
startnetsvc command 113
startsysdump command 114
starttrace command 115
stopnetsvc command 116
stoptrace command 117
syncvg command 118

T

topas command 119
traceroute command 128

U

unmirrorios command 130
unmount command 131
updateios command 132

X

xlplm
 command 133
xlplm command 133
xlpstat command 135



Printed in USA