

NAME

asmmenu - launch Advanced System Management menu

SYNOPSIS

asmmenu --ip *IP-address* [**--help**]

DESCRIPTION

asmmenu launches the Advanced System Management (ASM) interface for a managed system using the browser on the Hardware Management Console (HMC).

OPTIONS

--ip The IP address or host name of the service processor for the managed system for which to launch the ASM browser interface.

EXAMPLES

Launch the ASM browser interface for the managed system's service processor which has the IP address 192.168.131.25:

asmmenu --ip 192.168.131.25

ENVIRONMENT

None

BUGS

None

AUTHOR

IBM Austin

NAME

bkconsdata - back up console data

SYNOPSIS

```
bkconsdata -r { usb | ftp | sftp | nfs | none }  
[-h host-name] [-u user-ID] [--passwd password] [-k SSH-private-key]  
[-l mount-location] [-o "mount-command-options"]  
[-d remote-directory] [-x netcfg] [-i perfmon]  
[--help]
```

DESCRIPTION

bkconsdata backs up critical Hardware Management Console (HMC) data, which is HMC data that is stored on the HMC hard disk. This backup data may be used to restore the HMC if the HMC needs to be reinstalled from the HMC recovery CDs in the event of an HMC hard disk failure.

OPTIONS

-r The archive option for the backup data. Valid values are **usb** for a USB flash memory device, **ftp** for a remote FTP site, **sftp** for a remote secure FTP (SFTP) site, **nfs** for an NFS file system, or **none**. The **none** option is used to remove the backup task lock allowing this task to be re-executed immediately.

-h The host name or IP address of the remote server.

This option is required when backing up the data to a remote FTP site, SFTP site, or NFS file system. Otherwise, this option is not valid.

-u The user ID to use to log in to the remote FTP or SFTP site.

This option is required when backing up the data to a remote FTP or SFTP site. Otherwise, this option is not valid.

--passwd

The password to use to log in to the remote FTP or SFTP site.

If this option is omitted and the data is being backed up to a remote FTP site, you will be prompted to enter the password. If both this option and the **-k** option are omitted and the data is being backed up to a remote SFTP site, you will be prompted to enter the password. The **--passwd** and **-k** options are mutually exclusive.

This option is only valid when backing up the data to a remote FTP or SFTP site.

-k The name of the file that contains the SSH private key. If the file name is not fully qualified, the file must exist in the user's home directory on the HMC.

Use the **ssh-keygen** command to generate the public and private SSH key pair. The **ssh-keygen** command is not allowed to write to the **.ssh** directory in the user's home directory on the HMC, so when you run the command on the HMC, you must specify both the directory and the file name for the private key. If you generate a key with a passphrase, you will be prompted to enter the passphrase when you run any HMC command that uses the key.

If both this option and the **--passwd** option are omitted and the data is being backed up to a remote SFTP site, you will be prompted to enter the password. The **-k** and **--passwd** options are mutually exclusive.

This option is only valid when backing up the data to a remote SFTP site.

-l The mount location defined on the NFS server where the backup data will be written.

This option is required when backing up the data to an NFS file system. Otherwise, this option is

not valid.

- o** Options to be passed to the **mount** command used to mount the NFS file system where the backup data will be written. The options must be enclosed in double quotes.

The HMC supports NFS versions 3 and 4, and this command uses version 3 by default. To use version 4, include the option **vers=4**.

This option is only valid when backing up the data to an NFS file system.

- d** The directory on the remote server to which to write the backup data. If this option is not specified when backing up the data to a remote FTP or SFTP site, the backup data will be written to the user's home directory. If this option is not specified when backing up the data to an NFS file system, the backup data will be written to the *mount-location* on the NFS server.

This option is only valid when backing up the data to a remote FTP site, SFTP site, or NFS file system.

- x** The data to be excluded from the backup data. The only valid value is **netcfg** to exclude network configuration data.
- i** Additional data to be included in the backup data. The only valid value is **perfmon** to include performance monitoring data.
- help** Display the help text for this command and exit.

EXAMPLES

Back up critical HMC data to a USB flash memory device:

```
bkconsdata -r usb
```

Back up critical HMC data to a remote FTP site:

```
bkconsdata -r ftp -h ftpserver -u ftpuser --passwd ftppassword
```

Back up critical HMC data to a remote SFTP site using SSH keys for authentication and include performance monitoring data:

```
bkconsdata -r sftp -h sftpserver -u sftpuser -i perfmon  
-k /home/hmcuser/keys/id_rsa
```

Back up critical HMC data to an NFS file system using NFS version 4:

```
bkconsdata -r nfs -h 9.3.145.52 -l /home/hmc/backups -o "vers=4"
```

ENVIRONMENT

None

BUGS

None

AUTHOR

IBM Austin

NAME

bkprofdata - back up profile data

SYNOPSIS

bkprofdata -m *managed-system* **-f** *file* [**--force**] [**--help**]

DESCRIPTION

bkprofdata backs up profile data for the *managed-system*.

The **rstprofdata** command can be used to restore profile data for the *managed-system*.

The **rmprofdata** command can be used to remove a profile data backup file.

OPTIONS

-m The name of the managed system for which to back up profile data. The name may either be the user-defined name for the managed system, or be in the form *ttt-mmm*sssssss*, where *ttt* is the machine type, *mmm* is the model, and *sssssss* is the serial number of the managed system. The *ttt-mmm*sssssss* form must be used if there are multiple managed systems with the same user-defined name.

-f The name of the backup file where the profile data is to be written. If *file* is not fully qualified, *file* will be written to the */var/hsc/profiles/serial-number* directory on the HMC (*serial-number* is the serial number of the managed system).

If *file* already exists, this command will fail unless the **--force** option is specified.

To back up the profile data to removable media, the media must be present in the removable media device and the device must be mounted with the **mount** command before this command is issued. The **lsmediadev** command can be used to display all of the removable media devices on the HMC.

--force This option allows the specified backup file to be overwritten if it already exists. If the specified backup file already exists and this option is not specified, the backup will fail.

--help Display the help text for this command and exit.

EXAMPLES

Back up the profile data to a file called **myFile** on a USB flash memory device (a USB flash memory device must already be connected to the HMC):

lsmediadev (to obtain mount points)

mount /media/sdb1

bkprofdata -m mySystem -f /media/sdb1/myFile

Back up the profile data to a file called **backup1** (the file will be written to the */var/hsc/profiles/3413444* directory on the HMC):

bkprofdata -m 9406-570*3413444 -f backup1

Back up the profile data to a file called **backup1** (the file will be written to the */var/hsc/profiles/3413444* directory on the HMC). The file **backup1** will be overwritten if it already exists:

bkprofdata -m 9406-570*3413444 -f backup1 --force

ENVIRONMENT

None

BUGS

None

AUTHOR

IBM Austin

SEE ALSO

lsmediadev, rmprofdata, rstprofdata

NAME

chaccfg - change access control configuration

SYNOPSIS

```
chaccfg -t {resourcerole | taskrole}
{-f configuration-file | -i "configuration-data"}
[--help]
```

DESCRIPTION

chaccfg changes the configuration of an access control role.

OPTIONS

- t** The type of access control role to change. Valid values are **resourcerole** for managed resource role and **taskrole** for task role.
- f** The name of the file containing the configuration data needed to change the access control role. The configuration data consists of attribute name/value pairs, which are in comma separated value (CSV) format. These attribute name/value pairs form a configuration record. A line feed marks the end of a configuration record. There can only be one configuration record in the file.

The format of a configuration record is as follows:

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

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

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

When a list of values is specified, the attribute name/value pair must be enclosed in double quotes. Depending on the shell being used, nested double quote characters may need to be preceded by an escape character, which is usually a `'\'` character.

If `'+='` is used in the attribute name/value pair instead of `'='`, then the specified *value* is added to the existing value for the attribute if the attribute is numerical. If the attribute is a list, then the specified *value(s)* is added to the existing list.

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

The `'+='` and `'-='` operators can only be used when changing a managed resource role.

Attribute names for managed resource roles:

- name**
name of the managed resource role to change (required)
- resources**
comma separated list of managed resource objects (required)

Attribute names for task roles:

- name**
name of the task role to change (required)
- resources**
comma separated list of tasks (required)

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

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

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

- help** Display the help text for this command and exit.

EXAMPLES

Change the managed resource objects for the managed resource role **lpar_role** where **XXX** is a resource name returned from **lsaccfg -t resource --script**:

```
chaccfg -t resourcerole -i "name=lpar_role,resources=XXX"
```

Add a managed resource object to the managed resource role **mr1** where **XXX** is a resource name returned from **lsaccfg -t resource --script**:

```
chaccfg -t resourcerole -i "name=mr1,resources+=XXX"
```

Change the task role **tr1**:

```
chaccfg -t taskrole -i "name=tr1,"resources=  
cec:ChangeCECProperty+CECPowerOn+CECPowerOff,  
lpar:ChangeLPARProperty+ChangeProfileProperty""
```

Change a task role using the configuration data in the file **/tmp/cfgFile**:

```
chaccfg -t taskrole -f /tmp/cfgFile
```

ENVIRONMENT

None

BUGS

None

AUTHOR

IBM Austin

SEE ALSO

lsaccfg, **mkaccfg**, **rmaccfg**

NAME

chcod - change Capacity on Demand

SYNOPSIS

To enter a CoD code:

```
chcod -o e -m managed-system -k CoD-code
```

To activate or change the amount of On/Off CoD resources:

```
chcod -o a -m managed-system -c onoff -r {mem | proc}  
-q quantity-of-resources -d number-of-days
```

To activate or change the number of Utility CoD processors that are not unlimited Utility CoD processors:

```
chcod -o a -m managed-system -c utility -r proc  
-q quantity-of-processors
```

To deactivate all On/Off CoD or all Trial CoD resources:

```
chcod -o d -m managed-system -c {onoff | trial}  
-r {mem | proc}
```

To deactivate all Utility CoD processors that are not unlimited Utility CoD processors:

```
chcod -o d -m managed-system -c utility -r proc
```

To set or disable a Utility CoD processor minute usage limit:

```
chcod -o s -m managed-system -c utility -r proc  
-l number-of-processor-minutes
```

DESCRIPTION

chcod performs Capacity on Demand (CoD) operations on the *managed-system*.

chcod is used to enter a CoD code for the *managed-system*. It is also used to activate On/Off CoD or Utility CoD resources, or to deactivate On/Off CoD, Trial CoD, or Utility CoD resources. CoD resources are either memory or processors.

chcod is also used to set or disable a Utility CoD processor minute usage limit.

OPTIONS

-o The CoD operation to perform. Valid values are **e** to enter a CoD code, **a** to activate or change the number of On/Off CoD or Utility CoD resources, **d** to deactivate all On/Off CoD, all Trial CoD, or all Utility CoD resources, and **s** to set or disable a Utility CoD processor minute usage limit.

To change the number of On/Off CoD resources or days that are currently activated, specify an activate operation with this option, specify the total number of activated On/Off CoD resources that you want with the **-q** option, and specify the number of days that you want them for with the **-d** option. An activate operation is to be used regardless of whether the total number of activated On/Off CoD resources is being increased, decreased, or kept the same. A deactivate operation is to be used only when you no longer want to have any activated On/Off CoD resources.

To change the number of Utility CoD processors that are currently activated, specify an activate operation with this option, and specify the total number of activated Utility CoD processors that you want with the **-q** option. An activate operation is to be used regardless of whether the total number of activated Utility CoD processors is being increased or decreased. A deactivate operation is to be used only when you no longer want to have any activated Utility CoD processors in the shared processor pool.

Note that CUoD, Trial CoD, and unlimited Utility CoD resources are activated by entering a CoD code. Unlimited Utility CoD processors are also deactivated by entering a CoD code.

- m** The name of the managed system for which the CoD operation is to be performed. The name may either be the user-defined name for the managed system, or be in the form *tttt-mmm*sssssss*, where *tttt* is the machine type, *mmm* is the model, and *sssssss* is the serial number of the managed system. The *tttt-mmm*sssssss* form must be used if there are multiple managed systems with the same user-defined name.
- k** The CoD code (key) to enter. Letters may be entered in either upper case or lower case.
- c** The CoD type. Valid values are **onoff** for On/Off CoD, **trial** for Trial CoD, and **utility** for Utility CoD.
- r** The CoD resource type. Valid values are **mem** for memory and **proc** for processors.
- q** The quantity of On/Off CoD or Utility CoD resources requested. The value specified must be a whole number which is greater than 0.

When requesting On/Off CoD memory, the value specified must be in megabytes, and it must be a multiple of 1024 megabytes.
- d** The number of days for which the On/Off CoD resources are requested.
- l** The number of processor minutes to set as the Utility CoD processor minute usage limit.

To disable the Utility CoD processor minute usage limit, specify a 0 with this option.
- help** Display the help text for this command and exit.

EXAMPLES

Enter a CoD code:

```
chcod -m sys1 -o e -k AlphaNumericString12345
```

Activate 2 GB of On/Off CoD memory for 10 days:

```
chcod -m 9406-570*1001CA -o a -c onoff -r mem -q 2048 -d 10
```

Increase to 5 GB of On/Off CoD memory for 3 days:

```
chcod -m 9406-570*1001CA -o a -c onoff -r mem -q 5120 -d 3
```

Activate 4 On/Off CoD processors for 5 days:

```
chcod -m sys1 -o a -c onoff -r proc -q 4 -d 5
```

Decrease to 1 On/Off CoD processor for 5 days:

```
chcod -m sys1 -o a -c onoff -r proc -q 1 -d 5
```

Activate 1 Utility CoD processor:

```
chcod -m 9117-MMA*10BACEC -o a -c utility -r proc -q 1
```

Activate 2 more Utility CoD processors for a total of 3 Utility CoD processors:

```
chcod -m 9117-MMA*10BACEC -o a -c utility -r proc -q 3
```

Decrease the number of activated Utility CoD processors to 2:

chcod -m 9117-MMA*10BACEC -o a -c utility -r proc -q 2

Deactivate all On/Off CoD processors:

chcod -m sys1 -o d -c onoff -r proc

Deactivate all Utility CoD processors:

chcod -m sys1 -o d -c utility -r proc

Deactivate all Trial CoD memory and stop the trial:

chcod -m 9406-570*1001CA -o d -c trial -r mem

Set a Utility CoD processor minute usage limit of 100 processor minutes:

chcod -m sys1 -o s -c utility -r proc -l 100

Disable the Utility CoD processor minute usage limit:

chcod -m sys1 -o s -c utility -r proc -l 0

ENVIRONMENT

None

BUGS

None

AUTHOR

IBM Austin

SEE ALSO

lscod

NAME

chcodpool - change a CoD Power enterprise pool

SYNOPSIS

To update a Power enterprise pool:

```
chcodpool -o update {-p pool-name | --id pool-ID}  
[-f file] [-a "attributes"] [-v] [--help]
```

To add or remove Mobile CoD resources to or from a system in a Power enterprise pool:

```
chcodpool -o {add | remove} {-p pool-name | --id pool-ID}  
-m managed-system -r {mem | proc} -q quantity  
[--force] [--help]
```

To add, remove, or reauthenticate a managing Hardware Management Console (HMC):

```
chcodpool -o {add | remove | auth} {-p pool-name | --id pool-ID}  
--mc IP-address [-u user-ID [--passwd password]]  
[--force] [--forcenoconn] [--help]
```

To change the master HMC for a Power enterprise pool:

```
chcodpool -o setmaster {-p pool-name | --id pool-ID}  
--mc {this | IP-address} [--force] [--help]
```

To recover a Power enterprise pool after a clean install of the master HMC:

```
chcodpool -o recover -p pool-name -f file [-a "attributes"]  
[-v] [--help]
```

To synchronize Power enterprise pool information:

```
chcodpool -o sync {-p pool-name | --id pool-ID} [--help]
```

DESCRIPTION

chcodpool performs Capacity on Demand (CoD) Power enterprise pool operations.

OPTIONS

-o The operation to perform. Valid values are **update** to update a Power enterprise pool, **add** to add Mobile CoD resources to a system in a pool or to add a managing HMC to a pool, **remove** to remove Mobile CoD resources from a system in a pool or to remove a managing HMC from a pool, **auth** to reauthenticate a managing HMC in a pool, **setmaster** to change the master HMC for a pool, **recover** to recover a pool after a clean install of the master HMC for the pool, and **sync** to synchronize the pool information between the master HMC for a pool, the managing HMCs in the pool, and the systems in the pool.

-p The name of the Power enterprise pool for which to perform the operation.

Except when recovering a pool, you can either use this option to specify the name of the pool, or use the **--id** option to specify the pool's ID. The **--id** option must be used if there are multiple Power enterprise pools with the same name. The **-p** and the **--id** options are mutually exclusive.

When recovering a pool, you must use this option to specify a name for the pool being recovered.

--id The ID of the Power enterprise pool for which to perform the operation.

You can either use this option to specify the ID of the pool, or use the **-p** option to specify the pool's name. The **--id** option must be used if there are multiple Power enterprise pools with the same name. The **--id** and the **-p** options are mutually exclusive.

-m The name of the managed system to which to add or remove Mobile CoD resources. The name may either be the user-defined name for the managed system, or be in the form *tttt-mmmm*ssssss*, where *tttt* is the machine type, *mmm* is the model, and *ssssss* is the serial number of the managed

system. The *ttt-mmm*sssss* form must be used if there are multiple managed systems with the same user-defined name.

- r The type of Mobile CoD resources to add or remove. Valid values are **mem** for Mobile CoD memory, and **proc** for Mobile CoD processors.
- q The number of Mobile CoD processors or the amount of Mobile CoD memory to add or remove. The number of processors must be a whole number. The amount of memory must be specified in megabytes and must be a multiple of gigabytes.
- f The name of the XML file that contains the Power enterprise pool configuration data. If *file* is not fully qualified, *file* must exist in the user's home directory on the HMC.

If *file* exists on removable media, the media must be present in the removable media device and the device must be mounted with the **mount** command before this command is issued. The **lsmediadev** command can be used to display all of the removable media devices on the HMC.

This option must be specified when updating which systems are members of a Power enterprise pool, or when updating the amount of Mobile CoD resources in a pool. This option must also be specified when recovering a pool following the clean install of the master HMC for the pool.

- a The Power enterprise pool attributes to update or recover. The attribute data consists of attribute name/value pairs, which are in comma separated value (CSV) format.

The format of the attribute data is as follows:

attribute-name=value,attribute-name=value,...

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

"attribute-name=value,value,...",...

When a list of values is specified, the attribute name/value pair must be enclosed in double quotes. Depending on the shell being used, nested double quote characters may need to be preceded by an escape character, which is usually a `'\`` character.

Valid attribute names for updating a pool:

new_name

Valid attribute names for recovering a pool:

mcs

Comma separated list of managing HMCs for the pool, with each managing HMC having the following format:

host-name-or-IP-address/user-ID[password]

Both `'/'` characters must be present, but *password* may be omitted. If *password* is omitted, you will be prompted to enter the password.

Note: do not include the HMC which is running this command in the list.

For example:

hmc1/user1/ prompts the user for **user1**'s password then uses that password and the user ID **user1** to log into the HMC with host name **hmc1** to add that HMC to

the pool.

--mc Specify the HMC host name or IP address. When setting a new master HMC for the Power enterprise pool, you can specify **this** to set the HMC running this command as the new master.

When setting a new master HMC for the pool, if the current master HMC for the pool is running you can run this command from any managing HMC and set any managing HMC in the pool as the new master. If the current master HMC for the pool is not running or not reachable, you must run this command on the HMC that you want to set as the new master and you must specify the **--force** option.

-u Specify the user ID to use to log into the managing HMC being added or reauthenticated.

--passwd

Specify the password for *user-ID*. If this option is omitted, you will be prompted to enter the password.

--force Specify this option to force a Mobile CoD add or remove operation to occur for a system that is not currently reachable from the master HMC for the pool. If Mobile CoD resources are removed from a system that is not currently reachable from the master HMC, those resources will be considered unreturned until the system can be reached from the master HMC again. A grace period compliance timer will be started for that system as soon as the resources are removed. To add Mobile CoD resources to a system that is not currently reachable from the master HMC for the pool, the system must have unreturned Mobile CoD resources, and you cannot add more Mobile CoD resources to the system than are unreturned on that system.

When setting a new master HMC for a pool, you must specify this option to force the HMC to set itself as the new master HMC for a pool when the current master HMC for the pool is not running or not reachable.

When removing a managing HMC from the pool, you must specify this option if the HMC you want to remove is the only HMC managing one or more systems in the pool.

--forcenoconn

Specify this option to remove a managing HMC from a pool when the master HMC for the pool is not able to reach the managing HMC. After the HMC is removed from the pool, the pool data will remain on the managing HMC but will no longer be updated. To remove the pool data from the managing HMC, you must perform a clean install of the HMC.

-v Specify this option to enable verbose mode. When verbose mode is enabled, warning messages and informational messages are displayed for successful update pool and recover pool operations. Warning messages and informational messages are displayed for update pool and recover pool operations that partially succeed or fail, regardless of whether this option is specified.

--help Display the help text for this command and exit.

EXIT STATUS

This command has the following return codes:

0	Success
100	Partial success

This value is returned when at least one portion of an update pool or recover pool operation succeeded, and at least one failed.

Any other value means the command failed.

EXAMPLES

Update the pool configuration for Power enterprise pool **myPool** and change the pool's name to **pool1** (the configuration file **pool1032.xml** is in the user's home directory on the HMC):

```
chcodpool -o update -p myPool -f pool1032.xml -a "new_name=pool1" -v
```

Add 2 Mobile CoD processors to **sys4** in Power enterprise pool **myPool**:

```
chcodpool -o add -m sys4 -p myPool -r proc -q 2
```

Remove 4 GB of Mobile CoD memory from **sys4** in Power enterprise pool **myPool**:

```
chcodpool -o remove -m sys4 -p myPool -r mem -q 4096
```

Add the HMC with host name **hmc1** to the Power enterprise pool **myPool** (you will be prompted for the password):

```
chcodpool -o add --mc hmc1 -u user1 -p myPool
```

Remove the HMC with IP address **192.168.32.25** from the Power enterprise pool with ID **0235**:

```
chcodpool -o remove --mc 192.168.32.25 --id 0235
```

Set the HMC with host name **hmc1.company.com** as the new master HMC for Power enterprise pool **pool3**:

```
chcodpool -o setmaster --mc hmc1.company.com -p pool3
```

Set this HMC as the new master HMC for the Power enterprise pool with ID **1234** when the current master HMC is not reachable:

```
chcodpool -o setmaster --mc this --id 1234 --force
```

Recover the Power enterprise pool **pool1** using the configuration file **pool0195.xml** file on a USB flash memory device (the USB flash memory device must already be connected to the HMC). The managing HMCs for the pool are **hmc1** and **hmc2** (you will be prompted to enter the password for each HMC):

lsmediadev (to obtain mount points)

```
mount /media/sdb1
```

```
chcodpool -o recover -p pool1 -f /media/sdb1/pool0195.xml -v  
-a ""mcs=hmc1/myuserID/,hmc2/myuserID/""
```

Synchronize the Power enterprise pool information for the pool **myPool**:

```
chcodpool -o sync -p myPool
```

ENVIRONMENT

None

BUGS

None

AUTHOR

IBM Austin

SEE ALSO

lscodpool, mkcodpool, lsmediadev

NAME

chcomgmt - change co-management settings

SYNOPSIS

```
chcomgmt -m managed-system -o {setmaster | relmaster}  
[-t {norm | keep}] [--force] [--help]
```

DESCRIPTION

chcomgmt changes co-management settings for the *managed-system*.

This command is used to set or release this Hardware Management Console (HMC) as the PowerVM management master for the managed system. You must set this HMC as the PowerVM management master for a managed system when you want to perform an operation from this HMC that requires this HMC to be the PowerVM management master. In addition, you can choose to set this HMC as the PowerVM management master for a managed system when you want this HMC to be the only HMC that can perform PowerVM management for the managed system. Other HMCs would be able to perform PowerVM view operations only.

OPTIONS

- m** The name of the managed system for which to change co-management settings. The name may either be the user-defined name for the managed system, or be in the form *ttt-mmm*ssssss*, where *ttt* is the machine type, *mmm* is the model, and *ssssss* is the serial number of the managed system. The *ttt-mmm*ssssss* form must be used if there are multiple managed systems with the same user-defined name.
- o** The operation to perform. Valid values are **setmaster** to set this HMC as the PowerVM management master for the managed system, and **relmaster** to release this HMC as the PowerVM management master for the managed system.
- t** The type of PowerVM management master to set. Valid values are **norm** to set this HMC as a normal master, and **keep** to set this HMC as a persistent master.

If you set this HMC as the PowerVM management master for a managed system, it is recommended that you release it as master as soon as master mode is no longer needed, unless you want this HMC to remain the only HMC that can perform PowerVM management for the managed system. If you choose normal master mode and you do not release it, it will be gracefully released when another PowerVM manager, such as another HMC, sets itself as the PowerVM management master for the managed system. If you choose persistent master mode and you do not release it, no other PowerVM manager will be able to set itself as the PowerVM management master for the managed system without forcing the operation. It is recommended that you do not choose persistent master mode unless directed to do so by your hardware service representative.

To set this HMC as the persistent master for a managed system, you must first set this HMC as the normal master, unless there is no current PowerVM management master for the managed system. Once this HMC is the current master, then you can set this HMC as the persistent master. If there is no current PowerVM management master for the managed system, you can immediately set this HMC as the persistent master without setting it as the normal master first.

This option is required when setting the PowerVM management master. This option is not valid otherwise.

- force** Specify this option to force this HMC to be set as the PowerVM management normal master for the managed system. Use this option when you are not able to get the current PowerVM management master to release itself as the master. **Warning: use this option carefully. Forcing the set master operation can result in any operations in progress on the current PowerVM management master to fail, which can have unexpected impacts on managed system and partition configuration and states.**

If the current PowerVM management master for the managed system is in normal master mode, you must first try to set this HMC as the normal master without using the **--force** option. After doing so, if this HMC remains in pending master mode, then you can immediately set this HMC as the normal master by running this command again using the **--force** option. If the current PowerVM management master for the managed system is in persistent master mode, you can immediately set this HMC as the normal master by using the **--force** option.

This option is only valid when setting this HMC as the PowerVM management normal master for a managed system.

--help Display the help text for this command and exit.

EXAMPLES

Set this HMC as the PowerVM management normal master for managed system **sys1**:

```
chcomgmt -m sys1 -o setmaster -t norm
```

Release this HMC as the PowerVM management master for managed system **sys1**:

```
chcomgmt -m sys1 -o relmaster
```

ENVIRONMENT

None

BUGS

None

AUTHOR

IBM Austin

SEE ALSO

lscomgmt

NAME

chhmc - change Hardware Management Console (HMC) configuration information

SYNOPSIS

To change remote access settings:

```
chhmc -c {ssh | xntp | sol}
  -s {enable | disable | modify}
  [--sshprotocol {1 | 2 | all}]
  [--help]
```

To change remote web browser access:

```
chhmc -c remotewebui
  -s {enable | disable}
  -i interface
  [-a {IPv4-address | IPv6-address[/prefix-length]}]
  [-nm IPv4-network-mask]
  [--help]
```

To enable or disable TLS encrypted system logging over TCP:

```
chhmc -c syslog -t tls
  -s {enable | disable}
  [--force]
  [--help]
```

To add, modify, or remove a remote system logging destination:

```
chhmc -c syslog
  -s {add | modify | remove}
  {-a IP-address | -h host-name}
  [-t {tcp | tls | udp}]
  [-f input-data-file | --input "input-data"]
  [--help]
```

To add or remove an entry in the network time protocol configuration file:

```
chhmc -c xntp
  -s {add | remove}
  {-a {IPv4-address | IPv6-address[/prefix-length]} |
  -h host-name}
  [--ntpversion {1 | 2 | 3 | 4}]
  [-nm IPv4-network-mask]
  [-i interface]
  [--help]
```

To configure whether keyboard mapping configuration or display configuration will occur at the next reboot:

```
chhmc -c {kbdcfg | dispcfg}
  -s {enable | disable}
  [--help]
```

To configure an alternate disk partition on the HMC as a startup device:

```
chhmc -c altdiskboot
  -s {enable | disable}
  --mode {install | upgrade}
  [--help]
```

To configure firewall settings to permit or deny IP addresses from utilizing HMC services:

```
chhmc -c {ssh | service}
-s {add | remove}
-a {IPv4-address | IPv6-address[prefix-length]}
[-nm IPv4-network-mask]
[-i interface]
[--help]
```

If **-nm** is omitted, the default IPv4 network mask will be **255.255.255.255**.

If *prefix-length* is omitted, the default IPv6 prefix length will be **128**.

If **-i** is omitted, the rule will be applied to all interfaces.

To add or remove a static route:

```
chhmc -c netroute
-s {add | remove}
[--routetype {host | net}]
[-a {IPv4-address | IPv6-address[prefix-length]}]
[-nm IPv4-network-mask]
[-g gateway]
[-i interface]
[--position position]
```

Use the **chhmc -c network -s modify -g** command to modify the default route.

To add or remove entries from the DNS server search order or the domain suffix search order:

```
chhmc -c network
-s {add | remove}
[-ns DNS-server]
[-ds domain-suffix]
[--help]
```

To change network settings for a specific network interface:

```
chhmc -c network
-s modify
-i interface
[-a {none | list}]
[-nm IPv4-network-mask]
[--ipv6auto {on | off}]
[--ipv6privacy {on | off}]
[--ipv6dhcp {on | off}]
[--ipv4dhcp {on | off}]
[--lparcomm {on | off}]
[--tso {on | off}]
[--speed {auto | 10 | 100 | 1000}]
[--duplex {auto | half | full}]
[--jumboframe {on | off}]
[--help]
```

To change other network settings:

chhmc -c network**-s modify**[-h *host-name*][-d *network-domain-name*][-g *gateway* [-i *interface*]]

[--help]

To add or remove SLP service registration IP addresses (please allow at least 3 minutes for the change to take effect):

chhmc -c slp**-s {add | remove}****-a** {*IPv4-address* | *IPv6-address*[/*prefix-length*]}

[--help]

To change the locale for the HMC:

chhmc -c locale**-s modify****-l** *locale*

[--help]

To change the HMC date and time, or time zone:

chhmc -c date**-s modify**[--datetime *date-time*][--clock {**local** | **utc**}][--timezone *time-zone*]

[--help]

To change the HMC Kerberos configuration:

chhmc -c kerberos**-s {add | modify | remove}**[{-a *KDC-IPv4-address*[:*port*] |-h *KDC-host-name*[:*port*]}][--realm *realm*][--defaultrealm *realm*][--clockskew *clock-skew*][--ticketlifetime *ticket-lifetime*][--kpasswdadmin {*IPv4-address*[:*port*] | *host-name*[:*port*]}][--trace {**on** | **off**}][--weakcrypto {**on** | **off** | **default**}]

[--help]

To change the Integrated Management Module (IMM) settings:

chhmc -c imm**-s modify**[-a *IPv4-address* -nm *IPv4-network-mask* -g *gateway*][-u *user-ID* [--passwd *password*]][--mode {**ded** | **shared**}]

[--help]

To restart the IMM:

chhmc -c imm**-s restart**

[--help]

To change the cipher mode for the HMC chatlet interface use the following command. This setting must be consistent among all HMCs using the chatlet interface (port 9920). Enable sets the cipher mode to the legacy anonymous cipher mode and disable removes the use of anonymous ciphers on the chatlet interface. This setting is disabled when the HMC security mode is set to NIST SP 800-131A.

```
chhmc -c legacyhmccomm
  -s {enable | disable}
  [--help]
```

To change the HMC security mode (the HMC will automatically be restarted):

```
chhmc -c security
  -s modify
  --mode {legacy | nist_sp800_131a}
  [-r]
  [--help]
```

To enable or disable the Event Manager for Call Home:

```
chhmc -c emch
  -s {enable | disable}
  [--callhome {enable | disable}]
  [--help]
```

To enable, disable, or change the GRUB password:

```
chhmc -c grubpasswd
  -s {enable | disable | modify}
  [--passwd password]
  [--help]
```

DESCRIPTION

chhmc changes Hardware Management Console (HMC) configuration information, such as remote access settings and network settings.

OPTIONS

- c** The type of configuration to be modified. Valid values are **ssh**, **syslog**, **xntp**, **netroute**, **network**, **slp**, **kbdfg**, **dispcfg**, **altdiskboot**, **locale**, *service*, **date**, **remotewebui**, **kerberos**, **imm**, **sol** for Serial Over LAN, **legacyhmccomm**, **security**, **emch** for Event Manager for Call Home, and **grubpasswd**.
- s** The new state value of the configuration. Valid values are **enable**, **disable**, **add**, **modify**, **remove**, and **restart**.
- i** The interface to configure, such as **eth0** or **sl0**.

If this option is omitted when adding a static route, the route will be added for each interface.
- a** The network IP address.

For network interface configuration, this is the static IP address configuration. Valid values are:

```
none - no static IP address configured
list - specify a comma-separated list of static IP
addresses. The list can contain zero or one
static IPv4 address and zero or more static
IPv6 addresses in the following format:
static-address/prefix-length
```

For SLP service registration IP address configuration, specify a comma-separated list of zero or more IPv4 addresses and zero or more IPv6 addresses. IPv6 addresses must be specified in the following format: *IPv6-address[/prefix-length]*. If *prefix-length* is omitted, the default prefix

length will be 128.

For all other operations except Kerberos and IMM configuration, an IPv4 or IPv6 address can be specified.

For Kerberos configuration, use this option to specify the IPv4 address of the Key Distribution Center (KDC). An optional port number can be specified following the IPv4 address. If the KDC has an IPv6 address, then you cannot use this option to specify the IP address of the KDC. You must use the **-h** option to specify the host name of the KDC instead. Either this option or the **-h** option to specify the KDC is required for a Kerberos add or remove operation. This option is not valid for a Kerberos modify operation.

-nm The IPv4 network mask.

--ipv6auto

The IPv6 autoconfiguration setting for the network interface.

Valid values are:

- on** - autoconfigure IPv6 addresses
- off** - do not autoconfigure IPv6 addresses

--ipv6privacy

The IPv6 privacy extension setting for the network interface.

Valid values are:

- on** - use privacy extensions for autoconfiguration
- off** - do not use privacy extensions for autoconfiguration

--ipv6dhcp

The IPv6 DHCP setting for the network interface.

Valid values are:

- on** - obtain an IPv6 address automatically via DHCP
- off** - do not obtain an IPv6 address automatically via DHCP

--ipv4dhcp

The IPv4 DHCP setting for the network interface.

Valid values are:

- on** - obtain an IPv4 address automatically via DHCP
- off** - do not obtain an IPv4 address automatically via DHCP

This option cannot be used if a static IPv4 address has also been specified with the **-a** option.

--lparcomm

The partition communication setting for the network interface.

This option has been deprecated. Use the **chpsm** command to configure which network interfaces are enabled for partition communication.

--tso The TCP segmentation offload (TSO) setting for the network interface.

Valid values are:

- on** - enable TSO
- off** - disable TSO

--speed The speed setting for the network interface.

Valid values are:

auto - automatically detect and set speed
10 - 10Mbps
100 - 100Mbps
1000 - 1000Mbps

Specify **auto** unless you have a requirement to use a fixed speed setting.

--duplex

The duplex setting for the network interface.

Valid values are:

auto - automatically detect and set duplex
half - half duplex
full - full duplex

Specify **auto** unless you have a requirement to use a fixed duplex setting.

Valid combinations of **--speed** and **--duplex** are:

--speed auto --duplex auto
--speed 10 --duplex half
--speed 10 --duplex full
--speed 100 --duplex half
--speed 100 --duplex full
--speed 1000 --duplex full

--jumboframe

Enables or disables jumbo frames on the network interface. The jumbo frame size is 9000 bytes.

Valid values are:

on - enable jumbo frames
off - disable jumbo frames

Do not enable jumbo frames on the network interface that is used to connect to managed systems or managed frames. Managed systems and managed frames do not support jumbo frames.

-d The network domain name.

-h The host name.

For Kerberos configuration, use this option to specify the host name of the Key Distribution Center (KDC). An optional port number can be specified following the host name. Also, either this option or the **-a** option to specify the KDC is required for a Kerberos add or remove operation. This option is not valid for a Kerberos modify operation.

-g The default gateway IP address.

-ns The nameserver IP address to add or remove.

-ds The domain suffix to add or remove.

--routetype

The type of static route to add. Valid values are **host** to add a static route to a host and **net** to add a static route to a network.

--position

The position in the routing table of the static route to add or remove.

When adding a static route, if this option is omitted or if the position specified is greater than the position of the first default route entry, the route entry will be added to the routing table just before

the first default route entry.

The **lshmc --netroute** command can be used to display the position of all of the static route entries in the routing table.

--ntpversion

The NTP version. Versions 1 - 4 are supported.

--mode When configuring an alternate disk partition on the HMC as the startup device, specify the mode to use when starting up. Valid values are **install** and **upgrade**. The default value is **upgrade**.

When configuring the HMC IMM settings, specify the IMM network interface mode. Valid values are **ded** for dedicated and **shared**. The default value is **ded**.

When configuring the HMC security mode, valid values are **legacy** for no security mode, and **nist_sp800_131a** for NIST SP 800-131A (National Institute of Standards and Technology Special Publication 800-131A). The default value is **legacy**. The HMC will automatically be restarted after the security mode is changed.

-r When changing the HMC security mode, specify this option to cause the HMC to automatically be restarted without asking for confirmation.

-l The locale. For a list of all locales supported by the HMC, issue the **lshmc -L** command.

For the new locale to take effect for the local HMC console, you must log off the console and log back on. For the new locale to take effect for the HMC command line, you must reboot the HMC.

--datetime

The new date and time to set on the HMC's clock. *date-time* must be specified using the following format:

MMDDhhmm[[*CC*]*YY*][*.ss*]

where *MM* is the month, *DD* is the day, *hh* is the hour in 24 hour format, *mm* is the minutes, *CC* is the century, *YY* is the year, and *ss* is the seconds.

This option is required when the **--clock** option is specified.

You must reboot the HMC after setting the date and time.

--clock The BIOS clock type. Valid values are **local** and **utc**.

--timezone

The time zone to set for the HMC. *time-zone* must be specified in *continent/city* format.

You must reboot the HMC after setting the time zone.

--realm

The Kerberos realm name.

When the first Kerberos realm and KDC is added to the HMC, that realm is set as the default realm and Kerberos is automatically enabled on the HMC. When the last Kerberos realm and KDC is removed from the HMC, Kerberos is automatically disabled on the HMC. Then, if another Kerberos realm and KDC is added to the HMC, Kerberos will automatically be enabled again, and the default realm will be set to the new realm. There is no other way to enable or disable Kerberos on the HMC.

This option is required for a Kerberos add or remove operation. This option is not valid for a Kerberos modify operation.

--defaultrealm

The Kerberos default realm name.

When no realms exist on the HMC, the default realm is automatically set to the next realm that is added to the HMC.

This option is not valid for a Kerberos add or remove operation.

--clockskew

The Kerberos clock skew value in seconds. This is the maximum allowable amount of clock skew before Kerberos considers messages invalid. The clock skew is set to the default value of 120 seconds whenever Kerberos is enabled on the HMC, unless this option is specified.

This option is not valid for a Kerberos remove operation.

--ticketlifetime

The Kerberos ticket lifetime value (**lifetime for credentials**). The format of this value is a whole number followed by **s** for seconds, **m** for minutes, **h** for hours, or **d** for days. The ticket lifetime is set to the default value of **2d** (2 days) whenever Kerberos is enabled on the HMC, unless this option is specified.

This option is not valid for a Kerberos remove operation.

--kpasswdadmin

The Kerberos admin server host name or IP address. This option enables Kerberos users to change their own passwords using the **chhmcusr** command. If a realm is not specified, the default realm is used for the operation.

--trace Enables or disables Kerberos authentication logging. When enabled, trace messages are logged in the **/var/log/messages** file on the HMC.

Valid values are:

- on** - enable Kerberos authentication logging
- off** - disable Kerberos authentication logging (default value)

This option is only valid for a Kerberos modify operation.

--weakcrypto

Specifies whether or not Kerberos is allowed to use weak encryption types. This option is not allowed when the HMC security mode is set to NIST SP 800-131A.

Valid values are:

- on** - allow Kerberos to use weak encryption types
- off** - do not allow Kerberos to use weak encryption types
- default** - use the default setting on the HMC

This option is only valid for a Kerberos modify operation.

--sshprotocol

The SSH protocol to use. Valid values are **1** for version 1, **2** for version 2, or **all** for both versions 1 and 2. The default value is **2**.

-u The new IMM user ID.

--passwd

The new IMM user password or the new GRUB password. If this option is omitted, you will be prompted to enter the password.

- t** The type of connection to use for forwarding syslog messages to the remote destination. Valid values are **tcp** for unencrypted TCP, **tls** for TLS encrypted TCP, and **udp** for unencrypted UDP.

If this option is not specified, it defaults to **udp**.

The HMC does not support both encrypted and unencrypted remote system logging destinations simultaneously.

Before TLS encrypted system logging over TCP can be enabled, the **getfile** command must be run to deploy the rsyslog certificate files and private key file on the HMC.

- f** The name of the file containing the input data for this command. The input data consists of attribute name/value pairs, which are in comma separated value (CSV) format.

The format of the input data is as follows:

```
"attribute-name=value,value,...",...
```

When a list of values is specified, the attribute name/value pair must be enclosed in double quotes. Depending on the shell being used, nested double quote characters may need to be preceded by an escape character, which is usually a `'\'` character.

If `'+='` is used in the attribute name/value pair instead of `'='`, the specified value is added to the existing list.

If `'-='` is used in the attribute name/value pair instead of `'='`, the specified value is deleted from the existing list.

Valid attributes for this command:

filter_msg_contains_discard_strings

Defines a property-based syslog message filter that discards all syslog messages containing the specified strings in their message text

Input data for this command can be specified with this option or the **--input** option. The **-f** and the **--input** options are mutually exclusive.

This option is only valid for a syslog add or modify operation.

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

Input data for this command can be specified with this option or the **-f** option. The **--input** and the **-f** options are mutually exclusive.

This option is only valid for a syslog add or modify operation.

--callhome

When disabling the Event Manager for Call Home, use this option to specify whether automatic call home is to remain disabled or automatic call home is to be enabled. This option is required and only allowed when disabling the Event Manager for Call Home.

Valid values are:

disable - automatic call home is to remain disabled
enable - enable automatic call home

When the Event Manager for Call Home is enabled, automatic call home is automatically disabled.

- force** Specify this option to force TLS encrypted system logging over TCP to be disabled when there are remote system logging destinations configured. All configured remote system logging destinations will be removed.
- help** Display the help text for this command and exit.

EXAMPLES

To change the Hardware Management Console host name:

```
chhmc -c network -s modify -h mynewhost
```

To set the IP address and network mask for network interface eth0:

```
chhmc -c network -s modify -i eth0 -a 10.10.10.1  
-nm 255.255.255.0
```

To add a static route to the network **192.165.31.0** for interface **eth0** to the beginning of the routing table:

```
chhmc -c netroute -s add --routetype net -a 192.165.31.0  
-nm 255.255.255.0 -g 192.165.178.1 -i eth0 --position 1
```

To remove the static route entry in position 3 of the routing table:

```
chhmc -c netroute -s remove --position 3
```

To enable remote access via ssh:

```
chhmc -c ssh -s enable
```

To enable remote access via Serial Over LAN:

```
chhmc -c sol -s enable
```

To disable remote web browser access from all IP addresses over network interface eth0:

```
chhmc -c remotewebui -s disable -i eth0
```

To add a remote system logging destination that will receive syslog messages over UDP:

```
chhmc -c syslog -s add -a 10.10.10.2
```

To enable TLS encrypted system logging over TCP:

```
chhmc -c syslog -s enable -t tls
```

To add a remote system logging destination that will receive TLS encrypted syslog messages over TCP:

```
chhmc -c syslog -s add -t tls -h secure.ibm.com
```

To set a syslog filter for the server **secure.ibm.com**. The filter will cause all syslog messages that contain the string **informational** in their message text to be discarded:

```
chhmc -c syslog -s modify -t tls -h secure.ibm.com --input  
"filter_msg_contains_discard_strings=informational"
```

To remove the syslog filter for the server **secure.ibm.com**:

```
chhmc -c syslog -s modify -t tls -h secure.ibm.com --input  
"filter_msg_contains_discard_strings="
```

To enable keyboard mapping configuration to occur on the next reboot:

```
chhmc -c kbdcfg -s enable
```

To enable display configuration to occur on the next reboot:

```
chhmc -c dispcfg -s enable
```

To permit a single IP address to use the ssh service over network interface eth0:

```
chhmc -c ssh -s add -a 10.10.10.3 -nm 255.255.255.255  
-i eth0
```

To enable Network Time Protocol service:

```
chhmc -c xntp -s enable
```

To add a Network Time Protocol server to the configuration file:

```
chhmc -c xntp -s add -h mytimeserver.company.com
```

Specify the hostname will not change firewall rule settings. It is assumed that the user will use the Customize Network Settings to change firewall settings.

To add a Network Time Protocol server to the configuration file, using IP address and at the same time enable firewall access through network interface eth0:

```
chhmc -c xntp -s add -a 10.10.10.32 -i eth0
```

To remove a Network Time Protocol server from the configuration file:

```
chhmc -c xntp -s remove -h mytimeserver.company.com
```

To remove the HMC IP address 9.53.182.99 from the SLP service registration IP addresses:

```
chhmc -c slp -s remove -a 9.53.182.99
```

To set the current locale to Spanish:

```
chhmc -c locale -s modify -l es_ES
```

To set the alternate disk partition on HMC as a startup device on the next HMC boot:

```
chhmc -c altdiskboot -s enable --mode upgrade
```

To update the HMC clock to January 25, 2007 14:30:50 (the current year is 2007):

```
chhmc -c date -s modify --datetime 01251430.50 or
```

chhmc -c date -s modify --datetime 012514302007.50 or

chhmc -c date -s modify --datetime 0125143007.50

To update the HMC clock to December 8, 2008 09:45, local time:

chhmc -c date -s modify --datetime 120809452008 --clock local

To update the HMC time zone to United States, Central time:

chhmc -c date -s modify --timezone America/Chicago

To add a Kerberos realm and KDC to the HMC:

chhmc -c kerberos -s add --realm EXAMPLE.COM -a 10.10.0.20

To add a Kerberos admin server and allow Kerberos users to change their own passwords using the **chhmcusr** command:

chhmc -c kerberos -s add --kpasswdadmin 10.10.0.20 or

chhmc -c kerberos -s add --kpasswdadmin 10.10.0.20 --realm EXAMPLE.COM

To replace a Kerberos admin server:

chhmc -c kerberos -s modify --kpasswdadmin 10.10.0.20 or

chhmc -c kerberos -s modify --kpasswdadmin 10.10.0.20 --realm EXAMPLE.COM

To remove the Kerberos admin server:

chhmc -c kerberos -s remove --kpasswdadmin 10.10.0.20 --realm EXAMPLE.COM

To add a Kerberos realm and KDC and set the clock skew to 140 seconds and ticket lifetime to 1 day:

**chhmc -c kerberos -s add --realm EXAMPLE.COM
-h kdc.example.com:88 --clockskew 140
--ticketlifetime 1d**

To remove a Kerberos realm and KDC from the HMC:

**chhmc -c kerberos -s remove --realm EXAMPLE.COM
-h kdc.example.com**

To modify the Kerberos default realm on the HMC:

chhmc -c kerberos -s modify --defaultrealm EXAMPLE2.COM

To modify the Kerberos clock skew to 60 seconds on the HMC:

chhmc -c kerberos -s modify --clockskew 60

To modify the Kerberos ticket lifetime to 12 hours on the HMC:

```
chhmc -c kerberos -s modify --ticketlifetime 12h
```

To change the IMM network settings:

```
chhmc -c imm -s modify -a 9.3.99.100 -nm 255.255.255.0 -g 9.3.99.1 chhmc -c imm -s restart (to apply the new network settings)
```

To change the IMM user ID and password (the password must be entered when prompted):

```
chhmc -c imm -s modify -u immusername
```

To disable the use of anonymous ciphers on the HMC chatlet interface:

```
chhmc -c legacyhmccomm -s disable
```

To enable NIST SP 800-131A mode:

```
chhmc -c security -s modify --mode nist_sp800_131a
```

To enable the Event Manager for Call Home:

```
chhmc -c emch -s enable
```

To enable and set a GRUB password:

```
chhmc -c grubpasswd -s enable --passwd myPassword
```

To change the GRUB password (the password must be entered when prompted):

```
chhmc -c grubpasswd -s modify
```

ENVIRONMENT

None

BUGS

None

AUTHOR

IBM Austin

SEE ALSO

lshmc, chpsm, getfile

NAME

chhmcencr - change HMC encryption support

SYNOPSIS

```
chhmcencr -c { passwd | webui | ssh | sshmacc } -o { a | r | s }
-e encryption [--help]
```

DESCRIPTION

chhmcencr changes which encryption is used by the Hardware Management Console (HMC) to encrypt the passwords of locally authenticated HMC users. The new encryption will be used when a locally authenticated HMC user is created, or when the password for a locally authenticated HMC user is changed. The passwords of existing locally authenticated HMC users will not be affected by the encryption change until the passwords for those users are changed.

chhmcencr also changes which encryption ciphers can be used by the HMC Web user interface. The HMC must be restarted for any changes to HMC Web user interface encryption ciphers to take effect.

chhmcencr also changes which encryption ciphers and Message Authentication Code (MAC) algorithms can be used by the HMC Secure Shell (SSH) interface.

OPTIONS

-c The encryption configuration to change. Valid values are **passwd** to change which password encryption is used for locally authenticated HMC users, **webui** to change which encryption ciphers can be used by the HMC Web user interface, **ssh** to change which encryption ciphers can be used by the HMC SSH interface, or **sshmacc** to change which MAC algorithms can be used by the HMC SSH interface.

-o The operation to perform.

Specify **a** to add one or more encryption ciphers to the list of encryption ciphers currently supported by the HMC Web user interface or SSH interface. Also specify **a** to add one or more MAC algorithms to the list of MAC algorithms currently supported by the HMC SSH interface.

Specify **r** to remove one or more encryption ciphers from the list of encryption ciphers currently supported by the HMC Web user interface or SSH interface. Also specify **r** to remove one or more MAC algorithms from the list of MAC algorithms currently supported by the HMC SSH interface.

Specify **s** to set the encryption to use to encrypt the passwords of locally authenticated HMC users for all subsequent user creations and user password modifications.

-e The password encryption, encryption cipher, or MAC algorithm to add, remove, or set. When adding or removing encryption ciphers or MAC algorithms, multiple encryption ciphers or MAC algorithms can be specified and must be comma separated. Encryption ciphers and MAC algorithms are added to the end of their lists. The order of the encryption ciphers and MAC algorithms in their lists is the order the HMC attempts to use them.

A list of all of the HMC password encryptions, encryption ciphers, and MAC algorithms available on the HMC can be obtained from the **lshmcencr** command.

--help Display the help text for this command and exit.

EXAMPLES

Set the password encryption to be used for all subsequent HMC user creations or password modifications to SHA-512:

```
chhmcencr -c passwd -o s -e sha512
```

Remove **TLS_ECDH_ECDSA_WITH_3DES_EDE_CBC_SHA** from the encryption ciphers currently

supported by the HMC Web user interface:

```
chhmcencr -c webui -o r -e TLS_ECDH_ECDSA_WITH_3DES_EDE_CBC_SHA
```

Add **TLS_DHE_RSA_WITH_3DES_EDE_CBC_SHA** and **TLS_DHE_DSS_WITH_3DES_EDE_CBC_SHA** to the encryption ciphers currently supported by the HMC Web user interface:

```
chhmcencr -c webui -o a -e  
TLS_DHE_RSA_WITH_3DES_EDE_CBC_SHA,TLS_DHE_DSS_WITH_3DES_EDE_CBC_SHA
```

Remove **aes128-ctr** and **aes128-gcm@openssh.com** from the encryption ciphers currently supported by the HMC SSH interface:

```
chhmcencr -c ssh -o r -e aes128-ctr,aes128-gcm@openssh.com
```

Add **hmac-sha2-256** to the MAC algorithms currently supported by the HMC SSH interface:

```
chhmcencr -c sshmac -o a -e hmac-sha2-256
```

ENVIRONMENT

None

BUGS

None

AUTHOR

IBM Austin

SEE ALSO

lshmcencr

NAME

chhmcfs - free up space in HMC file systems

SYNOPSIS

```
chhmcfs -o f { -d days | -h hours | -s size }  
[-f file-system] [--help]
```

DESCRIPTION

chhmcfs frees up space in Hardware Management Console (HMC) file systems. Space is freed by removing temporary HMC files that are used for HMC and managed system firmware problem analysis from the HMC hard disk.

This command can only free up space in the following file systems: **/var**, **/dump**, **/extra**, and **/**. The temporary files that can be removed from the **/var** file system include HMC trace and log files. The temporary files that can be removed from the **/dump** file system include managed system dumps, managed frame dumps, and debug data collected using the HMC **pedbg** command. The temporary files that can be removed from the **/extra** file system include managed system dumps and managed frame dumps. The temporary files that can be removed from the **/** file system include HMC Java core dump and heap dump files, and HMC trace files for Power system firmware updates.

This command will not remove temporary HMC trace and log files that are in use.

OPTIONS

-o The operation to perform. The only valid value is **f** to free up file system disk space by removing temporary HMC files from the hard disk.

-d Remove temporary HMC files which have not been modified during the specified number of *days* prior to now.

If **0** is specified with this option, all temporary HMC files will be removed.

Specifying **-d 1** is equivalent to specifying **-h 24**.

The **-d**, **-h**, and **-s** options are mutually exclusive.

-h Remove temporary HMC files which have not been modified during the specified number of *hours* prior to now.

If **0** is specified with this option, all temporary HMC files will be removed.

The **-d**, **-h**, and **-s** options are mutually exclusive.

-s Remove temporary HMC files, starting with the oldest file, to free up to *size* megabytes in each file system.

The **-d**, **-h**, and **-s** options are mutually exclusive.

-f The file system from which to remove files. Valid values are **/var**, **/dump**, **/extra**, and **/**.

If this option is omitted, files will be removed from all file systems that have temporary HMC files that can be removed.

--help Display the help text for this command and exit.

EXAMPLES

Remove temporary HMC files which have not been modified during the last day (24 hours) from all file systems:

```
chhmcfs -o f -d 1
```

Remove all temporary HMC files from all file systems:

```
chhmcfcs -o f -d 0
```

Remove temporary HMC files which have not been modified during the last 36 hours from the **/var** file system:

```
chhmcfcs -o f -h 36 -f /var
```

Remove temporary HMC files from the **/dump** file system to free up to 100 MB:

```
chhmcfcs -o f -s 100 -f /dump
```

ENVIRONMENT

None

BUGS

None

AUTHOR

IBM Austin

SEE ALSO

lshmcfs

NAME

chhmcldap - change HMC LDAP configuration

SYNOPSIS

To set HMC LDAP configuration attributes:

```
chhmcldap -o s [--primary LDAP-server-URI]  
  [--backup LDAP-server-URI] [--basedn base-DN]  
  [--binddn bind-DN] [--bindpw bind-password]  
  [--timelimit time-limit]  
  [--bindtimelimit bind-time-limit]  
  [--automanage {0 | 1}]  
  [--auth {ldap | kerberos}]  
  [--loginattribute attribute]  
  [--hmcuserpropsattribute attribute]  
  [--hmcauthnameattribute attribute]  
  [--searchfilter LDAP-search-filter]  
  [--scope {one | sub}]  
  [--referrals {0 | 1}] [--starttls {0 | 1}]  
  [--hmcgroups group-names]
```

To remove an LDAP resource or the LDAP configuration from the HMC:

```
chhmcldap -o r  
  -r {backup | ldap | binddn | bindpw |  
  searchfilter | hmcgroups}
```

DESCRIPTION

chhmcldap changes the Hardware Management Console (HMC) Lightweight Directory Access Protocol (LDAP) client configuration.

By default, LDAP is not configured on the HMC. The HMC can only be configured as an LDAP client. To configure the HMC as an LDAP client, issue this command and specify both the primary LDAP server and the base Distinguished Name (DN) at a minimum.

The **getfile** command can be used to deploy an LDAP Certificate Authority (CA) certificate file on the HMC.

OPTIONS

- o** The operation to perform. Valid values are **s** to set an HMC LDAP configuration attribute, and **r** to remove an LDAP resource or the LDAP configuration from the HMC.
- r** The LDAP resource to remove. Valid values are **backup** to remove the backup LDAP server, **ldap** to unconfigure LDAP on the HMC, **binddn** to remove the bind Distinguished Name (DN), **bindpw** to remove the bind password, **searchfilter** to remove the LDAP search filter, and **hmcgroups** to remove the HMC user group configuration.

This option is required for a remove operation. This option is not valid for a set operation.

--primary

The primary LDAP server. The Uniform Resource Identifier (URI) scheme may be **ldap://** for LDAP over TCP, or **ldaps://** for LDAP over SSL. The URI format is **ldap://LDAP-server[:port-number]** or **ldaps://LDAP-server[:port-number]**. *LDAP-server* can either be the host name or the IP address of the LDAP server. The port number is optional. If the port number is omitted, port 389 is used for the **ldap://** scheme and port 636 is used for the **ldaps://** scheme. If Start TLS is enabled, the URI must use the **ldap://** scheme.

If LDAP is not configured on the HMC, both this option and the **--basedn** option must be specified to configure LDAP on the HMC.

This option is only valid for a set operation.

--backup

The backup LDAP server. The URI scheme may be **ldap://** for LDAP over TCP, or **ldaps://** for LDAP over SSL. The URI format is **ldap://LDAP-server[:port-number]** or **ldaps://LDAP-server[:port-number]**. *LDAP-server* can either be the host name or the IP address of the LDAP server. The port number is optional. If the port number is omitted, port 389 is used for the **ldap://** scheme and port 636 is used for the **ldaps://** scheme. If Start TLS is enabled, the URI must use the **ldap://** scheme.

This option is only valid for a set operation. Setting a backup LDAP server is optional.

--basedn

The default base DN to use when performing LDAP operations. The base DN must be specified in LDAP format (e.g. **dc=ldapclient,dc=example,dc=com**).

If LDAP is not configured on the HMC, both this option and the **--primary** option must be specified to configure LDAP on the HMC.

This option is only valid for a set operation.

--binddn

The DN to use for binding to the LDAP server when the LDAP server is configured to use non-anonymous binding. The bind DN must be specified in LDAP format (e.g. **cn=admin,dc=yourorg,dc=com**).

This option is only valid for a set operation.

--bindpw

The password to use when binding to the LDAP server when the LDAP server is configured to use non-anonymous binding.

This option is only valid for a set operation.

--timelimit

The LDAP search time limit in seconds. A value of 0 means there is no time limit. If this option is not specified when LDAP is configured, the search time limit is set to 30 seconds.

This option is only valid for a set operation.

--bindtimelimit

The LDAP server bind time limit in seconds. If this option is not specified when LDAP is configured, the bind time limit is set to 30 seconds.

This option is only valid for a set operation.

--automanage

Use this option to indicate whether the HMC should automatically manage remotely authenticated LDAP users. Valid values are **0** to disable automatic management, or **1** to enable automatic management.

When automatic management is enabled, an LDAP user can log in to the HMC. An HMC user will be automatically created for the LDAP user if the HMC user does not exist when the LDAP user logs in. If the HMC user already exists, it will be updated with the current user definition retrieved from the LDAP server when the LDAP user logs in.

If this option is not specified when LDAP is configured, automatic management is disabled.

This option is only valid for a set operation.

--auth The type of authentication to use for automatically managed LDAP users. Valid values are **ldap** to use LDAP authentication, or **kerberos** to use Kerberos authentication.

If **kerberos** is specified and Kerberos is not enabled on the HMC, LDAP users will not be able to log in to the HMC because the HMC will not be able to retrieve the user data from the LDAP server.

If this option is not specified when LDAP is configured, LDAP authentication will be used.

This option is only valid for a set operation.

--loginattribute

The login attribute to use to authenticate LDAP users on the HMC. The value of the login attribute stored on the LDAP server is compared with the HMC user name. If a match is found on the LDAP server, then authentication with the LDAP server is attempted.

If this option is not specified when LDAP is configured, the login attribute is set to **uid**.

This option is only valid for a set operation.

--hmcuserpropsattribute

The attribute to use to retrieve the user roles and properties from the LDAP server. These user roles and properties are used when the HMC user is created or updated for an automatically managed LDAP user.

If this option is not specified when LDAP is configured, this attribute is set to **ibm-aixAdminPolicyEntry**.

This option is only valid for a set operation.

--hmcauthnameattribute

The attribute to use to retrieve the remote user ID used in Kerberos authentication. When the HMC user is created or updated for an automatically managed LDAP user, the HMC user is assigned this remote user ID if Kerberos authentication is used.

If this option is not specified when LDAP is configured, this attribute is set to **userPrincipalName**.

This option is only valid for a set operation.

--searchfilter

The filter to use for limiting the search of the LDAP server for user information. This is useful when you have a large number of LDAP users and you want to decrease the amount of time for retrieving the user information.

If this option is not specified when LDAP is configured, then no filter is set.

This option is only valid for a set operation.

--scope The search scope starting from base DN. Valid values are **one** to search one level or **sub** to search the subtree. This is useful when you want to search for users in different organizations.

If this option is not specified when LDAP is configured, the search scope is set to one level.

This option is only valid for a set operation.

--referrals

Specifies whether automatic referral chasing is to be enabled or disabled. Valid values are **0** to disable automatic referral chasing, and **1** to enable it. If this option is not specified when LDAP is configured, automatic referral chasing is enabled.

This option is only valid for a set operation.

--starttls

Specifies whether Start Transport Layer Security (TLS) is to be enabled or disabled. Valid values are **0** to disable Start TLS, and **1** to enable it. This option can only be enabled if the LDAP server supports the Start TLS protocol. When Start TLS is enabled, a connection is opened using the normal LDAP port. Then, when needed (for example when binding), the connection is converted to a SSL connection using the Start TLS protocol. This functionality exists, for example, in OpenLDAP. When Start TLS is enabled, the URI scheme **ldap://** must be used.

If this option is not specified when LDAP is configured, Start TLS is disabled.

This option is only valid for a set operation.

--hmcgroups

Specifies the name of one or more user groups allowed to log in to this HMC. Multiple group names must be comma separated. The user groups specified with this option will replace the current user groups configured for this HMC.

If this option is not specified when LDAP is configured, no HMC user groups will be configured and all LDAP users will be allowed to log in to this HMC.

This option is only valid for a set operation.

--help Display the help text for this command and exit.

EXAMPLES

Configure LDAP on the HMC for the first time and enable Start TLS:

```
chhmclldap -o s --primary ldap://ldap.example.com --basedn dc=example,dc=com --starttls 1
```

Set a backup LDAP server using the **ldap://** URI scheme:

```
chhmclldap -o s --backup ldap://ldap2.example.com
```

Set the primary LDAP server using the **ldaps://** URI scheme and port number **625**:

```
chhmclldap -o s --primary ldaps://ldap.example.com:625
```

Set the bind DN and bind password:

```
chhmclldap -o s --binddn cn=admin,dc=yourorg,dc=com --bindpw secret
```

Enable automatic management of LDAP users and set the authentication type to Kerberos:

```
chhmclldap -o s --automanage 1 --auth kerberos
```

Set the LDAP search filter:

```
chhmclldap -o s --searchfilter "mail=*@yourcompany.com"
```

Set the LDAP search scope to subtree:

```
chhmcdap -o s --scope sub
```

Set the login attribute to **notesshortname**:

```
chhmcdap -o s --loginattribute notesshortname
```

Set both the time limit and bind time limit to **40** seconds and disable referrals:

```
chhmcdap -o s --timelimit 40 --bindtimelimit 40 --referrals 0
```

Configure this HMC to only allow users in groups **group1** and **group2** to log in:

```
chhmcdap -o s --hmcgroups group1,group2
```

Remove the backup LDAP server:

```
chhmcdap -o r -r backup
```

Unconfigure LDAP on the HMC:

```
chhmcdap -o r -r ldap
```

ENVIRONMENT

None

BUGS

None

AUTHOR

IBM Austin

SEE ALSO

lshmcdap, **getfile**, **rmfile**

NAME

chhmcusr - change a Hardware Management Console user

SYNOPSIS

```
chhmcusr -u user-name
-t {assign | desc | name | passwd | pwage | taskrole |
auth | remoteuser}
[-o {a | r}] [-r {resource | resourcerole}]
[-v new-attribute-value]
[--remoteuser remote-user-name] [--localpasswd password]
[--help]
```

or

```
chhmcusr { -f input-data-file | -i "input-data" } [--help]
```

DESCRIPTION

chhmcusr changes attributes of a Hardware Management Console (HMC) user.

OPTIONS

-u The user name of the HMC user to change.

You can either use this option, or use the **name** attribute with the **-f** or **-i** option, to specify the user name. The **-u**, **-f**, and **-i** options are mutually exclusive.

-t The user attribute to change. Valid values are **assign** for the user's access control managed resource object assignment or managed resource role assignment, **desc** for the user's description, **name** for the user's user name, **passwd** for the user's password, **pwage** for the number of days until the user's password expires, **taskrole** for the user's access control task role, **auth** for the user's authentication type, and **remoteuser** for the user's remote user ID used for remote Kerberos authentication.

Only users that have the **hmcsuperadmin** task role, or that have the **ManageAllUserPasswords** task in their task role, are authorized to change other locally authenticated user's passwords. The password for a remotely authenticated Kerberos user can be changed only by that user. Passwords for remotely authenticated LDAP users cannot be changed.

Password expirations can be changed for locally authenticated users only.

You can either use this option, or use the **-f** or **-i** option, to specify the user attribute(s) to change. The **-t**, **-f**, and **-i** options are mutually exclusive.

-o The managed resource object or role assignment operation to perform. Valid values are **a** to add a managed resource object or role to the user and **r** to remove a managed resource object or role from the user.

This option is required when changing the user's managed resource object assignment.

You can either use this option, or use the **resourcerole** attribute with the **-f** or **-i** option, to change the user's managed resource role assignment. The **-o**, **-f**, and **-i** options are mutually exclusive.

-r The type of access control assignment to change. Valid values are **resource** for managed resource object assignment and **resourcerole** for managed resource role assignment.

This option is required when the **-o** option is used to change the user's managed resource object assignment or managed resource role assignment. This option is not valid otherwise.

-v The new value for the attribute being changed.

When changing the user's managed resource object assignment, specify the managed resource object to be added or removed.

When changing the user's managed resource role assignment, specify the managed resource role to be added or removed.

When changing the user's description, specify the new description with this option. The new description can be any string.

When changing the user's user name, specify the new user name with this option. The new user name must not be longer than 32 characters, and it must begin with a letter.

When changing the user's password, you can either specify the new password with this option, or you can omit this option and you will be prompted to enter the password. The new password must be at least 7 characters in length.

When changing the number of days until the user's password expires, specify the new number of days with this option.

When changing the user's access control task role, specify the new task role with this option. Valid values are **hmcsuperadmin**, **hmcoperator**, **hmcviewer**, **hmcpe**, **hmcserVICerep**, or a user-defined task role.

When changing the user's authentication type, specify the new authentication type with this option. Valid values are **local**, **kerberos**, or **ldap**.

When changing the user's remote user ID used for remote Kerberos authentication, specify the new remote user ID with this option.

This option is required when the **-t** option is specified to change any user attribute other than the user's password.

You can either use this option, or use the **-f** or **-i** option, to specify the new user attribute value(s). The **-v**, **-f**, and **-i** options are mutually exclusive.

--remoteuser

The remote user ID used for remote Kerberos authentication for this user. This is the user's Kerberos principal. The format of a typical Kerberos principal is *primary/instance@REALM*.

The remote user ID must be specified when changing the user's authentication type to remote Kerberos authentication.

This option is only valid when the **-t** option is specified to change the user's authentication type to remote Kerberos authentication.

You can either use this option, or use the **remote_user_name** attribute with the **-f** or **-i** option, to change the remote user ID. The **--remoteuser**, **-f**, and **-i** options are mutually exclusive.

--localpasswd

The password for this user. The password must be at least 7 characters in length.

This option is only valid when the **-t** option is specified to change the user's authentication type to local authentication.

You can either use this option, or use the **passwd** attribute with the **-f** or **-i** option, to specify the password for this user when changing the user's authentication type to local authentication. If this option is omitted or the **-f** or **-i** option is specified and the **passwd** attribute is omitted, you will be prompted to enter the password. The **--localpasswd**, **-f**, and **-i** options are mutually exclusive.

- f** The name of the file containing the input data for this command. The input data consists of attribute name/value pairs, which are in comma separated value (CSV) format.

The format of the input data is as follows:

attribute-name=value,attribute-name=value,...

Valid attribute names for this command:

name

[**new_name**]

[**taskrole**]

Valid values are **hmcsuperadmin**, **hmcoperator**, **hmcviewer**, **hmcpe**, **hmcservicerep**, or a user-defined task role

[**resourcerole**]

[**description**]

[**passwd**]

Local and Kerberos users only

[**current_passwd**]

When changing the password for a Kerberos user, use this attribute to specify the user's current password.

If this attribute is omitted, you will be prompted to enter the current password.

[**pwage**]

Local users only

number of days

[**min_pwage**]

Local users only

number of days

[**authentication_type**]

Valid values are:

local - local authentication

kerberos - remote Kerberos authentication

ldap - remote LDAP authentication

[**session_timeout**]

number of minutes

[**verify_timeout**]

number of minutes

[**idle_timeout**]

number of minutes

[**inactivity_expiration**]

number of days

[**remote_webui_access**]

Valid values are:

0 - do not allow this user to log in remotely to the HMC Web user interface

1 - allow this user to log in remotely to the HMC Web user interface

[**remote_ssh_access**]

Valid values are:

- 0** - do not allow this user to log in remotely to the HMC using SSH
 - 1** - allow this user to log in remotely to the HMC using SSH
- [remote_user_name]**
Kerberos users only

Input data for this command can be specified with this option, the **-i** option, or any of the other command options. The **-f** and the **-i** options are mutually exclusive, and they cannot be specified if any of the other command options are specified.

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

Input data for this command can be specified with this option, the **-f** option, or any of the other command options. The **-i** and the **-f** options are mutually exclusive, and they cannot be specified if any of the other command options are specified.

- help** Display the help text for this command and exit.

EXAMPLES

Change the password for the user **tester** (the new password must be entered when prompted):

```
chhmcusr -u tester -t passwd
```

Change the password for the user **tester** without prompting:

```
chhmcusr -u tester -t passwd -v secretpassword
```

or

```
chhmcusr -i "name=tester,passwd=secretpassword"
```

Change the number of days until the password expires for the user **hmcuser1** to be **180**:

```
chhmcusr -u hmcuser1 -t pwage -v 180
```

or

```
chhmcusr -i "name=hmcuser1,pwage=180"
```

Change the task role for the user **tester** to **hmcoperator**:

```
chhmcusr -u tester -t taskrole -v hmcoperator
```

or

```
chhmcusr -i "name=tester,taskrole=hmcoperator"
```

Change the remote user ID for the user **krbuser** to **krbuser@EXAMPLE.ORG**:

```
chhmcusr -u krbuser -t remoteuser -v krbuser@EXAMPLE.ORG
```

or

```
chhmcusr -i "name=krbuser,remote_user_name=krbuser@EXAMPLE.ORG"
```

Change the remotely authenticated Kerberos user **user1** to a locally authenticated user (the password must be entered when prompted):

```
chhmcusr -u user1 -t auth -v local
```

or

```
chhmcusr -i "name=user1,authentication_type=local"
```

Change the locally authenticated user **user2** to a remotely authenticated Kerberos user and set the remote user ID to **user2@EXAMPLE.ORG**:

```
chhmcusr -u user2 -t auth -v kerberos --remoteuser  
user2@EXAMPLE.ORG
```

or

```
chhmcusr -i "name=user2,authentication_type=kerberos,  
remote_user_name=user2@EXAMPLE.ORG"
```

Change the locally authenticated user **user3** to a remotely authenticated LDAP user:

```
chhmcusr -u user3 -t auth -v ldap
```

or

```
chhmcusr -i "name=user3,authentication_type=ldap"
```

Change the remotely authenticated LDAP user **user4** to a locally authenticated user:

```
chhmcusr -u user4 -t auth -v local --localpasswd jk3ds00b
```

or

```
chhmcusr -i "name=user4,authentication_type=local,passwd=jk3ds00b"
```

ENVIRONMENT

None

BUGS

None

AUTHOR

IBM Austin

SEE ALSO

lshmcusr, mkhmcusr, rmhmcusr, lsaccfg

NAME

chhwres - change hardware resources

SYNOPSIS

To add, remove, or move a physical I/O slot:

```
chhwres -r io -m managed-system -o {a | r | m}
  {-p partition-name | --id partition-ID}
  [{-t target-partition-name |
  --tid target-partition-ID}]
  -l slot-DRC-index [-a "attributes"]
  [-w wait-time] [-d detail-level] [--force]
```

To set physical I/O attributes:

```
chhwres -r io -m managed-system -o s
  {-p partition-name | --id partition-ID}
  --rsubtype {iopool | taggedio}
  -a "attributes"
```

To clear a physical I/O slot or bus:

```
chhwres -r io -m managed-system -o c
  [--rsubtype {bus | slot}] [-l slot-DRC-index]
```

To add or remove a virtual I/O adapter:

```
chhwres -r virtualio -m managed-system -o {a | r}
  {-p partition-name | --id partition-ID}
  [--rsubtype {eth | fc | scsi | serial | vnic}]
  [-s virtual-slot-number] [-a "attributes"]
  [-w wait-time] [-d detail-level] [--force] [-v]
```

To change a virtual ethernet adapter or virtual NIC assigned to a partition:

```
chhwres -r virtualio -m managed-system -o s
  --rsubtype {eth | vnic}
  {-p partition-name | --id partition-ID}
  -s virtual-slot-number -a "attributes"
  [-w wait-time] [-d detail-level] [--force]
  [--allownetdisrupt]
```

To enable or disable a virtual ethernet adapter or virtual NIC:

```
chhwres -r virtualio -m managed-system -o {d | e}
  --rsubtype {eth | vnic}
  {-p partition-name | --id partition-ID}
  -s virtual-slot-number
```

To set virtual I/O attributes:

```
chhwres -r virtualio -m managed-system -o s
  [{-p partition-name | --id partition-ID}]
  --rsubtype {eth | hsl | virtualopti}
  -a "attributes"
```

To set virtual NIC backing device attributes:

```
chhwres -r virtualio -m managed-system -o s
  --rsubtype vnicbkdev
  {-p partition-name | --id partition-ID}
  -s virtual-slot-number --logport logical-port-ID
  -a "attributes"
```

To make a virtual NIC backing device active or clear a virtual NIC backing device error:

```
chhwres -r virtualio -m managed-system -o {act | c}
--subtype vnicbkdev
{-p partition-name | --id partition-ID}
-s virtual-slot-number --logport logical-port-ID
```

To add or remove a virtual network or set virtual network attributes:

```
chhwres -r virtualio -m managed-system -o {a | r | s}
--subtype vnetwork --vnetwork virtual-network
[-a "attributes"]
```

To add or remove a virtual switch:

```
chhwres -r virtualio -m managed-system -o {a | r}
--subtype vswitch --vswitch virtual-switch
```

To set virtual switch attributes:

```
chhwres -r virtualio -m managed-system -o s
--subtype vswitch --vswitch virtual-switch
[-a "attributes"] [--force]
```

To synchronize the virtual switch mode:

```
chhwres -r virtualio -m managed-system -o sync
--subtype vswitch --vswitch virtual-switch
```

To add, remove, or move memory:

```
chhwres -r mem -m managed-system -o {a | r | m | s}
{-p partition-name | --id partition-ID}
[{-t target-partition-name |
--tid target-partition-ID}]
[-q quantity] [--entitled quantity]
[-w wait-time] [-d detail-level] [--force]
```

To set memory attributes for the managed system:

```
chhwres -r mem -m managed-system -o s
-a "attributes"
```

To set memory attributes for a partition:

```
chhwres -r mem -m managed-system -o s
{-p partition-name | --id partition-ID}
-a "attributes"
[-w wait-time] [-d detail-level] [--force]
```

To switch over a running shared memory partition to its redundant paging Virtual I/O Server (VIOS) partition:

```
chhwres -r mem -m managed-system -o so
{-p partition-name | --id partition-ID}
```

To reset I/O entitled memory statistics for a partition:

```
chhwres -r mem -m managed-system -o rs
{-p partition-name | --id partition-ID}
[-w wait-time] [-d detail-level]
```

To add or remove a shared memory pool, or set shared memory pool attributes:

```
chhwres -r mempool -m managed-system -o {a | r | s}
```

[-a "attributes"]

To add or remove a device in the shared memory pool (on managed systems that support partition suspend and resume operations, this adds or removes a device in the reserved storage device pool):

```
chhwres -r mempool -m managed-system -o {a | r}
  {-p paging-VIOS-partition-name |
    --id paging-VIOS-partition-ID}
  --subtype pgdev --device device-name
```

To add or remove a reserved storage device pool, or set reserved storage device pool attributes (only supported on managed systems that support partition suspend and resume operations):

```
chhwres -r rspool -m managed-system -o {a | r | s}
  [-a "attributes"]
```

To add or remove a device in the reserved storage device pool (only supported on managed systems that support partition suspend and resume operations):

```
chhwres -r rspool -m managed-system -o {a | r}
  {-p VIOS-partition-name | --id VIOS-partition-ID}
  --subtype rsdev --device device-name [--manual]
```

To add, remove, or move processing resources:

```
chhwres -r proc -m managed-system -o {a | r | m | s}
  {-p partition-name | --id partition-ID}
  [{-t target-partition-name |
    --tid target-partition-ID}]
  [--procs quantity] [--procunits quantity]
  [-w wait-time] [-d detail-level] [--force]
```

To set processing attributes for the managed system:

```
chhwres -r proc -m managed-system -o s -a "attributes"
```

To set processing attributes for a partition:

```
chhwres -r proc -m managed-system -o s
  {-p partition-name | --id partition-ID}
  -a "attributes"
```

To set shared processor pool attributes on a managed system that supports multiple shared processor pools:

```
chhwres -r procpool -m managed-system -o s
  {--poolname pool-name | --poolid pool-ID}
  -a "attributes"
```

To set shared processor pool attributes for a partition on a managed system that supports multiple shared processor pools:

```
chhwres -r procpool -m managed-system -o s
  {-p partition-name | --id partition-ID}
  -a "attributes"
```

To add, remove, or move a Host Ethernet Adapter (HEA) logical port:

```
chhwres -r hea -m managed-system -o {a | r | m}
  {-p partition-name | --id partition-ID}
  [{-t target-partition-name |
    --tid target-partition-ID}]
  -l HEA-adapter-ID
  [--physport physical-port-ID]
```

```
-g port-group --logport logical-port-ID
[-a "attributes"]
[-w wait-time] [-d detail-level] [--force]
```

To set HEA attributes:

```
chhwres -r hea -m managed-system -o s
-l HEA-adapter-ID
[--physport physical-port-ID]
-g port-group -a "attributes"
```

To switch an SR-IOV adapter to shared mode:

```
chhwres -r sriov -m managed-system --subtype adapter
-o a -a "attributes"
```

To switch an SR-IOV adapter to dedicated mode:

```
chhwres -r sriov -m managed-system --subtype adapter
-o r -a "attributes"
```

To move the configuration of a failed SR-IOV adapter to a new adapter:

```
chhwres -r sriov -m managed-system --subtype adapter
-o m -a "attributes"
```

To set SR-IOV physical port attributes:

```
chhwres -r sriov -m managed-system --subtype physport
-o s -a "attributes"
```

To add or remove an SR-IOV logical port, or to set SR-IOV logical port attributes:

```
chhwres -r sriov -m managed-system --subtype logport
-o {a | r | s}
{-p partition-name | --id partition-ID }
-a "attributes"
[-w wait-time] [-d detail-level] [--force]
```

To reset statistics for an SR-IOV logical or physical port:

```
chhwres -r sriov -m managed-system -o rs
--subtype {logport | physport}
-a "attributes"
```

DESCRIPTION

chhwres changes the hardware resource configuration of the *managed-system*. **chhwres** is used to perform dynamic logical partitioning (DLPAR) operations.

When **chhwres** is used to add a virtual I/O adapter to a Virtual I/O Server (VIOS) partition, **chhwres** issues the VIOS **cfgdev** command to configure the device in the VIOS. When **chhwres** is used to remove a virtual I/O adapter from a VIOS partition, **chhwres** issues the VIOS **rmdev** command to remove the device from the VIOS partition.

When **chhwres** is used to remove memory or processing resources from a shutdown partition, if the amount of memory or processing resources falls below the minimum for the partition, the minimum, assigned, and maximum values for the partition are all set to 0. Also, if the memory values for a partition are set to 0, the processing resource values for the partition will also be set to 0. Likewise, if the processing resource values for a partition are set to 0, the memory values for the partition will also be set to 0.

OPTIONS

-r The type of hardware resources to change. Valid values are **io** for physical I/O, **virtualio** for virtual I/O, **mem** for memory, **mempool** for shared memory pool, **rspool** for reserved storage device

pool, **proc** for processing resources, **procpool** for shared processor pool, **hea** for Host Ethernet Adapter (HEA) resources, and **sriov** for Single Root I/O Virtualization (SR-IOV) resources.

Reserved storage device pool operations are only supported on managed systems that support partition suspend and resume operations.

--subtype

The subtype of hardware resources to change. Valid physical I/O resource subtypes are **slot** for I/O slots, **bus** for I/O buses, **iopool** for I/O pools, and **taggedio** for tagged I/O resources. Valid virtual I/O resource subtypes are **eth** for virtual ethernet, **fc** for virtual fibre channel, **scsi** for virtual SCSI, **serial** for virtual serial, **hsl** for High Speed Link (HSL) OptiConnect, **virtualopti** for virtual OptiConnect, **vnetwork** for virtual network, **vnic** for virtual NIC, **vnickdev** for virtual NIC backing device, and **vswitch** for virtual switch resources. Valid SR-IOV resource subtypes are **adapter** for SR-IOV adapter resources, **logport** for SR-IOV logical port resources, and **physport** for SR-IOV physical port resources.

The only valid shared memory pool resource subtype is **pgdev** for paging space devices. The only valid reserved storage device pool resource subtype is **rsdev** for reserved storage devices. On a managed system that supports Active Memory Sharing and partition suspend and resume operations, the command **chhwres -m managed-system -r mempool --subtype pgdev -o {a | r} -p VIOS-partition-name --device device-name** adds or removes a device in the reserved storage device pool, and is equivalent to the command **chhwres -m managed-system -r rspool --subtype rsdev -o {a | r} -p VIOS-partition-name --device device-name**.

This option is required for physical I/O and virtual I/O set operations, for virtual I/O add operations, for virtual network and virtual switch operations, and for SR-IOV operations. This option is not valid for memory, processor, shared processor pool, or HEA operations.

-m The name of the managed system for which the hardware resource configuration is to be changed. The name may either be the user-defined name for the managed system, or be in the form *ttt-**mmm*****sssss***, where *ttt* is the machine type, *mmm* is the model, and *sssss* is the serial number of the managed system. The *ttt-**mmm*****sssss*** form must be used if there are multiple managed systems with the same user-defined name.

-o The operation to perform. Valid values are **a** to add hardware resources, **r** to remove hardware resources, **m** to move hardware resources from one partition to another, **s** to set hardware resource related attributes, **d** to disable a virtual ethernet adapter or virtual NIC, **e** to enable a virtual ethernet adapter or virtual NIC, **rs** to reset I/O entitled memory statistics for a partition or to reset statistics for an SR-IOV logical or physical port, **so** to switch over a running shared memory partition to its redundant paging VIOS partition, **sync** to synchronize the virtual switch mode, **c** to clear a physical I/O slot or bus, or clear a virtual NIC backing device error, or **act** to make a virtual NIC backing device active.

The **c** operation is useful for clearing stale physical I/O slot or bus information after a physical I/O adapter or drawer has been concurrently removed but not replaced.

-p The name of the partition for which to perform the operation. For a move operation, this is the source partition (the partition the resources will be moved from) for the operation.

When adding or removing a device in the shared memory pool or the reserved storage device pool, use this option to specify the name of the VIOS partition which has the device. If the device is accessed by more than one VIOS partition, then only one of the VIOS partitions must be specified, and the device name specified with the **--device** option must be the name of the device on that VIOS partition.

You can either use this option to specify the name of the partition for which to perform the operation, or use the **--id** option to specify the partition's ID. The **-p** and the **--id** options are mutually

exclusive.

- id** The ID of the partition for which to perform the operation. For a move operation, this is the source partition (the partition the resources will be moved from) for the operation.

When adding or removing a device in the shared memory pool or the reserved storage device pool, use this option to specify the ID of the VIOS partition which has the device. If the device is accessed by more than one VIOS partition, then only one of the VIOS partitions must be specified, and the device name specified with the **--device** option must be the name of the device on that VIOS partition.

You can either use this option to specify the ID of the partition for which to perform the operation, or use the **-p** option to specify the partition's name. The **--id** and the **-p** options are mutually exclusive.

- t** The name of the target partition for a move operation.

You can either use this option to specify the name of the target partition, or use the **--tid** option to specify the ID of the partition. The **-t** and the **--tid** options are mutually exclusive.

A target partition is required to be specified with this option or the **--tid** option for a move operation. This option is not valid for any other operation.

- tid** The ID of the target partition for a move operation.

You can either use this option to specify the ID of the target partition, or use the **-t** option to specify the name of the target partition. The **--tid** and the **-t** options are mutually exclusive.

A target partition is required to be specified with this option or the **-t** option for a move operation. This option is not valid for any other operation.

--poolname

The name of the shared processor pool for which to set attributes.

You can either use this option to specify the name of the shared processor pool, or use the **--poolid** option to specify the ID of the shared processor pool. The **--poolname** and the **--poolid** options are mutually exclusive.

A shared processor pool is required to be specified with this option or the **--poolid** option when setting the attributes for a shared processor pool. This option is not valid for any other operation.

--poolid

The ID of the shared processor pool for which to set attributes.

You can either use this option to specify the ID of the shared processor pool, or use the **--poolname** option to specify the name of the shared processor pool. The **--poolid** and the **--poolname** options are mutually exclusive.

A shared processor pool is required to be specified with this option or the **--poolname** option when setting the attributes for a shared processor pool. This option is not valid for any other operation.

- l** When adding, removing, moving, or clearing a physical I/O slot, use this option to specify the DRC index of the slot.

When performing an HEA operation, use this option to specify the adapter ID of the HEA for which the operation is to be performed.

This option is not valid for any other operation.

- s The virtual slot number of the virtual I/O adapter to add, change, disable, enable, or remove. When performing an operation on a virtual NIC backing device, this is the virtual slot number of the virtual NIC.

When adding a virtual I/O adapter, if this option is not specified then the next available virtual slot number will be assigned to the virtual I/O adapter.

When changing, disabling, enabling, or removing a virtual I/O adapter, or when performing an operation on a virtual NIC backing device, this option is required.

- q The quantity of memory to add, remove, move, or set. The *quantity* specified must be in megabytes and it must be a multiple of the memory region size for the *managed-system*.

For a partition using shared memory, this is the quantity of logical memory to add, remove, or set. Memory move operations are not supported for partitions using shared memory.

When adding memory to a partition or removing memory from a partition, you can either specify an add or remove operation and specify the amount of memory you want to add or remove, or you can specify a set operation and specify the total amount of memory you want the partition to have when the operation completes.

When the current, pending, and runtime memory values for a partition are not the same and an add or remove operation is being performed, the *quantity* of memory specified with this option is added to or removed from the current value.

--entitled

The quantity of I/O entitled memory to add, remove, or set. The *quantity* specified must be in megabytes, or *quantity* can be **auto**. If *quantity* is a number, then automatic I/O entitled memory management will no longer be provided for the partition after the *quantity* of I/O entitled memory is added, removed, or set. If *quantity* is **auto**, then the partition will be set to have automatic I/O entitled memory management, and if necessary, I/O entitled memory will be added or removed at this time so that the partition has the amount of memory required with automatic I/O entitled memory management.

When adding I/O entitled memory to a partition or removing I/O entitled memory from a partition, you can either specify an add or remove operation and specify the amount of I/O entitled memory you want to add or remove, or you can specify a set operation and specify the total amount of I/O entitled memory you want the partition to have when the operation completes.

This option is only valid for partitions using shared memory.

- procs** When adding or removing processing resources to or from a partition using dedicated processors, or when moving processing resources from a partition using dedicated processors to another partition using dedicated processors, use this option to specify the quantity of dedicated processors to add, remove, move, or set.

When adding or removing processing resources to or from a partition using shared processors, or when moving processing resources from a partition using shared processors to another partition using shared processors, use this option to specify the quantity of virtual processors to add, remove, move, or set.

When moving processing resources from a partition using dedicated processors to a partition using shared processors, use this option to specify the quantity of dedicated processors to be moved from the source partition and added as shared processors to the target partition.

This option is not valid when moving processing resources from a partition using shared processors to a partition using dedicated processors. The **--procunits** option must be used instead.

The *quantity* of processing resources specified with this option must be a whole number.

When adding processing resources to a partition or removing processing resources from a partition, you can either specify an add or remove operation and specify the number of processing resources you want to add or remove, or you can specify a set operation and specify the total number of processing resources you want the partition to have when the operation completes.

When the current, pending, and runtime processing resource values for a partition are not the same and an add or remove operation is being performed, the *quantity* of processing resources specified with this option is added to or removed from the current value.

--procunits

When adding or removing processing resources to or from a partition using shared processors, or when moving processing resources from a partition using shared processors to another partition using shared processors, use this option to specify the quantity of processing units to add, remove, move, or set.

When moving processing resources from a partition using shared processors to a partition using dedicated processors, use this option to specify the quantity of shared processors to be moved from the source partition and added as dedicated processors to the target partition.

This option is not valid when moving processing resources from a partition using dedicated processors to a partition using shared processors. The **--procs** option must be used instead.

When moving processing resources from a partition using shared processors to a partition using dedicated processors, the *quantity* of processing units specified with this option must be a whole number. Otherwise, the *quantity* of processing units specified with this option can have up to 2 decimal places.

When adding processing resources to a partition or removing processing resources from a partition, you can either specify an add or remove operation and specify the number of processing resources you want to add or remove, or you can specify a set operation and specify the total number of processing resources you want the partition to have when the operation completes.

When the current, pending, and runtime processing resource values for a partition are not the same and an add or remove operation is being performed, the *quantity* of processing resources specified with this option is added to or removed from the current value.

--physport

The ID of the HEA physical port. This option is required when adding an HEA logical port to a partition. This option is also required when setting HEA physical port attributes. This option is not valid for any other operation.

-g The HEA port group. This option is required for all HEA operations, and is not valid for any other operation.

--logport

The ID of the HEA logical port to add, remove, or move, or the ID of the SR-IOV logical port of the virtual NIC backing device.

This option is required for an HEA add, remove, or move operation and for a virtual NIC backing device operation. This option is not valid for any other operation.

--vnetwork

The virtual network name.

This option is required for all virtual network operations. This option is not valid for any other operation.

--vswitch

The virtual switch name.

This option is required for all virtual switch operations. This option is not valid for any other operation.

--device

The name of the device to add or remove.

This option is required when adding or removing a device in the shared memory pool or the reserved storage device pool. This option is not valid for any other operation.

--manual

When adding a device to the reserved storage device pool, specify this option to indicate that the device will only be available for manual assignment to partitions by the user. A reserved storage device that is added to the pool as a manual device cannot be automatically assigned to a partition by the Hardware Management Console (HMC).

If this option is not specified when adding a device to the reserved storage device pool, the device will only be available for automatic assignment to partitions by the HMC.

This option is only valid when adding a device to the reserved storage device pool.

-w

The maximum time, in minutes, to wait for partition operating system commands issued by the HMC to complete. If a partition operating system command does not complete within the time specified, the operation will be stopped.

wait-time must be a whole number. If *wait-time* is 0, partition operating system commands will not be timed out.

If this option is not specified, a default value of 5 minutes is used.

This option is valid for add, remove, move, and memory and processor set operations for AIX, Linux, and VIOS partitions. In addition, this option is valid for I/O entitled memory statistics reset operations for AIX and Linux partitions. This option is also valid for memory add, remove, move, and set operations for IBM i partitions, and when changing the additional VLAN IDs for a virtual ethernet adapter assigned to a VIOS partition.

-d

The level of detail to request from partition operating system commands issued by the HMC. Valid values are **1** (lowest) through **4** (highest).

If this option is not specified, a default value of 1 is used.

This option is valid for add, remove, move, and memory and processor set operations for AIX, Linux, and VIOS partitions. In addition, this option is valid for I/O entitled memory statistics reset operations for AIX and Linux partitions. This option is also valid when changing the additional VLAN IDs for a virtual ethernet adapter assigned to a VIOS partition.

--force

This option allows you to force a remove or move operation to be performed for a physical I/O slot that is currently in use (varied on) by an IBM i partition.

This option also allows you to force an add, remove, move, or memory or processor set operation

to be performed for an AIX, Linux, or VIOS partition that does not have an RMC connection to the HMC. If this command completes successfully, you will need to restart your operating system for the change to take affect. You should only use this option if you intentionally configured your LAN to isolate the HMC from the operating system of your partition.

This option also allows you to force the additional VLAN IDs for a virtual ethernet adapter to be changed for a VIOS partition that does not have an RMC connection to the HMC.

This option also allows you to force a virtual switch to be set to VEPA mode when the physical switch does not support the IEEE 802.1Qbg standard.

--allownetdisrupt

This option must be specified when adding one or more backing devices to a virtual NIC that currently has only one backing device and its MAC address is the same as the virtual NIC's MAC address. The HMC will remove the existing backing device and recreate it with a different MAC address after adding the specified backing devices. The recreate will cause a short network disruption for the virtual NIC.

- a** The configuration data for the command. The configuration data consists of attribute name/value pairs, which are in comma separated value (CSV) format. The configuration data must be enclosed in double quotes.

The format of the configuration data is as follows:

attribute-name=value,attribute-name=value,...

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

"attribute-name=value,value,...",...

When a list of values is specified, the attribute name/value pair must be enclosed in double quotes. Depending on the shell being used, nested double quote characters may need to be preceded by an escape character, which is usually a `'\'` character.

If `'+='` is used in the attribute name/value pair instead of `'='`, then the specified *value* is added to the existing value for the attribute if the attribute is numerical. If the attribute is a list, then the specified *value(s)* is added to the existing list.

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

Valid attribute names for attributes that can be set when adding, removing, or moving a physical I/O slot:

slot_io_pool_id

Valid attribute names for setting I/O pool attributes:

lpar_io_pool_ids

comma separated

Valid attribute names for setting tagged I/O resources (IBM i partitions only):

load_source_slot

DRC index of I/O slot, DRC index of SR-IOV logical port, or virtual slot number

alt_restart_device_slot

DRC index of I/O slot, DRC index of HEA or SR-IOV

logical port, or virtual slot number

console_slot

DRC index of I/O slot, DRC index of HEA or SR-IOV

logical port, virtual slot number, or the value **hmc**

alt_console_slot

DRC index of I/O slot

op_console_slot

DRC index of I/O slot

Valid attribute names for adding a virtual ethernet adapter:

ieee_virtual_eth

Valid values:

0 - not IEEE 802.1Q compatible

1 - IEEE 802.1Q compatible

Required

port_vlan_id

Required

addl_vlan_ids

is_trunk

Valid values:

0 - no

1 - yes

trunk_priority

Valid values are integers between **1**

and **15**, inclusive

Required for a trunk adapter

vswitch

mac_addr

12 hexadecimal characters

If not specified, a unique MAC address will be automatically generated for the adapter. It is highly recommended that you use an automatically generated MAC address.

allowed_os_mac_addrs

1 to 4 comma separated MAC addresses, each specified as 12 hexadecimal characters. Other valid values:

all - all OS defined MAC addresses are allowed

none - no OS defined MAC addresses are allowed

qos_priority

Valid values are **none** and integers between **0** and **7**, inclusive

vsi_manager_id

vsi_type_id

vsi_type_version

Valid attribute names for adding a virtual fibre channel adapter:

adapter_type

Valid values are **client** or **server**

Required

remote_lpar_id | **remote_lpar_name**

One of these attributes is required

remote_slot_num

Required

wwpns

Optional for a **client** adapter, and not valid for a **server** adapter. If this attribute is not specified, WWPNs will be automatically generated for a **client** adapter. It is highly recommended that you do not specify WWPNs so that they will be automatically generated. If you do specify WWPNs, you must specify exactly two, and they must be comma separated.

Valid attribute names for adding a virtual SCSI adapter:

adapter_type

Valid values are **client** or **server**

Required

remote_lpar_id | **remote_lpar_name**

One of these attributes is required for a **client** adapter

remote_slot_num

Required for a **client** adapter

Valid attribute names for adding a virtual serial adapter:

adapter_type

Valid values are **client** or **server**

Required

remote_lpar_id | **remote_lpar_name**

One of these attributes is required for a **client** adapter

remote_slot_num

Required for a **client** adapter

supports_hmc

The only valid value is **0** for no.

Valid attribute names for adding a virtual NIC:

backing_devices

Required

Comma separated list of virtual NIC backing devices, with each backing device having the following format:

```
sriov/vios-lpar-name/vios-lpar-ID/  
sriov-adapter-ID/sriov-physical-port-ID/  
[capacity][/failover-priority]
```

The first 5 '/' characters must be present. The 6th '/' character is optional, but it must be present if *failover-priority* is specified.

Optional values may be omitted. Optional values are *capacity*, either *vios-lpar-name* or *vios-lpar-ID*, and *failover-priority*.

Specify *capacity* as a percentage with up to 2 decimal places.

Specify *failover-priority* as a value between **1** and **100** inclusive, with **1** being the highest

priority. It defaults to **50** if not specified.

For example:

sriov/vios1//1/3/5.5 specifies a virtual NIC backing device on VIOS **vios1** and physical port **3** of SR-IOV adapter **1** with a capacity of **5.5%**.

port_vlan_id

Port VLAN ID or **0** to disable VLAN tag insertions for untagged frames (default value)

pvid_priority

Valid values are:

An integer between **0** and **7**, inclusive.

Default value is **0**.

allowed_vlan_ids

Comma separated list of VLAN IDs, or one of the following values:

all - all VLAN IDs are allowed (default value)

none - no VLAN IDs are allowed

mac_addr

12 hexadecimal characters

If not specified, a unique MAC address will be automatically generated. It is highly recommended that you use an automatically generated MAC address.

allowed_os_mac_addrs

Comma separated list of MAC addresses, each specified as 12 hexadecimal characters. Other valid values are:

all - all OS defined MAC addresses are allowed (default value)

none - no OS defined MAC addresses are allowed

auto_priority_failover

Valid values:

0 - automatic priority failover is disabled

1 - automatic priority failover is enabled

(default value)

Valid attribute names for changing a virtual ethernet adapter assigned to a running partition:

ieee_virtual_eth

Valid values:

0 - not IEEE 802.1Q compatible

1 - IEEE 802.1Q compatible

addl_vlan_ids

qos_priority

Valid values are **none** and integers between **0** and **7**, inclusive

vsi_manager_id

To remove a VSI profile, specify **none**

vsi_type_id

To remove a VSI profile, specify **none**

vsi_type_version

To remove a VSI profile, specify **none**

Valid attribute names for changing a virtual NIC assigned to a partition:

backing_devices

Comma separated list of virtual NIC backing devices,

with each backing device having the following format:

```
sriov/vios-lpar-name/vios-lpar-ID/  
sriov-adapter-ID/sriov-physical-port-ID/  
[capacity][/failover-priority]
```

The first 5 '/' characters must be present. The 6th '/' character is optional, but it must be present if *failover-priority* is specified.

Optional values may be omitted. Optional values are *capacity*, either *vios-lpar-name* or *vios-lpar-ID*, and *failover-priority*.

Specify *capacity* as a percentage with up to 2 decimal places.

Specify *failover-priority* as a value between **1** and **100** inclusive, with **1** being the highest priority. It defaults to **50** if not specified.

For example:

sriov/vios1//1/0/6.25/10 specifies a virtual NIC backing device on VIOS **vios1** and physical port **0** of SR-IOV adapter **1** with a capacity of **6.25%** and a failover priority of **10**.

port_vlan_id

Port VLAN ID or **0** to disable VLAN tag insertions for untagged frames

pvid_priority

Valid values are:

An integer between **0** and **7**, inclusive.

allowed_vlan_ids

Comma separated list of VLAN IDs

allowed_os_mac_addr

Comma separated list of MAC addresses, each specified as 12 hexadecimal characters

auto_priority_failover

Valid values:

0 - automatic priority failover is disabled

1 - automatic priority failover is enabled

Valid attribute names for changing the backing device for a virtual NIC assigned to a partition:

failover_priority

Valid values are:

An integer between **1** and **100** inclusive, with **1** being the highest priority.

Valid attribute names for setting virtual ethernet attributes for the *managed-system*:

mac_prefix

Valid attribute names for setting HSL OptiConnect attributes (IBM i partitions only):

hsl_pool_id

Valid values are:

- 0** - HSL OptiConnect is disabled
- 1** - HSL OptiConnect is enabled

Valid attribute names for setting virtual OptiConnect attributes (IBM i partitions only):

virtual_opti_pool_id

Valid values are:

- 0** - virtual OptiConnect is disabled
- 1** - virtual OptiConnect is enabled

Valid attribute names for adding a virtual network:

vlan_id

Required

vswitch

Required

is_tagged

Valid values are:

- 0** - do not use IEEE 802.1Q VLAN tagging
- 1** - use IEEE 802.1Q VLAN tagging

Required

Valid attribute names for setting virtual network attributes:

new_name

Valid attribute names for setting virtual switch attributes:

new_name

switch_mode

Valid values are:

- VEB** - Virtual Ethernet Bridge mode
- VEPA** - Virtual Ethernet Port Aggregator mode

Valid attribute names for setting memory attributes for the *managed-system*:

requested_num_sys_huge_pages

mem_mirroring_mode

Valid values are:

- none** - disable memory mirroring
- sys_firmware_only** - enable memory mirroring for system firmware memory only

pend_mem_region_size

Memory region size of the *managed-system* after it is restarted. The valid values for the *managed-system* are returned by the **lshwres -m managed-system -r mem --level sys -F possible_mem_region_sizes** command.

Valid attribute names for setting memory attributes for a partition:

mem_weight

mem_expansion

Valid values are from **1.00** to **10.00**

hardware_mem_encryption

Valid values are:

- 0** - disable hardware-accelerated encryption
- 1** - enable hardware-accelerated encryption

hardware_mem_expansion

Valid values are:

- 0** - disable hardware-accelerated Active Memory

Expansion
1 - enable hardware-accelerated Active Memory
 Expansion

Valid attribute names for creating the shared memory pool or setting shared memory pool attributes:

pool_mem
 megabytes
max_pool_mem
 megabytes
paging_vios_names | **paging_vios_ids**
 specify one or two VIOS partitions
mem_dedup
 Valid values are:
0 - disable Active Memory Deduplication
1 - enable Active Memory Deduplication
dedup_table_ratio
 Ratio of the deduplication table size to the maximum memory for the pool. Valid values are displayed by the **lshwres -r mem -m managed-system --level sys -F possible_dedup_table_ratios** command.

Valid attribute names for creating the reserved storage device pool or setting reserved storage device pool attributes:

vios_names | **vios_ids**
 specify one or two VIOS partitions

Valid attribute names for setting processing attributes for the *managed-system*:

max_curr_procs_per_lpar
 The valid values for the *managed-system* are returned by the **lshwres -r proc --level sys -m managed-system -F possible_max_curr_procs_per_lpar** command.

Valid attribute names for setting processing attributes for a partition:

sharing_mode
 Valid values for partitions using dedicated processors are:
keep_idle_procs - never share processors
share_idle_procs - share processors only when partition is inactive
share_idle_procs_active - share processors only when partition is active
share_idle_procs_always - always share processors

Valid values for partitions using shared processors are:
cap - capped
uncap - uncapped
uncap_weight

Valid attribute names for setting shared processor pool attributes:

new_name
max_pool_proc_units
reserved_pool_proc_units

Valid attribute names for setting shared processor pool attributes for a partition:
shared_proc_pool_name | **shared_proc_pool_id**

Valid attribute names when adding an HEA logical port:

vlan_id_list

comma separated

lhea_capabilities

Comma separated list of Logical Host Ethernet adapter (LHEA) capabilities, with each capability having one of the following formats:

capability

or

5lieqnieqlqplcq/mr

where *ieq* (interruptible event queues), *nieq* (non-interruptible event queues), *qp* (queue pairs), *cq* (completion queues), and *mr* (memory regions) each specify the resource amount in addition to the base minimum.

Valid values for *capability*:

0 - base minimum

1 - low

2 - medium

3 - high

4 - dedicated

For example:

5/22/128/1021/1019/63424

allowed_os_mac_addr

1 to 4 comma separated MAC addresses, each specified as 12 hexadecimal characters. Other valid values:

all - all OS defined MAC addresses are allowed

none - no OS defined MAC addresses are allowed

Valid attribute names for setting HEA physical port attributes:

conn_speed

Valid values are:

auto - system selects automatically

10 - 10 Mbps

100 - 100 Mbps

1000 - 1 Gbps

10000 - 10 Gbps

duplex

Valid values are:

auto - system selects automatically
full - full duplex
max_rcv_packet_size
 Valid values are:
1500 - 1500 bytes per frame (non-jumbo)
9000 - 9000 bytes per frame (jumbo)
flow_control
 Valid values are:
0 - disable flow control
1 - enable flow control
promisc_lpar_name | promisc_lpar_id

Valid attribute names for setting HEA port group attributes:

pend_port_group_mcs_value

Valid attribute names for switching an SR-IOV adapter to shared mode:

slot_id

Required

adapter_id

Valid values are:

1 - 32, inclusive

If the adapter ID is not specified, the next available adapter ID will be assigned.

Valid attribute names for switching an SR-IOV adapter to dedicated mode:

slot_id

Required

Valid attribute names for moving the configuration of an SR-IOV adapter:

slot_id

Required

DRC index of the I/O slot of the failed SR-IOV adapter

target_slot_id

Required

DRC index of the I/O slot of the new SR-IOV adapter

Valid attribute names for changing an SR-IOV physical port:

Warning! When an attribute for an SR-IOV physical port is changed, a short network interruption may occur for all partitions that share the physical port.

adapter_id

Required

phys_port_id

Required

conn_speed

Possible valid values are:

auto

10 - 10 Mbps

100 - 100 Mbps

1000 - 1 Gbps

10000 - 10 Gbps

40000 - 40 Gpbs
100000 - 100 Gpbs
max_eth_logical_ports
 An integer value less than or equal to the maximum number of Ethernet logical ports allowed on any physical port on the adapter
max_rcv_packet_size
 Valid values are:
1500 - 1500 bytes
9000 - 9000 bytes (Jumbo frames)
phys_port_label
 1-16 characters
 Specify **none** to clear the physical port label
phys_port_sub_label
 1-8 characters
 Specify **none** to clear the physical port sublabel
rcv_flow_control
 Valid values are:
0 - disable
1 - enable
trans_flow_control
 Valid values are:
0 - disable
1 - enable
veb_mode
 Valid values are:
0 - disable Virtual Ethernet Bridge mode
1 - enable Virtual Ethernet Bridge mode
vepa_mode
 Valid values are:
0 - disable Virtual Ethernet Port Aggregator mode
1 - enable Virtual Ethernet Port Aggregator mode

Valid attribute names for adding an SR-IOV logical port:

adapter_id
 Required
phys_port_id
 Required
logical_port_type
 Required
 Valid values are:
eth - ethernet logical port
allowed_os_mac_addr
 Comma separated list of MAC addresses, each specified as 12 hexadecimal characters. Other valid values are:
all - all OS defined MAC addresses are allowed (default value)
none - no OS defined MAC addresses are allowed
allowed_priorities
 Comma separated list of integers between **0** and **7**, inclusive, or one of the following values:
all - all supported priorities are allowed
none - no priorities are allowed (default value)
allowed_vlan_ids

Comma separated list of VLAN IDs, or one of the following values:

all - all VLAN IDs are allowed (default value)

none - no VLAN IDs are allowed

capacity

Must be a multiple of **min_eth_capacity_granularity** of the physical port, which is displayed by the

lshwres -m managed-system -r sriov

--subtype physport --level {eth | ethc}

-F min_eth_capacity_granularity command.

The minimum value and the default value is the value of **min_eth_capacity_granularity**, the maximum value is **100**.

config_id

If not specified, the management console will assign the next available value.

diag_mode

Valid values are:

0 - disable (default value)

1 - enable

Warning! If diagnostics mode is enabled, no other logical port may be attached to the physical port. If logical ports are already in use on the physical port, you will have to power off partitions or use Dynamic Logical Partitioning (DLPAR) to remove the logical ports that are attached to the physical port.

huge_dma_window_mode

Valid values are:

0 - disable (default value)

1 - enable

logical_port_id

DRC index of an SR-IOV logical port. If not specified, an SR-IOV logical port will be automatically assigned.

mac_addr

12 hexadecimal characters

If not specified, a unique MAC address will be automatically generated for the logical port. It is highly recommended that you use an automatically generated MAC address.

port_vlan_id

Port VLAN ID or **0** to disable VLAN tag insertions for untagged frames (default value)

promisc_mode

Valid values are:

0 - disable (default value)

1 - enable

pvid_priority

Valid values are:

An integer between **0** and **7**, inclusive.

Default value is **0**.

Valid attribute names for changing an SR-IOV logical port:

adapter_id

Required

logical_port_id

Required

allowed_os_mac_addrs

Comma separated list of MAC addresses, each specified as 12 hexadecimal characters

allowed_prioritiesComma separated list of integers between **0** and **7**, inclusive**allowed_vlan_ids**

Comma separated list of VLAN IDs

diag_mode

Valid values are:

0 - disable (default value)**1** - enable

Warning! If diagnostics mode is enabled, no other logical port may be attached to the physical port. If logical ports are already in use on the physical port, you will have to power off partitions or use Dynamic Logical Partitioning (DLPAR) to remove the logical ports that are attached to the physical port.

port_vlan_idPort VLAN ID or **0** to disable VLAN tag insertions for untagged frames**pvid_priority**

Valid values are:

An integer between **0** and **7**, inclusive.

Valid attribute names for removing an SR-IOV logical port:

adapter_id

Required

logical_port_id

Required

Valid attribute names for resetting statistics for an SR-IOV physical port:

adapter_id

Required

phys_port_id

Required

Valid attribute names for resetting statistics for an SR-IOV logical port:

adapter_id

Required

logical_port_id

Required

- v** Specify this option to enable verbose mode. When verbose mode is enabled, warning messages and informational messages are displayed for successful virtual NIC operations. Warning messages and informational messages are displayed for virtual NIC operations that fail, regardless of whether this option is specified.
- help** Display the help text for this command and exit.

EXAMPLES

Add the I/O slot with DRC index **21010001** to partition **p1** and set the I/O pool ID for the slot to **3**:

```
chhwres -r io -m sys1 -o a -p p1 -l 21010001
-a "slot_io_pool_id=3"
```

Add I/O pools **2** and **3** to the I/O pools in which partition **p1** is participating:

```
chhwres -r io --subtype iopool -m 9406-520*1234321A -o s
-p p1 -a ""lpar_io_pool_ids+=2,3""
```

Clear the physical I/O slot with DRC index **21010207** after the physical I/O adapter in that slot was concurrently removed but not replaced:

```
chhwres -r io -m sys1 -o c -l 21010207
```

Add a virtual ethernet adapter to the partition with ID **3**:

```
chhwres -r virtualio -m 9406-520*1234321A -o a --id 3
--subtype eth -a "ieee_virtual_eth=1,
port_vlan_id=4,"addl_vlan_ids=5,6",is_trunk=1,
trunk_priority=1"
```

Add a virtual fibre channel client adapter to the partition **p1**:

```
chhwres -r virtualio -m sys1 -o a -p p1 --subtype fc
-a "adapter_type=client,remote_lpar_name=vios,remote_slot_num=16"
```

Add a virtual NIC to the partition **p1**:

```
chhwres -r virtualio -m sys1 -o a -p p1 --subtype vnic -v
-a "port_vlan_id=3,backing_devices=sriov/vios1//1/2/"
```

Remove the virtual adapter in slot **3** from partition **p1**:

```
chhwres -r virtualio -m sys1 -o r -p p1 -s 3
```

Add an additional VLAN ID to the virtual ethernet adapter in slot **4** of the partition **AIX**:

```
chhwres -r virtualio -m sys1 -o s --subtype eth -p AIX -s 4
-a "addl_vlan_ids+=5"
```

Add additional VLAN IDs to the virtual NIC in slot **6** of the partition **lp4** and change its port VLAN ID priority:

```
chhwres -r virtualio -m sys1 -o s --subtype vnic -p lp4 -s 6
-a ""allowed_vlan_ids+=10,20,30",pvid_priority=5"
```

Add additional backing devices to the virtual NIC in slot **3** of the partition with ID **4**:

```
chhwres -r virtualio -m sys1 -o s --subtype vnic --id 4 -s 3
-a ""backing_devices+=sriov/vios1//1/0/4/50,sriov/vios2//2/0/4/62""
--allownetdisrupt
```

Remove a backing device from the virtual NIC in slot **3** of the partition **lp4**:

```
chhwres -r virtualio -m sys1 -o s --rsubtype vnic -p lp4 -s 3  
-a "backing_devices==sriov///1/0///"
```

Change the failover priority of the backing device, with logical port ID **27004001**, for the virtual NIC in slot **3** of the partition **lp4**:

```
chhwres -m sys1 -r virtualio -o s --rsubtype vnicbkdev -p lp4 -s 3  
--logport 27004001 -a "failover_priority=75"
```

Make the backing device, with logical port ID **27004001**, for the virtual NIC in slot **3** of the partition **lp4** active:

```
chhwres -m sys1 -r virtualio -o act --rsubtype vnicbkdev -p lp4 -s 3  
--logport 27004001
```

Disable the virtual ethernet adapter in slot **2** of the partition **p5**:

```
chhwres -r virtualio -m sys1 -o d --rsubtype eth -p p5 -s 2
```

Enable HSL OptiConnect for the IBM i partition **i_p1**:

```
chhwres -r virtualio -m sys1 -o s -p i_p1  
--rsubtype hsl -a "hsl_pool_id=1"
```

Add the virtual network **vnet5**:

```
chhwres -r virtualio --rsubtype vnetwork -m sys1 -o a --vnetwork vnet5  
-a "vswitch=ETHERNET0,vlan_id=5,is_tagged=1"
```

Add the virtual switch **vs1**:

```
chhwres -r virtualio --rsubtype vswitch -m sys1 -o a --vswitch vs1
```

Add **128** MB of memory to the partition with ID **1**, and time out after **10** minutes:

```
chhwres -r mem -m sys1 -o a --id 1 -q 128 -w 10
```

Remove **512** MB of memory from the AIX partition **aix_p1** and return a detail level of **5**:

```
chhwres -r mem -m 9406-520*1234321A -o r -p aix_p1 -q 512 -d 5
```

Set the total amount of memory for the partition **p1** to **8192** MB:

```
chhwres -r mem -m sys1 -o s -p p1 -q 8192
```

Remove **128** MB of logical memory and **100** MB of I/O entitled memory from the partition **smp1**:

```
chhwres -r mem -m sys1 -o r -p smp1 -q 128 --entitled 100
```

Set the partition **smp1** to automatic I/O entitled memory management mode:

```
chhwres -r mem -m sys1 -o a -p smp1 --entitled auto
```

Reset I/O entitled memory statistics for partition **smp1**:

```
chhwres -r mem -m sys1 -o rs -p smp1
```

Create the shared memory pool with redundant VIOS partitions (on a managed system that supports partition suspend and resume operations, this command also creates the reserved storage device pool if it does not already exist):

```
chhwres -r mempool -m sys1 -o a -a "pool_mem=4096,max_pool_mem=8192,"paging_vios_names=vios_p1,vios_p2""
```

Increase the size of the shared memory pool by **256** MB:

```
chhwres -r mempool -m sys1 -o s -a "pool_mem+=256"
```

Add the device **hdisk1** on VIOS partition **vios_p1** to the shared memory pool (on a managed system that supports partition suspend and resume operations, this command adds the device to the reserved storage device pool):

```
chhwres -r mempool -m sys1 --rsubtype pgdev -o a -p vios_p1 --device hdisk1
```

Create the reserved storage device pool with a single VIOS partition:

```
chhwres -r rspool -m sys1 -o a -a "vios_names=vios_p2"
```

Add the VIOS partition with ID **1** to the reserved storage device pool:

```
chhwres -r rspool -m sys1 -o s -a "vios_ids+=1"
```

Remove the device **hdisk1** on VIOS partition **vios_p1** from the reserved storage device pool:

```
chhwres -r rspool -m sys1 --rsubtype rsdev -o r -p vios_p1 --device hdisk1
```

Switch over the running shared memory partition **smp1** to its redundant paging VIOS partition:

```
chhwres -r mem -m sys1 -p smp1 -o so
```

Set the number of pages of huge page memory requested for the managed system to **2** (the managed system must be powered off):

```
chhwres -r mem -m sys1 -o s -a "requested_num_sys_huge_pages=2"
```

Move **1** processor from partition **p1** to partition **p2** (both partitions are using dedicated processors):

```
chhwres -r proc -m 9406-520*1234321A -o m -p p1 -t p2 --procs 1
```

Move **.5** processing units from the partition with ID **1** to the partition with ID **2** (both partitions are using shared processors):

```
chhwres -r proc -m sys1 -o m --id 1 --tid 2 --procunits .5
```

Add **.25** processing units to the IBM i partition **i_p1**:

```
chhwres -r proc -m sys1 -o a -p i_p1 --procunits .25
```

Set the number of processors for partition **p1** to 2:

```
chhwres -r proc -m sys1 -o s -p p1 --procs 2
```

Configure the shared processor pool with ID 1:

```
chhwres -r procpool -m sys1 -o s --poolid 1 -a "new_name=pool1,max_pool_proc_units=2,reserved_pool_proc_units=.5"
```

Move the partition **sharedlpar1** to shared processor pool **pool1**:

```
chhwres -r procpool -m sys1 -o s -p sharedlpar1 -a "shared_proc_pool_name=pool1"
```

Add logical port 4 for physical port 0 belonging to port group 2 of the HEA with an adapter ID of 23000020 to partition **p1**. Also set the LHEA capability level to low:

```
chhwres -r hea -m mySystem -o a -p p1 -l 23000020 --physport 0 -g 2 --logport 4 -a "lhea_capabilities=1"
```

Remove logical port 1 belonging to port group 2 of the HEA with an adapter ID of 23000020 from the partition with ID 8:

```
chhwres -r hea -m 9117-MMA*123432C -o r --id 8 -l 23000020 -g 2 --logport 1
```

Set physical port attributes for port group 2 of physical port 1 of the HEA with an adapter ID of 23000020:

```
chhwres -r hea -m mySystem -o s -l 23000020 -g 2 --physport 1 -a "conn_speed=auto,duplex=auto,flow_control=1"
```

Set port group attributes for port group 1 of the HEA with an adapter ID of 23000030:

```
chhwres -r hea -m sys1 -o s -l 23000030 -g 1 -a "pend_port_group_mcs_value=4"
```

Switch an SR-IOV adapter to shared mode:

```
chhwres -r sriov -m sys1 --subtype adapter -o a -a "slot_id=21010202"
```

Switch an SR-IOV adapter to dedicated mode:

```
chhwres -r sriov -m sys1 --subtype adapter -o r -a "slot_id=21010202"
```

Set the connection speed for SR-IOV physical port 0 to 100 Gbps:

```
chhwres -r sriov -m sys1 --subtype physport -o s -a "adapter_id=1,phys_port_id=0,conn_speed=100000"
```

Add an SR-IOV ethernet logical port (using defaults) to partition **lpar1**:

```
chhwres -r sriov -m sys1 --subtype logport -o a -p lpar1
-a "adapter_id=1,phys_port_id=1,logical_port_type=eth"
```

Remove an SR-IOV ethernet logical port from partition **lpar1**:

```
chhwres -r sriov -m sys1 --subtype logport -o r -p lpar1
-a "adapter_id=1,logical_port_id=27004001"
```

Change the port VLAN ID for an SR-IOV ethernet logical port in partition **lpar1**:

```
chhwres -r sriov -m sys1 --subtype logport -o s -p lpar1
-a "adapter_id=1,logical_port_id=27004001,port_vlan_id=2"
```

Reset the statistics for an SR-IOV physical port:

```
chhwres -r sriov -m sys1 --subtype physport -o rs
-a "adapter_id=1,phys_port_id=0"
```

ENVIRONMENT

None

BUGS

None

AUTHOR

IBM Austin

SEE ALSO

lshwres, **rsthwres**, **chsyscfg**, **lssyscfg**

NAME

chipsec - change IPsec configuration

SYNOPSIS

To create and start an IPsec connection to a managed system:

```
chipsec -m managed-system --left IPv6-address --passkey passphrase
```

To create and start a user-defined IPsec connection:

```
chipsec -f connection-file --passkey passphrase  
[--cert certificate-file]
```

To start, stop, or remove an IPsec connection:

```
chipsec { --up connection-name | --down connection-name |  
  -r connection-name }
```

To start or stop the IPsec service on the Hardware Management Console:

```
chipsec { --start | --stop }
```

DESCRIPTION

chipsec changes the Internet Protocol Security (IPsec) configuration on the Hardware Management Console (HMC). It can also be used to manage IPsec connections and the IPsec service.

An IPsec connection is automatically started when it is created. The IPsec service is automatically started when the first IPsec connection is created.

The HMC uses the strongSwan IPsec implementation.

OPTIONS

-m The name of the managed system to which to create an IPsec connection. The name may either be the user-defined name for the managed system, or be in the form *tttt-mmm*ssssss*, where *tttt* is the machine type, *mmm* is the model, and *ssssss* is the serial number of the managed system. The *tttt-mmm*ssssss* form must be used if there are multiple managed systems with the same user-defined name.

--left The HMC IPv6 address to be used for the IPsec connection to the *managed-system*.

--passkey

The passphrase.

-f The name of the file that contains the configuration information for the user-defined IPsec connection to be created. The configuration information in the file must be the **conn** section for the **ipsec.conf** file. The file can contain the configuration information for one connection only.

If the file is on removable media, the media must be present in the removable media device and the device must be mounted with the **mount** command before this command is issued. The **lsmediadev** command can be used to display all of the removable media devices on the HMC.

--cert The name of the file that contains the certificate for the user-defined IPsec connection to be created. The file name must end with the **.pem** suffix.

If the file is on removable media, the media must be present in the removable media device and the device must be mounted with the **mount** command before this command is issued. The **lsmediadev** command can be used to display all of the removable media devices on the HMC.

--up Specify this option to start the IPsec connection *connection-name*. If the IPsec service is not running, it will be started.

An IPsec connection is automatically started when it is created.

- down** Specify this option to stop the IPsec connection *connection-name*. If no other IPsec connections are active, the IPsec service will be stopped.
- r** Specify this option to remove the IPsec connection *connection-name*.
- An active IPsec connection is stopped before it is removed.
- start** Specify this option to start the IPsec service on the HMC.
- The IPsec service is automatically started when the first IPsec connection is created or started.
- stop** Specify this option to stop the IPsec service on the HMC.
- The IPsec service is automatically stopped when the last IPsec connection is removed or stopped.
- help** Display the help text for this command and exit.

EXAMPLES

Create and start an IPsec connection to managed system **j21**:

```
chipsec -m j21 --left fe80:0:0:0:221:5eff:fe46:57ce  
--passkey hmctest1234!
```

Create and start a user-defined IPsec connection:

```
chipsec -f connfile.txt --passkey passphrase1
```

Create and start a user-defined IPsec connection with certificate. Both files are on a USB flash memory device (the USB flash memory device must already be connected to the HMC):

```
lsmediadev (to obtain mount points)
```

```
mount /media/sdb1
```

```
chipsec -f /media/sdb1/user_file.txt --cert /media/sdb1/user_cert.pem  
--passkey hmctest1234!
```

Start the IPsec connection **j21**:

```
chipsec --up j21
```

Stop the IPsec connection **j21**:

```
chipsec --down j21
```

Remove the IPsec connection **j21**:

```
chipsec -r j21
```

Start the IPsec service on the HMC:

```
chipsec --start
```

Stop the IPsec service on the HMC:

```
chipsec --stop
```

ENVIRONMENT

None

BUGS

None

AUTHOR

IBM Austin

SEE ALSO

lsipsec, lsmediadev

NAME

chkmedia - check media readiness

SYNOPSIS

chkmedia -r usbflashmem [--help]

DESCRIPTION

chkmedia allows a user to test for media readiness on the Hardware Management Console (HMC). Media devices that can be tested for readiness are USB data storage devices. In order to test for media readiness the media device is mounted, then a small amount of data is written to, read from, then deleted from the media. Finally, the media device is unmounted. No other diagnostic actions are performed on the media or the media device. The media device cannot already be mounted when this command is run.

OPTIONS

-r The media device to be tested for readiness. The only valid value is **usbflashmem** for a flash memory device or other data storage device that is inserted into a USB port on the HMC.

--help Display the help text for this command and exit.

EXAMPLES

Check a USB flash memory device for media readiness:

chkmedia -r usbflashmem

ENVIRONMENT

None

BUGS

None

AUTHOR

IBM Austin

NAME

chled - change LED

SYNOPSIS

```
chled -r sa -t {phys | virtuallpar | virtualsys}  
-m managed-system -o {on | off}  
[-p partition-name | --id partition-ID] [--help]
```

DESCRIPTION

chled changes the state of an LED on the *managed-system*.

OPTIONS

- r** The type of LED resource to change. The only valid value is **sa** for System Attention (SA) LED.
- t** The type of System Attention (SA) LED to change. Valid values are **phys** for the physical SA LED for the *managed-system*, **virtuallpar** for a virtual partition SA LED, or **virtualsys** for the virtual SA LED for the *managed-system*.
- m** The name of the managed system on which to change the LED. The name may either be the user-defined name for the managed system, or be in the form *tttt-mmm*sssssss*, where *tttt* is the machine type, *mmm* is the model, and *sssssss* is the serial number of the managed system. The *tttt-mmm*sssssss* form must be used if there are multiple managed systems with the same user-defined name.
- o** The operation to perform on the LED. Valid values are **on** to activate the LED, and **off** to deactivate the LED.

The value **on** is not valid when changing the physical SA LED for the *managed-system*.

- p** The name of the partition for which the virtual partition SA LED is to be changed.

To change a virtual partition SA LED, you must either use this option to specify the name of the partition for which the LED is to be changed, or use the **--id** option to specify the partition's ID. The **-p** and the **--id** options are mutually exclusive.

This option is not valid when changing the physical SA LED or the virtual SA LED for the *managed-system*.

- id** The ID of the partition for which the virtual partition SA LED is to be changed.

To change a virtual partition SA LED, you must either use this option to specify the ID of the partition for which the LED is to be changed, or use the **-p** option to specify the partition's name. The **--id** and the **-p** options are mutually exclusive.

This option is not valid when changing the physical SA LED or the virtual SA LED for the *managed-system*.

- help** Display the help text for this command and exit.

EXAMPLES

Deactivate the physical SA LED for the system:

```
chled -m 9406-570*100103C -r sa -t phys -o off
```

Activate the virtual partition SA LED for partition **lpar3**:

```
chled -m system1 -r sa -t virtuallpar -o on -p lpar3
```

Deactivate the virtual partition SA LED for the partition with ID **3**:

chled -m system1 -r sa -t virtuallpar -o off --id 3

ENVIRONMENT

None

BUGS

None

AUTHOR

IBM Austin

SEE ALSO

lsled

NAME

chlpstate - change partition state

SYNOPSIS

To suspend a partition or validate to suspend a partition:

```
chlpstate -m managed-system -o {suspend | validate} [-t suspend]  
  {-p partition-name | --id partition-ID}  
  [{-f input-data-file | -i "input-data"}]  
  [--protectstorage {1 | 2}]  
  [-w wait-time] [-d detail-level] [-v] [--nodetails]
```

To resume a partition or validate to resume a partition:

```
chlpstate -m managed-system -o {resume | validate} [-t resume]  
  {-p partition-name | --id partition-ID}  
  [--mpio {1 | 2}] [--vlanbridge {1 | 2}]  
  [--vsi {1 | 2}]  
  [-w wait-time] [-d detail-level] [-v] [--nodetails]
```

To cancel a partition suspend or resume operation:

```
chlpstate -m managed-system -o cancel -t {suspend | resume}  
  {-p partition-name | --id partition-ID}  
  [-w wait-time] [-d detail-level]
```

To recover a partition suspend or resume operation:

```
chlpstate -m managed-system -o recover -t {suspend | resume}  
  {-p partition-name | --id partition-ID}  
  [-w wait-time] [-d detail-level] [-v] [--nodetails]  
  [--force]
```

To shut down a suspended partition:

```
chlpstate -m managed-system -o shutdown  
  {-p partition-name | --id partition-ID}  
  [--mpio {1 | 2}] [--vlanbridge {1 | 2}]  
  [-w wait-time] [-d detail-level] [-v] [--nodetails]  
  [--immed] [--force]
```

To shut down or restart a partition that is not suspended:

```
chlpstate -m managed-system -o shutdown  
  {-p partition-name | --id partition-ID}  
  [--immed] [--restart] [--force]
```

DESCRIPTION

chlpstate changes the state of a partition on the *managed-system*.

chlpstate is used to perform partition suspend and resume operations. Partition suspend and resume operations can be performed for AIX, Linux, and IBM i partitions.

See the **chsysstate** command for additional partition state change operations.

chlpstate can also be used to perform partition shutdown operations. Either **chlpstate** or **chsysstate** can be used to shut down a partition that is not suspended.

The recommended procedure for shutting down a suspended partition is to resume the partition first, then shut it down. However, if you want to shut down a partition while it is suspended, then it is recommended that **chlpstate**, not **chsysstate**, be used to shut down the partition because **chlpstate** supports additional options that **chsysstate** does not.

OPTIONS

- o** The operation to perform. Valid values are **suspend** to validate then suspend a partition if validation succeeds, **resume** to validate then resume a suspended partition if validation succeeds, **validate** to validate a partition suspend or resume operation, **cancel** to stop a partition suspend or resume operation, **recover** to recover from a failed partition suspend or resume operation, and **shutdown** to shut down a partition.
- t** To validate, cancel, or recover a partition suspend operation, specify **suspend** with this option.
To validate, cancel, or recover a partition resume operation, specify **resume** with this option.
- m** The name of the managed system on which to perform the partition state change operation. The name may either be the user-defined name for the managed system, or be in the form *tttt-
mmm*ssssss*, where *tttt* is the machine type, *mmm* is the model, and *ssssss* is the serial number of the managed system. The *tttt-
mmm*ssssss* form must be used if there are multiple managed systems with the same user-defined name.
- p** The name of the partition for which to perform the operation.
You must either use this option to specify the name of the partition, or use the **--id** option to specify the partition's ID. The **-p** and the **--id** options are mutually exclusive.
- id** The ID of the partition for which to perform the operation.
You must either use this option to specify the ID of the partition, or use the **-p** option to specify the partition's name. The **--id** and the **-p** options are mutually exclusive.
- f** The name of the file containing the input data for this command. The data consists of attribute name/value pairs, which are in comma separated value (CSV) format.

The format of the data is as follows:

attribute-name=value,attribute-name=value,...

Valid attribute names for this command:

primary_rs_vios_name | **primary_rs_vios_id**
secondary_rs_vios_name | **secondary_rs_vios_id**
primary_vios_name | **primary_vios_id**

These attributes are deprecated. Use the

primary_rs_vios_name | **primary_rs_vios_id**
attributes instead.

secondary_vios_name | **secondary_vios_id**

These attributes are deprecated. Use the

secondary_rs_vios_name | **secondary_rs_vios_id**
attributes instead.

Input data for this command can be specified with this option or the **-i** option. The **-f** and the **-i** options are mutually exclusive.

Input data can be specified when suspending a partition or validating a partition suspend operation only.

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

Input data for this command can be specified with this option or the **-f** option. The **-i** and the **-f** options are mutually exclusive.

Input data can be specified when suspending a partition or validating a partition suspend operation only.

--protectstorage

When suspending a partition or validating a partition suspend operation, use this option to specify whether the partition's virtual storage devices are required to be protected to prevent accidental reassignment of the devices while the partition is suspended. To protect a suspended partition's virtual storage devices, all Virtual I/O Server (VIOS) partitions hosting the devices must be capable of reporting virtual storage device usage of suspended partitions (VIOS version 2.2.1.4 or later).

Valid values for this option are **1** if the partition's virtual storage devices are required to be protected while the partition is suspended, or **2** if the partition's virtual storage devices are not required to be protected while the partition is suspended, but should be protected if possible. If this option is not specified, this option defaults to the value **1**.

If this option has a value of **1** and not all of the VIOS partitions hosting the partition's virtual storage devices are capable of reporting virtual storage device usage of suspended partitions, an error will occur and the operation will fail since not all of the partition's virtual storage devices can be protected. The operation can succeed if this option has a value of **2**, but you would be responsible for ensuring the integrity of the partition's virtual storage devices while the partition is suspended.

This option is valid when suspending a partition or validating a partition suspend operation only.

--mpio When resuming a partition, validating a partition resume operation, or shutting down a suspended partition, use this option to specify whether the management console is required to maintain an equivalent multipath I/O (MPIO) configuration of the partition's virtual SCSI and virtual fibre channel adapters. Valid values are **1** if the management console is required to maintain an equivalent MPIO configuration, or **2** if the management console is not required to maintain an equivalent MPIO configuration, but should do so if possible. If this option is not specified, then this option defaults to the value **1**.

If this option has a value of **1** and the management console cannot guarantee that an equivalent MPIO configuration of all of the partition's virtual SCSI and virtual fibre channel adapters can be maintained, then an error will occur and the operation will fail.

This option is valid when resuming a partition, validating a partition resume operation, or shutting down a suspended partition only.

--vlanbridge

When resuming a partition, validating a partition resume operation, or shutting down a suspended partition, use this option to specify whether each of the partition's virtual ethernet adapters is required to be configured so that it is bridged on the same VLAN to an external network. Valid values are **1** if the partition's virtual ethernet adapters are required to be bridged, or **2** if the partition's virtual ethernet adapters are not required to be bridged, but should be bridged if possible. If this option is not specified, then this option defaults to the value **1**.

If this option has a value of **1** and the management console cannot guarantee that all of the partition's virtual ethernet adapters can be bridged, then an error will occur and the operation will fail.

This option is valid when resuming a partition, validating a partition resume operation, or shutting down a suspended partition only.

--vsi When resuming a partition or validating a partition resume operation, use this option to specify whether the partition's virtual ethernet adapter Virtual Station Interface (VSI) profiles are required to be configured. Valid values are **1** if VSI profiles are required to be configured, or **2** if VSI profiles are not required to be configured, but should be configured if possible. If this option is not

specified, then this option defaults to the value of the **--vlanbridge** option.

If this option has a value of **1** and the management console cannot configure all of the partition's VSI profiles, an error will occur and the operation will fail.

A value of **1** cannot be specified with this option when a value of **2** is specified with the **--vlanbridge** option.

This option is valid when resuming a partition or validating a partition resume operation only.

- w** The maximum time, in minutes, to wait for operating system commands issued by the management console during the partition state change operation to complete. If an operating system command does not complete within the time specified, the partition state change operation will be stopped.

wait-time must be a whole number. If wait-time is 0, the partition state change operation will not be timed out.

If this option is not specified, a default value of 3 minutes is used.

- d** The level of detail requested from operating system commands issued by the management console during the partition state change operation. Valid values are **1** (lowest) through **5** (highest).

If this option is not specified, a default value of 1 is used.

- v** Specify this option to enable verbose mode for the partition state change operation. When verbose mode is enabled, detail messages and warning messages are displayed for a successful partition state change operation. Detail messages and warning messages are always displayed for a partition state change operation that fails, regardless of whether this option is specified.

When the **--nodetails** option is also specified, no detail messages are displayed.

This option can be specified for a partition suspend, resume, or recover operation. This option can also be specified when shutting down a suspended partition.

--nodetails

Specify this option to suppress the display of all detail messages.

--immed

Specify this option on a partition shutdown or restart operation to perform an immediate shutdown or restart.

If this option is specified, but the **--restart** option is not, then an immediate shutdown (operator panel function 8) is performed. If both the **--immed** and the **--restart** options are specified, an immediate restart (operator panel function 3) is performed. If neither the **--immed** nor the **--restart** option is specified, a delayed shutdown is performed.

When shutting down a suspended partition, this option is ignored.

--restart

Specify this option on a partition shutdown operation to restart the partition.

If this option is specified, but the **--immed** option is not, a dump restart operation (operator panel function 22) is performed. If both the **--immed** and the **--restart** options are specified, an immediate restart (operator panel function 3) is performed.

This option is not valid when shutting down a suspended partition.

--force When performing a recover operation, use this option to force the recover operation to proceed when errors are encountered.

When performing a shutdown operation for a suspended partition, use this option to immediately power off the partition without restoring the configuration of its virtual I/O. You will not be able to restart the partition without manually restoring the virtual I/O configuration.

You can also use this option to force a migrating partition to be shut down. Use this option carefully as manual intervention may be required to recover the failed migration after the partition is shut down.

--help Display the help text for this command and exit.

EXAMPLES

Suspend partition **aix1**:

```
chlpstate -m sys1 -o suspend -p aix1
```

Resume partition **aix1**:

```
chlpstate -m sys1 -o resume -p aix1
```

Resume partition **aix1**, and do not fail if an equivalent MPIIO configuration of all of the partition's virtual SCSI and virtual fibre channel adapters cannot be maintained:

```
chlpstate -m sys1 -o resume -p aix1 --mpio 2
```

Validate the operation to suspend the partition with ID **8**:

```
chlpstate -m 8233-E8B*7654321 -o validate -t suspend --id 8
```

Stop the partition suspend operation that is just starting for partition **mylpar**:

```
chlpstate -m systemP -o cancel -t suspend -p mylpar
```

Recover the failed partition resume operation for partition **mylpar**:

```
chlpstate -m systemP -o recover -t resume -p mylpar
```

Shut down the partition with ID **1**:

```
chlpstate -m 9406-570*1234567 -o shutdown --id 1
```

Immediately restart the partition **p1**:

```
chlpstate -m sys1 -o shutdown -p p1 --immed --restart
```

ENVIRONMENT

None

BUGS

None

AUTHOR

IBM Austin

SEE ALSO

chsysstate, lssyscfg

NAME

chlpoutil - change utilization data collection settings

SYNOPSIS

```
chlpoutil -r config [-m managed-system] -s sample-rate
[--help]
```

DESCRIPTION

chlpoutil changes the Hardware Management Console (HMC) settings for utilization data collection for managed systems.

OPTIONS

- r** The type of utilization data collection resources to change. The only valid value is **config** for configuration settings.
- m** The name of the managed system for which the settings for utilization data collection are to be changed. The name may either be the user-defined name for the managed system, or be in the form *tttt-mmm*sssssss*, where *tttt* is the machine type, *mmm* is the model, and *sssssss* is the serial number of the managed system. The *tttt-mmm*sssssss* form must be used if there are multiple managed systems with the same user-defined name, or if the HMC does not currently have a connection to the managed system.

If this option is not specified, then the HMC settings for utilization data collection for all of the systems currently managed by this HMC, and for all of the systems for which utilization data has been previously collected by this HMC will be changed. Even though utilization data collection is enabled for a managed system, the HMC will only collect utilization data for that managed system if the HMC currently has a connection to that managed system.

This command only changes the settings for utilization data collection for this HMC. Utilization data collection by other HMCs for the same managed systems is not affected.

- s** The rate, in seconds, at which to sample the utilization data. Sample rates of **0**, **30** (30 seconds), **60** (60 seconds), **300** (5 minutes), **1800** (30 minutes), and **3600** (1 hour) are supported. A sample rate of **0** disables utilization data collection.
- help** Display the help text for this command and exit.

EXAMPLES

Enable the collection of utilization data with a sample rate of 1 hour for all managed systems:

```
chlpoutil -r config -s 3600
```

Enable the collection of utilization data with a sample rate of 60 seconds for the managed system **mySystem**:

```
chlpoutil -r config -s 60 -m mySystem
```

Disable the collection of utilization data for the managed system with the type, model, and serial number **9406-520*98765432**:

```
chlpoutil -r config -s 0 -m 9406-520*98765432
```

ENVIRONMENT

None

BUGS

None

AUTHOR

IBM Austin

SEE ALSO

lsparutil, rmlparutil

NAME

chnportlogin - N_Port log in and log out

SYNOPSIS

```
chnportlogin -o {login | logout} -m managed-system  
{-p partition-name | --id partition-ID} [-n profile-name]  
[-w wait-time] [-d detail-level] [-v] [--help]
```

DESCRIPTION

chnportlogin performs N_Port login and logout operations for virtual Fibre Channel client adapters that are configured in a partition or a partition profile.

When performing a login operation, all inactive WWPNs will be activated, including the second WWPN in the pair assigned to each virtual Fibre Channel client adapter. When performing a logout operation, all WWPNs not in use will be deactivated.

To successfully log in a virtual Fibre Channel client adapter, the corresponding virtual Fibre Channel server adapter must exist and it must be mapped.

The primary intent of the login operation is to allow the system administrator to allocate, log in and zone WWPNs before the client partition is activated. With best practices, the WWPNs should be logged out after they are zoned on the Storage Area Network (SAN) and before the partition is activated. If a partition is activated with WWPNs still logged in, the WWPNs used for client access are automatically logged out so they can be logged in by the client.

The login operation can also be used to zone the inactive WWPNs in preparation for a partition mobility operation. If the login operation is performed when a partition is already active, only the inactive WWPNs are activated to the "constant login" state similar to physical Fibre Channel adapters. The WWPNs that are already in use by the virtual Fibre Channel client adapters remain in control of the virtual Fibre Channel clients and are not under the control of this command. This means that active client virtual Fibre Channel WWPNs do not achieve a "constant login" state similar to physical Fibre Channel adapters.

The login operation can interfere with partition mobility operations. Best practice is to perform a logout operation for a partition before attempting to migrate the partition to another server. If a mobility operation is attempted with WWPNs still logged in, the firmware will attempt to automatically log out the WWPNs. However, in some cases, the logouts may not complete in time and may therefore cause the mobility operation to fail.

OPTIONS

-o The operation to perform. Valid values are **login** to log in the WWPNs assigned to the virtual Fibre Channel client adapters, or **logout** to log out the WWPNs assigned to the virtual Fibre Channel adapters.

-m The name of the managed system on which to perform the operation. The name may either be the user-defined name for the managed system, or be in the form *tttt-mmm*sssssss*, where *tttt* is the machine type, *mmm* is the model, and *sssssss* is the serial number of the managed system. The *tttt-mmm*sssssss* form must be used if there are multiple managed systems with the same user-defined name.

-p The name of the partition for which to perform the operation.

You must either use this option to specify the name of the partition, or use the **--id** option to specify the partition's ID. The **-p** and the **--id** options are mutually exclusive.

--id The ID of the partition for which to perform the operation.

You must either use this option to specify the ID of the partition, or use the **-p** option to specify the partition's name. The **--id** and the **-p** options are mutually exclusive.

- n** The name of the partition profile for which to perform the operation. The operation will be performed for all of the virtual Fibre Channel client adapters configured in the partition profile. If this option is not specified, then the operation will be performed for all of the virtual Fibre Channel client adapters in the current configuration of the partition.
- w** The maximum time, in minutes, to wait for each Virtual I/O Server (VIOS) command issued by the HMC to complete. If a VIOS command does not complete within the time specified, the login or logout operation for the virtual Fibre Channel client adapter will be stopped.

wait-time must be a whole number. If wait-time is 0, the operation will not be timed out.

If this option is not specified, a default value of 3 minutes is used.

- d** The level of detail requested from VIOS commands issued by the HMC. Valid values are **0** (none) through **5** (highest).

If this option is not specified, a default value of 1 is used.

- v** Specify this option to enable verbose mode. When verbose mode is enabled and this command succeeds, all warning messages are displayed. When this command partially succeeds or fails, all warning and error messages are displayed regardless of whether this option is specified.

- help** Display the help text for this command and exit.

EXIT STATUS

This command has the following return codes:

- 0 Success
- 100 Partial success

Any other value means the command failed.

EXAMPLES

Log in all of the inactive WWPNs assigned to the virtual Fibre Channel client adapters in a partition's current configuration:

```
chnportlogin -m 9117-MMB*1234567 -o login -p clientPartition
```

Log in all of the WWPNs assigned to the virtual Fibre Channel client adapters in a partition profile:

```
chnportlogin -m sys1 -o login -p clientPartition -n defaultProf
```

Log out all of the WWPNs assigned to the virtual Fibre Channel client adapters in a partition profile:

```
chnportlogin -m sys1 -o logout -p clientPartition -n defaultProf
```

ENVIRONMENT

None

BUGS

None

AUTHOR

IBM Austin

SEE ALSO

lsnportlogin

NAME

chproxy - change proxy settings

SYNOPSIS

```
chproxy -t sslproxy -o {enable | disable}  
[-h host-name -p port [-u user-ID [--passwd password]]]  
 [--help]
```

DESCRIPTION

chproxy changes the Secure Sockets Layer (SSL) proxy settings used by the Hardware Management Console when performing call-home functions.

OPTIONS

- t** The type of proxy settings to change. The only valid value is **sslproxy** to change the SSL proxy settings used when performing call-home functions.
- o** The operation to perform. Valid values are **enable** to enable the use of an SSL proxy server for call-home functions, and **disable** to disable the use of an SSL proxy server for call-home functions.
- h** The host name or IP address of the SSL proxy server.

This option is required when enabling the use of an SSL proxy server. Otherwise, this option is not valid.
- p** The port number to use on the SSL proxy server.

This option is required when enabling the use of an SSL proxy server. Otherwise, this option is not valid.
- u** The user ID to use for SSL proxy server authentication.

This option is only valid when enabling the use of an SSL proxy server.
- passwd** The password to use for SSL proxy server authentication. If this option is omitted, you will be prompted to enter the password.

This option is only valid when enabling the use of an SSL proxy server.
- help** Display the help text for this command and exit.

EXAMPLES

Enable the use of an SSL proxy server without authentication for call-home:

```
chproxy -t sslproxy -o enable -h 9.3.12.345 -p 30000
```

Enable the use of an SSL proxy server with authentication for call-home (you will be prompted for the password):

```
chproxy -t sslproxy -o enable -h server1 -p 30000 -u adminuser
```

Disable the use of an SSL proxy server for call-home:

```
chproxy -t sslproxy -o disable
```

ENVIRONMENT

None

BUGS

None

AUTHOR

IBM Austin

SEE ALSO

lsproxy

NAME

chpsm - change Power systems management console configuration

SYNOPSIS

chpsm -o s -a "attributes" [--help]

DESCRIPTION

chpsm changes Power systems management console configuration settings.

You must reboot the management console for the changes to take effect.

OPTIONS

- o** The operation to perform. The only valid value is **s** for a set operation.
- a** The configuration data to set. The configuration data consists of attribute name/value pairs, which are in comma separated value (CSV) format. The configuration data must be enclosed in double quotes.

The format of the configuration data is as follows:

attribute-name=value,attribute-name=value,...

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

"attribute-name=value,value,...",...

When a list of values is specified, the attribute name/value pair must be enclosed in double quotes. Depending on the shell being used, nested double quote characters may need to be preceded by an escape character, which is usually a `'\'` character.

Valid attribute names:

lpar_rmc_comm_ifs

Comma separated list of management console network interfaces, in priority order, to enable for partition communication. The following special values are also allowed:

all - enable all management console network interfaces for partition communication (default value)

none - disable all management console network interfaces for partition communication

- help** Display the help text for this command and exit.

EXAMPLES

Enable only the **eth0** interface for partition communication:

```
chpsm -o s -a "lpar_rmc_comm_ifs=eth0"
```

Enable the **eth0** and **eth1** interfaces for partition communication, giving priority to **eth0**:

```
chpsm -o s -a ""lpar_rmc_comm_ifs=eth0,eth1""
```

ENVIRONMENT

None

BUGS

None

AUTHOR

IBM Austin

SEE ALSO

lpsm

NAME

chpwdpolicy - change password policies

SYNOPSIS

To activate a password policy:

```
chpwdpolicy -o a [-n policy-name]
```

To disable password policies:

```
chpwdpolicy -o d
```

To modify a password policy:

```
chpwdpolicy -o m {-f input-data-file | -i "input-data"}
```

DESCRIPTION

chpwdpolicy activates and disables password policies on the Hardware Management Console (HMC). **chpwdpolicy** also modifies password policies.

Password policies are enforced for locally authenticated HMC users only. All HMC user passwords that are created or changed after a password policy is activated must conform to the activated policy.

After password policies are disabled, the only password rules that are enforced are the minimum and maximum password ages that are configured for each user.

Only user-defined password policies can be modified. The active password policy cannot be modified.

OPTIONS

- o** The operation to perform. Valid values are **a** to activate a password policy, **d** to disable password policies, and **m** to modify a password policy.
- n** The name of the password policy to activate. If this option is omitted, the pre-defined **HMC Medium Security Password Policy** password policy will be activated.
- f** The name of the file containing the data to modify a password policy. The data consists of attribute name/value pairs, which are in comma separated value (CSV) format.

The format of the data is as follows:

```
attribute-name=value,attribute-name=value,...
```

Valid attribute names for this command:

name

The name of the password policy to modify.

[new_name]

The new name of the password policy.

[description]

[min_pwage]

The number of days that must elapse before a password can be changed.

[pwage]

The number of days that can elapse before a password expires and must be changed. A value of **99999** indicates no password expiration.

[warn_pwage]

The number of days prior to password expiration when a warning message will begin to be displayed.

[min_length]

The minimum password length.

[hist_size]

The number of times a password must be changed before a password can be reused. This value cannot exceed 50.

[min_digits]

The minimum number of digits that a password must contain.

[min_uppercase_chars]

The minimum number of uppercase characters that a password must contain.

[min_lowercase_chars]

The minimum number of lowercase characters that a password must contain.

[min_special_chars]

The minimum number of special characters that a password must contain. Special characters include symbols, punctuation, and white space characters.

Brackets around an attribute name indicate that the attribute is optional.

Comments are allowed in the input file. Comments must begin with the '#' character.

The data to modify a password policy is required to be specified with this option or the **-i** option. The **-f** and the **-i** options are mutually exclusive.

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

The data to modify a password policy is required to be specified with this option or the **-f** option. The **-i** and the **-f** options are mutually exclusive.

- help** Display the help text for this command and exit.

EXAMPLES

Activate the password policy **xyzPolicy**:

```
chpwdpolicy -o a -n xyzPolicy
```

Disable password policies:

```
chpwdpolicy -o d
```

Modify the password policy **xyzPolicy**:

```
chpwdpolicy -o m -i "name=xyzPolicy,min_pwage=5,pwage=90"
```

ENVIRONMENT

None

BUGS

None

AUTHOR

IBM Austin

SEE ALSO

lspwdpolicy, mkpwdpolicy, rmpwdpolicy

NAME

chpwrmgmt - change power management settings

SYNOPSIS

To enable or disable power saver mode:

```
chpwrmgmt -m managed-system -r sys -o {enable | disable}
[-t {static | dynamic_favor_perf | dynamic_favor_power |
fixed_max_frequency | idle}] [--help]
```

To set or reset dynamic or idle power saver parameters:

```
chpwrmgmt -m managed-system -r sys -o {set | reset}
-t {dynamic_parms | idle_parms} [-a "attributes"]
[--help]
```

DESCRIPTION

chpwrmgmt configures the power management settings for the *managed-system*.

OPTIONS

- m** The name of the managed system for which to configure power management settings. The name may either be the user-defined name for the managed system, or be in the form *ttt-mmm*ssssss*, where *ttt* is the machine type, *mmm* is the model, and *ssssss* is the serial number of the managed system. The *ttt-mmm*ssssss* form must be used if there are multiple managed systems with the same user-defined name.
- r** The type of resource for which to configure power management settings. The only valid value is **sys** for managed system.
- o** The operation to perform. Valid values are **enable** to enable power saver mode, **disable** to disable power saver mode, **set** to set dynamic or idle power saver parameters, and **reset** to reset dynamic or idle power saver parameters to their default values.

Warning: enabling any of the power saver modes will cause changes in processor frequencies, changes in processor utilization, and changes in power consumption, which can cause performance of the managed system to vary.

Warning: changing dynamic or idle power saver parameters can result in unexpected behavior and performance impacts. It is recommended that you consult your system provider before changing dynamic or idle power saver parameters.

- t** The type of power saver mode to enable or the type of power saver parameters to set or reset. Valid power saver mode types are **static** for static power saver mode, **dynamic_favor_perf** for dynamic power saver mode favoring performance, **dynamic_favor_power** for dynamic power saver mode favoring power, **fixed_max_frequency** for fixed maximum frequency power saver mode, and **idle** for idle power saver mode. Valid power saver parameter types are **dynamic_parms** for dynamic power saver tuning parameters and **idle_parms** for idle power saver parameters.

When enabling power saver mode, static power saver mode will be enabled if this option is not specified. With the exception of idle power saver mode, only one power saver mode can be enabled at a time. Idle power saver mode can be enabled independently of all other modes.

To disable idle power saver mode, you must specify **idle** with this option. Otherwise, this option is not allowed when disabling power saver mode, as the current power saver mode will be disabled.

- a** The power saver parameters to set. The data consists of attribute name/value pairs, which are in comma separated value (CSV) format.

The format of the attribute data is as follows:

attribute-name=value,attribute-name=value,...

The valid attribute names for setting dynamic power saver tuning parameters are listed below. Issue the **lspwrmgmt -m managed-system -r sys -t dynamic_parms** command to display the minimum and maximum values allowed for each parameter.

util_thresh_incr_freq
percentage with up to 1 decimal place

util_thresh_decr_freq
percentage with up to 1 decimal place

num_samples_util_stats
integer

freq_step_up_size
percentage with up to 1 decimal place

freq_step_down_size
percentage with up to 1 decimal place

active_cores_delta
percentage

active_slack_cores_util_thresh
percentage with up to 1 decimal place

core_freq_delta_state
Valid values:
0 - disable
1 - enable

core_max_freq_delta
percentage

The valid attribute names for setting idle power saver parameters are listed below. Issue the **lspwrmgmt -m managed-system -r sys -t idle_parms** command to display the minimum and maximum values allowed for each parameter.

enter_delay_seconds
integer

exit_delay_seconds
integer

enter_util_thresh
percentage

exit_util_thresh
percentage

This option is required when setting dynamic or idle power saver parameters. Otherwise this option is not valid.

--help Display the help text for this command and exit.

EXAMPLES

Enable static power saver mode on managed system **system1**:

```
chpwrmgmt -m system1 -r sys -o enable
```

or

```
chpwrmgmt -m system1 -r sys -o enable -t static
```

Enable dynamic power saver mode favoring performance on managed system **system1**:

```
chpwrmgmt -m system1 -r sys -o enable -t dynamic_favor_perf
```

Enable idle power saver mode on managed system **system1**:

```
chpwrmgmt -m system1 -r sys -o enable -t idle
```

Disable the current power saver mode on managed system **9117-MMA*1234567**:

```
chpwrmgmt -m 9117-MMA*1234567 -r sys -o disable
```

Disable idle power saver mode on managed system **system1**:

```
chpwrmgmt -m system1 -r sys -o disable -t idle
```

Set the utilization thresholds for increasing and decreasing frequency on managed system **system1** to **95.5%**:

```
chpwrmgmt -m system1 -r sys -o set -t dynamic_parms  
-a "util_thresh_incr_freq=95.5,util_thresh_decr_freq=95.5"
```

Reset the dynamic power saver tuning parameters on managed system **system1** to their default values:

```
chpwrmgmt -m system1 -r sys -o reset -t dynamic_parms
```

ENVIRONMENT

None

BUGS

None

AUTHOR

IBM Austin

SEE ALSO

lspwrmgmt

NAME

chsacfg - change Service Agent configuration

SYNOPSIS

To add or remove authorized users:

```
chsacfg -t authuser -o {add | remove} -a email-address
```

To enable or disable call-home for a managed system or this HMC:

```
chsacfg -t callhome -o {enable | disable}
  -m {managed-system | hmc} | --all
```

To change call-home settings:

```
chsacfg -t callhomeserver -o set
  -f input-data-file | -i "input-data"
```

To test an existing Internet connection for performing call-home functions over an encrypted SSL connection (this operation may take a long time to complete):

```
chsacfg -t callhomeserver -o test
```

To change customer information:

```
chsacfg -t custinfo -o set
  -f input-data-file | -i "input-data"
```

To change customer email notification settings:

```
chsacfg -t email
  -o {enable | disable | setsmtplib | add |
  remove | test | addfrom | rmfrom}
  [-a {email-address | all}]
  [--events {all | callhome | dpo}]
  [-h host-name [-p port-number]]
```

To test FTP server and firewall settings for the offload of service information or reset those settings to their default values:

```
chsacfg -t ftp -o {test | reset}
```

To change FTP firewall settings:

```
chsacfg -t ftpfirewall
  [-o {add | remove | enable | disable}]
  [-h host-name] [-p port-number]
  [-u user-ID] [--passwd password]
  [--passive passive-mode]
  [--fwtype firewall-type] [-n name]
```

To change FTP offload server settings:

```
chsacfg -t ftpoffload
  [-o {enable | disable}]
  [-h host-name] [-p port-number]
  [-u user-ID] [--passwd password]
  [--passive passive-mode] [-d directory]
```

To change SNMP trap configuration:

```
chsacfg -t snmp
  -o {add | remove}
  -a {SNMP-address | all} [-p port-number]
  [--events SNMP-trap-list]
```

To set SNMP options:

```
chsacfg -t snmpopt -o set
  {-f input-data-file | -i "input-data" }
```

DESCRIPTION

chsacfg changes the Service Agent configuration.

OPTIONS

- t** The type of Service Agent configuration information to change. Valid values are **authuser** for the users that are authorized to access information collected by Electronic Service Agent for the systems managed by this Hardware Management Console (HMC), **callhome** for the call-home state for a managed system or this HMC, **callhomeserver** for the call-home server settings for this HMC, **custinfo** for customer information settings, **email** for customer email notification settings, **ftp** to test FTP server and firewall settings for the offload of service information or to reset those settings to their default values, **ftpfirewall** for FTP firewall settings, **ftpoffload** for FTP offload server settings, **snmp** for SNMP trap configuration, and **snmpopt** for SNMP options.
- o** The operation to perform. Valid values are:
 - enable** - enable call-home for a managed system or this HMC, customer email notification, FTP firewall settings, or FTP offload of service information
 - disable** - disable call-home for a managed system or this HMC, customer email notification, FTP firewall settings, or FTP offload of service information
 - setsmtp** - set the SMTP server information for customer notification email
 - add** - add an authorized user, a notification email address, a name to the exclusion list for the FTP firewall server, or an address for SNMP trap notifications
 - remove** - remove an authorized user, a notification email address, a name from the exclusion list for the FTP firewall server, or an address from the SNMP trap notification list
 - set** - set call-home server settings, customer information, or SNMP options
 - test** - test an existing Internet connection for performing call-home functions, send a test email to a notification email address, or test FTP server and firewall settings used to transmit service information
 - reset** - reset FTP server and firewall settings to their default values
 - addfrom** - set the email address from which customer notification emails are sent
 - rmfrom** - reset the email address from which customer notification emails are sent to the default value
- a** The email address to add or remove as an authorized user, the email address to add, remove, or test for customer notification email, the email address to set as the email address from which customer notification emails are sent, or the SNMP address to add or remove for SNMP trap notifications. SNMP addresses are specified in the form *Community-Name@IP-Address*.

When adding or removing an authorized user, multiple email addresses can be specified with this option. The email addresses must be comma separated.

When performing a remove or test operation for a customer notification email address, **all** can be specified with this option instead of a specific email address. If **all** is specified, then the remove or test operation will be performed for all notification email addresses.

When performing a remove operation for an SNMP address, **all** can be specified with this option instead of a specific SNMP address. If **all** is specified, then the remove operation will be performed for all SNMP trap notification addresses.

This option is required when adding or removing authorized users, when adding, removing, or testing a customer notification email address, when setting the email address from which customer notification emails are sent, or when adding or removing an SNMP trap notification address. This option is not valid for any other operation.

- m** The name of the managed system for which to enable or disable call-home. The name may either be the user-defined name for the managed system, or be in the form *tttt-mmm*ssssss*, where *tttt* is the machine type, *mmm* is the model, and *ssssss* is the serial number of the managed system. The *tttt-mmm*ssssss* form must be used if there are multiple managed systems with the same user-defined name.

Specify **hmc** with this option to enable or disable call-home for this HMC.

You can either use this option to enable or disable call-home for this HMC or a single managed system, or use the **--all** option to enable or disable call-home for this HMC and all systems managed by this HMC. The **-m** and the **--all** options are mutually exclusive.

- all** Specify this option to enable or disable call-home for this HMC and all systems managed by this HMC.

You can either use this option to enable or disable call-home for this HMC and all systems managed by this HMC, or use the **-m** option to enable or disable call-home for this HMC or a single managed system. The **--all** and the **-m** options are mutually exclusive.

- f** The name of the file containing the input data for this command. The input data consists of attribute name/value pairs, which are in comma separated value (CSV) format.

The format of the input data is as follows:

```
attribute-name=value,attribute-name=value,...
```

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

```
"attribute-name=value,value,...",...
```

When a list of values is specified, the attribute name/value pair must be enclosed in double quotes. Depending on the shell being used, nested double quote characters may need to be preceded by an escape character, which is usually a `'\'` character.

If `'+='` is used in the attribute name/value pair instead of `'='`, then the specified *value* is added to the existing list.

If `'-='` is used in the attribute name/value pair instead of `'='`, then the specified *value* is subtracted from the existing list.

Valid attribute names for setting call-home server settings:

is_hmc_callhome_server

Valid values:

0 - disable this HMC as a call-home server
1 - enable this HMC as a call-home server
use_internet_conn
 Valid values:
0 - do not use an existing Internet connection to perform call-home functions
1 - use an existing Internet connection to perform call-home functions
use_ssl_proxy
 Valid values:
0 - do not use an SSL proxy to access the Internet
1 - use an SSL proxy to access the Internet
use_authentication
 Valid values:
0 - do not authenticate with the SSL proxy
1 - authenticate with the SSL proxy
host
 Host name or IP address of the SSL proxy
port
 Port number of the SSL proxy
user
 User name used to authenticate with the SSL proxy
password
 Password used to authenticate with the SSL proxy. If omitted, you will be prompted to enter the password.
protocol
 Internet protocol used by the SSL proxy
 Valid values:
ipv4 - IPv4
ipv6 - IPv6
ipv4,ipv6 - try IPv6 then IPv4
servers
 Comma separated list of host names or IP addresses or the value **discover**

Valid attribute names for setting customer information:

admin_company_name
admin_name
admin_email
admin_phone
admin_alt_phone
admin_fax
admin_alt_fax
admin_addr
admin_addr2
admin_city
admin_country
 ISO 3166 2-letter country code
admin_state
 2-letter abbreviation
admin_postal_code
sys_use_admin_info
 Valid values:
0 - do not use the administrator's information as

the system location
1 - use the administrator's information as the system location
sys_phone
sys_addr
sys_addr2
sys_city
sys_country
 ISO 3166 2-letter country code
sys_state
 2-letter abbreviation
sys_postal_code
acct_customer_num
acct_enterprise_num
acct_sales_office
acct_service_office
acct_area

Valid attribute names for setting SNMP options:

enhanced

Valid values:

- 0** - disable enhanced information for SNMP traps
- 1** - enable enhanced information for SNMP traps

Input data for this command can be specified with this option or the **-i** option. The **-f** and the **-i** options are mutually exclusive.

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

Input data for this command can be specified with this option or the **-f** option. The **-i** and the **-f** options are mutually exclusive.

--events

When adding a customer email notification address, specify the type of problem events for which the email address is to be notified. Valid values are **all** for all problem events, **callhome** for only call home problem events, and **dpo** for only Dynamic Platform Optimization events. If this option is not specified, then the email address will be notified for all problem events.

When adding an address for SNMP trap notifications, specify the list of SNMP traps for which the SNMP address is to be notified. If there is more than one SNMP trap, the SNMP traps must be separated by blanks. The defined SNMP traps can be listed by issuing the command **lssacfg -t snmptrapnames**.

This option is only valid when adding a customer email notification address or when adding an SNMP address for SNMP trap notifications.

- h** When enabling customer email notification or setting the SMTP server information for customer notification email, use this option to specify the host name or IP address of the SMTP server.

When changing the FTP firewall settings, use this option to specify the host name or IP address of the firewall server.

When changing the FTP offload server settings, use this option to specify the host name or IP address of the FTP offload server.

This option is required when setting the SMTP server information for customer notification email.

- p** When setting the SMTP server information for customer notification email, use this option to specify the port number to use on the SMTP server. The default port number is 25.

When changing the FTP firewall settings, use this option to specify the port number to use on the firewall server. The default port number is 21.

When changing the FTP offload server settings, use this option to specify the port number to use on the FTP offload server. The default port number is 21.

When adding an SNMP trap notification address, use this option to specify the port number to listen for the SNMP trap. The default port number is 162.

- u** When changing the FTP firewall settings, use this option to specify the user ID to use to log in to the firewall server.

When changing the FTP offload server settings, use this option to specify the user ID to use to log in to the FTP offload server.

--passwd

When changing the FTP firewall settings, use this option to specify the password to use to log in to the firewall server. If this option is omitted, you will be prompted to enter the password.

When changing the FTP offload server settings, use this option to specify the password to use to log in to the FTP offload server. If this option is omitted, you will be prompted to enter the password.

- d** The directory to use on the FTP offload server.

--passive

When changing the FTP firewall settings, use this option to specify the passive mode to use. Valid values are **on** for always use PASV, **off** for always use PORT, and **optional** for try PASV then PORT.

When changing the FTP offload server settings, use this option to specify the passive mode to use. Valid values are **on** for always use PASV and **off** for always use PORT.

--fwtype

The FTP firewall type to use when offloading service information. Valid values are:

- 0** - do not use a firewall
- 1** - connect to firewall server, but send "USER user@real.host.name"
- 2** - connect to firewall server, log in with "USER fwuser" and "PASS fwpassword", and then "USER user@real.host.name"
- 3** - connect to and log in to firewall server, and then use "SITE real.host.name" followed by "USER user" and "PASS password"
- 4** - connect to and log in to firewall server, and then use "OPEN real.host.name" followed by "USER user" and "PASS password"
- 5** - connect to firewall server, but send "USER user@fwuser@real.host.name" and "PASS pass@fwpass" to log in
- 6** - connect to firewall server, but send "USER fwuser@real.host.name" and "PASS fwpass" followed by "USER user" and

- "PASS password" to complete the login
- 7 - connect to firewall server, but send
"USER user@real.host.name fwuser" and
"PASS pass" followed by "ACCT fwpass" to
complete the login

This option is only valid when changing the FTP firewall settings.

- n** The host name or domain name to add to or remove from the exclusion list for the firewall server. The exclusion list is a list of domains or hosts where the firewall should not be used.

This option is required when adding or removing a name in the exclusion list for the FTP firewall server. This option is not valid for any other operation.

- help** Display the help text for this command and exit.

EXAMPLES

Add users that are authorized to access information collected by Electronic Service Agent for the systems managed by this HMC:

```
chsacfg -t authuser -o add -a user1@company.com,user2@company.com
```

Enable call-home for this HMC and all systems managed by this HMC:

```
chsacfg -t callhome -o enable --all
```

Disable call-home for this HMC:

```
chsacfg -t callhome -o disable -m hmc
```

Configure the call-home server settings for this HMC:

```
chsacfg -t callhomeserver -o set -i "is_hmc_callhome_server=1,  
use_internet_conn=1,use_ssl_proxy=1,use_authentication=1,  
host=9.53.2.5,port=8080,user=joeuser,protocol=ipv4,  
servers=sys1.company.com,sys2.company.com"
```

Set customer information:

```
chsacfg -t custinfo -o set -i "admin_company_name=My Company,  
admin_name=Joe Customer,admin_email=joe@my.company.com,  
admin_phone=512-555-1212,admin_addr=1234 Company Ave.,  
admin_city=Austin,admin_country=US,admin_state=TX,  
admin_postal_code=78759,sys_use_admin_info=1"
```

Enable customer email notification for problem events:

```
chsacfg -t email -o enable
```

Set the SMTP server information for customer notification email:

```
chsacfg -t email -o setsmtpt -h mycompany.com -p 30
```

Add an email address to be notified when call home problem events occur:

```
chsacfg -t email -o add -a addr1@mycompany.com
```

--events callhome

Send a test email to all customer notification email addresses:

```
chsacfg -t email -o test -a all
```

Set the email address from which customer notification emails are sent:

```
chsacfg -t email -o addfrom -a myuser@mycompany.com
```

Enable the use of FTP to offload service information and set the FTP server information:

```
chsacfg -t ftpoffload -o enable -h ftphost -u ftpuser  
--passwd userpwd
```

Set the FTP firewall server information to be used when offloading service information:

```
chsacfg -t ftpfirewall --fwtype 1 -h ftpfirewall  
-u ftpuser --passwd userpwd
```

Add a name to the exclusion list for the FTP firewall server:

```
chsacfg -t ftpfirewall -o add -n excludedhost
```

Disable the use of an FTP firewall when offloading service information:

```
chsacfg -t ftpfirewall --fwtype 0
```

Add an SNMP notification to community **mySNMP** at IP address **target.com** for SNMP traps 1, 2 and 7:

```
chsacfg -t snmp -o add -a mySNMP@target.com  
--events Trap1 Trap2 Trap7
```

Remove all SNMP trap notification addresses:

```
chsacfg -t snmp -o remove -a all
```

Enable enhanced information for SNMP traps:

```
chsacfg -t snmpopt -o set -i "enhanced=1"
```

ENVIRONMENT

None

BUGS

None

AUTHOR

IBM Austin

SEE ALSO

lssacfg

NAME

chstat - change statistics collection settings

SYNOPSIS

chstat -r vlan [-m *managed-system*] -s *sample-rate* [--help]

DESCRIPTION

chstat changes the Hardware Management Console (HMC) settings for statistics collection for managed systems.

OPTIONS

- r** The type of resources for which to change the settings for statistics collection. The only valid value is **vlan** for Virtual Local Area Network (VLAN) statistics.
- m** The name of the managed system for which to change the settings for statistics collection. The name may either be the user-defined name for the managed system, or be in the form *ttt-*mmm***ssssss**, where *ttt* is the machine type, *mmm* is the model, and *ssssss* is the serial number of the managed system. The *ttt-*mmm***ssssss** form must be used if there are multiple managed systems with the same user-defined name.

If this option is not specified, then the settings for statistics collection for all of the systems that support statistics collection and are currently managed by this HMC will be changed.

This command only changes the settings for statistics collection for this HMC. Statistics collection by other HMCs for the same managed systems is not affected.
- s** The rate, in seconds, at which to sample the statistics. Sample rates of **0**, **60** (1 minute), and **120** (2 minutes) are supported. A sample rate of **0** disables statistics collection.
- help** Display the help text for this command and exit.

EXAMPLES

Enable VLAN statistics collection with a sample rate of 1 minute for all managed systems that support VLAN statistics collection:

```
chstat -r vlan -s 60
```

Enable VLAN statistics collection with a sample rate of 2 minutes for managed system **mySystem**:

```
chstat -r vlan -s 120 -m mySystem
```

Disable VLAN statistics collection for the managed system with the type, model, and serial number **8233-E8B*7654321**:

```
chstat -r vlan -s 0 -m 8233-E8B*7654321
```

ENVIRONMENT

None

BUGS

None

AUTHOR

IBM Austin

SEE ALSO

chsvc, **lsstat**

NAME

chsvc - change a service

SYNOPSIS

To start or stop monitoring managed systems using SNMP:

```
chsvc -s snmp -o {start | stop} [-m managed-system] [--help]
```

To add or remove a network management system to or from the SNMP trap notification list:

```
chsvc -s snmp -o {add | remove} --ip IP-address  
[--traps trap-list] [--help]
```

DESCRIPTION

chsvc changes the configuration of a service, such as SNMP, on the Hardware Management Console (HMC).

OPTIONS

-s The service for which to change the configuration. The only valid value is **snmp** for Simple Network Management Protocol (SNMP) version 3.

-o The operation to perform.

Specify **start** to start monitoring managed systems on the HMC using SNMP. If the SNMPv3 agent is not already running on the HMC, it will be started. In addition, Virtual Local Area Network (VLAN) statistics collection for the managed systems will be enabled, and the sample rate will be set to 1 minute. The **chstat** command can be used to change the sample rate.

Specify **stop** to stop monitoring managed systems on the HMC using SNMP. Also, VLAN statistics collection for the managed systems will be disabled. In addition, if there are no remaining managed systems on the HMC that are being monitored, the SNMPv3 agent on the HMC will be stopped.

Specify **add** to add a network management system to the SNMP trap notification list.

Specify **remove** to remove a network management system from the SNMP trap notification list.

-m The name of the managed system to start or stop being monitored by the SNMPv3 agent on the HMC. The name may either be the user-defined name for the managed system, or be in the form *tttt-mmm*ssssss*, where *tttt* is the machine type, *mmm* is the model, and *ssssss* is the serial number of the managed system. The *tttt-mmm*ssssss* form must be used if there are multiple managed systems with the same user-defined name.

If this option is not specified for a start operation, then a maximum of five of the systems currently managed by the HMC will be monitored. If this option is not specified for a stop operation, then monitoring for all managed systems will be stopped.

--ip The IP address or host name of the network management system to add to or remove from the SNMP trap notification list.

--traps A list of traps to send to the network management system. If multiple traps are specified, they must be comma separated.

For a list of all of the traps supported by the SNMPv3 agent on the HMC, run the **lssvc -s snmp -r trapnames** command.

This option is required when adding a network management system to the SNMP trap notification list. This option is not valid otherwise.

--help Display the help text for this command and exit.

EXAMPLES

Start the SNMPv3 agent on the HMC if it is not already running, and start monitoring managed system **sys1**:

```
chsvc -s snmp -o start -m sys1
```

Stop monitoring managed system **sys1** and stop the SNMPv3 agent on the HMC if no other managed systems are being monitored:

```
chsvc -s snmp -o stop -m sys1
```

Stop monitoring all managed systems and stop the SNMPv3 agent on the HMC:

```
chsvc -s snmp -o stop
```

Add notification of the network management system with IP address **9.3.99.100** of the **linkUp** and **linkDown** traps:

```
chsvc -s snmp -o add --ip 9.3.99.100 --traps linkUp,linkDown
```

Remove notification of the network management system with IP address **9.3.99.100**:

```
chsvc -s snmp -o remove --ip 9.3.99.100
```

ENVIRONMENT

None

BUGS

None

AUTHOR

IBM Austin

SEE ALSO

chstat, **lsstat**, **lssvc**

NAME

`chsvcevent` - updates serviceable events on the HMC

SYNOPSIS

`chsvcevent -o {close | closeall | approve}`
`[-p problem-number -h analyzing-HMC] [--help]`

DESCRIPTION

`chsvcevent` closes serviceable events on the Hardware Management Console (HMC).

When the Event Manager for Call Home is enabled on the HMC, `chsvcevent` can also be used to approve and initiate call home of a serviceable event.

OPTIONS

-o The operation to perform. Valid values are **close** to close the specified serviceable event, **closeall** to close all serviceable events on this HMC, and **approve** to approve and initiate call home of the specified serviceable event.

-p The problem number of the serviceable event to close or approve. This option is required when closing a single serviceable event or approving a serviceable event. This option is not allowed when closing all serviceable events.

The problem number can be obtained using the `lssvcevents` command.

-h The analyzing HMC for the serviceable event to close or approve. This option is required when closing a single serviceable event or approving a serviceable event. This option is not allowed when closing all serviceable events.

The analyzing HMC can be obtained using the `lssvcevents` command.

--help Display the help text for this command and exit.

EXAMPLES

To close the serviceable event with problem number **1048** and analyzing HMC **7315-C02/KLRMT0L**:

```
chsvcevent -p 1048 -h 7315-C02/KLRMT0L -o close
```

To close all serviceable events on this HMC:

```
chsvcevent -o closeall
```

To approve the serviceable event with problem number **56** and analyzing HMC **myHmc**:

```
chsvcevent -p 56 -h myHMC -o approve
```

ENVIRONMENT

None

BUGS

None

AUTHOR

IBM Austin

SEE ALSO

`lssvcevents`

NAME

chsyscfg - change system resources

SYNOPSIS

To change attributes:

```
chsyscfg -r {lpar | prof | sys | sysprof | frame}
          {-m managed-system | -e managed-frame}
          {-f configuration-file | -i "configuration-data" }
          [--force]
          [--help]
```

To change partition attributes by applying a partition profile:

```
chsyscfg -r lpar -m managed-system
          -o apply {-p partition-name | --id partition-ID}
          [-n profile-name]
          [--help]
```

DESCRIPTION

chsyscfg changes the attributes of partitions, partition profiles, or system profiles for the *managed-system*. It can also change the attributes of the *managed-system*.

chsyscfg can also change the attributes of the *managed-frame*.

USAGE IN A POWERVM MANAGEMENT MASTER MODE ENVIRONMENT

When the *managed-system* is in PowerVM management master mode, both the **chsyscfg -r lpar** and the **chsyscfg -r prof** commands may be used to change the current configuration of a partition. The following attributes are not allowed with the **chsyscfg -r lpar** command: **profile_name** and **sync_curr_profile**. The following attributes are not allowed with the **chsyscfg -r prof** command: **name** and **new_name**. The partition will be configured immediately with the specified attributes.

When running **chsyscfg -r lpar -o apply** the **-n** option is not allowed. The last valid configuration partition profile will always be applied.

The Hardware Management Console (HMC) must be the current PowerVM management master for the *managed-system* when running the **chsyscfg -r lpar** and **chsyscfg -r prof** commands.

The **chsyscfg -r sysprof** command is not allowed.

OPTIONS

-r The type of resources to change. Valid values are **lpar** for partitions, **prof** for partition profiles, **sys** for the managed system, **sysprof** for system profiles, and **frame** for the managed frame.

-m The name of either the managed system to be changed, or the managed system which has the system resources to be changed. The name may either be the user-defined name for the managed system, or be in the form *ttt-mmm*sssssss*, where *ttt* is the machine type, *mmm* is the model, and *sssssss* is the serial number of the managed system. The *ttt-mmm*sssssss* form must be used if there are multiple managed systems with the same user-defined name.

This option is required when changing the attributes of partitions, partition profiles, system profiles, or the managed system. This option is not valid otherwise.

-e The name of the managed frame to be changed. The name may either be the user-defined name for the managed frame, or be in the form *ttt-mmm*sssssss*, where *ttt* is the type, *mmm* is the model, and *sssssss* is the serial number of the managed frame. The *ttt-mmm*sssssss* form must be used if there are multiple managed frames with the same user-defined name.

This option is required when changing the attributes of the managed frame. This option is not valid otherwise.

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

The format of a configuration record is as follows:

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

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

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

When a list of values is specified, the attribute name/value pair must be enclosed in double quotes. Depending on the shell being used, nested double quote characters may need to be preceded by an escape character, which is usually a backslash character.

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

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

Attribute names for partitions:

name | lpar_id

name or ID of the partition to change (required)

new_name

default_profile

ipl_source

IBM i only

Valid values are **a**, **b**, **c**, or **d**

allow_perf_collection

Valid values are:

0 - do not allow

1 - allow

This attribute replaces the

shared_proc_pool_util_auth attribute.

Setting this attribute also sets the

shared_proc_pool_util_auth attribute

to the same value.

shared_proc_pool_util_auth

Valid values are:

0 - do not allow authority

1 - allow authority

This attribute has been deprecated. Use

the **allow_perf_collection** attribute

instead. Setting this attribute also

sets the **allow_perf_collection**

attribute to the same value.

work_group_id

Specify **none** if you are not using Enterprise Workload Manager (EWM)

lpar_avail_priority

Valid values are **0 - 255**.

electronic_err_reporting

IBM i only
Valid values are:

0 - disable

1 - enable

misp

Virtual I/O server only
Valid values are:
0 - the partition is not a mover service partition
1 - the partition is a mover service partition

time_ref

Valid values are:
0 - the partition is not a time reference partition
1 - the partition is a time reference partition

suspend_capable

AIX, Linux, and IBM i only
Valid values are:
0 - the partition cannot be suspended
1 - the partition can be suspended

simplified_remote_restart_capable

AIX, Linux, and IBM i only
Valid values are:
0 - remote restart without using a reserved storage device cannot be performed for the partition
1 - remote restart without using a reserved storage device can be performed for the partition

remote_restart_capable

AIX, Linux, and IBM i only
Valid values are:
0 - remote restart using a reserved storage device cannot be performed for the partition
1 - remote restart using a reserved storage device can be performed for the partition

primary_rs_vios_name | primary_rs_vios_id

Remote restart capable partitions only

secondary_rs_vios_name | secondary_rs_vios_id

Remote restart capable partitions only
To set no secondary reserved storage VIOS specify **none** for **secondary_rs_vios_id**

rs_device_name

Remote restart capable partitions only
Name of the device on the primary reserved storage VIOS partition

vtpm_enabled

AIX, Linux, and VIOS only

Valid values are:

0 - virtual Trusted Platform Module (vTPM) is not enabled for this partition

1 - vTPM is enabled for this partition

os400_restricted_io_mode

IBM i only

Valid values are:

0 - disable IBM i restricted I/O mode

1 - enable IBM i restricted I/O mode

sync_curr_profile

Specifies whether this partition's current configuration is to be synchronized with its current active profile. Valid values are:

0 - disable synchronization

1 - enable synchronization

2 - suspend synchronization until the next time a profile is activated or applied

powervm_mgmt_capable

Linux only

Valid values are:

0 - do not allow this partition to provide PowerVM management functions

1 - enable this partition to provide PowerVM management functions

migration_disabled

AIX, Linux, and IBM i only

Valid values are:

0 - do not disable partition migration for this partition

1 - disable partition migration for this partition

Attribute names for partition profiles:

name

name of the partition profile to change (required)

lpar_name | lpar_id

name or ID of the partition for which the profile to change is defined (required)

new_name

min_mem

megabytes

desired_mem

megabytes

max_mem

megabytes

min_num_huge_pages

AIX and Linux only

desired_num_huge_pages

AIX and Linux only

max_num_huge_pages

AIX and Linux only

mem_mode

Valid values are:

ded - dedicated memory

shared - shared memory

desired_io_entitled_mem

Specify the number of megabytes or specify **auto** for automatic I/O entitled memory management mode

mem_weight**primary_paging_vios_name** | **primary_paging_vios_id****secondary_paging_vios_name** | **secondary_paging_vios_id**

To set no secondary paging VIOS specify **none** for

secondary_paging_vios_id**mem_expansion**

Valid values are:

0 - disable Active Memory Expansion

1.00-10.00 - expansion factor

hpt_ratio

Ratio of hardware page table size to the maximum memory for the partition. Valid values are displayed

by the **lshwres -r mem -m managed-system**

--level sys -F possible_hpt_ratios command.

bsr_arrays**proc_mode**

Valid values are:

ded - dedicated processors

shared - shared processors

min_procs**desired_procs****max_procs****min_proc_units****desired_proc_units****max_proc_units****sharing_mode**

Valid values for partitions using dedicated processors are:

keep_idle_procs - never share processors

share_idle_procs - share processors only when partition is inactive

share_idle_procs_active - share processors only when partition is active

share_idle_procs_always - always share processors

Valid values for partitions using shared processors are:

cap - capped

uncap - uncapped

uncap_weight**shared_proc_pool_name** | **shared_proc_pool_id****affinity_group_id**

Memory and processor affinity group in which the partition will participate. Valid values are

none and **1 - 255**.

io_slots

Comma separated list of I/O slots, with each I/O slot having the following format:

slot-DRC-index[*slot-IO-pool-ID*]/
is-required

Both '/' characters must be present, but optional values may be omitted. Optional values are *slot-IO-pool-ID*.

Valid values for *is-required*:

0 - no

1 - yes

For example:

21010003/3/1 specifies an I/O slot with a DRC index of **21010003**, it is assigned to I/O pool **3**, and it is a required slot.

lpar_io_pool_ids

comma separated

load_source_slot

IBM i only

DRC index of I/O slot, DRC index of SR-IOV logical port, or virtual slot number

alt_restart_device_slot

IBM i only

DRC index of I/O slot, DRC index of HEA or SR-IOV logical port, or virtual slot number

console_slot

IBM i only

DRC index of I/O slot, DRC index of HEA or SR-IOV logical port, virtual slot number, or the value **hmc**

alt_console_slot

IBM i only

DRC index of I/O slot

op_console_slot

IBM i only

DRC index of I/O slot

auto_start

Valid values are:

0 - off

1 - on

boot_mode

AIX, Linux, and virtual I/O server only

Valid values are:

norm - normal

dd - diagnostic with default boot list

ds - diagnostic with stored boot list

of - Open Firmware OK prompt

sms - System Management Services

power_ctrl_lpar_ids | power_ctrl_lpar_names

comma separated

conn_monitoring

Valid values are:

0 - off

1 - on

hsl_pool_id

IBM i only

Valid values are:

0 - HSL OptiConnect is disabled

1 - HSL OptiConnect is enabled

virtual_opti_pool_id

IBM i only

Valid values are:

0 - virtual OptiConnect is disabled

1 - virtual OptiConnect is enabled

max_virtual_slots**virtual_eth_adapters**

Comma separated list of virtual ethernet adapters, with each adapter having the following format:

```
virtual-slot-number/is-IEEE/port-vlan-ID/  
[additional-vlan-IDs]/[trunk-priority]/  
is-required[/[virtual-switch]/[MAC-address]/  
[allowed-OS-MAC-addresses]/[QoS-priority]]
```

The first 5 '/' characters must be present. The 6th '/' character is optional, but it must be present if *virtual-switch* or any of the values following *virtual-switch* are specified. The last 3 '/' characters are optional, but all 3 must be present if *MAC-address*, *allowed-OS-MAC-addresses*, or *QoS-priority* is specified.

Optional values may be omitted. Optional values are *additional-vlan-IDs*, *trunk-priority*, *virtual-switch*, *MAC-address*, *allowed-OS-MAC-addresses*, and *QoS-priority*.

Valid values for *is-IEEE* and *is-required*:

0 - no

1 - yes

If values are specified for *additional-vlan-IDs*, they must be comma separated.

Valid values for *trunk-priority*:

0 - this adapter is not a trunk adapter (default value)

1 - 15 - this adapter is a trunk adapter with the specified priority

If *MAC-address* is not specified, a unique MAC address will be automatically generated for the adapter. If it is specified, it must be specified as 12 hexadecimal characters. It is highly recommended that

you do not specify *MAC-address* so that the MAC address will be automatically generated.

allowed-OS-MAC-addresses can be a comma separated list of 1 to 4 MAC addresses, where each MAC address is specified as 12 hexadecimal characters, or it can be one of the following values:

all - all OS defined MAC addresses are allowed

(default value)

none - no OS defined MAC addresses are allowed

Valid values for *QoS-priority*:

none - do not use Quality of Service (QoS) (default value)

0 - 7 - the QoS priority level to use

For example:

13/1/5/6,7/2/1

specifies a virtual ethernet adapter with a virtual slot number of **13**, it is IEEE 802.1Q compatible, it has a port virtual LAN ID of **5**, additional virtual LAN IDs of **6** and **7**, it is a trunk adapter with a trunk priority of **2**, and it is required. It also uses the default virtual switch, it uses an automatically generated MAC address, it allows all OS defined MAC addresses, and it does not use QoS.

4/1/3//0/1///fecd537a910b,fecd537a910c/2

specifies a virtual ethernet adapter with a virtual slot number of **4**, it is IEEE 802.1Q compatible, it has a port virtual LAN ID of **3**, no additional virtual LAN IDs, it is not a trunk adapter, it is required, it uses the default virtual switch, it uses an automatically generated MAC address, it allows the OS defined MAC addresses FE:CD:53:7A:91:0B and FE:CD:53:7A:91:0C, and it uses a QoS priority level of **2**.

virtual_eth_vsi_profiles

Comma separated list of virtual ethernet adapter Virtual Station Interface (VSI) profile attributes:

virtual-slot-number/VSI-manager-ID/VSI-type-ID/VSI-type-version

All 3 '/' characters must be present, and all values are required.

For example:

15/2/1193040/1

specifies a VSI profile for the virtual ethernet adapter in slot **15**. The VSI profile has a VSI manager ID of **2**, a VSI type ID of **1193040**, and a VSI type version of **1**.

virtual_fc_adapters

Comma separated list of virtual fibre channel adapters, with each adapter having the following format:

```
virtual-slot-number/client-or-server/
[remote-lpar-ID]/[remote-lpar-name]/
remote-slot-number/[wwpns]/is-required
```

All 6 '/' characters must be present, but optional values may be omitted. Optional values are *remote-lpar-ID* or *remote-lpar-name* (one of those values is required, but not both).

wwpns is optional for a client adapter, and is not allowed for a server adapter. If *wwpns* is not specified, WWPNs will be automatically generated for a client adapter. It is highly recommended that you do not specify *wwpns* so that WWPNs will be automatically generated, unless you are changing an existing virtual fibre channel adapter and you want to retain its existing WWPNs. If you do specify *wwpns*, you must specify exactly two WWPNs, and they must be comma separated.

Valid values for *client-or-server*:

client

server

Valid values for *is-required*:

0 - no

1 - yes

For example:

4/client/vios_p1/16/1

specifies a virtual fibre channel **client** adapter with a virtual slot number of **4**, a remote (server) partition name of **vios_p1**, a remote (server) slot number of **16**, and it is required. WWPNs will be automatically generated for this client adapter.

virtual_scsi_adapters

Comma separated list of virtual SCSI adapters, with each adapter having the following format:

```
virtual-slot-number/client-or-server/
[remote-lpar-ID]/[remote-lpar-name]/
[remote-slot-number]/is-required
```

All 5 '/' characters must be present, but optional values may be omitted. Optional values for server adapters are *remote-lpar-ID*, *remote-lpar-name*, and *remote-slot-number*. Optional values for client adapters are *remote-lpar-ID* or

remote-lpar-name (one of those values is required, but not both).

Valid values for *client-or-server*:

client
server

Valid values for *is-required*:

0 - no
1 - yes

For example:

14/client/2//13/0

specifies a virtual SCSI **client** adapter with a virtual slot number of **14**, a remote (server) partition ID of **2**, a remote (server) slot number of **13**, and it is not required.

virtual_serial_adapters

Comma separated list of virtual serial adapters, with each adapter having the following format:

```
virtual-slot-number/client-or-server/
[supports-HMC]/[remote-lpar-ID]/
[remote-lpar-name]/[remote-slot-number]/
is-required
```

All 6 '/' characters must be present, but optional values may be omitted. Optional values for server adapters are *supports-HMC*, *remote-lpar-ID*, *remote-lpar-name*, and *remote-slot-number*. Optional values for client adapters are *remote-lpar-ID* or *remote-lpar-name* (one of those values is required, but not both), and the *supports-HMC* value is not allowed.

Valid values for *client-or-server*:

client
server

Valid values for *supports-HMC*:

0 - no

Valid values for *is-required*:

0 - no
1 - yes

For example:

14/server/0///0

specifies a virtual serial **server** adapter with a virtual slot number of **14**, it does

not support an HMC connection, any client adapter is allowed to connect to it, and it is not required.

vnic_adapters

AIX, Linux, and IBM i only

Comma separated list of virtual NICs, with each virtual NIC consisting of a list of colon separated property name/value pairs, which can be specified in any order:

property-name=value:property-name=value:...

Property names:

slot_num

Required

port_vlan_id

Port VLAN ID or **0** to disable VLAN tag insertions for untagged frames (default value)

pvid_priority

Valid values:

An integer between **0** and **7**, inclusive.

Default value is **0**.

allowed_vlan_ids

Comma separated list of VLAN IDs, or one of the following values:

all - all VLAN IDs are allowed (default value)

none - no VLAN IDs are allowed

mac_addr

12 hexadecimal characters

If not specified, a unique MAC address will be automatically generated for the virtual NIC.

It is highly recommended that you use an automatically generated MAC address.

allowed_os_mac_addr

Comma separated list of MAC addresses, each specified as 12 hexadecimal characters. Other valid values:

all - all OS defined MAC addresses are allowed (default value)

none - no OS defined MAC addresses are allowed

is_required

Valid values:

0 - no

1 - yes (default value)

auto_priority_failover

Valid values:

0 - automatic priority failover is disabled

1 - automatic priority failover is enabled (default value)

backing_devices

Required

Comma separated list of virtual NIC backing devices, with each backing device having the

following format:

```
sriov/vios-lpar-name/vios-lpar-ID/  
sriov-adapter-ID/sriov-physical-port-ID/  
[capacity][/failover-priority]
```

The first 5 '/' characters must be present. The 6th '/' character is optional, but it must be present if *failover-priority* is specified.

Optional values may be omitted. Optional values are *capacity*, either *vios-lpar-name* or *vios-lpar-ID*, and *failover-priority*.

Specify *capacity* as a percentage with up to 2 decimal places.

Specify *failover-priority* as a value between **1** and **100** inclusive, with **1** being the highest priority. It defaults to **50** if not specified.

For example:

slot_num=5:backing_devices=sriov/vios1//1/0/1,sriov/vios2//2/0/10 specifies a virtual NIC with a virtual slot number of **5** and 2 backing devices. The first backing device is on VIOS **vios1** and physical port **0** of SR-IOV adapter **1** and has a failover priority of **1**. The second backing device is on VIOS **vios2** and physical port **0** of SR-IOV adapter **2** and has a failover priority of **10**. All the other virtual NIC properties are configured with default values.

hca_adapters

AIX, Linux, and virtual I/O server only
Comma separated list of Host Channel adapters (HCA), with each adapter having the following format:

```
adapter-ID/GUID/capability
```

All 3 values must be specified for each adapter.

Valid values for *capability*:

- 1** - low
- 2** - medium
- 3** - high
- 4** - dedicated

For example:

23000cff/2550000000609/3 specifies an HCA with an adapter ID of **23000cff**, a Globally Unique ID (GUID) of **2:55:00:00:00:06:09**, and a capability setting of high.

lhea_logical_ports

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

```
adapter-ID/port-group/physical-port-ID/
logical-port-ID[allowed-VLAN-IDs][/
allowed-OS-MAC-addresses]
```

The first 4 '/' characters must be present. The last '/' character is optional, but it must be present if *allowed-OS-MAC-addresses* is specified. Optional values may be omitted. Optional values are *allowed-VLAN-IDs* and *allowed-OS-MAC-addresses*.

If values are specified for *allowed-VLAN-IDs*, they must be comma separated, or one of the following values can be specified:

all - allowed to participate in all VLANs
none - not allowed to participate in any VLAN
 (default value)

allowed-OS-MAC-addresses can be a comma separated list of 1 to 4 MAC addresses, where each MAC address is specified as 12 hexadecimal characters, or it can be one of the following values:

all - all OS defined MAC addresses are allowed
 (default value)
none - no OS defined MAC addresses are allowed

For example:

23000000/2/0/1/3 specifies logical port **1** for physical port **0** belonging to port group **2** of the Host Ethernet adapter (HEA) with an adapter ID of **23000000**. This logical port is allowed to participate in the VLAN with an ID of **3** only. It also allows all OS defined MAC addresses.

23000000/2/1/5/all/fecd537a910b,fecd537a910c specifies logical port **5** for physical port **1** belonging to port group **2** of the HEA with an adapter ID of **23000000**. This logical port is allowed to participate in all VLANs, and it allows the OS defined MAC addresses FE:CD:53:7A:91:0B and FE:CD:53:7A:91:0C.

lhea_capabilities

Comma separated list of LHEA capabilities, with each capability having one of the following formats:

```
adapter-ID/capability
```

or

```
adapter-ID/5/ieq/nielq/qplcq/mr
```

where *ieq* (interruptible event queues), *nieq* (non-interruptible event queues), *qp* (queue pairs), *cq* (completion queues), and *mr* (memory regions) each specify the resource amount in addition to the base minimum.

Valid values for *capability*:

- 0** - base minimum
- 1** - low
- 2** - medium
- 3** - high
- 4** - dedicated

For example:

23000000/3 sets the LHEA capability for the HEA with an adapter ID of **23000000** to high.

sni_device_ids

AIX, Linux, and virtual I/O server only
Comma separated list of Switch Network Interface (SNI) adapter device IDs

work_group_id

Specify **none** if you are not using Enterprise Workload Manager (EWM)

redundant_err_path_reporting

Valid values are:

- 0** - disable
- 1** - enable

electronic_err_reporting

IBM i only
Valid values are:

- 0** - disable
- 1** - enable

lpar_proc_compat_mode

The valid values for the *managed-system* are displayed by the **lssyscfg -r sys -m managed-system -F lpar_proc_compat_modes** command.

sriov_eth_logical_ports

Comma separated list of SR-IOV ethernet logical ports, with each logical port consisting of a list of colon separated property name/value pairs, which can be specified in any order:

property-name=value:property-name=value:...

Property names:

adapter_id

Required when using the = or += operator

phys_port_id

Required when using the = or += operator

allowed_os_mac_addr

Comma separated list of MAC addresses, each

specified as 12 hexadecimal characters. Other valid values:

all - all OS defined MAC addresses are allowed (default value)

none - no OS defined MAC addresses are allowed

allowed_priorities

Comma separated list of integers between **0** and **7**, inclusive, or one of the following values:

all - all supported priorities are allowed

none - no priorities are allowed (default value)

allowed_vlan_ids

Comma separated list of VLAN IDs, or one of the following values:

all - all VLAN IDs are allowed (default value)

none - no VLAN IDs are allowed

capacity

Must be a multiple of **min_eth_capacity_granularity** of the physical port, which is displayed by the **lshwres -m managed-system -r sriov --subtype physport --level {eth | ethc} -F min_eth_capacity_granularity** command. The minimum value and the default value is the value of **min_eth_capacity_granularity**, the maximum value is **100**.

config_id

If not specified, the management console will assign the next available value.

Required when using the **-=** operator

diag_mode

Valid values:

0 - disable (default value)

1 - enable

Warning! If diagnostics mode is enabled, no other logical port may be attached to the physical port. If logical ports are already in use on the physical port, you will have to power off partitions or use Dynamic Logical Partitioning (DLPAR) to remove the logical ports that are attached to the physical port.

huge_dma_window_mode

Valid values:

0 - disable (default value)

1 - enable

logical_port_id

DRC index of an SR-IOV logical port. If not specified, an SR-IOV logical port will be automatically assigned.

mac_addr

12 hexadecimal characters

If not specified, a unique MAC address will be automatically generated for the logical port.

It is highly recommended that you use an automatically generated MAC address.

port_vlan_id

Port VLAN ID or **0** to disable VLAN tag insertions for untagged frames (default value)

promisc_mode

Valid values:

0 - disable (default value)

1 - enable

pvid_priority

Valid values:

An integer between **0** and **7**, inclusive.

Default value is **0**.

For example:

adapter_id=1:phys_port_id=0 specifies that an unconfigured SR-IOV logical port on physical port **0** of the SR-IOV adapter with the adapter ID of **1** will be selected and configured with default ethernet logical port values.

Attribute names for a managed system:

new_name**power_off_policy**

Valid values are:

0 - power off after all partitions are shut down

1 - do not power off after all partitions are shut down

service_lpar_id

Specify **none** to remove the service partition assignment

service_lpar_name**power_on_lpar_start_policy**

Valid values are **autostart**, **userinit**, or **autorecovery**.

This attribute replaces the **power_on_option** attribute.

Setting this attribute also sets the **power_on_option** attribute to the equivalent value.

power_on_option

Valid values are **autostart** or **standby**.

This attribute has been deprecated. Use the

power_on_lpar_start_policy attribute instead. Setting this attribute also sets the

power_on_lpar_start_policy attribute to the equivalent value.

power_on_side

Valid values are **perm** or **temp**

power_on_speed

Valid values are **slow** or **fast**.

This attribute can be set for POWER6 and

POWER7 (firmware level 710 only) servers only.

power_on_speed_override

Valid values are **none**, **slow**, or **fast**.

This attribute can be set for POWER6 and

POWER7 (firmware level 710 only) servers only.

sp_failover_enabled

Valid values are:

0 - disable service processor failover

1 - enable service processor failover

addr_broadcast_perf_policy

Valid values are:

chip_affinity

no_affinity

node_affinity

max_lpars_per_hca

The managed system must be restarted for this change to take effect.

hca_bandwidth_capabilities

Host Channel adapter (HCA) bandwidth capability settings. Specify *highmediumlow*, where *high*, *medium*, and *low* are each expressed as a percentage of the total bandwidth, or specify **default** to set all 3 settings to their default values.

If *highmediumlow* is specified, both *'/'* characters must be present, but any setting may be omitted. A setting that is omitted is not changed.

The default value for each setting is:

high - 25%

medium - 12.5%

low - 6.25%

New HCA bandwidth capability settings are applied only when partitions are activated.

For example:

50/25.5/10.25 specifies a high setting of **50%**, a medium setting of **25.5%**, and a low setting of **10.25%**.

33.3// specifies a high setting of **33.3%**, and does not change the medium and low settings.

turbocore

Valid values are:

0 - disable TurboCore mode

1 - enable TurboCore mode

The managed system must be restarted for this change to take effect.

sys_time

Managed system date and time in the following format:

MM/DD/YYYY hh:mm:ss

where *MM* is the month, *DD* is the day, *YYYY* is the year, *hh* is the hour in 24 hour format, *mm* is the minute, and *ss* is the seconds.

This attribute can be set only when the managed system is powered off.

mfg_default_boot_mode

Valid values are:

norm - normal

dd - diagnostic with default boot list

ds - diagnostic with stored boot list

of - Open Firmware OK prompt

sms - System Management Services

mfg_default_ipl_source

Valid values are **a**, **b**, **c**, or **d**

Attribute names for system profiles:

name

name of the system profile to change (required)

new_name

lpar_names | lpar_ids

comma separated

profile_names

comma separated

Attribute names for a managed frame:

new_name

frame_num

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

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

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

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

- o** The operation to perform. The only valid value is **apply** to apply a partition profile to a shutdown partition. Applying a partition profile is equivalent to activating a partition with that partition profile except the partition is not powered on.
- p** The name of the partition to which to apply the partition profile.

You can either use this option to specify the name of the partition, or use the **--id** option to specify the partition's ID. The **-p** and the **--id** options are mutually exclusive.

- id** The ID of the partition to which to apply the partition profile.

You can either use this option to specify the ID of the partition, or use the **-p** option to specify the partition's name. The **--id** and the **-p** options are mutually exclusive.

- n** The name of the partition profile to apply. If this option is omitted, the last valid configuration partition profile will be applied.
- force** When changing a partition profile, specify this option to allow a partition profile that is synchronized with the partition's current configuration to be changed. After the profile is changed, synchronization of a profile with the partition's current configuration will be suspended until the next time the partition is activated with a profile or a profile is applied to the partition.

When disabling the remote restart capability of a partition, specify this option to allow the capability to be disabled when the reserved storage device assigned to the partition cannot be read. Note that this will cause some partition settings, such as the partition's boot or IPL settings and the partition's time of day, to be reset to their default values.

--help Display the help text for this command and exit.

EXAMPLES

Change the managed system's user-defined name and power off policy:

```
chsyscfg -r sys -m 9406-570*89320051 -i "new_name=sys1,
power_off_policy=1"
```

Change partitions using the configuration data in the file **/tmp/lparfile**:

```
chsyscfg -r lpar -m sys1 -f /tmp/lparfile
```

Change the attributes of a shutdown partition by applying the profile **prof1**:

```
chsyscfg -r lpar -m system1 -o apply -p p1 -n prof1
```

Change a partition profile's memory amounts (reduce the profile's current memory amounts each by 256 MB), and number of desired processors:

```
chsyscfg -r prof -m sys1 -i "name=profile1,
lpar_name=partition3,min_mem-=256,desired_mem-=256,
max_mem-=256,desired_procs=2"
```

Change a system profile (add 2 new partition profiles):

```
chsyscfg -r sysprof -m sys1 -i "name=sysprof1,
'lpar_names+=partition3,partition4',
'profile_names+=3_prof1,4_defaultProf'"
```

Change the managed frame's user-defined name and frame number:

```
chsyscfg -r frame -e 9119-59*1234567 -i "new_name=frame1,
frame_num=1"
```

ENVIRONMENT

None

BUGS

None

AUTHOR

IBM Austin

SEE ALSO

lssyscfg, mksyscfg, rmsyscfg, chhwres

NAME

chsyspwd - change system password

SYNOPSIS

```
chsyspwd -t {access | admin | general}
          {-m managed-system | -e managed-frame}
          [--passwd current-password] [--newpasswd new-password]
          [--help]
```

DESCRIPTION

chsyspwd changes a password for the *managed-system* or the *managed-frame*.

chsyspwd can also be used to enter the current password for the *managed-system* or the *managed-frame* when it is in Failed Authentication state.

OPTIONS

-t The password to change. Valid values are **access** for the Hardware Management Console (HMC) Access password, **admin** for the Advanced System Management (ASM) Administrator password, and **general** for the ASM General password.

-m The name of the managed system for which to change the password. The name may either be the user-defined name for the managed system, or be in the form *ttt-mmm*sssssss*, where *ttt* is the machine type, *mmm* is the model, and *sssssss* is the serial number of the managed system. The *ttt-mmm*sssssss* form must be used if there are multiple managed systems with the same user-defined name.

The **-m** and the **-e** options are mutually exclusive.

-e The name of the managed frame for which to change the password. The name may either be the user-defined name for the managed frame, or be in the form *ttt-mmm*sssssss*, where *ttt* is the type, *mmm* is the model, and *sssssss* is the serial number of the managed frame. The *ttt-mmm*sssssss* form must be used if there are multiple managed frames with the same user-defined name.

The **-e** and the **-m** options are mutually exclusive.

--passwd

The current password. If this option is omitted, you will be prompted to enter the current password.

--newpasswd

The new password. If this option is omitted, you will be prompted to enter the new password.

The new password must be at least 4 characters in length, and cannot be longer than 63 characters.

--help Display the help text for this command and exit.

EXAMPLES

Change the HMC Access password for the managed system **system1**:

```
chsyspwd -m system1 -t access --passwd sys1pw --newpasswd
sys1newpw
```

Change the ASM Administrator password (the current and new passwords must be entered when prompted) for the managed system **9406-520*10013CA**:

```
chsyspwd -m 9406-520*10013CA -t admin
```

Change the HMC Access password for the managed frame **frame1**:

```
chsyspwd -e frame1 -t access --passwd frame1pw --newpasswd  
frame1newpw
```

Enter the current password for the managed system **system1**, which is in Failed Authentication state:

```
chsyspwd -m system1 -t access --newpasswd sys1pw
```

ENVIRONMENT

None

BUGS

None

AUTHOR

IBM Austin

NAME

chsysstate - change partition state or system state

SYNOPSIS

To power on a managed system:

```
chsysstate {-m managed-systems | --all} -r sys  
-o {on | onstandby | onstartpolicy | onsysprof |  
onhwdisc}  
[-f system-profile-name] [-k keylock-position]
```

To power off a managed system:

```
chsysstate {-m managed-systems | --all} -r sys  
-o off [--immed]
```

To restart a managed system:

```
chsysstate {-m managed-systems | --all} -r sys  
-o off --immed --restart
```

To rebuild a managed system or a managed frame:

```
chsysstate {-m managed-systems | -e managed-frame | --all}  
-r {sys | frame} -o rebuild
```

To recover partition data for a managed system:

```
chsysstate -m managed-system -r sys -o recover
```

To initiate service processor failover for a managed system:

```
chsysstate -m managed-system -r sys -o spfailover
```

To set the keylock position for a managed system or a partition:

```
chsysstate -m managed-system -r {sys | lpar}  
-o chkey -k keylock-position  
[-n partition-name | --id partition-ID]
```

To activate a partition with a profile:

```
chsysstate -m managed-system -r lpar -o on  
{-n partition-name | --id partition-ID}  
-f partition-profile-name  
[-k keylock-position]  
[-b boot-mode] [-i IPL-source]  
[--novsi] [--force]
```

To activate a partition with its current configuration:

```
chsysstate -m managed-system -r lpar -o on  
{-n partition-names | --id partition-IDs | --all}  
[-b boot-mode]
```

To activate and perform a network install of an IBM i partition:

```
chsysstate -m managed-system -r lpar -o on  
{-n partition-name | --id partition-ID}  
-f partition-profile-name  
[-k keylock-position] [-i IPL-source]  
--ip IP-address [--netmask network-mask]  
--gateway gateway-IP-address  
--serverip server-IP-address  
--serverdir server-directory
```

```
[--speed {auto | 1 | 10 | 100 | 1000}]
[--duplex {auto | half | full}]
[--mtu {1500 | 9000}]
[--vlan VLAN-ID]
```

To shut down or restart a partition (it is recommended that the **chlparstate** command be used to shut down a suspended partition):

```
chsysstate -m managed-system -r lpar
-o {shutdown | osshutdown | dumprestart |
retrydump}
{-n partition-name | --id partition-ID | --all |
--allnonvios}
[--immed] [--restart] [--force]
```

To perform an operator panel service function on a partition:

```
chsysstate -m managed-system -r lpar
-o {dston | remotestoff | remoteston |
consoleservice | iopreset | iopdump}
{-n partition-name | --id partition-ID}
```

To validate or activate a system profile:

```
chsysstate -m managed-system -r sysprof
-n system-profile-name
[-o on] [--continue] [--test]
```

To power off all of the unowned I/O units in a managed frame:

```
chsysstate -e managed-frame -r frame -o unownediooff
```

DESCRIPTION

chsysstate changes the state of a partition, the *managed-system*, or the *managed-frame*.

See the **chlparstate** command for additional partition state change operations.

USAGE IN A POWERVM MANAGEMENT MASTER MODE ENVIRONMENT

When the *managed-system* is in PowerVM management master mode, the **-f** option is not allowed when activating a partition or performing a network install of an IBM i partition. For both of those operations partitions will be activated with their current configurations. The **chsyscfg -r lpar** and **chsyscfg -r prof** commands can be used to modify the current configuration of a partition before activating it.

The Hardware Management Console (HMC) must be the current PowerVM management master for the *managed-system* when running the **chsysstate -r lpar** command.

The **chsysstate -r sysprof** and **chsysstate -r sys -o onsysprof** commands are not allowed.

OPTIONS

-m The name of the managed system on which to perform the operation. The name may either be the user-defined name for the managed system, or be in the form *tttt-mmm*sssssss*, where *tttt* is the machine type, *mmm* is the model, and *sssssss* is the serial number of the managed system. The *tttt-mmm*sssssss* form must be used if there are multiple managed systems with the same user-defined name.

When powering on managed systems (**-o on** or **-o onstandby**), powering off or restarting managed systems (**-o off**), or rebuilding managed systems (**-o rebuild**), multiple managed system names can be specified with this option, or the **--all** option can be specified to perform the operation on all managed systems. Multiple managed system names must be comma separated.

This option is required when performing a partition or a system profile operation. When performing a managed system operation, either this option or the **--all** option is required. This option is not valid for any other operations.

The **-m** and **--all** options are mutually exclusive for managed system operations.

- e** The name of the managed frame on which to perform the operation. The name may either be the user-defined name for the managed frame, or be in the form *tttt-mmm*sssssss*, where *tttt* is the type, *mmm* is the model, and *sssssss* is the serial number of the managed frame. The *tttt-mmm*sssssss* form must be used if there are multiple managed frames with the same user-defined name.

This option is required when performing a managed frame operation. This option is not valid otherwise.

- r** The type of resource on which to perform the operation. Valid values are **lpar** for partition, **sys** for managed system, **sysprof** for system profile, and **frame** for managed frame.
- o** The operation to perform. Valid values are:
 - on** - activates a partition or a system profile, or powers on the *managed-system*. When powering on the *managed-system*, partitions that are marked as auto start are activated.
 - onstandby** - powers on the *managed-system* to Standby state.
 - onstartpolicy** - powers on the *managed-system* as defined by the next partition start policy.
 - onsysprof** - powers on the *managed-system* then activates a system profile. Only those partitions in the system profile are activated.
 - onhwdisc** - powers on the *managed-system* and runs the hardware discovery process. When the hardware discovery process completes, the *managed-system* will be in Operating state, but no partitions will be automatically started.
 - off** - powers off the *managed-system*. If the **--immed** option is specified, a fast power off (operator panel function 8) is performed, otherwise a normal power off is performed. If both the **--immed** and the **--restart** options are specified, a restart (operator panel function 3) of the *managed-system* is performed.
 - rebuild** - rebuilds the *managed-system* or the *managed-frame*.
 - recover** - recovers partition data for the *managed-system* by restoring the data from the backup file on the HMC.
 - spfailover** - initiates service processor failover for the *managed-system*.
 - chkey** - sets the keylock position for a partition or the *managed-system*.
 - shutdown** - shuts down a partition. If the **--immed** option but not the **--restart**

option is specified, an immediate shutdown (operator panel function 8) is performed. If the **--restart** option but not the **--immed** option is specified, a **dumpprestart** operation (operator panel function 22) is performed. If both the **--immed** and the **--restart** options are specified, an immediate restart (operator panel function 3) of the partition is performed. If neither the **--immed** nor the **--restart** option is specified, a delayed shutdown is performed.

It is recommended that the **chlpstate** command be used to shut down a suspended partition.

- osshutdown** - issues the operating system shutdown command to shut down a partition. If the **--immed** option is specified, the immediate option is included on the operating system shutdown command to immediately shut down the partition. If the **--restart** option is specified, the restart option is included on the operating system shutdown command to restart the partition. Normal operating system shutdown IBM i, with or without the **--restart** option, will take longer to complete.
- dumpprestart** - initiates a dump on the partition and restarts the partition when the dump is complete (operator panel function 22).
- retrydump** - retries the dump on the partition and restarts the partition when the dump is complete (operator panel function 34). This operation is valid for IBM i partitions only.
- dston** - activates dedicated service tools for the partition (operator panel function 21). This operation is valid for IBM i partitions only.
- remotedstoffs** - disables a remote service session for the partition (operator panel function 65). This operation is valid for IBM i partitions only.
- remotedston** - enables a remote service session for the partition (operator panel function 66). This operation is valid for IBM i partitions only.
- consoleservice** - enables console service functions for the partition (operator panel function 65 followed by 21). This operation is valid for IBM i partitions only.
- iopreset** - resets or reloads the failed IOP (operator panel function 67). This operation is valid for IBM i partitions

only.

iopdump - allows use of the IOP control storage dump (operator panel function 70). This operation is valid for IBM i partitions only.

unownediooff - powers off all of the unowned I/O units in a managed frame.

- f** When activating a partition with a profile, use this option to specify the name of the partition profile to use.

If the *managed-system* is in the manufacturing default configuration and this option is not specified, the partition will be activated with no configuration changes and the *managed-system* will remain in the manufacturing default configuration. If the *managed-system* is in the manufacturing default configuration and this option is specified, the partition will be activated with the configuration specified by *partition-profile-name* and the manufacturing default configuration will be permanently exited.

When powering on a managed system with a system profile, use this option to specify the name of the system profile to use. This option is required.

- k** The keylock position to set. Valid values are **manual** and **norm** for normal.

This option is required when setting the keylock position for a partition or a managed system. This option is optional when powering on a managed system or activating a partition.

--immed

If this option is specified when powering off a managed system, a fast power off is performed.

This option must be specified when restarting a managed system.

If this option is specified when shutting down or restarting a partition, an immediate shutdown or restart is performed.

--restart

If this option is specified, the partition or managed system will be restarted.

- n** When performing a system profile operation, use this option to specify the name of the system profile on which to perform the operation.

When performing a partition operation, use this option to specify the name of the partition on which to perform the operation, the **--id** option to specify the partition's ID, the **--all** option to perform the operation on all partitions, or the **--allnonvios** option to perform the operation on all partitions except Virtual I/O Server (VIOS) partitions. The **--all** option is only allowed when activating partitions with their current configurations (**-o on**), or when shutting down or restarting partitions (**-o shutdown**). The **--allnonvios** option is only allowed when shutting down or restarting partitions (**-o shutdown**).

When activating partitions with their current configurations, multiple partition names can be specified with this option. The partition names must be comma separated.

The **-n**, **--id**, **--all**, and **--allnonvios** options are mutually exclusive for partition operations.

- id** When performing a partition operation, use this option to specify the ID of the partition on which to perform the operation, the **-n** option to specify the partition's name, the **--all** option to perform the operation on all partitions, or the **--allnonvios** option to perform the operation on all partitions except VIOS partitions. The **--all** option is only allowed when activating partitions with their current configurations (**-o on**), or when shutting down or restarting partitions (**-o shutdown**). The

--allnonvios option is only allowed when shutting down or restarting partitions (**-o shutdown**).

When activating partitions with their current configurations, multiple partition IDs can be specified with this option. The partition IDs must be comma separated.

The **--id**, **-n**, **--all**, and **--allnonvios** options are mutually exclusive for partition operations.

--all When performing a managed system operation, specify this option to perform the operation on all of the systems currently managed by this HMC. This option is only allowed for the following managed system operations: power on (**-o on** or **-o onstandby**), power off or restart (**-o off**), or rebuild (**-o rebuild**).

When performing a partition operation, specify this option to perform the operation on all of the partitions in the *managed-system*. This option is only allowed for the following partition operations: activate with current configuration (**-o on**), or shut down or restart (**-o shutdown**).

--allnonvios

When performing a partition shutdown or restart operation (**-o shutdown**), specify this option to perform the operation on all of the partitions in the *managed-system* except VIOS partitions.

-b The boot mode to use when activating an AIX, Linux, or Virtual I/O Server partition. Valid values are **norm** for normal, **dd** for diagnostic with default boot list, **ds** for diagnostic with stored boot list, **of** for Open Firmware OK prompt, or **sms** for System Management Services.

-i The IPL source to use when activating an IBM i partition. Valid values are **a**, **b**, **c**, or **d**.

--ip The IP address of the Ethernet adapter, HEA logical port, or SR-IOV logical port assigned as the alternate restart device of the IBM i partition for which to perform the network install operation.

IP-address can be either an IPv4 or an IPv6 address. If an IPv6 address is specified, it must be in the following format: *IPv6-address/prefix-length*.

--netmask

If the IP address specified with the **--ip** option is an IPv4 address, this option is required and is used to specify the network mask.

If the IP address specified with the **--ip** option is an IPv6 address, this option is not valid.

--gateway

The IP address of the gateway to use when performing a network install of an IBM i partition. Either an IPv4 or an IPv6 address can be specified.

--serverip

The IP address of the server that contains the network install images for an IBM i partition. Either an IPv4 or an IPv6 address can be specified.

--serverdir

The directory on the server that contains the network install images for an IBM i partition.

--speed The speed setting to use when performing a network install of an IBM i partition. Valid values are:

- auto** - automatically detect and set speed
- 1** - 1 Mbps
- 10** - 10 Mbps
- 100** - 100 Mbps
- 1000** - 1000 Mbps

If this option is not specified, it defaults to **auto**.

--duplex

The duplex setting to use when performing a network install of an IBM i partition. Valid values are:

auto - automatically detect and set duplex

half - half duplex

full - full duplex

If this option is not specified, it defaults to **auto**.

--mtu The maximum transmission unit, in bytes, to use when performing a network install of an IBM i partition. Valid values are **1500** or **9000** bytes. If this option is not specified, it defaults to **1500** bytes.

--vlan The VLAN ID to use when performing a network install of an IBM i partition. Ethernet packets will be tagged with the specified VLAN ID. If this option is not specified, Ethernet packets will be untagged.

--novsi When activating a partition that uses virtual ethernet adapter Virtual Station Interface (VSI) profiles, use this option to allow the partition to be activated without VSI profiles. Note that when this option is specified, the partition will still be activated with VSI profiles if possible.

--force This option allows you to shut down a Virtual I/O Server (VIOS) partition which is the only VIOS partition providing access to the paging device or reserved storage device for one or more partitions.

This option also allows you to activate a shared memory partition under any of the following conditions:

- the partition is configured to use redundant paging VIOS partitions, but only one paging VIOS partition is currently available and that paging VIOS partition has access to an available paging space device in the shared memory pool,

- the partition is configured to use redundant paging VIOS partitions, but there is no available paging space device in the shared memory pool that can be accessed by both paging VIOS partitions. However, there is an available paging space device in the shared memory pool that can be accessed by one of the paging VIOS partitions,

- the partition is configured to use a single paging VIOS partition which is not currently available or does not have access to an available paging space device in the shared memory pool, but there is another paging VIOS partition in the shared memory pool that is available and that has access to an available paging space device in the shared memory pool.

--test If this option is specified when performing a system profile operation, the system profile is validated.

--continue

If this option is specified when activating a system profile, remaining partitions will continue to be activated after a partition activation failure occurs.

--help Display the help text for this command and exit.

EXIT STATUS

This command has the following return codes:

0 Success

100 Partial success

This value is returned when performing an operation on multiple managed systems or multiple partitions and at least one operation succeeded, and at least one failed.

Any other value means the command failed.

EXAMPLES

Power on a managed system and auto start partitions:

```
chsysstate -m 9406-520*10110CA -r sys -o on
```

Power on a managed system with a system profile:

```
chsysstate -m sys1 -r sys -o onsysprof -f mySysProf
```

Power off a managed system normally:

```
chsysstate -m sys1 -r sys -o off
```

Power off a managed system fast:

```
chsysstate -m sys1 -r sys -o off --immed
```

Restart a managed system:

```
chsysstate -m 9406-570*12345678 -r sys -o off --immed  
--restart
```

Rebuild a managed system:

```
chsysstate -m 9406-570*12345678 -r sys -o rebuild
```

Recover partition data for a managed system:

```
chsysstate -m sys1 -r sys -o recover
```

Initiate service processor failover for a managed system:

```
chsysstate -m myServer -r sys -o spfailover
```

Set the keylock position for a managed system:

```
chsysstate -m sys1 -r sys -o chkey -k manual
```

Activate IBM i partition **p1** using partition profile **p1_prof1** and IPL source **b**:

```
chsysstate -m sys1 -r lpar -o on -n p1 -f p1_prof1 -i b
```

Activate partitions **p1**, **p2**, and **p3** with their current configurations:

```
chsysstate -m sys1 -r lpar -o on -n p1,p2,p3
```

Activate and perform a network install of the IBM i partition **iLpar**:

```
chsysstate -m mySys -r lpar -o on -n iLpar -f iProf -i d --ip 9.1.2.33  
--netmask 255.255.255.0 --gateway 9.1.0.1 --serverip 9.5.12.34  
--serverdir /IBMi/611
```

Shut down the partition with ID **1**:

```
chsysstate -m 9406-570*12345678 -r lpar -o shutdown --id 1
```

Issue the operating system shutdown command to immediately shut down the partition **p1**:

```
chsysstate -m 9406-570*12345678 -r lpar -o osshutdown  
-n p1 --immed
```

Issue the operating system shutdown command to shut down then restart the partition **p1**:

```
chsysstate -m 9406-570*12345678 -r lpar -o osshutdown  
-n p1 --restart
```

Immediately restart the partition with ID **1**:

```
chsysstate -m 9406-570*12345678 -r lpar -o shutdown --id 1  
--immed --restart
```

Enable a remote service session for the IBM i partition **mylpar**:

```
chsysstate -m sys1 -r lpar -o remoteston -n mylpar
```

Validate system profile **sp1**:

```
chsysstate -m sys1 -r sysprof -n sp1 --test
```

Validate then activate system profile **sp1**:

```
chsysstate -m sys1 -r sysprof -n sp1 -o on --test
```

Activate system profile **mySysProf** and continue activating remaining partitions if a partition activation failure occurs:

```
chsysstate -m 9406-570*12345678 -r sysprof -n mySysProf  
-o on --continue
```

Rebuild a managed frame:

```
chsysstate -e myFrame -r frame -o rebuild
```

ENVIRONMENT

None

BUGS

None

AUTHOR

IBM Austin

SEE ALSO

chlparstate, **lssyscfg**

NAME

chtskey - change trusted system key

SYNOPSIS

```
chtskey -m managed-system -o {change | restore}  
--newkey file [--oldkey file] [--force] [--help]
```

DESCRIPTION

chtskey changes the trusted system key for the *managed-system*. The trusted system key is used for virtual Trusted Platform Module (vTPM) data encryption.

If the trusted system key for a managed system is not set, the managed system automatically generates the default trusted system key when the first partition that uses vTPM is activated.

OPTIONS

-m The name of the managed system for which to set the trusted system key. The name may either be the user-defined name for the managed system, or be in the form *ttt-mmm*ssssss*, where *ttt* is the machine type, *mmm* is the model, and *ssssss* is the serial number of the managed system. The *ttt-mmm*ssssss* form must be used if there are multiple managed systems with the same user-defined name.

-o The operation to perform. Valid values are **change** to set the trusted system key and re-encrypt vTPM data, and **restore** to set the trusted system key without re-encrypting vTPM data.

The restore operation is only allowed when all of the partitions that are using vTPM are shut down.

Re-encryption of vTPM data may take a long time to complete.

--newkey

The name of the binary file that contains the new key.

If the file is on removable media, the media must be present in the removable media device and the device must be mounted with the **mount** command before this command is issued. The **lsmediadev** command can be used to display all of the removable media devices on the HMC.

--oldkey

The name of the binary file that contains the current key. This option is not required the first time a user sets the trusted system key for the *managed-system*. After the first time, this option is required.

If the file is on removable media, the media must be present in the removable media device and the device must be mounted with the **mount** command before this command is issued. The **lsmediadev** command can be used to display all of the removable media devices on the HMC.

--force Specify this option to force the trusted system key to be changed when vTPM data is still being re-encrypted due to a previous trusted system key change operation. vTPM data loss may occur.

--help Display the help text for this command and exit.

EXAMPLES

Set the trusted system key for the first time. The key file exists in the user's home directory on the HMC:

```
chtskey -m sys1 -o change --newkey keyfile
```

Set the trusted system key after it has been set at least once. The key files are on a USB flash memory device which is connected to the HMC:

```
lsmediadev (to obtain mount points)
```

mount /media/sdb1

**chtskey -m 8233-E8B*1000ABC -o change --newkey /media/sdb1/newkeyfile
--oldkey /media/sdb1/oldkeyfile**

ENVIRONMENT

None

BUGS

None

AUTHOR

IBM Austin

SEE ALSO

lstskey

NAME

chusrta - displays welcome text at local console login or displays text when remotely login using SSH.

SYNOPSIS

```
chusrta [-o { a | r | ab | rb }]  
[-f welcome-file] [-b banner-file]  
[-c] [--help]
```

DESCRIPTION

chusrta provides a method to deploy text file(s) containing welcome or warning message and present to all users who login locally on the Hardware Management Console (HMC). The text file must be named **WelcomeFile.txt**. A banner file can also be specified, to have its contents displayed prior to login with SSH. The banner file must first be copied to the HMC, then it will be copied to a location accessible by SSH daemon. The banner file must be named **BannerFile.txt**.

OPTIONS

- o** Adds (**a**) or removes (**r**) the text file containing the welcome text. Add banner (**ab**) and remove banner (**rb**) can only be specified when deploying a banner file.
- f** The text file whose contents will be used as welcome, when login locally on the HMC, must be named **WelcomeFile.txt**.
- b** The text file whose contents will be used as banner, when remotely login using SSH, must be named **BannerFile.txt**.
- c** Removes the welcome or the banner file from the HMC filesystem. This option requires the **-f** or **-b** flags.
- help** Displays the help text for this command and exit.

EXAMPLES

Adds welcome text. The file is on floppy:

```
mount /mnt/floppy
```

```
chusrta -o a -f /mnt/floppy/WelcomeFile.txt
```

Use scp to copy the **WelcomeFile.txt** file to /tmp on the HMC, then adds welcome text then removes it:

```
scp someID@someHost:/tmp/WelcomeFile.txt /tmp
```

```
chusrta -o a -f /tmp/WelcomeFile.txt -c
```

Removes welcome text from the HMC. The welcome text will no longer be displayed:

```
chusrta -o r
```

Removes banner text from the HMC. The banner text will no longer be displayed:

```
chusrta -o rb
```

Use scp to copy the banner file to /tmp on the HMC, then deploy it:

```
scp someID@someHost:/tmp/BannerFile.txt /tmp
```

```
chusrta -o ab -b /tmp/BannerFile.txt -c
```

Removes banner text file from the HMC:

chusrta -c -b /tmp/BannerFile.txt

ENVIRONMENT

None

BUGS

None

AUTHOR

IBM Austin

NAME

chvet - activate Capacity on Demand advanced functions

SYNOPSIS

chvet -o e -m *managed-system* -k *activation-code* [--help]

DESCRIPTION

chvet performs activation of Capacity on Demand (CoD) advanced functions on the *managed-system*. CoD advanced functions include PowerVM and Enterprise Enablement.

CoD advanced functions are sometimes referred to as Virtualization Engine systems technologies.

OPTIONS

- o** The operation to perform. The only valid value is **e** to enter an activation code.
- m** The name of the managed system on which to enter the activation code. The name may either be the user-defined name for the managed system, or be in the form *ttt-*mmm***sssssss**, where *ttt* is the machine type, *mmm* is the model, and *sssssss* is the serial number of the managed system. The *ttt-*mmm***sssssss** form must be used if there are multiple managed systems with the same user-defined name.
- k** The activation code (key) to enter. Letters may be entered in either upper case or lower case.
- help** Display the help text for this command and exit.

EXAMPLES

Enter an activation code:

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

ENVIRONMENT

None

BUGS

None

AUTHOR

IBM Austin

SEE ALSO

lsvet, lssyscfg

NAME

cpdump - copy dumps

SYNOPSIS

```
cpdump -r {ftp | usb} -f "file1 file2 ...fileN"  
[-h host-name] [-u user-ID] [--passwd password]  
[-d directory] [--help]
```

DESCRIPTION

cpdump copies managed system dumps and managed frame dumps from the Hardware Management Console (HMC) to a remote FTP site or a USB data storage device.

Use the **lsdump** command to list the managed system dumps and managed frame dumps that reside on the HMC.

OPTIONS

-r The storage location to where the dumps will be copied. Valid values are **ftp** for a remote FTP site and **usb** for a USB data storage device.

When copying dumps to a USB data storage device, the device must be inserted into a USB port on the HMC and the device must be mounted with the **mount** command before running this command. The **lsmediadev** command can be used to display the mount points for the removable media devices on the HMC.

-f The file name of the dump to be copied. If multiple dumps are to be copied, the file names must be separated by blanks and enclosed in double quotes.

-h The host name or IP address of the remote FTP site.

This option is required when copying dumps to a remote FTP site. Otherwise, this option is not valid.

-u The user ID to use to log in to the remote FTP site.

This option is required when copying dumps to a remote FTP site. Otherwise, this option is not valid.

--passwd

The password to use to log in to the remote FTP site. If this option is omitted, you will be prompted to enter the password.

This option is only valid when copying dumps to a remote FTP site.

-d When copying dumps to a remote FTP site, specify the directory on the FTP server to where the dumps will be copied. If this option is omitted, the dumps will be copied to the home directory of the user.

When copying dumps to a USB data storage device, specify the directory on the USB device to where the dumps will be copied.

--help Display the help text for this command and exit.

EXAMPLES

Copy a dump file to a remote FTP site:

```
cpdump -r ftp -f "FSPDUMP.100072A.01000019.20040629205455"  
-h ftpserver -u ftpuser --passwd ftppassword -d /tmp/dumpfiles/
```

Copy dump files to a USB flash drive:

lsmediadev (to obtain mount points)

mount /media/sdb1

**cpdump -r usb -f "FSPDUMP.100072A.01000019.20040629205455
SYSDUMP.100072A.01000019.20040630183112" -d /media/sdb1**

ENVIRONMENT

None

BUGS

None

AUTHOR

IBM Austin

SEE ALSO

lsdump, lsmediadev, rmdump

NAME

cpfile - copy file

SYNOPSIS

cpfile -t modelmap -l {l | r} -f file-name -o {import | export}
[-h host-name -u user-ID [--passwd password]] [--help]

DESCRIPTION

cpfile copies a file to or from the Hardware Management Console (HMC) and the local HMC file system, removable media, or a remote system using secure File Transfer Protocol (FTP).

This command allows a user to copy a file to or from a directory on the HMC to which the user does not have write access. The type of file specified determines the directory.

OPTIONS

-t The type of file to copy. The only valid value is **modelmap** for an XML file which indicates the HMC is managing an IBM Smart Analytics Solution. This file describes an appliance model ID and its associated server model ID and is used during a call-home event to associate the failing server model to the appliance model.

-l The location of the file. For an import operation, this is the source location of the file. For an export operation, this is the destination location of the file. Valid values are **l** for the local HMC file system or removable media, or **r** for a remote secure FTP server.

When copying a file to or from removable media, the media must be present in the removable media device and the device must be mounted with the **mount** command before this command is issued. The **lsmediadev** command can be used to display all of the removable media devices on the HMC.

-f The fully qualified name of the file to copy. If the file already exists in the destination location, it will be overwritten.

-o The direction of the copy. Valid values are **import** to copy the file from the specified location to the destination location on the HMC, or **export** to copy the file from the source location on the HMC to the specified location.

-h The host name or IP address of the secure FTP server.

This option is required when copying a file to or from a secure FTP server. Otherwise, this option is not valid.

-u The user ID to use to log in to the secure FTP server.

This option is required when copying a file to or from a secure FTP server. Otherwise, this option is not valid.

--passwd

The password to use to log in to the secure FTP server. If this option is omitted, you will be prompted to enter the password.

This option is only valid when copying a file to or from a secure FTP server.

--help Display the help text for this command and exit.

EXAMPLES

Copy the model mapping file called **hmcModelMap.xml** on a USB flash memory device (the USB flash memory device must already be connected to the HMC) to the HMC:

lsmediadev (to obtain mount points)

mount /media/sdb1

```
cpfile -t modelmap -l l -o import -f /media/sdb1/hmcModelMap.xml
```

Use **scp** to copy the model mapping file called **myfile.xml** to the user's home directory on the HMC. Then use **cpfile** to copy the file into the model mapping file directory on the HMC:

```
scp remoteuser@remotehostname:/directory/myfile.xml /home/user
```

```
cpfile -t modelmap -l l -o import -f /home/user/myfile.xml
```

Copy the model mapping file on the HMC to the file **/home/maps/map.xml** on a remote SFTP server (you will be prompted to enter your password):

```
cpfile -t modelmap -l r -o export -f /home/maps/map.xml  
-h sftpserver -u sftpuser
```

ENVIRONMENT

None

BUGS

None

AUTHOR

IBM Austin

SEE ALSO

lsmediadev

NAME

cpsysplan - copy system plan

SYNOPSIS

```
cpsysplan -r {ftp | media | homedir} -f file-name  
-o {import | export}  
[-h host-name] [-p ftp-port-number] [-u user-ID]  
[--passwd password] [-d directory]  
[--check] [--help]
```

DESCRIPTION

cpsysplan copies a system plan file between the system plan file directory on the Hardware Management Console (HMC) and a remote host, removable media, or the user's home directory on the HMC.

OPTIONS

-r The location where the system plan file will be copied to or from. Valid values are **ftp** for a remote FTP site, **media** for removable media, and **homedir** for the user's home directory on the HMC.

If **media** is specified, the media must be present in the removable media device and the device must be mounted with the **mount** command before this command is issued. The **lsmediadev** command can be used to display all of the removable media devices on the HMC.

-f Specifies the name of the system plan file that is to be copied. The file name is restricted to the alphanumeric (upper and lower) as well as a set of special characters (comma, period, hyphen, underscore, and space). The file name cannot begin with a period or hyphen and must end with the **.sysplan** suffix.

cpsysplan performs a raw file copy. It does not parse or validate the system plan file content.

When copying from the HMC, the file will be searched for in the system plan file directory on the HMC. When copying to the HMC, the file will be written to this directory. If the file already exists, it will be overwritten unless the **--check** option is specified.

-o Specifies if the file is to be copied to or from the system plan file directory on the HMC. Valid values are **export** to copy the file from the HMC to the resource location specified or **import** to copy the file from the resource location specified to the system plan file directory on the HMC.

-h The host name or IP address of the remote FTP site.

This option is required when copying a file to or from a remote FTP site. Otherwise, this option is not valid.

-p The port number to use for the remote FTP site. If this option is omitted, the default FTP port 21 will be used.

This option is only valid when copying a file to or from a remote FTP site.

-u The user ID to use to log in to the remote FTP site.

This option is required when copying a file to or from a remote FTP site. Otherwise, this option is not valid.

--passwd

The password to use to log in to the remote FTP site. If this option is omitted, you will be prompted to enter the password.

This option is only valid when copying a file to or from a remote FTP site.

-d The optional directory to use when copying a file to or from a remote FTP site.

This option is required when copying a system plan file to or from media. The directory is the name of the mounted removable media where the file is read from when **-o import** is specified or written to when **-o export** is specified.

--check

Specifies that when importing a file, this command should check the system plan directory and fail with an error if the file with the name specified with the **-f** option already exists.

This option is only valid when **-o import** is specified.

--help Display the help text for this command and exit.

EXAMPLES

Copy the system plan file **sysplan.sysplan** to a remote FTP site:

```
cpsysplan -r ftp -f sysplan.sysplan -o export -h ftpserver -u ftpuser --passwd ftppassword
```

Copy the system plan file **sysplan.sysplan** to a remote FTP site:

```
cpsysplan -r ftp -f sysplan.sysplan -o export -h 9.3.145.52 -d /home/user/systemplans
```

Copy the system plan file **sysplan.sysplan** from a remote FTP site:

```
cpsysplan -r ftp -f sysplan.sysplan -o import -h ftpserver -p 9922 -u ftpuser --passwd ftppassword
```

Copy the system plan file **sysplan.sysplan** from CD-ROM (the CD must be present in the drive). Don't overwrite **sysplan.sysplan** if the file already exists in the system plan directory:

```
mount /media/cdrom
```

```
cpsysplan -r media -f sysplan.sysplan -o import -d /media/cdrom --check
```

Copy the system plan file **sysplan.sysplan** to a USB flash memory device (the USB flash memory device must already be connected to the HMC):

```
mount /media/sda1
```

```
cpsysplan -r media -f sysplan.sysplan -o export -d /media/sda1
```

ENVIRONMENT

None

BUGS

None

AUTHOR

IBM Austin

SEE ALSO

deploysysplan, lssysplan, mksysplan, rmsysplan, lsmediadev

NAME

deploysysplan - deploy system plan

SYNOPSIS

```
deploysysplan -f file-name -o { dv | v | d }  
[-m managed-system] [--force]  
[--plan plan-number] [-v]  
[--help]
```

DESCRIPTION

deploysysplan deploys a managed system's partition plan from a system plan file.

OPTIONS

- f** Specifies the name of the file that contains the system plan that is to be deployed. The file must exist in the system plan file directory on the Hardware Management Console (HMC). This file will not be changed by the **deploysysplan** command.
- The file name must end with the **.sysplan** suffix.
- o** Specifies the deployment option to be used when deploying the system plan.
- Use the **dv** option to validate the system plan on the managed system, and if it is valid, to deploy it.
- Use the **v** option to validate the system plan on the managed system but not deploy it.
- Use the **d** option to skip validation of the system plan against the managed system, and just deploy the plan-specified partitions and profiles on the managed system.
- m** Specifies the managed system's name where the plan should be deployed. The system plan optionally specifies the managed system's name where that system plan was intended to be deployed. If a managed system name is specified with this option, that name would override the plan-provided name. If no managed system name was specified in the system plan file then this option is required.
- The name may either be the user-defined name for the managed system, or be in the form *tttt-mmm*sssssss*, where *tttt* is the machine type, *mmm* is the model, and *sssssss* is the serial number of the managed system. The *tttt-mmm*sssssss* form must be used if there are multiple managed systems with the same user-defined name.
- force** If the managed system has existing partitions that the system plan contains that are in an active state, this option allows this command to proceed. Otherwise, if active partitions are found, this command will fail with an error message.
- plan** System plan files can contain more than one system plan. If the file contains more than one plan, this option specifies the **name** of the system plan to deploy. If this **name** is not specified, then the first system plan in the file is deployed. If there is only one system plan in the file, this option is ignored.
- v** When **-o v** is specified, this option specifies that a list of identifiers that identify the steps of the validated system plan should be written to stdout. These steps are part of the plan but not yet implemented on the managed system. Otherwise, this list will not be written out.
- When **-o dv** or **-o d** is specified, this option specifies that a list of log entries that record the steps performed as part of deploying the system plan should be written to stdout. Otherwise, this list will not be written out.
- help** Display the help text for this command and exit.

EXAMPLES

Deploy a system plan from the file **mysysplan.sysplan** onto the managed system configured in the file. Validate the plan against the managed system before deploying it:

```
deploysysplan -f mysysplan.sysplan -o dv
```

Deploy a system plan from the file **mysysplan.sysplan** onto managed system **mySystem**. Validate the plan against the managed system before deploying it. Deploy the plan, if valid, even though there might be active partitions on the managed system:

```
deploysysplan -f mysysplan.sysplan -o dv -m mySystem --force
```

Deploy a system plan from the file **mysysplan.sysplan** onto the managed system identified by the machine type, model, and serial number **9406-570*34134888**. Validate the plan against the managed system before deploying it. Write any errors found in validation or deployment to stdout:

```
deploysysplan -f mysysplan.sysplan -o dv -m 9406-570*34134888 -v
```

Validate the ability to deploy a system plan from the file **multisysplan.sysplan** onto managed system **mySystem**. Use the third system plan that appears in this file (the system plan file has at least three system plans). Validate the plan specified against the managed system to determine if that plan can be deployed there but do not actually deploy the plan:

```
deploysysplan -f multisysplan.sysplan -o v -m mySystem --plan 3
```

Validate the ability to deploy a system plan from the file **mysysplan.sysplan** onto managed system **mySystem**. Validate the plan against the managed system to determine if that plan can be deployed there but do not actually deploy the plan. Write out the steps in the system plan that are part of the plan but not yet implemented to stdout. Write any errors found in validation or deployment to stdout:

```
deploysysplan -f mysysplan.sysplan -o v -m mySystem -v
```

Deploy a system plan from the file **mysysplan.sysplan** onto managed system **mySystem**. Do not validate the plan against the managed system to determine if that plan can be deployed there. Write information about the steps performed in deployment to stdout. Write any errors found in validation or deployment to stdout. Deploy the plan, if valid, even though there might be active partitions on the managed system:

```
deploysysplan -f mysysplan.sysplan -o d -m mySystem -v --force
```

ENVIRONMENT

None

BUGS

None

AUTHOR

IBM Austin

SEE ALSO

cpsysplan, lssysplan, mksysplan, rmsysplan

NAME

diagrmc - diagnose RMC connection issues

SYNOPSIS

To check the management console for general RMC connection issues:

```
diagrmc [--autocorrect] [-v] [--help]
```

To check the management console for general RMC connection issues and also check for RMC connection issues to the specified partition:

```
diagrmc -m managed-system --ip partition-IP-address  
{ -p partition-name | --id partition-ID }  
[--autocorrect] [-v] [--help]
```

To check the management console for general RMC connection issues and also check for RMC connection issues to the specified management console:

```
diagrmc --mcip MC-IP-address  
[--autocorrect] [-v] [--help]
```

DESCRIPTION

diagrmc checks the management console for possible problems that will prevent the management console from establishing Resource Monitoring and Control (RMC) connections to partitions. Some of the things that are checked include the network setup, firewall permissions for RMC ports, the RMC subsystem, duplicate partition NodeIDs, and recycled IP addresses.

diagrmc also checks the management console for possible problems that will prevent it from establishing RMC connections to other management consoles in order to manage a Power enterprise pool.

OPTIONS

-m The name of the managed system which has the partition for which to check for RMC connection issues. The name may either be the user-defined name for the managed system, or be in the form *ttt-mmm*sssss*, where *ttt* is the machine type, *mmm* is the model, and *sssss* is the serial number of the managed system. The *ttt-mmm*sssss* form must be used if there are multiple managed systems with the same user-defined name.

--ip The IP address of the partition for which to check for RMC connection issues.

-p The name of the partition for which to check for RMC connection issues.

You must either use this option to specify the name of the partition, or use the **--id** option to specify the partition's ID. The **-p** and the **--id** options are mutually exclusive.

--id The ID of the partition for which to check for RMC connection issues.

You must either use this option to specify the ID of the partition, or use the **-p** option to specify the partition's name. The **--id** and the **-p** options are mutually exclusive.

--mcip The IP address of the management console for which to check for RMC connection issues.

--autocorrect

Specify this option to allow the management console to automatically correct any RMC connection issues that are found that the management console can correct. Note that the management console will not be automatically restarted.

-v Specify this option to enable verbose mode. When verbose mode is enabled, additional details are displayed.

--help Display the help text for this command and exit.

EXAMPLES

Check for general RMC connection issues:

diagrmc

Check for general RMC connection issues and for RMC connection issues to partition **mylpar** with IP address **9.53.22.86**:

```
diagrmc -m system1 -p mylpar --ip 9.53.22.86
```

Check for general RMC connection issues and for RMC connection issues to partition ID **5** with IP address **9.41.161.33**, and automatically correct the issues found:

```
diagrmc -m system1 --id 5 --ip 9.41.161.33 --autocorrect
```

Check for general RMC connection issues and for RMC connection issues to the management console with IP address **9.41.161.45** that may be preventing the management consoles from communicating to manage a Power enterprise pool:

```
diagrmc --mcip 9.41.161.45
```

ENVIRONMENT

None

BUGS

None

AUTHOR

IBM Austin

NAME

dump - perform dump operation

SYNOPSIS

```
dump -m managed-system -t sys -o set -a "attributes"  
[--help]
```

DESCRIPTION

dump sets the system dump parameters for the *managed-system*. This operation is only supported for POWER6 and later servers.

OPTIONS

- m** The name of the managed system for which to set the system dump parameters. The name may either be the user-defined name for the managed system, or be in the form *tttt-mmm*sssssss*, where *tttt* is the machine type, *mmm* is the model, and *sssssss* is the serial number of the managed system. The *tttt-mmm*sssssss* form must be used if there are multiple managed systems with the same user-defined name.
- t** The type of dump for which to set the dump parameters. Valid values are **sys** for system dump.
- o** The dump operation to perform. Valid values are **set** to set the system dump parameters.
- a** The configuration data needed to set the system dump parameters. The configuration data consists of attribute name/value pairs, which are in comma separated value (CSV) format. The configuration data must be enclosed in double quotes.

The format of the configuration data is as follows:

```
attribute-name=value,attribute-name=value,...
```

Valid attribute names:

as_needed

Specifies when to collect system dumps.

Valid values:

0 - do not collect system dumps

1 - collect system dumps as necessary

Setting **as_needed** to **0** will remove the ability of service personnel to diagnose server errors. Do not set **as_needed** to **0** unless directed to do so by your hardware service representative.

sys_hardware

Specifies how much hardware data to collect when a system dump occurs.

Valid values:

auto - collect hardware data based on the system dump event that occurred

max - collect as much hardware data as possible

Setting **sys_hardware** to **max** will cause system dumps to take longer to complete. Do not set **sys_hardware** to **max** unless directed to do so by your hardware service representative.

sys_firmware

Specifies how much system firmware data to collect when a system dump occurs.

Valid values:

- auto** - collect the primary data in memory owned by system firmware
- max** - collect the primary data in memory owned by system firmware plus all TCEs
- io** - collect the primary data in memory owned by system firmware plus all PCI TCEs
- virtualio** - collect the primary data in memory owned by system firmware plus all virtual I/O TCEs
- hps** - collect the primary data in memory owned by system firmware plus all High Performance Switch (HPS) TCEs
- hcaio** - collect the primary data in memory owned by system firmware plus all Host Channel Adapter (HCA) TCEs

Setting **sys_firmware** to **max**, **io**, **virtualio**, **hps**, or **hcaio** will cause system dumps to take longer to complete. Do not set **sys_firmware** to one of these values unless directed to do so by your hardware service representative.

--help Display the help text for this command and exit.

EXAMPLES

Set the system dump parameters for managed system **sys1**:

```
dump -m sys1 -t sys -o set -a "sys_hardware=auto,  
sys_firmware=auto"
```

ENVIRONMENT

None

BUGS

None

AUTHOR

IBM Austin

SEE ALSO

lsdump

NAME

expdata - export data

SYNOPSIS

expdata -f file [--force] [--help]

DESCRIPTION

expdata exports Hardware Management Console (HMC) configuration data to a file. The exported data includes HMC user ID data, HMC task role and managed resource role data, call-home data, and managed system IP addresses.

OPTIONS

-f The name of the file where the HMC configuration data is to be written. If *file* is not fully qualified, *file* will be written to the **/dump/expdata** directory on the HMC.

If *file* already exists, this command will fail unless the **--force** option is specified.

To export the data to removable media, the media must be present in the removable media device and the device must be mounted with the **mount** command before this command is issued. The **lsmediadev** command can be used to display all of the removable media devices on the HMC.

--force This option allows the specified file to be overwritten if it already exists.

--help Display the help text for this command and exit.

EXAMPLES

Export HMC configuration data to a file called **cfgdata** on a USB flash memory device (the USB flash memory device must already be connected to the HMC):

```
mount /media/sdb1
```

```
expdata -f /media/sdb1/cfgdata
```

Export HMC configuration data to a file called **hmcdata** (the file will be written to the **/dump/expdata** directory on the HMC). The file **hmcdata** will be overwritten if it already exists:

```
expdata -f hmcdata --force
```

ENVIRONMENT

None

BUGS

None

AUTHOR

IBM Austin

SEE ALSO

lsmediadev

NAME

getdump - get dump

SYNOPSIS

```
getdump {-m managed-system | -e managed-frame}  
-t {pss | sp | sys | resource} [-s {a | b | p | s}] [--help]
```

DESCRIPTION

getdump offloads a dump from the *managed-system* or the *managed-frame* to the Hardware Management Console (HMC). The dump is written to a file in the **/dump** or **/extra** directory on the HMC.

The **lsdump** command can be used to list the dumps that are available on the *managed-system* or the *managed-frame*.

OPTIONS

-m The name of the managed system from which to offload the dump. The name may either be the user-defined name for the managed system, or be in the form *tttt-mmm*sssssss*, where *tttt* is the machine type, *mmm* is the model, and *sssssss* is the serial number of the managed system. The *tttt-mmm*sssssss* form must be used if there are multiple managed systems with the same user-defined name.

This option is required when offloading a service processor dump, a system dump, or a resource dump. This option is not valid otherwise.

-e The name of the managed frame from which to offload the dump. The name may either be the user-defined name for the managed frame, or be in the form *tttt-mmm*sssssss*, where *tttt* is the type, *mmm* is the model, and *sssssss* is the serial number of the managed frame. The *tttt-mmm*sssssss* form must be used if there are multiple managed frames with the same user-defined name.

This option is required when offloading a power subsystem dump. This option is not valid otherwise.

-t The dump to offload. Valid values are **pss** for a power subsystem dump, **sp** for a service processor dump, **sys** for a system dump, or **resource** for a resource dump.

-s When offloading a power subsystem dump, use this option to specify the side of the managed frame's bulk power assembly (BPA) from which to offload the dump. Valid values are **a** for side A and **b** for side B.

When offloading a service processor dump, use this option to specify the service processor from which to offload the dump. Valid values are **p** for the primary service processor and **s** for the secondary service processor. If this option is not specified, the service processor dump will be offloaded from the primary service processor.

--help Display the help text for this command and exit.

EXAMPLES

Offload the system dump from managed system **sys1**:

```
getdump -m sys1 -t sys
```

Offload the service processor dump from the secondary service processor on managed system **9117-570*12345678**:

```
getdump -m 9117-570*12345678 -t sp -s s
```

Offload the power subsystem dump from side B of the BPA for the managed frame **9119-59*000012C**:

getdump -e 9119-59*000012C -t pss -s b

Offload the resource dump from managed system **sys1**:

getdump -m sys1 -t resource

ENVIRONMENT

None

BUGS

None

AUTHOR

IBM Austin

SEE ALSO

lsdump, rmdump, startdump

NAME

getfile - get file

SYNOPSIS

```
getfile -t {krbkeyfile | ldapcert | rsyslogcert | rsysloghmc-cert |  
          rsysloghmckey}  
-l {l | s} -f file  
[-h host-name -u user-ID [--passwd password] [-k SSH-private-key]]  
[--help]
```

DESCRIPTION

getfile obtains and deploys a file of the specified type on the Hardware Management Console (HMC).

OPTIONS

-t The type of file to get and deploy. Valid values are **krbkeyfile** for the Kerberos service key (keytab) file, **ldapcert** for the LDAP Certificate Authority (CA) certificate file, **rsyslogcert** for the rsyslog CA certificate file, **rsysloghmc-cert** for the rsyslog HMC certificate file, and **rsysloghmckey** for the rsyslog HMC private key file.

The Kerberos service key file is generated on a Key Distribution Center (KDC) server. The service key file contains the host principal that identifies the HMC to the KDC server. When this command is run, the file obtained will become the `/etc/krb5.keytab` file on the HMC. This file will be used to verify the identity of the KDC and the HMC during remote Kerberos authentication of HMC users. Kerberos must be configured on the HMC before this command is run, and the HMC must be rebooted after the successful completion of this command for the change to take effect.

The LDAP CA certificate file contains the certificate for the Certificate Authority the HMC recognizes. This file will be used to provide a secure connection between the HMC and the LDAP server during remote LDAP authentication of HMC users. LDAP must be configured on the HMC before this command is run.

The rsyslog CA certificate file, rsyslog HMC certificate file, and rsyslog HMC private key file must be deployed on the HMC before TLS encrypted system logging over TCP can be enabled. These 3 files will automatically be removed from the HMC when TLS encrypted system logging over TCP is disabled.

-l The location where the file to get resides.

Specify **l** with this option if the file has already been copied to the HMC or resides on removable media. If the file resides on removable media, the media must be present in the removable media device and the device must be mounted with the **mount** command before this command is issued. The **lsmediadev** command can be used to display all of the removable media devices on the HMC.

Specify **s** with this option if the file resides on a remote site. Secure FTP (SFTP) will be used to obtain the file from the remote site.

-f The fully qualified name of the file on the HMC, removable media, or remote site.

-h The host name or IP address of the SFTP server from which to get the file.

This option is required when *file* resides on a remote site. Otherwise, this option is not valid.

-u The user ID to use to log in to the remote site.

This option is required when *file* resides on a remote site. Otherwise, this option is not valid.

--passwd

The password to use to log in to the remote site.

If both this option and the **-k** option are omitted and *file* resides on a remote site, you will be prompted to enter the password. The **--passwd** and **-k** options are mutually exclusive.

This option is only valid when *file* resides on a remote site.

-k The name of the file that contains the SSH private key. If the file name is not fully qualified, the file must exist in the user's home directory on the HMC.

Use the **ssh-keygen** command to generate the public and private SSH key pair. The **ssh-keygen** command is not allowed to write to the **.ssh** directory in the user's home directory on the HMC, so when you run the command on the HMC, you must specify both the directory and the file name for the private key. If you generate a key with a passphrase, you will be prompted to enter the passphrase when you run any HMC command that uses the key.

If both this option and the **--passwd** option are omitted and *file* resides on a remote site, you will be prompted to enter the password. The **-k** and **--passwd** options are mutually exclusive.

This option is only valid when *file* resides on a remote site.

--help Display the help text for this command and exit.

EXAMPLES

Get the Kerberos service key file from the remote site **sftpserver** using SFTP, then deploy the file as **/etc/krb5.keytab** on the HMC (you will be prompted to enter the user's password):

```
getfile -t krbkeyfile -l s -h sftpserver -u sftpuser  
-f /home/sftpuser/krb5.keytab
```

Get the rsyslog CA certificate file from the remote site **sftpserver** using SFTP and SSH keys for authentication, then deploy the file on the HMC:

```
getfile -t rsyslogcacert -l s -h sftpserver -u sftpuser  
-f /home/sftpuser/rsyslog/cacert -k /home/hmcuser/keys/id_rsa
```

Get the LDAP CA certificate file locally from the user's home directory, then deploy the file on the HMC:

```
getfile -t ldapcacert -l l -f /home/user/certs/ldapcert
```

ENVIRONMENT

None

BUGS

None

AUTHOR

IBM Austin

SEE ALSO

chhmc, **chhmcldap**, **lsmediadev**, **rmfile**

NAME

getupgfiles - get upgrade files

SYNOPSIS

```
getupgfiles -r {ftp | sftp | nfs}
-h host-name [-u user-ID] [--passwd password] [-k SSH-private-key]
[-l mount-location] [-o mount-command-options] [-d remote-directory]
[--help]
```

DESCRIPTION

getupgfiles obtains required files for a Hardware Management Console (HMC) upgrade from a remote server. The files are transferred onto a special partition on the HMC hard disk. After the files have been transferred, the **chhmc** command can be used to allow the HMC to boot from this partition and perform the upgrade.

OPTIONS

- r** The type of remote server where the upgrade files are located. Valid values are **ftp** for an FTP server, **sftp** for a secure FTP (SFTP) server, and **nfs** for an NFS server.
- h** The host name or IP address of the remote server where the upgrade files are located.
- u** The user ID to use to log in to the remote FTP or SFTP server.

This option is required when the upgrade files are located on a remote FTP or SFTP server. Otherwise, this option is not valid.

--passwd

The password to use to log in to the remote FTP or SFTP server.

If this option is omitted and the upgrade files are located on a remote FTP server, you will be prompted to enter the password. If both this option and the **-k** option are omitted and the upgrade files are located on a remote SFTP server, you will be prompted to enter the password. The **--passwd** and **-k** options are mutually exclusive.

This option is only valid when the upgrade files are located on a remote FTP or SFTP server.

- k** The name of the file that contains the SSH private key. If the file name is not fully qualified, the file must exist in the user's home directory on the HMC.

Use the **ssh-keygen** command to generate the public and private SSH key pair. The **ssh-keygen** command is not allowed to write to the **.ssh** directory in the user's home directory on the HMC, so when you run the command on the HMC, you must specify both the directory and the file name for the private key. If you generate a key with a passphrase, you will be prompted to enter the passphrase when you run any HMC command that uses the key.

If both this option and the **--passwd** option are omitted and the upgrade files are located on a remote SFTP server, you will be prompted to enter the password. The **-k** and **--passwd** options are mutually exclusive.

This option is only valid when the upgrade files are located on a remote SFTP server.

- l** The mount location defined on the NFS server where the upgrade files are located.

This option is required when the upgrade files are located on a remote NFS server. Otherwise, this option is not valid.

- o** The options to be passed to the **mount** command used to mount the NFS file system where the upgrade files are located.

This command uses NFS version 4 by default.

This option is only valid when the upgrade files are located on a remote NFS server.

- d** The directory on the remote server that contains the upgrade files. If all of the required files are not present in this directory, this command will fail.

This option is required when the upgrade files are located on a remote FTP or SFTP server.

- help** Display the help text for this command and exit.

EXAMPLES

Get HMC upgrade files from a remote FTP server (you will be prompted to enter the user's password):

```
getupgfiles -r ftp -h ftpserver -u ftpuser -d /home/ftpuser/images
```

Get HMC upgrade files from a remote SFTP server and use SSH keys for authentication:

```
getupgfiles -r sftp -h 9.3.123.25 -u sftpuser  
-k /home/hmcuser/keys/id_rsa -d /home/sftpuser/images
```

Get HMC upgrade files from a remote NFS server:

```
getupgfiles -r nfs -h nfsserver -l /hmc/upgrades -d 830
```

ENVIRONMENT

None

BUGS

None

AUTHOR

IBM Austin

SEE ALSO

chhmc

NAME

hmcshutdown - shut down the Hardware Management Console

SYNOPSIS

hmcshutdown -t {**now** | *number-of-minutes*} [**-r**] [**--help**]

DESCRIPTION

hmcshutdown shuts down the Hardware Management Console (HMC). If requested, the HMC will be rebooted after the shutdown completes.

OPTIONS

- t** The number of minutes to wait before starting the shutdown. If **now** is specified, the shutdown will be started immediately.
- r** Reboot the HMC after the shutdown. If this option is omitted, the HMC will be halted after the shutdown.
- help** Display the help text for this command and exit.

EXAMPLES

Reboot the HMC after 3 minutes:

```
hmcshutdown -t 3 -r
```

Halt the HMC immediately:

```
hmcshutdown -t now
```

ENVIRONMENT

None

BUGS

None

AUTHOR

IBM Austin

NAME

hmcwin - Provides screen capture on HMC

SYNOPSIS

hmcwin { -o c | v |r } -f file-name [--help]

DESCRIPTION

hmcwin Provide screen capture on HMC.

OPTIONS

- o c captures the screen and save it in file-name. The file will be saved under the user's \$HOME/.screen_capture/ directory. v displays the image stored in file file-name from the user's \$HOME/.screen_capture/ directory. When using the v option, left mouse click on the displayed window or enter ctrl-c to dismiss it. r removes the file file-name under the user's \$HOME/.screen_capture directory.
- f The file name to be used to create or remove the screen captured file under the user's \$HOME/.screen_capture directory.
- help Prints the help message and exits.

ENVIRONMENT

None

BUGS

None

AUTHOR

IBM Austin

NAME

hwdbg - hardware debug tools

SYNOPSIS

```
hwdbg -m managed-system -r sriov -o s
{-p partition-name | --id partition-ID}
-a "attributes" [--force] [--help]
```

DESCRIPTION

hwdbg provides hardware debug tools for hardware service personnel.

This command is used to set SR-IOV hardware debug mode for an SR-IOV logical port.

OPTIONS

- m** The name of the managed system on which to perform the operation. The name may either be the user-defined name for the managed system, or be in the form *ttt-mmm*sssss*, where *ttt* is the machine type, *mmm* is the model, and *sssss* is the serial number of the managed system. The *ttt-mmm*sssss* form must be used if there are multiple managed systems with the same user-defined name.
- r** The type of hardware resources to debug. Valid values are **sriov** for Single Root I/O Virtualization (SR-IOV) resources.
- o** The operation to perform. Valid values are **s** to set hardware debug attributes.
- p** The name of the partition for which to perform the operation.

You must either use this option to specify the name of the partition, or use the **--id** option to specify the partition's ID. The **-p** and the **--id** options are mutually exclusive.

- id** The ID of the partition for which to perform the operation.

You must either use this option to specify the ID of the partition, or use the **-p** option to specify the partition's name. The **--id** and the **-p** options are mutually exclusive.

- a** The configuration data needed to set hardware debug attributes. The configuration data consists of attribute name/value pairs, which are in comma separated value (CSV) format. The configuration data must be enclosed in double quotes.

The format of the configuration data is as follows:

```
attribute-name=value,attribute-name=value,...
```

Valid attribute names (all are required):

adapter_id

debug_mode

Valid values:

0 - disable

1 - enable

Warning! Enabling SR-IOV debug mode is a destructive operation which may be disruptive to all users of the adapter. An adapter reset may be required to recover.

logical_port_id

- force** This option allows you to force the operation to be performed for an AIX, Linux, or VIOS partition that does not have an RMC connection to the management console.
- help** Display the help text for this command and exit.

EXAMPLES

Enable SR-IOV hardware debug mode for an SR-IOV logical port:

```
hwdbg -m 9117-MMB*1234567 -r sriov -p lpar1 -o s  
-a "adapter_id=1,logical_port_id=27004001,debug_mode=1"
```

Disable SR-IOV hardware debug mode for an SR-IOV logical port:

```
hwdbg -m 9117-MMB*1234567 -r sriov -p lpar1 -o s  
-a "adapter_id=1,logical_port_id=27004001,debug_mode=0"
```

ENVIRONMENT

None

BUGS

None

AUTHOR

IBM Austin

SEE ALSO

lshwres

NAME

installios - Network install a Virtual I/O Server Logical Partition

SYNOPSIS

```
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] [-A net_interface]  
[-E bosinst_pvid] [-T destination_dir]  
[-V vlan_tag] [-Y vlan_priority] [-C]  
[-o nfs_option] [-R label_name]  
| -u  
| -q  
| [-F] -e -R label_name
```

DESCRIPTION

installios installs the Virtual I/O Server. It must be run from the Hardware Management Console (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.

Please note:

If the managed system is in co-management mode (set with *is_master=1*) when listed with the **lscomgmt** command, the partition profile value will need to be omitted.

OPTIONS

- 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. Valid *path* types are:

/dev/cdrom - for installation images on DVD

localpath - for a system backup of the Virtual I/O Server created by the **backupios** command -- or for an ISO image of Virtual I/O Server install media -- that resides in the directory *localpath* on the HMC. When using an ISO image, *localpath* should include the image filename of the first volume. If the ISO image span multiple volumes, the **installios** will automatically locate the remaining volumes, provided they reside within the same directory as the first.

nfs_server:remote_path_to_image - for a system backup of the Virtual I/O Server, created by the **backupios** command, that resides in the *remote_path_to_image* location on the NFS server, *nfs_server*. Installation from an ISO image is not supported when an NFS path is used.

- 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, 1000, *or* auto, and is 100 by default if this flag is not specified. For newer machines, a speed value of auto is desired although this value is determined by the network configuration of the environment and the network adapter of the installing machine.
- D** Specifies *duplex* (optional). The *duplex* setting with which to configure the client's network interface. This value can be full, half *or* auto and is set to full by default if this flag is not specified. For newer machines, a duplex value of auto is desired although this value is determined by the network configuration of the environment and the network adapter of the installing machine.
- 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.
- A** Specifies *net_interface* (optional): The specified *net_interface* will be used to network install the partition. The network interface specified must be connected to a network in which the partition's IP address is reachable. If no interface is specified, a default public interface will be chosen automatically, based on the HMC's current hostname.
- E** Specifies *bosinst_pvid* (optional): The specified *bosinst_pvid* value signifies the physical volume identity for a disk targeted for installation. The *bosinst_pvid* value can be obtained by running the *lspv* command on the machine being installed.
- T** Specifies *destination_dir* (optional): The specified *destination_dir* points to a path where the source will be copied to on the HMC. The default destination path is the */extra* filesystem.
- V** Specifies the VLAN tag identifier to use for tagging Ethernet frames during network install for virtual network communication. Valid value is from **0** to **4094** (optional).
- Y** Specifies the VLAN tag priority to use for tagging Ethernet frames during network install for virtual network communication. Valid value is from **0** to **7** (optional).
- C** Specifies the option to save the label name and its image for later use (optional). To remove the saved labels and its image, use the **-e** flag to erase the label from */etc/nimol.conf*.
- R** Specifies a label name for **installios** to use instead of creating a default label name (optional).
- 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. However, if the **-C** flag is specified, **installios** will not unconfigure itself. If a label exists in the **installios** configuration file */etc/nimol.conf*, the labels will need to be removed first in order to unconfigure **installios**.
- q** Query **installios** for labels created (optional). The **installios** command creates labels when it rips installable images from the source image to disk. The labels are stored in */etc/nimol.conf* and are linked to clients when they are being used for installation. The query flag will list out the labels known to **installios**.
- e** Erase a label known to **installios** (optional). When executing the **installios** command with the **-R** label, it creates a label when it rips the installable images from the source image to disk. The label is stored in */etc/nimol.conf* and its reference to install resources are linked to a client when it is used for installation. The query flag will list out the labels known to **installios**.

- F Force the operation to continue and to bypass certain checking that may be preventing **installios** to run (optional).
- o Specifies the NFS option to use when mounting install resources from a remote server (optional).

EXAMPLES

To use the **installios** wizard which will prompt the user for partition and network information to install a partition, type:

installios

To create resources from the CDROM media for installing a partition with the following information:

```
Managed System : wampeter
Partition Name  : shire02
Partition Profile : default
Partition IP    : 9.3.6.234
Partition Subnet : 255.255.255.0
Partition Gateway : 9.3.6.1
```

installios -d /dev/cdrom -i 9.3.6.234 -g 9.3.6.1 -S 255.255.255.0 -s wampeter -p shire02 -r default

To create resources from a directory at /mnt for installing a partition with the following information:

```
Managed System : bokonon
Partition Name  : mordor02
Partition Profile : full_resource
Partition IP    : 9.3.6.134
Partition Subnet : 255.255.255.0
Partition Gateway : 9.3.6.1
MAC address    : f2:d4:60:00:d0:03
```

installios -d /mnt -i 9.3.6.134 -g 9.3.6.1 -S 255.255.255.0 -m f2:d4:60:00:d0:03 -s bokonon -p mordor02 -r full_resource

To create resources from an ISO image at the location /extra/GOLD_BUILD, and saving the image for later execution in installing a partition with the following information:

```
Managed System : bokonon
Partition Name  : mordor05
Partition Profile : vios
Partition IP    : 9.3.6.250
Partition Subnet : 255.255.255.0
Partition Gateway : 9.3.6.1
MAC address    : f2:d8:90:00:e0:10
```

installios -d /extra/GOLD_BUILD/dvdimage.v1.iso -i 9.3.6.250 -g 9.3.6.1 -S 255.255.255.0 -m f2:d8:90:00:e0:10 -s bokonon -p mordor05 -r vios -R GOLD_BUILD -C

To create resources from a remote server that will be serving a **nim_resources.tar** image at the location **server1:/export/nim_resources.tar**, using an NFS option of "**-o vers=3 -o nolock**", with the following information:

```
Managed System : bokonon
Partition Name  : mordor05
Partition Profile : vios
Partition IP    : 9.3.6.250
Partition Subnet : 255.255.255.0
```

Partition Gateway : 9.3.6.1
MAC address : f2:d8:90:00:e0:10

```
installios -d server1:/export/nim_resources.tar -o "-o vers=3 -o nolock" -i 9.3.6.250 -g 9.3.6.1 -S 255.255.255.0 -m f2:d8:90:00:e0:10 -s bokonon -p mordor05 -r vios
```

To clean up tasks from a previous `installios` execution, type:

```
installios -u
```

To query `installios` for a list of label names and their image locations, type:

```
installios -q
```

To erase a label name and the image saved from `installios` where the label name is `GOLD_BUILD`, type:

```
installios -e -R GOLD_BUILD
```

ENVIRONMENT

INSTALLIOS_DEBUG

This variable is used to print debug messages from `lpar_netboot` when `installios` executes the `lpar_netboot` command. This variable is set with `INSTALLIOS_DEBUG=yes`. When executing the `installios` command in the wizard mode, this variable must not be set as the debug output from `lpar_netboot` will interfere with the MAC address discovery resulting in a failed execution. This variable must only be set when executing `installios` on the command line with flags.

INSTALLIOS_PRIVATE_IF

This variable is used by `installios` to enable DHCP service when responding to BOOTP request from an installing partition. This variable must only be used in a situation where the FSP (Flexible Service Processor) and the installing partition communicate with the HMC through a private network. If the private network the HMC is using is `eth0`, then this variable is set with `INSTALLIOS_PRIVATE_IF=eth0`.

INSTALLIOS_DESTINATION_DIRECTORY

This variable is used by `installios` to specify the destination directory to copy the source to. This variable is set with `INSTALLIOS_DESTINATION_DIRECTORY=/info`. This variable is overridden by the `-T destination_dir` flag. In certain situations where calling `installios` fails due to space limitations when copying over the source image, setting this variable or the `-T` flag to an alternate location with space capable of storing the content of the source location will resolve the situation.

BUGS

None

AUTHOR

IBM Austin

SEE ALSO

`chsysstate`, `lssyscfg`, `mkvterm`, `rmvterm`, `lpar_netboot`

NAME

logssh - log ssh command usage

SYNOPSIS

logssh

DESCRIPTION

logssh is a script that can be added to a user's **authorized_keys2** file to restrict a user from being able to open a pseudo-tty using ssh. Furthermore, it will log all commands executed by the user in syslog. This command is not intended to be run interactively.

Note: The logssh `${SSH_ORIGINAL_COMMAND}` argument has been deprecated and will be ignored if passed to the script.

EXAMPLES

Use the command **mkauthkeys** to update a user's **authorized_keys2** file:

```
mkauthkeys -a -u john 'command="logssh"  
ssh-rsa AAAAB3NzaC1yc2EAAAABIwAAAIEAzxTNjoX  
AvyZBw390oJ27uj90PxZNtUWhYVN1/kaAfilSIr3z5Hhm7BdaaarU  
ru94qhiM0xds6cgQpNUQUy6GByoWDrNhdEIdAzXj3uaPscY6wKkNi  
a0llTJPUoBDBsadaa4oEc0/4poNG/X3uYrsdnbbMNkt/jmnEilSXI  
gOEmWk= john@somehost'
```

The above command updates user **john**'s **authorized_keys2** file on the HMC, so that all commands executed by **john** will be logged in syslog.

```
ssh john@somehmc lshmc -V
```

The above command will result in the following entries being logged in `/var/log/messages`:

```
Nov 28 15:08:14 somehmc hmc_ssh: john login on Tue Nov 28  
15:08:14 CST 2006 from 192.168.128.125  
Nov 28 15:08:14 somehmc hmc_ssh: john runs lshmc -V on  
Tue Nov 28 15:08:14 CST 2006 from 192.168.128.125
```

ENVIRONMENT

None

BUGS

None

AUTHOR

IBM Austin

NAME

lpar_netboot - retrieve MAC address and physical location code from network adapters for a partition or instruct a partition to network boot

SYNOPSIS

To retrieve MAC address and physical location code:

```
lpar_netboot -M -n [-v] [-x] [-f] [-i] [-E environment [-E ...]]  
  [-A] -t ent [-T {on|off}] [-D -s speed -d duplex]  
  -S server -G gateway -C client [-K subnetmask]  
  [-V vlan_tag] [-Y vlan_priority]  
  partition-name partition-profile managed-system
```

To perform network boot:

```
lpar_netboot [-v] [-x] [-f] [-i] [-E environment [-E ...]]  
  [-g args] [{-A -D | [-D] -I physical-location-code |  
  [-D] -m MAC-address}] -t ent [-T {on|off}]  
  -s speed -d duplex -S server -G gateway -C client  
  [-K subnetmask] [-V vlan_tag] [-Y vlan_priority]  
  [[-a -B tftp_image_filename] |  
  -B bootp_image_filename]  
  partition-name partition-profile managed-system
```

To retrieve MAC address and physical location code on a system supporting a full system partition:

```
lpar_netboot -M -n [-v] [-x] [-f] [-i] [-E environment [-E ...]]  
  [-A] -t ent [-T {on|off}] [-D -s speed -d duplex]  
  -S server -G gateway -C client [-K subnetmask]  
  [-V vlan_tag] [-Y vlan_priority]  
  managed-system managed-system
```

To perform network boot on a system supporting a full system partition:

```
lpar_netboot [-v] [-x] [-f] [-i] [-E environment [-E ...]]  
  [-g args] [{-A -D | [-D] -I physical-location-code |  
  [-D] -m MAC-address}] -t ent [-T {on|off}]  
  -s speed -d duplex -S server -G gateway -C client  
  [-K subnetmask] [-V vlan_tag] [-Y vlan_priority]  
  [[-a -B tftp_image_filename] |  
  -B bootp_image_filename]  
  managed-system managed-system
```

DESCRIPTION

lpar_netboot instructs a logical partition to network boot by having it send out a bootp request to a server specified with the **-S** option. The server can be an AIX NIM server serving SPOT resources or any server serving network boot images. If specified with the **-M** and **-n** options, **lpar_netboot** will return the Media Access Control (MAC) address and the physical location code for a network adapter of the type specified with the **-t** option. When the **-m** option is specified, **lpar_netboot** will boot a partition using the network adapter which has the specified MAC address. When the **-I** option is specified, **lpar_netboot** will boot a partition using the network adapter which has the specified physical location code. The MAC address and physical location code of a network adapter is dependent upon the hardware resource allocation in the partition profile the partition was booted with. The **lpar_netboot** command requires arguments for partition name, partition profile, and the name of the managed system which has the partition.

Please note:

If the managed system is in co-management mode (set with *is_master=1*) when listed with the **lscomgmt** command, the partition profile value will need to be blanked out with "". The **lpar_netboot** command will also ignore the partition profile value if its managed system is in co-management mode.

OPTIONS

- A** Return all adapters of the type specified with the **-t** option.
 - B** Network boot image filename, required option if IPv6 addresses.
 - C** The IP address of the partition to network boot.
 - D** Perform a ping test and use the adapter that successfully pings the server specified with the **-S** option.
 - E** Set environment variable setting. The **-E LPAR_NETBOOT_DEBUG=1** is the same as **export LPAR_NETBOOT_DEBUG=1**. See **ENVIRONMENT**.
 - G** The gateway IP address of the partition specified with the **-C** option.
 - K** Subnetmask IP address.
 - M** Discover network adapter MAC address and physical location code.
 - S** The IP address of the machine from which to retrieve the network boot image during network boot.
 - T** Enable or disable firmware spanning tree discovery. Valid values are **on**, **off**.
 - V** Specifies the VLAN tag identifier to use for tagging Ethernet frames during network install for virtual network communication. Valid value is from **0** to **4094**.
 - Y** Specifies the VLAN tag priority to use for tagging Ethernet frames during network install for virtual network communication. Valid value is from **0** to **7**.
 - a** Network ip addresses for server, client and gateway are IPv6.
 - d** The duplex setting of the partition specified with the **-C** option. Valid values are **full**, **half**, and **auto**.
 - f** Force close the virtual terminal session for the partition.
 - g** Specify generic arguments for booting the partition.
 - i** Force immediate shutdown of the partition. If this option is not specified, a delayed shutdown will be performed.
 - l** The physical location code of the network adapter to use for network boot.
 - m** The MAC address of the network adapter to use for network boot.
 - n** Instruct the partition to not network boot.
 - s** The speed setting of the partition specified with the **-C** option. Valid values are **10**, **100**, **1000**, and **auto**.
 - t** The type of adapter for MAC address or physical location code discovery or for network boot. The only valid value is **ent** for ethernet.
 - v** Display additional information during command execution.
 - x** Display debug output during command execution.
- partition-name**
The name of the partition.
- partition-profile**
The name of the partition profile.
- managed-system**
The name of the managed system which has the partition.
- help** Display the help text for this command and exit.

EXAMPLES

To retrieve the MAC address and physical location code for partition **machA** with partition profile **machA_prof** on managed system **test_sys**:

```
lpar_netboot -M -n -t ent "machA" "machA_prof" "test_sys"
```

To network boot the partition **machA** with partition profile **machA_prof** on managed system **test_sys**:

```
lpar_netboot -t ent -s auto -d auto -S 9.3.6.49 -G 9.3.6.1 -C 9.3.6.234 "machA" "machA_prof" "test_sys"
```

To network boot the partition **machA** using the network adapter with a MAC address of 00:09:6b:dd:02:e8 with partition profile **machA_prof** on managed system **test_sys**:

```
lpar_netboot -t ent -m 00096bdd02e8 -s auto -d auto -S 9.3.6.49 -G 9.3.6.1 -C 9.3.6.234 "machA" "machA_prof" "test_sys"
```

To network boot the partition **machA** using the network adapter with a physical location code of U1234.121.A123456-P1-T6 with partition profile **machA_prof** on managed system **test_sys**:

```
lpar_netboot -t ent -l U1234.121.A123456-P1-T6 -s auto -d auto -S 9.3.6.49 -G 9.3.6.1 -C 9.3.6.234 "machA" "machA_prof" "test_sys"
```

To perform a ping test along with a network boot of the partition **machA** with partition profile **machA_prof** on managed system **test_sys**:

```
lpar_netboot -t ent -D -s auto -d auto -S 9.3.6.49 -G 9.3.6.1 -C 9.3.6.234 "machA" "machA_prof" "test_sys"
```

To perform ping test along with a network boot of the partition **machA** with partition profile **machA_prof** on managed system **test_sys** and disable firmware spanning tree discovery:

```
lpar_netboot -t ent -T off -D -s auto -d auto -S 9.3.6.49 -G 9.3.6.1 -C 9.3.6.234 "machA" "machA_prof" "test_sys"
```

ENVIRONMENT

INSTALLIOS_DEBUG

The installios command used this environment variable to print out lpar_netboot debug.

LPAR_NETBOOT_3EXEC

Executing OF "dev /" and "ls" returns nothing, retry reboot and "ls".

LPAR_NETBOOT_DEBUG

Setting this variable will enable lpar_netboot debug. Similar to -x flag.

LPAR_NETBOOT_DEBUG_BOOT

Setting this variable will initiate the firmware boot command with -s trap.

LPAR_NETBOOT_ADD_TIMEOUT

Extend timeout value by 5 seconds. Ex. LPAR_NETBOOT_ADD_TIMEOUT=5

LPAR_NETBOOT_SUB_TIMEOUT

Lower timeout value by 8 seconds. Ex. LPAR_NETBOOT_SUB_TIMEOUT=8

LPAR_NETBOOT_SPANNING_TREE

LPAR_NETBOOT_SPANNING_TREE=on enable spanning tree, LPAR_NETBOOT_SPANNING_TREE=off disable spanning tree. Similar to the -T flag.

OPEN_DEV_DEBUG

Show firmware OPEN_DEV debug by setting value to yes.

FIRMWARE_DUMP

Show firmware dump for firmware debugging by setting value to yes.

BUGS

None

AUTHOR

IBM Austin

SEE ALSO

chsysstate, lssyscfg, mkvterm, rmvterm

NAME

lpcfgop - perform a partition configuration image operation

SYNOPSIS

lpcfgop -m *managed-system* **-o** { **clear** | **dump** } [**--help**]

DESCRIPTION

lpcfgop clears or dumps partition configuration data on the *managed-system*.

OPTIONS

-m The name of the managed system on which to perform the operation. The name may either be the user-defined name for the managed system, or be in the form *ttt-mmm*sssssss*, where *ttt* is the machine type, *mmm* is the model, and *sssssss* is the serial number of the managed system. The *ttt-mmm*sssssss* form must be used if there are multiple managed systems with the same user-defined name.

-o The operation to perform. Valid values are **clear** or **dump**.

The **clear** operation clears all partition configuration data on the *managed-system* and sets the configuration back to its original state. This operation can only be performed when the *managed-system* is in the Standby state.

The **dump** operation dumps all partition configuration data on the *managed-system* to a file. The file is written to the **/tmp** directory on the Hardware Management Console (HMC).

--help Display the help text for this command and exit.

EXAMPLES

Clear the partition configuration data (confirmation will be required):

```
lpcfgop -m sys1 -o clear
```

Dump the partition configuration data to a file in **/tmp**:

```
lpcfgop -m 9406-520*100132A -o dump
```

ENVIRONMENT

None

BUGS

None

AUTHOR

IBM Austin

NAME

lsaccfg - list access control configuration information

SYNOPSIS

```
lsaccfg -t {resource | resourcerole | taskrole}
[--script] [--filter "filter-data"]
[-F [attribute-names] [--header]] [--help]
```

DESCRIPTION

lsaccfg lists managed resource objects, managed resource roles, or task roles.

OPTIONS

-t The type of access control objects to list. Valid values are **resource** for managed resource objects, **resourcerole** for managed resource roles, and **taskrole** for task roles.

--script

List managed resource objects in a format that can be used as input to the **chaccfg**, **chhmcusr**, and **mkaccfg** commands.

This option is only valid when listing managed resource objects.

--filter

The filter to apply to the access control objects to be listed. A filter is used to select which access control objects of the specified type are to be listed. If a filter is not used, then all of the access control objects of the specified type will be listed. For example, specific task roles can be listed by using a filter to specify the names of the task roles to list. Otherwise, if no filter is used, then all of the task roles on this Hardware Management Console (HMC) will be listed.

The filter data consists of filter name/value pairs, which are in comma separated value (CSV) format. The filter data must be enclosed in double quotes.

The format of the filter data is as follows:

```
"filter-name=value,filter-name=value,..."
```

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

```
""filter-name=value,value,...","..."
```

When a list of values is specified, the filter name/value pair must be enclosed in double quotes. Depending on the shell being used, nested double quote characters may need to be preceded by an escape character, which is usually a `'\'` character.

Multiple values can be specified for each filter.

Valid filter names for managed resource objects:

resource_type

Specify resource type(s):

cec - managed system objects

lpar - partition objects

Valid filter names for managed resource roles:

resourceroles

Specify managed resource role name(s)

Valid filter names for task roles:

taskroles

Specify task role name(s)

-F A delimiter separated list of attribute names for the desired attribute values to be displayed for each access control object. If no attribute names are specified, then values for all of the attributes for each access control object will be displayed.

When this option is specified, only attribute values will be displayed. No attribute names will be displayed. The attribute values displayed will be separated by the delimiter which was specified with this option.

This option is useful when only attribute values are desired to be displayed, or when the values of only selected attributes are desired to be displayed.

--header

Display a header record, which is a delimiter separated list of attribute names for the attribute values that will be displayed. This header record will be the first record displayed. This option is only valid when used with the **-F** option.

--help Display the help text for this command and exit.

EXAMPLES

List all managed resource objects on this HMC:

lsaccfg -t resource

List only managed partition objects on this HMC:

lsaccfg -t resource --filter "resource_type=lpar"

List all managed resource roles on this HMC:

lsaccfg -t resourcerole

List the managed resource role lpar_role:

lsaccfg -t resourcerole --filter "resourceroles=lpar_role"

List only the names and parents of the task roles on this HMC, and separate the output values with a colon:

lsaccfg -t taskrole -F name:parent

List the task roles tr1 and tr2:

lsaccfg -t taskrole --filter ""taskroles=tr1,tr2""

ENVIRONMENT

None

BUGS

None

AUTHOR

IBM Austin

SEE ALSO

chaccfg, mkaccfg, rmaccfg, chhmcusr, lshmcusr

NAME

lscod - list Capacity on Demand information

SYNOPSIS

```
lscod -t { bill | cap | code | hist | util }  
-m managed-system  
[-c { cuod | mobile | onoff | trial | trialexc | trialstd |  
      utility | utilityen | utilityunl | updateacc }]  
[-r { mem | proc }]  
[-F [attribute-names] [--header]] [--help]
```

DESCRIPTION

lscod lists Capacity on Demand (CoD) information for the *managed-system*.

OPTIONS

-t The type of CoD information to list. Valid values are **bill** for On/Off CoD billing information, **cap** for CoD capacity information, **code** for information used to generate CoD codes, **hist** for the CoD history log, and **util** for shared processor pool utilization information.

The time stamps shown when listing the CoD history log are managed system time stamps.

-m The name of the managed system for which CoD information is to be listed. The name may either be the user-defined name for the managed system, or be in the form *tttt-mmm*sssssss*, where *tttt* is the machine type, *mmm* is the model, and *sssssss* is the serial number of the managed system. The *tttt-mmm*sssssss* form must be used if there are multiple managed systems with the same user-defined name.

-c The CoD type. Valid values are **cuod** for CUoD (permanent), **mobile** for Mobile CoD, **onoff** for On/Off CoD, **trial** for Trial CoD, **utility** for Utility CoD or Utility CoD reporting codes, **utilityen** for Utility CoD enablement codes, **utilityunl** for Utility CoD unlimited codes, **trialexc** for Trial CoD exception request codes, **trialstd** for Trial CoD standard request codes, and **updateacc** for update access key codes.

This option is required when listing CoD capacity information, CoD code generation information, or shared processor pool utilization information. This option is not valid when listing the CoD history log.

To list Trial CoD capacity information, specify **trial** with this option. To list Trial CoD code generation information, specify **trialexc** or **trialstd** with this option.

To list Utility CoD capacity information, specify **utility** with this option. To list Utility CoD code generation information, specify **utility** for Utility CoD reporting codes, **utilityen** for Utility CoD enablement codes, or **utilityunl** for Utility CoD unlimited codes.

mobile and **updateacc** can only be specified when listing CoD code generation information.

-r The CoD resource type. Valid values are **mem** for memory and **proc** for processors.

This option is required when listing On/Off CoD billing information, CoD capacity information, CoD code generation information, or shared processor pool utilization information. This option is not valid when listing the CoD history log.

-F A delimiter separated list of attribute names representing the desired CoD attribute values to display. If this option is specified without any attribute names, then all of the CoD attributes for the type of CoD listing specified will be displayed.

When this option is specified, only attribute values will be displayed. No attribute names will be displayed. The attribute values displayed will be separated by the delimiter which was specified with this option.

This option is useful when only attribute values are desired to be displayed, or when the values of only selected attributes are desired to be displayed.

--header

Display a header record, which is a delimiter separated list of attribute names for the attribute values that will be displayed. This header record will be the first record displayed. This option is only valid when used with the **-F** option.

--help Display the help text for this command and exit.

EXAMPLES

Display On/Off CoD processor billing information:

lscod -m sys1 -t bill -r proc -c onoff

Display CUoD processor capacity information:

lscod -m 9406-570*98127634 -t cap -r proc -c cuod

Display Trial CoD memory capacity information:

lscod -m 9406-570*98127634 -t cap -r mem -c trial

Display the current number of unreturned On/Off CoD processors:

**lscod -m sys1 -t cap -r proc -c onoff -F
unreturned_onoff_procs**

Display Utility CoD capacity information:

lscod -m sys1 -t cap -r proc -c utility

Display CUoD processor activation code generation information:

lscod -m sys1 -t code -r proc -c cuod

Display Utility CoD reporting code generation information:

lscod -m 9117-MMA*1234321 -t code -r proc -c utility

Display the CoD history log:

lscod -m 9406-570*98127634 -t hist

Display shared processor pool utilization information for Utility CoD:

lscod -m sys1 -t util -r proc -c utility

ENVIRONMENT

None

BUGS

None

AUTHOR

IBM Austin

SEE ALSO
chcod, lshwres

NAME

lscodpool - list CoD Power enterprise pool information

SYNOPSIS

To list Power enterprise pool information:

```
lscodpool --level pool [{-p pool-name | --id pool-ID}]  
[-F attribute-names] [--header] [--help]
```

To list the systems in a Power enterprise pool:

```
lscodpool --level sys {-p pool-name | --id pool-ID}  
[--filter "filter-data"]  
[-F attribute-names] [--header] [--help]
```

To list the Hardware Management Consoles (HMC)s managing a Power enterprise pool:

```
lscodpool --level mc {-p pool-name | --id pool-ID}  
[-F attribute-names] [--header] [--help]
```

To list a CoD history log:

```
lscodpool -t hist  
{-p pool-name | --id pool-ID | -m managed-system}  
[-F attribute-names] [--header] [--help]
```

To list CoD code generation information:

```
lscodpool -t code -m managed-system  
-r {mem | proc | member | convertmem | convertproc}  
[-F attribute-names] [--header] [--help]
```

DESCRIPTION

lscodpool lists Capacity on Demand (CoD) Power enterprise pool information.

OPTIONS

--level The level of the Power enterprise pool information to list. Valid values are **pool** to list pool information, **sys** to list the systems in a pool, and **mc** to list the HMCs managing a pool.

-p The name of the Power enterprise pool for which to list information.

You can either use this option to specify the name of the pool, or use the **--id** option to specify the pool's ID. The **--id** option must be used if there are multiple Power enterprise pools with the same name. The **-p** and the **--id** options are mutually exclusive.

--id The ID of the Power enterprise pool for which to list information.

You can either use this option to specify the ID of the pool, or use the **-p** option to specify the pool's name. The **--id** option must be used if there are multiple Power enterprise pools with the same name. The **--id** and the **-p** options are mutually exclusive.

-m The name of the managed system for which to list information. The name may either be the user-defined name for the managed system, or be in the form *tttt-mmm*ssssss*, where *tttt* is the machine type, *mmm* is the model, and *ssssss* is the serial number of the managed system. The *tttt-mmm*ssssss* form must be used if there are multiple managed systems with the same user-defined name.

This option is only valid when listing the CoD history log for a managed system or listing CoD code generation information.

-t The type of information to list. Specify **code** to list CoD code generation information for Power enterprise pool membership codes, Mobile CoD processor or memory activation codes, or permanent to Mobile CoD processor or memory conversion codes. Specify **hist** to list the CoD history log for a managed system or for a Power enterprise pool.

When listing the CoD history log for a managed system, the time stamps shown are managed system time stamps. When listing the CoD history log for a Power enterprise pool, the time stamps shown are time stamps from the HMCs that logged the entries.

- r** The resource type of the CoD code for which to list code generation information. Valid values are **mem** for Mobile CoD memory, **proc** for Mobile CoD processor, **member** for Power enterprise pool membership, **convertmem** for permanent to Mobile CoD memory conversion, and **convertproc** for permanent to Mobile CoD processor conversion.
- filter** The filter to apply to the systems in the Power enterprise pool to list. The filter is used to select which systems in the pool to list. If a filter is not specified, then all systems in the pool will be listed.

The filter data consists of filter name/value pairs, which are in comma separated value (CSV) format. The filter data must be enclosed in double quotes.

The format of the filter data is as follows:

```
"filter-name=value"
```

Note that the filters accept a comma separated list of values, as follows:

```
""filter-name=value,value,...""
```

When a list of values is specified, the filter name/value pair must be enclosed in double quotes. Depending on the shell being used, nested double quote characters may need to be preceded by an escape character, which is usually a `\` character.

Multiple values can be specified for each filter.

Valid filter names:

names | mtms

This option is only valid when listing the systems in a Power enterprise pool.

- F** A delimiter separated list of attribute names for the desired attribute values to be displayed. If no attribute names are specified, then values for all of the attributes will be displayed.

When this option is specified, only attribute values will be displayed. No attribute names will be displayed. The attribute values displayed will be separated by the delimiter which was specified with this option.

This option is useful when only attribute values are desired to be displayed, or when the values of only selected attributes are desired to be displayed.

- header** Display a header record, which is a delimiter separated list of attribute names for the attribute values that will be displayed. This header record will be the first record displayed. This option is only valid when used with the **-F** option.
- help** Display the help text for this command and exit.

EXAMPLES

List all of the Power enterprise pools managed by this HMC:

```
lscodpool --level pool
```

List all of the systems in the Power enterprise pool **myPool**:

lscodpool --level sys -p myPool

List only systems **sys1** and **sys2** in the Power enterprise pool with ID **0095**:

lscodpool --level sys --id 0095 --filter ""names=sys1,sys2""

List all of the HMCs managing the Power enterprise pool **myPool**:

lscodpool --level mc -p myPool

List the CoD history log for the Power enterprise pool **pool1**:

lscodpool -p pool1 -t hist

List the CoD history log for the managed system **mySys**:

lscodpool -m mySys -t hist

List the Power enterprise pool membership CoD code generation information for the managed system **s1**:

lscodpool -m s1 -t code -r member

ENVIRONMENT

None

BUGS

None

AUTHOR

IBM Austin

SEE ALSO

chcodpool, mkcodpool

NAME

lscomgmt - list co-management settings

SYNOPSIS

lscomgmt -m *managed-system* [-F [*attribute-names*] [--header]] [--help]

DESCRIPTION

lscomgmt lists co-management settings for the *managed-system*.

OPTIONS

-m The name of the managed system for which to list co-management settings. The name may either be the user-defined name for the managed system, or be in the form *ttt-mmm*ssssss*, where *ttt* is the machine type, *mmm* is the model, and *ssssss* is the serial number of the managed system. The *ttt-mmm*ssssss* form must be used if there are multiple managed systems with the same user-defined name.

-F A delimiter separated list of attribute names for the desired attribute values to be displayed. If no attribute names are specified, then values for all of the attributes will be displayed.

When this option is specified, only attribute values will be displayed. No attribute names will be displayed. The attribute values displayed will be separated by the delimiter which was specified with this option.

This option is useful when only attribute values are desired to be displayed, or when the values of only selected attributes are desired to be displayed.

--header

Display a header record, which is a delimiter separated list of attribute names for the attribute values that will be displayed. This header record will be the first record displayed. This option is only valid when used with the **-F** option.

--help Display the help text for this command and exit.

EXAMPLES

List the co-management settings for managed system **sys1**:

```
lscomgmt -m sys1
```

ENVIRONMENT

None

BUGS

None

AUTHOR

IBM Austin

SEE ALSO

chcomgmt

NAME

lsdump - list dump information

SYNOPSIS

To list the dumps that are available on a managed system or managed frame:

```
lsdump {-m managed-system | -e managed-frame}  
[-r avail] [-s {a | b | p | s}]  
[-F [attribute-names] [--header]]
```

To list the managed system dumps and managed frame dumps that are available on the HMC:

```
lsdump -h [-r avail] [--filter "filter-data"]  
[-F [attribute-names] [--header]]
```

To list the system dump parameters:

```
lsdump -m managed-system -r parm  
[-F [attribute-names] [--header]]
```

DESCRIPTION

lsdump lists the dumps that are available on the *managed-system* or the *managed-frame*. **lsdump** can also list the managed system dumps and the managed frame dumps that are available on the Hardware Management Console (HMC).

lsdump also lists the system dump parameters for the *managed-system*.

The **getdump** command can be used to offload an available dump from the *managed-system* or the *managed-frame* to the HMC.

The **cpdump** command can be used to copy a dump from the HMC to a remote FTP site or removable media.

The **rmddump** command can be used to remove a dump from the HMC.

OPTIONS

-m The name of the managed system for which available dumps or system dump parameters are to be listed. The name may either be the user-defined name for the managed system, or be in the form *ttt-mmm*ssssss*, where *ttt* is the machine type, *mmm* is the model, and *ssssss* is the serial number of the managed system. The *ttt-mmm*ssssss* form must be used if there are multiple managed systems with the same user-defined name.

Either this option, the **-e** option, or the **-h** option is required. The **-m**, **-e**, and the **-h** options are mutually exclusive.

-e The name of the managed frame for which available dumps are to be listed. The name may either be the user-defined name for the managed frame, or be in the form *ttt-mmm*ssssss*, where *ttt* is the type, *mmm* is the model, and *ssssss* is the serial number of the managed frame. The *ttt-mmm*ssssss* form must be used if there are multiple managed frames with the same user-defined name.

Either this option, the **-m** option, or the **-h** option is required. The **-e**, **-m**, and the **-h** options are mutually exclusive.

-h List the managed system dumps and managed frame dumps that are available on the HMC.

Either this option, the **-m** option, or the **-e** option is required. The **-h**, **-m**, and the **-e** options are mutually exclusive.

-r The type of dump resources to list. Valid values are **avail** for available dumps, and **parm** for system dump parameters. If this option is not specified, available dumps will be listed.

- s When listing dumps that are available on a managed frame, use this option to specify the side of the managed frame's bulk power assembly (BPA) for which available dumps are to be listed. Valid values are **a** for side A and **b** for side B.

When listing dumps that are available on a managed system, use this option to specify the service processor for which available dumps are to be listed. Valid values are **p** for the primary service processor and **s** for the secondary service processor. If this option is not specified, available dumps on the primary service processor will be listed. System dumps and resource dumps are only available from the primary service processor.

This option is not valid when listing the managed system dumps and managed frame dumps that are available on the HMC.

- filter The filter to apply to the available dumps to be listed. The filter is used to select the type of dump that is to be listed. If no filter is specified, then all of the available dumps will be listed.

The filter data consists of filter name/value pairs, which are in comma separated value (CSV) format. The filter data must be enclosed in double quotes.

The format of the filter data is as follows:

```
"filter-name=value,filter-name=value,..."
```

Valid filter names for this command:

dump_type

Only one value may be specified.

Valid values are:

pss - power subsystem dumps

sp - service processor dumps

sys - system dumps

resource - resource dumps

This option is not valid when listing the dumps that are available on a managed system or managed frame, or when listing the system dump parameters.

- F A delimiter separated list of attribute names representing the desired dump attribute values to display. If this option is specified without any attribute names, then all of the dump attributes will be displayed.

When this option is specified, only attribute values will be displayed. No attribute names will be displayed. The attribute values displayed will be separated by the delimiter which was specified with this option.

This option is useful when only attribute values are desired to be displayed, or when the values of only selected attributes are desired to be displayed.

--header

Display a header record, which is a delimiter separated list of attribute names for the attribute values that will be displayed. This header record will be the first record displayed. This option is only valid when used with the **-F** option.

- help Display the help text for this command and exit.

EXAMPLES

List the dumps that are currently available from the primary service processor on the managed system **9406-520*100128A**:

lsdump -m 9406-520*100128A

List the dumps that are currently available from the secondary service processor on the managed system **sys1**:

lsdump -m sys1 -s s

List the dumps that are currently available on side A of the BPA for the managed frame **frame1**:

lsdump -e frame1 -s a

List the managed system dumps and managed frame dumps that are currently available on the HMC:

lsdump -h

List only the managed system system dumps that are currently available on the HMC:

lsdump -h --filter "dump_type=sys"

List the system dump parameters for the managed system **sys1**:

lsdump -m sys1 -r parm

ENVIRONMENT

None

BUGS

None

AUTHOR

IBM Austin

SEE ALSO

cpdump, dump, getdump, rmdump, startdump

NAME

lsfru - list FRU information

SYNOPSIS

lsfru -t {**nsp** | **sp**} **-m** *managed-system*
[**-F** [*attribute-names*] [**--header**]] [**--help**]

DESCRIPTION

lsfru lists selected service processor field-replaceable unit (FRU) information for the *managed-system*.

This command is only supported for POWER6 and later servers.

OPTIONS

- t** The type of FRUs to list. Valid values are **nsp** for node service processors, and **sp** for primary and secondary service processors.
- m** The name of the managed system for which to list FRU information. The name may either be the user-defined name for the managed system, or be in the form *tttt-mmm*sssssss*, where *tttt* is the machine type, *mmm* is the model, and *sssssss* is the serial number of the managed system. The *tttt-mmm*sssssss* form must be used if there are multiple managed systems with the same user-defined name.
- F** A delimiter separated list of attribute names for the desired attribute values to be displayed for each FRU. If no attribute names are specified, then values for all of the attributes for each FRU will be displayed.

When this option is specified, only attribute values will be displayed. No attribute names will be displayed. The attribute values displayed will be separated by the delimiter which was specified with this option.

This option is useful when only attribute values are desired to be displayed, or when the values of only selected attributes are desired to be displayed.

--header

Display a header record, which is a delimiter separated list of attribute names for the attribute values that will be displayed. This header record will be the first record displayed. This option is only valid when used with the **-F** option.

- help** Display the help text for this command and exit.

EXAMPLES

List the node service processor FRU information for managed system **system1**:

```
lsfru -t nsp -m system1
```

List the primary and secondary service processor FRU information for managed system **9117-MMA*1234321**:

```
lsfru -t sp -m 9117-MMA*1234321
```

ENVIRONMENT

None

BUGS

None

AUTHOR

IBM Austin

NAME

lshmc - list Hardware Management Console configuration information

SYNOPSIS

```
lshmc {-b | -e | -h | -i | -l | -L | -n | -r | -v | -V | --firewall |  
--netroute | --syslog}  
[-F [attribute-names]] [--header] [--help]
```

DESCRIPTION

lshmc lists Hardware Management Console (HMC) configuration information. This command can be used to list the BIOS level, settings for the Event Manager for Call Home, hardware information, Integrated Management Module (IMM) settings, current locale, supported locales, network settings, remote access settings, Vital Product Data (VPD) information, version information, firewall settings, network routing information, and syslog server configuration for the HMC.

Only one of the following options can be specified at a time: **-b**, **-e**, **-h**, **-i**, **-l**, **-L**, **-n**, **-r**, **-v**, **-V**, **--firewall**, **--netroute**, or **--syslog**.

OPTIONS

- b** List the BIOS level of the HMC.
- e** List the settings for the HMC's Event Manager for Call Home.
- h** List HMC hardware information.
- i** List the Integrated Management Module (IMM) settings for the HMC.
- l** List the current locale for the HMC.
- L** List all of the locales supported for the HMC.
- n** List the network settings for the HMC.

The **lshmcldap** command can be used to display information about LDAP configuration settings.

- r** List the remote access settings for the HMC.
- v** List the VPD information for the HMC.
- V** List the version information for the HMC.

--firewall
List the firewall settings for the HMC.

--netroute
List the network routing information for the HMC.

--syslog
List the syslog server configuration for the HMC.

-F A delimiter separated list of attribute names representing the desired attribute values to display. If this option is specified without any attribute names, then all of the attributes for the type of HMC configuration information specified will be displayed.

When this option is specified, only attribute values will be displayed. No attribute names will be displayed. The attribute values displayed will be separated by the delimiter which was specified with this option.

This option is useful when only attribute values are desired to be displayed, or when the values of only selected attributes are desired to be displayed.

--header
Display a header record, which is a delimiter separated list of attribute names for the attribute values that will be displayed. This header record will be the first record displayed. This option is only valid when used with the **-F** option.

--help Display the help text for this command and exit.

EXAMPLES

Display the HMC's BIOS level:

lshmc -b

Display the settings for the HMC's Event Manager for Call Home:

lshmc -e

Display HMC hardware information:

lshmc -h

Display the HMC's IMM settings:

lshmc -i

Display the HMC's current locale:

lshmc -l

Display all locales supported by the HMC:

lshmc -L

Display the HMC's network settings:

lshmc -n

Display the HMC's host name and IP address, and separate the output values with a colon:

lshmc -n -F hostname:ipaddr

Display the HMC's remote access settings:

lshmc -r

Display the HMC's VPD information:

lshmc -v

Display the HMC's version information:

lshmc -V

Display the firewall settings for the HMC:

lshmc --firewall

Display the network routing information for the HMC:

lshmc --netroute

Display the syslog server configuration for the HMC:

lshmc --syslog

ENVIRONMENT

None

BUGS

None

AUTHOR

IBM Austin

SEE ALSO

chhmc

NAME

lshmcencr - list HMC encryptions

SYNOPSIS

lshmcencr -c { **passwd** | **webui** | **ssh** | **sshmac** } **-t** { **c** | **a** }
[-F [*attribute-names*] **--header**] **--help**]

DESCRIPTION

lshmcencr lists the encryptions which are available and the one which is currently used by the Hardware Management Console (HMC) to encrypt the passwords of locally authenticated HMC users. The current encryption is used when a new locally authenticated HMC user is created, or when the password for a locally authenticated HMC user is changed.

lshmcencr also lists the encryption ciphers which are available and those which can currently be used by the HMC Web user interface.

lshmcencr also lists the encryption ciphers and Message Authentication Code (MAC) algorithms which are available and those which can currently be used by the HMC Secure Shell (SSH) interface.

OPTIONS

- c** The encryption configuration to list. Valid values are **passwd** to list password encryptions for locally authenticated HMC users, **webui** to list encryption ciphers for the HMC Web user interface, **ssh** to list encryption ciphers for the HMC SSH interface, or **sshmac** to list MAC algorithms for the HMC SSH interface.
- t** The type of encryptions to list. Valid values are **c** for currently supported encryption(s), or **a** for all available encryptions.
- F** A delimiter separated list of attribute names for the desired attribute values to be displayed. If no attribute names are specified, then values for all of the attributes will be displayed.

When this option is specified, only attribute values will be displayed. No attribute names will be displayed. The attribute values displayed will be separated by the delimiter which was specified with this option.

This option is useful when only attribute values are desired to be displayed, or when the values of only selected attributes are desired to be displayed.

--header

Display a header record, which is a delimiter separated list of attribute names for the attribute values that will be displayed. This header record will be the first record displayed. This option is only valid when used with the **-F** option.

- help** Display the help text for this command and exit.

EXAMPLES

List the encryption which is currently used by the HMC to encrypt user passwords when locally authenticated HMC users are created, or when the passwords for locally authenticated HMC users are changed:

```
lshmcencr -c passwd -t c
```

List all available HMC password encryptions:

```
lshmcencr -c passwd -t a
```

List the encryption ciphers which can currently be used by the HMC Web user interface:

```
lshmcencr -c webui -t c
```

List all available encryption ciphers for the HMC Web user interface:

lshmcencr -c webui -t a

List the encryption ciphers which can currently be used by the HMC SSH interface:

lshmcencr -c ssh -t c

List all available encryption ciphers for the HMC SSH interface:

lshmcencr -c ssh -t a

List the MAC algorithms which can currently be used by the HMC SSH interface:

lshmcencr -c sshmac -t c

List all available MAC algorithms for the HMC SSH interface:

lshmcencr -c sshmac -t a

ENVIRONMENT

None

BUGS

None

AUTHOR

IBM Austin

SEE ALSO

chhmcencr

NAME

lshmcfs - list HMC file system information

SYNOPSIS

lshmcfs [-o **c** {-d *days* | -h *hours* | -s *size*}]
[-F [*attribute-names*] [--header]] [--help]

DESCRIPTION

lshmcfs lists Hardware Management Console (HMC) file system disk space usage information. Disk space usage information is listed for the HMC file systems that can contain temporary HMC files used for HMC and managed system firmware problem analysis.

This command lists information for the following file systems: **/var**, **/dump**, **/extra**, and **/**. The temporary files in the **/var** file system include HMC trace and log files. The temporary files in the **/dump** file system include managed system dumps, managed frame dumps, and debug data collected using the HMC **pedbg** command. The temporary files in the **/extra** file system include managed system dumps and managed frame dumps. The temporary files in the **/** file system include HMC Java core dump and heap dump files, and HMC trace files for Power system firmware updates.

All size and free space values displayed by this command are in megabytes.

OPTIONS

-o The operation to perform. The only valid value is **c** to calculate and list the amount of disk space that can be freed in each file system by removing temporary HMC files from the hard disk.

This command will not remove any files. The **chhmcfs** command can be used to remove the files.

If this option is omitted, current disk space usage information is listed.

-d List the amount of space that can be freed if files, which have not been modified during the specified number of *days* prior to now, were removed.

To list the amount of space that can be freed if all temporary HMC files, other than the trace and log files that are in use, were removed, specify **0** with this option.

Specifying **-d 1** is equivalent to specifying **-h 24**.

When **-o c** is specified to calculate and list the amount of disk space that can be freed, this option, the **-h** option, or the **-s** option must be specified to indicate which files should be considered for removal. The **-d**, **-h**, and **-s** options are mutually exclusive.

-h List the amount of space that can be freed if files, which have not been modified during the specified number of *hours* prior to now, were removed.

To list the amount of space that can be freed if all temporary HMC files, other than the trace and log files that are in use, were removed, specify **0** with this option.

When **-o c** is specified to calculate and list the amount of disk space that can be freed, this option, the **-d** option, or the **-s** option must be specified to indicate which files should be considered for removal. The **-d**, **-h**, and **-s** options are mutually exclusive.

-s List the amount of space that can be freed if files, starting with the oldest file, were removed from each file system to free up to *size* megabytes in each file system. The trace and log files that are in use would not be included.

When **-o c** is specified to calculate and list the amount of disk space that can be freed, this option, the **-d** option, or the **-h** option must be specified to indicate which files should be considered for removal. The **-d**, **-h**, and **-s** options are mutually exclusive.

-F A delimiter separated list of attribute names for the desired attribute values to be displayed for each file system. If no attribute names are specified, then values for all of the file system attributes will be displayed.

When this option is specified, only attribute values will be displayed. No attribute names will be displayed. The attribute values displayed will be separated by the delimiter which was specified with this option.

This option is useful when only attribute values are desired to be displayed, or when the values of only selected attributes are desired to be displayed.

--header

Display a header record, which is a delimiter separated list of attribute names for the attribute values that will be displayed. This header record will be the first record displayed. This option is only valid when used with the **-F** option.

--help Display the help text for this command and exit.

EXAMPLES

List current HMC file system disk space usage information:

lshmcfs

List HMC file system disk space usage information if temporary HMC files which have not been modified during the last 2 days (48 hours) were removed:

lshmcfs -o c -d 2

List HMC file system disk space usage information if all temporary HMC files, except the trace and log files that are in use, were removed:

lshmcfs -o c -d 0

List HMC file system disk space usage information if temporary HMC files were removed to free up to 100 MB in each file system:

lshmcfs -o c -s 100

ENVIRONMENT

None

BUGS

None

AUTHOR

IBM Austin

SEE ALSO

chhmcfs

NAME

lshmcldap - list LDAP information

SYNOPSIS

```
lshmcldap -r { config | user } [-v] [--filter "filter-data"]  
[-F [attribute-names] [--header]] [--help]
```

DESCRIPTION

lshmcldap lists Hardware Management Console (HMC) Lightweight Directory Access Protocol (LDAP) configuration information. **lshmcldap** can also list LDAP user information from the LDAP server.

OPTIONS

- r** The LDAP resources to list. Valid values are **config** to list HMC LDAP configuration information, and **user** to list LDAP user information from the LDAP server.
- v** When this option is specified, command execution details are output to stderr.
- filter** The filter to apply to the LDAP users to be listed. The filter is used to select which LDAP users to list. If no filter is used, then all LDAP users will be listed.

The filter data consists of a filter name/value pair, which is in comma separated value (CSV) format. The filter data must be enclosed in double quotes.

The format of the filter data is as follows:

```
""filter-name=value,value,...","..."
```

When a list of values is specified, the filter name/value pair must be enclosed in double quotes. Depending on the shell being used, nested double quote characters may need to be preceded by an escape character, which is usually a ``` character.

Valid filter names:

names

Specify one or more user names.

LDAP users whose login attribute values match the names specified will be listed. Which login attribute to use can be configured using the **chmcldap** command.

This option is only valid when listing LDAP user information.

- F** A delimiter separated list of attribute names for the desired attribute values to be displayed. If no attribute names are specified, then values for all of the attributes will be displayed.

When this option is specified, only attribute values will be displayed. No attribute names will be displayed. The attribute values displayed will be separated by the delimiter which was specified with this option.

This option is useful when only attribute values are desired to be displayed, or when the values of only selected attributes are desired to be displayed.

--header

Display a header record, which is a delimiter separated list of attribute names for the attribute values that will be displayed. This header record will be the first record displayed. This option is only valid when used with the **-F** option.

- help** Display the help text for this command and exit.

EXAMPLES

List the HMC LDAP configuration:

lshmcldap -r config

List all LDAP users on the LDAP server:

lshmcldap -r user

List LDAP users **uname1** and **uname2** and output command execution details to stderr:

```
lshmcldap -r user --filter ""names=uname1,uname2"" -v
```

ENVIRONMENT

None

BUGS

None

AUTHOR

IBM Austin

SEE ALSO

chhmldap, **lshmcusr**

NAME

lshmcusr - list Hardware Management Console user information

SYNOPSIS

```
lshmcusr [--script] [--filter "filter-data"]  
[-F [attribute-names] [--header]] [--help]
```

DESCRIPTION

lshmcusr lists Hardware Management Console (HMC) user information.

OPTIONS

--script

List managed resource objects in a format that can be used as input to the **chhmcusr**, **chaccfg**, and **mkaccfg** commands.

--filter

The filter to apply to the HMC users to be listed. A filter is used to select which HMC users are to be listed. If a filter is not used, then all HMC users will be listed.

The filter data consists of filter name/value pairs, which are in comma separated value (CSV) format. The filter data must be enclosed in double quotes.

The format of the filter data is as follows:

```
"filter-name=value,filter-name=value,..."
```

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

```
""filter-name=value,value,...","..."
```

When a list of values is specified, the filter name/value pair must be enclosed in double quotes. Depending on the shell being used, nested double quote characters may need to be preceded by an escape character, which is usually a `'\`` character.

Multiple values can be specified for each filter.

Valid filter names for this command:

names

Specify user name(s)

resources

Specify managed resource object(s)

resourceroles

Specify managed resource role name(s)

taskroles

Specify task role name(s)

password_encryptions

Specify password encryption(s)

-F

A delimiter separated list of attribute names for the desired attribute values to be displayed for each HMC user. If no attribute names are specified, then values for all of the attributes for each HMC user will be displayed.

When this option is specified, only attribute values will be displayed. No attribute names will be displayed. The attribute values displayed will be separated by the delimiter which was specified with this option.

This option is useful when only attribute values are desired to be displayed, or when the values of only selected attributes are desired to be displayed.

--header

Display a header record, which is a delimiter separated list of attribute names for the attribute values that will be displayed. This header record will be the first record displayed. This option is only valid when used with the **-F** option.

--help Display the help text for this command and exit.

EXAMPLES

List all HMC users:

lshmcusr

List only the user names and managed resource roles for all HMC users, and separate the output values with a colon:

lshmcusr -F name:resourcerole

List the HMC users hscroot and user1:

lshmcusr --filter ""names=hscroot,user1""

List the HMC users with the task role hmcviewer and the managed resource role mr1:

lshmcusr --filter "taskroles=hmcviewer,resourceroles=mr1"**ENVIRONMENT**

None

BUGS

None

AUTHOR

IBM Austin

SEE ALSO

chhmcusr, mkhmcusr, rmhmcusr, lsaccfg

NAME

lshwinfo - list environmental information

SYNOPSIS

```
lshwinfo -r { frame | sys } -e managed-frame [-s { a | b }]
[-F [attribute-names] [--header]] [--help]
```

DESCRIPTION

lshwinfo lists environmental information, such as input power levels, for the *managed-frame*. The **lshwinfo** command also lists environmental information, such as ambient air temperatures, for systems contained in the *managed-frame*.

OPTIONS

- r** The type of resource for which to list environmental information. Valid values are **frame** for managed frame and **sys** for systems contained in a managed frame.
- e** The name of the managed frame for which to list environmental information. The name may either be the user-defined name for the managed frame, or be in the form *ttt-mmm*sssssss*, where *ttt* is the type, *mmm* is the model, and *sssssss* is the serial number of the managed frame. The *ttt-mmm*sssssss* form must be used if there are multiple managed frames with the same user-defined name.
- s** Use this option to specify the side of the managed frame's bulk power assembly (BPA) for which to list environmental information. Valid values are **a** for side A and **b** for side B. If this option is omitted, environmental information for side A will be listed.

This option is only valid when listing environmental information for a managed frame.

- F** A delimiter separated list of attribute names for the desired attribute values to be displayed. If no attribute names are specified, then values for all of the attributes will be displayed.

When this option is specified, only attribute values will be displayed. No attribute names will be displayed. The attribute values displayed will be separated by the delimiter which was specified with this option.

This option is useful when only attribute values are desired to be displayed, or when the values of only selected attributes are desired to be displayed.

--header

Display a header record, which is a delimiter separated list of attribute names for the attribute values that will be displayed. This header record will be the first record displayed. This option is only valid when used with the **-F** option.

- help** Display the help text for this command and exit.

EXAMPLES

Display environmental information for side A of the BPA for the managed frame **myFrame**:

```
lshwinfo -r frame -e myFrame
```

Display environmental information for side B of the BPA for the managed frame **myFrame**:

```
lshwinfo -r frame -e myFrame -s b
```

Display environmental information for the systems contained in the managed frame **9119-595*020000A**:

```
lshwinfo -r sys -e 9119-595*020000A
```

ENVIRONMENT

None

BUGS

None

AUTHOR

IBM Austin

NAME

lshwres - list hardware resources

SYNOPSIS

To list physical I/O resources:

```
lshwres -r io --rsubtype {unit | bus | slot |  
iopool | taggedio | slotchildren}  
-m managed-system  
[--level {pool | sys}] [-R]  
[--filter "filter-data"]  
[-F [attribute-names] [--header]] [--help]
```

To list virtual I/O resources:

```
lshwres -r virtualio  
--rsubtype {eth | fc | hsl | virtualopti | scsi | serial |  
slot | vnetwork | vnic | vnicbkdev | vswitch}  
-m managed-system  
[--level {lpar | slot | sys}]  
[--filter "filter-data"]  
[-F [attribute-names] [--header]] [--help]
```

To list memory resources:

```
lshwres -r mem -m managed-system  
--level {lpar | sys} [-R]  
[--maxmem quantity [--hpratio ratio]]  
[--filter "filter-data"]  
[-F [attribute-names] [--header]] [--help]
```

To list I/O entitled memory statistics for a partition:

```
lshwres -r mem -m managed-system --level lpar  
--filter "filter-data" --stat  
[-F [attribute-names] [--header]] [--help]
```

To list shared memory pool resources:

```
lshwres -r mempool -m managed-system [--rsubtype pgdev]  
[--filter "filter-data"]  
[-F [attribute-names] [--header]] [--help]
```

To list reserved storage device pool resources (only supported on managed systems that support partition suspend and resume operations):

```
lshwres -r rspool -m managed-system [--rsubtype rsdev]  
[--filter "filter-data"]  
[-F [attribute-names] [--header]] [--help]
```

To list processing resources:

```
lshwres -r proc -m managed-system  
--level {lpar | pool | sys} [-R]  
[--filter "filter-data"]  
[-F [attribute-names] [--header]] [--help]
```

To list shared processor pool resources on a managed system that supports multiple shared processor pools:

```
lshwres -r procpool -m managed-system  
[--filter "filter-data"]  
[-F [attribute-names] [--header]] [--help]
```

To list Host Channel adapter (HCA) resources:

```
lshwres -r hca -m managed-system  
--level {lpar | sys}  
[-filter "filter-data"]  
[-F [attribute-names] [--header] [--help]]
```

To list Host Ethernet adapter (HEA) resources:

```
lshwres -r hea -m managed-system  
--subtype {logical | phys}  
--level {port | port_group | sys}  
[-R] [--stat]  
[-filter "filter-data"]  
[-F [attribute-names] [--header] [--help]]
```

To list Switch Network Interface (SNI) adapter resources:

```
lshwres -r sni -m managed-system  
[-filter "filter-data"]  
[-F [attribute-names] [--header] [--help]]
```

To list Single Root I/O Virtualization (SR-IOV) adapter resources:

```
lshwres -r sriov -m managed-system  
--subtype {adapter | logport | physport}  
[-level {eth | ethc}] [-R]  
[-filter "filter-data"]  
[-F [attribute-names] [--header] [--help]]
```

DESCRIPTION

lshwres lists the hardware resources of the *managed-system*, including physical I/O, virtual I/O, memory, shared memory pool, processing, shared processor pool, Host Channel adapter (HCA), Host Ethernet adapter (HEA), Switch Network Interface (SNI) adapter, and Single Root I/O Virtualization (SR-IOV) resources.

OPTIONS

-r The type of hardware resources to list. Valid values are **io** for physical I/O, **virtualio** for virtual I/O, **mem** for memory, **mempool** for shared memory pool, **rspool** for reserved storage device pool, **proc** for processing, **procpool** for shared processor pool, **hca** for HCA, **hea** for HEA, **sni** for SNI adapter, and **sriov** for SR-IOV resources.

The **procpool** resource type is only supported for managed systems that support multiple shared processor pools.

The **rspool** resource type is only supported for managed systems that support partition suspend and resume operations.

--subtype

The subtype of hardware resources to list. Valid physical I/O resource subtypes are **unit** for I/O units, **bus** for I/O buses, **slot** for I/O slots, **iopool** for I/O pools, **taggedio** for tagged I/O, and **slotchildren** for I/O slot children resources. Valid virtual I/O resource subtypes are **eth** for virtual ethernet, **fc** for virtual fibre channel, **hsl** for High Speed Link (HSL) OptiConnect, **virtualopti** for virtual OptiConnect, **scsi** for virtual SCSI, **serial** for virtual serial, **slot** for virtual slot, **vnetwork** for virtual network, **vnic** for virtual NIC, **vnickbdev** for virtual NIC backing devices, and **vswitch** for virtual switch resources. Valid HEA resource subtypes are **logical** for logical HEA resources, and **phys** for physical HEA resources. Valid SR-IOV resource subtypes are **adapter** for SR-IOV adapter resources, **logport** for SR-IOV logical port resources, and **physport** for SR-IOV physical port resources.

The only valid shared memory pool resource subtype is **pgdev** for paging space devices. The only valid reserved storage device pool subtype is **rsdev** for reserved storage devices. On a managed system that supports Active Memory Sharing and partition suspend and resume operations, the command **lshwres -m managed-system -r mempool --rsubtype pgdev** lists the devices in the reserved storage device pool, and is equivalent to the command **lshwres -m managed-system -r rspool --rsubtype rsdev**.

This option is required when listing physical I/O, virtual I/O, HEA, or SR-IOV resources. This option is not valid when listing memory, processing, shared processor pool, HCA, or SNI adapter resources.

- m** The name of the managed system which has the hardware resources to list. The name may either be the user-defined name for the managed system, or be in the form *tttt-mmm*ssssss*, where *tttt* is the machine type, *mmm* is the model, and *ssssss* is the serial number of the managed system. The *tttt-mmm*ssssss* form must be used if there are multiple managed systems with the same user-defined name.
- level** The level of information to list. Valid values are **lpar** for partition, **pool** for pool, **port** for port, **port_group** for port group, **slot** for slot, **sys** for system, **eth** for ethernet, and **ethc** for converged ethernet.

This option is required when listing I/O pool resources, virtual ethernet, fibre channel, serial, or slot resources, memory, processing, HCA, or HEA resources, or SR-IOV physical ports.

Valid levels for I/O pool resources are **pool** or **sys**. Valid levels for virtual ethernet resources are **lpar** or **sys**. Valid levels for virtual fibre channel resources are **lpar** or **sys**. Valid levels for virtual serial resources are **lpar**. Valid levels for virtual slot resources are **lpar**, **slot**, or **sys**. Valid levels for memory resources are **lpar** or **sys**. Valid levels for processing resources are **lpar**, **pool**, or **sys**. Valid levels for HCA resources are **lpar** or **sys**. Valid levels for physical HEA resources are **port**, **port_group**, or **sys**. Valid levels for logical HEA resources are **port** or **sys**. Valid levels for SR-IOV logical ports are **eth**. Valid levels for SR-IOV physical ports are **eth** or **ethc**.

- R** Only list information for partitions with hardware resources that can be restored due to a dynamic logical partitioning (DLPAR) operation failure.

The **rsthwres** command can be used to restore those hardware resources.

This option is only valid for listing physical I/O slots, partition level memory or processing resources, logical HEA resources, or SR-IOV logical ports.

--maxmem

When this option is specified, the required minimum memory amount needed for partitions to support the maximum memory *quantity* and Hardware Page Table *ratio* specified is listed. All memory quantities are in megabytes, and are a multiple of the memory region size for the *managed-system*.

This information is useful for specifying memory amounts in partition profiles.

The required minimum memory amount listed is the minimum memory amount required for a partition by the managed system's system firmware. An operating system installed on the partition may require more memory than the amount listed.

This option is only valid when listing system level memory resources.

--hptratio

The ratio of the Hardware Page Table (HPT) size to the maximum memory for a partition. Valid values for the *managed-system* are displayed by the **lshwres -r mem -m managed-system --level**

sys -F possible_hpt_ratios command.

If this option is not used to specify an HPT ratio, the default HPT ratio for the *managed-system* is used.

This option is only valid when used with the **--maxmem** option.

--stat When this option is specified when listing HEA physical port resources, port counter statistics for HEA physical ports are listed.

When this option is specified when listing partition level memory resources, I/O entitled memory statistics are listed.

This option is only valid when listing HEA physical port resources or partition level memory resources.

--filter The filter(s) to apply to the hardware resources to be listed. Filters are used to select which hardware resources of the specified type are to be listed. If no filters are used, then all of the hardware resources of the specified type will be listed. For example, all of the physical I/O slots on a specific I/O unit and bus can be listed by using a filter to specify the I/O unit and the bus which has the slots to list. Otherwise, if no filter is used, then all of the physical I/O slots in the managed system will be listed.

The filter data consists of filter name/value pairs, which are in comma separated value (CSV) format. The filter data must be enclosed in double quotes.

The format of the filter data is as follows:

```
"filter-name=value,filter-name=value,..."
```

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

```
""filter-name=value,value,...",..."
```

When a list of values is specified, the filter name/value pair must be enclosed in double quotes. Depending on the shell being used, nested double quote characters may need to be preceded by an escape character, which is usually a ``\`` character.

Unless otherwise indicated, multiple values can be specified for each filter.

When listing I/O entitled memory statistics for a partition, this parameter is required, and exactly one partition name or ID must be specified with the **lpar_names** or **lpar_ids** filter.

Valid filter names for this command:

adapter_ids

Specify HCA, HEA, or SR-IOV adapter ID(s)

buses

Specify I/O bus ID(s)

logical_port_ids

Specify SR-IOV logical port ID(s)

lpar_ids

Specify partition ID(s)

lpar_names

Specify partition user-defined name(s)

phys_port_ids

Specify SR-IOV physical port ID(s)

pool_names
Specify pool name(s)

pools
Specify pool ID(s)

port_groups
Specify HEA port group(s)

slots
Specify physical I/O slot DRC index(ices)
or virtual I/O slot number(s)

sni_device_ids
Specify SNI adapter device ID(s)

units
Specify I/O unit physical location
code(s)

vlans
Specify virtual LAN ID(s)

vswitches
Specify virtual switch name(s)

Valid filters with **-r io --rsubtype unit**:
units

Valid filters with **-r io --rsubtype bus**:
buses, units

Valid filters with **-r io --rsubtype slot**:
**buses, lpar_ids | lpar_names, pools, slots,
units**

Valid filters with **-r io --rsubtype iopool --level pool**:
lpar_ids | lpar_names, pools

Valid filters with **-r io --rsubtype taggedio**:
lpar_ids | lpar_names

Valid filters with **-r io --rsubtype slotchildren**:
lpar_ids | lpar_names, slots

Valid filters with **-r virtualio --rsubtype eth --level lpar**:
lpar_ids | lpar_names, slots, vlans, vswitches

Valid filters with **-r virtualio --rsubtype fc --level lpar**:
lpar_ids | lpar_names, slots

Valid filters with **-r virtualio --rsubtype hsl**:
lpar_ids | lpar_names, pools

Valid filters with **-r virtualio --rsubtype virtualopti**:
lpar_ids | lpar_names, pools

Valid filters with **-r virtualio --rsubtype scsi**:
lpar_ids | lpar_names, slots

Valid filters with **-r virtualio --rsubtype serial --level lpar**:

lpar_ids | lpar_names, slots

Valid filters with **-r virtualio --subtype slot --level lpar:**
lpar_ids | lpar_names

Valid filters with **-r virtualio --subtype slot --level slot:**
lpar_ids | lpar_names, slots

Valid filters with **-r virtualio --subtype vnic:**
lpar_ids | lpar_names, slots

Valid filters with **-r virtualio --subtype vnicbkdev:**
lpar_ids | lpar_names

Valid filters with **-r virtualio --subtype vswitch:**
vlans, vswitches

Valid filters with **-r mem --level lpar:**
lpar_ids | lpar_names

Valid filters with **-r mempool --subtype pgdev:**
lpar_ids | lpar_names
Specify **none** for **lpar_ids** to list all
devices that are not assigned to a partition

Valid filters with **-r rspool --subtype rsdev:**
lpar_ids | lpar_names
Specify **none** for **lpar_ids** to list all
devices that are not assigned to a partition

Valid filters with **-r proc --level lpar:**
lpar_ids | lpar_names, pools | pool_names

Valid filters with **-r procpool:**
pools | pool_names

Valid filters with **-r hca --level sys:**
adapter_ids

Valid filters with **-r hca --level lpar:**
adapter_ids, lpar_ids | lpar_names

Valid filters with **-r hea --subtype logical --level sys:**
adapter_ids, lpar_ids | lpar_names

Valid filters with **-r hea --subtype logical --level port:**
adapter_ids, lpar_ids | lpar_names,
port_groups

Valid filters with **-r hea --subtype phys --level sys:**
adapter_ids

Valid filters with **-r hea --subtype phys --level port:**
adapter_ids, port_groups

Valid filters with **-r hea --rsubtype phys --level port_group**:
adapter_ids, port_groups

Valid filters with **-r sni**:
lpar_ids | lpar_names, sni_device_ids

valid filters with **-r sriov --rsubtype adapter**:
adapter_ids

valid filters with **-r sriov --rsubtype logport**:
adapter_ids, logical_port_ids

valid filters with **-r sriov --rsubtype logport --level**:
adapter_ids, logical_port_ids, phys_port_ids,
lpar_ids | lpar_names

valid filters with **-r sriov --rsubtype logport -R**:
adapter_ids, logical_port_ids,
lpar_ids | lpar_names

Valid filters with **-r sriov --rsubtype physport**:
adapter_ids, phys_port_ids

-F A delimiter separated list of attribute names for the desired attribute values to be displayed for each hardware resource. If no attribute names are specified, then values for all of the attributes for each hardware resource will be displayed.

When this option is specified, only attribute values will be displayed. No attribute names will be displayed. The attribute values displayed will be separated by the delimiter which was specified with this option.

This option is useful when only attribute values are desired to be displayed, or when the values of only selected attributes are desired to be displayed.

Descriptions of selected command attributes:

topology

A comma separated list of the end-to-end device mappings for each virtual device. This is displayed for virtual SCSI and virtual fibre channel adapters and for SR-IOV logical ports when specified with the **-F** option.

For virtual SCSI, each device mapping has the following format:

*virtual-device/state/virtual-SCSI-client-device/
virtual-SCSI-server-adapter/storage-pool/
backing-device*

where *virtual-device* is the name of the virtual device on the client operating system (will be an empty string for virtual SCSI server adapters), *state* is the state of the virtual device on the client operating system (will be an empty string for virtual SCSI server adapters), *virtual-SCSI-client-device* is the

name of the virtual SCSI client device on the client operating system (will be an empty string for virtual SCSI server adapters), *virtual-SCSI-server-adapter* is the name of the virtual SCSI server adapter on the Virtual I/O Server (VIOS), *storage-pool* is the name of the storage pool containing the backing device on the VIOS (will be an empty string if the backing device is not in a storage pool), and *backing-device* is the name of the backing device on the VIOS.

For virtual fibre channel, each device mapping has the following format:

```
virtual-device/state/virtual-FC-client-device/  
virtual-FC-server-adapter/physical-port
```

where *virtual-device* is the name of the virtual device on the client operating system (will be an empty string for virtual fibre channel server adapters), *state* is the state of the virtual device on the client operating system (will be an empty string for virtual fibre channel server adapters), *virtual-FC-client-device* is the name of the virtual fibre channel client device on the client operating system (will be an empty string for virtual fibre channel server adapters), *virtual-FC-server-adapter* is the name of the virtual fibre channel server adapter on the VIOS, and *physical-port* is the name of the physical fibre channel port to which this device is backed.

For SR-IOV logical ports, each device mapping has the following format:

```
device-name/adaptor-id/physical-port-id/  
logical-port-id
```

where *device-name* is the name of the OS device, *adaptor-id* is the SR-IOV adapter ID, *physical-port-id* is the SR-IOV physical port ID, and *logical-port-id* is the SR-IOV logical port ID.

--header

Display a header record, which is a delimiter separated list of attribute names for the attribute values that will be displayed. This header record will be the first record displayed. This option is only valid when used with the **-F** option.

--help Display the help text for this command and exit.

EXAMPLES

List all I/O units on the managed system:

```
lshwres -r io --rsubtype unit -m system1
```

List all buses on I/O unit **U787A.001.0395036**:

```
lshwres -r io --rsubtype bus -m 9406-570*12345678
```

```
--filter "units=U787A.001.0395036"
```

List only the DRC index, description, and the owning partition for each physical I/O slot on buses **2** and **3** of I/O unit **U787A.001.0395036**:

```
lshwres -r io --subtype slot -m system1 --filter  
"units=U787A.001.0395036,"buses=2,3"" -F drc_index,  
description,lp_name
```

List all I/O pools and the partitions and slots assigned to each I/O pool:

```
lshwres -r io --subtype iopool -m system1 --level pool
```

List the tagged I/O devices for the IBM i partition that has an ID of 1:

```
lshwres -r io --subtype taggedio -m 9406-520*100103A  
--filter "lp_ids=1"
```

List the I/O slot children for the I/O slot with DRC index **21010207**:

```
lshwres -r io --subtype slotchildren -m 9117-MMA*107791F  
--filter "slots=21010207"
```

List all virtual ethernet adapters on the managed system:

```
lshwres -r virtualio --subtype eth --level lp -m system1
```

List all virtual SCSI adapters on the managed system, and only display attribute values for each adapter, following a header of attribute names:

```
lshwres -r virtualio --subtype scsi -m system1 -F --header
```

List all virtual slots for partition lp1:

```
lshwres -r virtualio --subtype slot -m system1 --level  
slot --filter "lp_names=lp1"
```

List all virtual networks in the managed system:

```
lshwres -r virtualio --subtype vnetwork -m mySys
```

List all virtual NICs assigned to partition lp1:

```
lshwres -r virtualio --subtype vnic -m mySys  
--filter "lp_names=lp1"
```

List all virtual NIC backing devices in the managed system:

```
lshwres -r virtualio --subtype vnicbkdev -m mySys
```

List all virtual switches in the managed system:

```
lshwres -r virtualio --subtype vswitch -m system1
```

List system level memory information:

```
lshwres -r mem -m 9406-570*98765432 --level sys
```

List recoverable memory information:

```
lshwres -r mem -m 9406-570*98765432 --level lpar -R
```

List memory information for partitions lpar1 and lpar2:

```
lshwres -r mem -m system1 --level lpar --filter  
"'lpar_names=lpar_1,lpar_2'"
```

List only the installed and configurable processors on the system, and separate the output values with a colon:

```
lshwres -r proc -m 9406-570*98765432 --level sys -F  
installed_sys_proc_units:configurable_sys_proc_units
```

List I/O entitled memory statistics for partition lpar1:

```
lshwres -r mem -m system1 --level lpar --filter "lpar_names=lpar_1"  
--stat
```

List the shared memory pool:

```
lshwres -r mempool -m system1
```

List all of the devices in the shared memory pool (on managed systems that support partition suspend and resume operations, this lists all of the devices in the reserved storage device pool):

```
lshwres -r mempool -m system1 --rsubtype pgdev
```

List the reserved storage device pool:

```
lshwres -r rspool -m 8233-E8B*1234321
```

List all of the devices in the reserved storage device pool:

```
lshwres -r rspool -m system1 --rsubtype rsdev
```

List processing resources for all partitions:

```
lshwres -r proc -m system1 --level lpar
```

List the physical processor pool:

```
lshwres -r proc -m system1 --level pool
```

List all configured shared processor pools in the managed system:

```
lshwres -r procpool -m system1
```

List all HCAs on the managed system:

lshwres -r hca -m 9117-570*12345678 --level sys

List all HCAs assigned to partition AIX1:

**lshwres -r hca -m 9117-570*12345678 --level lpar
--filter "lpar_names=AIX1"**

List all physical HEAs on the managed system:

lshwres -r hea -m mySys --subtype phys --level sys

List all port groups for all HEAs on the managed system:

lshwres -r hea -m mySys --subtype phys --level port_group

List all physical ports belonging to port group 2 for the HEA with adapter ID 23000010:

**lshwres -r hea -m 9117-MMA*1234ABC --subtype phys
--level port --filter "adapter_ids=23000010,
port_groups=2"**

List all Logical Host Ethernet adapters (LHEA) on the managed system:

lshwres -r hea -m mySys --subtype logical --level sys

List all HEA logical ports assigned to partition p1:

**lshwres -r hea -m mySys --subtype logical --level port
--filter "lpar_names=p1"**

List all SNI adapters on the managed system:

lshwres -r sni -m system1

List all SR-IOV adapters on the managed system:

lshwres -r sriov -m system1 --subtype adapter

List all unconfigured SR-IOV logical ports:

lshwres -r sriov -m system1 --subtype logport

List all recoverable SR-IOV logical ports:

lshwres -r sriov -m system1 --subtype logport -R

List all SR-IOV ethernet logical ports:

lshwres -r sriov -m system1 --subtype logport --level eth

List all SR-IOV ethernet physical ports:

lshwres -r sriov -m system1 --subtype physport --level eth

List all SR-IOV converged ethernet physical ports:

lshwres -r sriov -m system1 --subtype physport --level ethc

ENVIRONMENT

None

BUGS

None

AUTHOR

IBM Austin

SEE ALSO

chhwres, rsthwres, lssyscfg

NAME

lsiotopo - list I/O topology

SYNOPSIS

```
lsiotopo -m managed-system [-r {slot | hea}] [--filter "filter-data"]  
[-F [attribute-names] [--header]] [--help]
```

DESCRIPTION

lsiotopo lists the physical I/O slot and logical Host Ethernet Adapter (HEA) topology of the *managed-system*.

lsiotopo is useful for displaying the I/O topology of multi-node managed systems.

OPTIONS

- m** The name of the managed system for which to list the I/O topology. The name may either be the user-defined name for the managed system, or be in the form *ttt-mmm*sssssss*, where *ttt* is the machine type, *mmm* is the model, and *sssssss* is the serial number of the managed system. The *ttt-mmm*sssssss* form must be used if there are multiple managed systems with the same user-defined name.
- r** The type of I/O resources to list. Valid values are **slot** to list physical I/O slot topology only, and **hea** to list logical HEA topology only. If this option is omitted, both the physical I/O slot and logical HEA topology will be listed.
- filter** The filters to apply to the I/O resources to be listed. Filters are used to select which I/O resources are to be listed. For example, to list only the physical I/O slots in a specific system enclosure and fabric bus, use the **sys_enclosures** and **fabric_bus_ids** filters.

The filter data consists of filter name/value pairs, which are in comma separated value (CSV) format. The filter data must be enclosed in double quotes.

The format of the filter data is as follows:

```
"filter-name=value,filter-name=value,..."
```

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

```
""filter-name=value,value,...",..."
```

When a list of values is specified, the filter name/value pair must be enclosed in double quotes. Depending on the shell being used, nested double quote characters may need to be preceded by an escape character, which is usually a ``` character.

Multiple values can be specified for each filter.

Valid filter names for this command:

sys_enclosures

Specify system enclosure location codes

slot_enclosures

Specify I/O resource enclosure location codes

fabric_bus_ids

Specify fabric bus IDs

owner_lpar_names

Specify I/O resource owner partition user-defined names

owner_lpar_ids

Specify I/O resource owner partition IDs

-F A delimiter separated list of attribute names for the desired attribute values to be displayed for each I/O resource. If no attribute names are specified, then values for all of the attributes for each I/O resource will be displayed.

When this option is specified, only attribute values will be displayed. No attribute names will be displayed. The attribute values displayed will be separated by the delimiter which is specified with this option.

This option is useful when only attribute values are desired to be displayed, or when the values of only selected attributes are desired to be displayed.

--header

Display a header record, which is a delimiter separated list of attribute names for the attribute values that will be displayed. This header record will be the first record displayed. This option is only valid when used with the **-F** option.

--help Display the help text for this command and exit.

EXAMPLES

List the physical I/O slot topology of the managed system:

lsiotopo -m system1 -r slot

List the physical I/O slot and logical HEA topology of the managed system:

lsiotopo -m system1

List only the system enclosure, slot enclosure, owner partition ID, and description, separated by a space, for each physical I/O slot on fabric buses **1920** and **1672** of system enclosure **U78A2.001.992002Y-P9**:

**lsiotopo -m system1 -r slot -F sys_enclosure slot_enclosure
owner_lpar_id description --filter "sys_enclosures=
U78A2.001.992002Y-P9,"fabric_bus_ids=1920,1672""**

ENVIRONMENT

None

BUGS

None

AUTHOR

IBM Austin

SEE ALSO

lshwres

NAME

lsipsec - list IPsec connection information

SYNOPSIS

```
lsipsec {-l | -c connection-name | --status connection-name |  
--statusall}  
[-F [attribute-names] [--header]] [--help]
```

DESCRIPTION

lsipsec lists Internet Protocol Security (IPsec) connection information for the Hardware Management Console.

OPTIONS

- l** Specify this option to list all existing IPsec connection names.

This option cannot be specified with the **-c**, **--status**, or **--statusall** options.
- c** Specify this option to list connection configuration details for *connection-name*.

This option cannot be specified with the **-l**, **--status**, or **--statusall** options.
- status** Specify this option to list connection status for *connection-name*.

This option cannot be specified with the **-l**, **-c**, or **--statusall** options.
- statusall** Specify this option to list IPsec daemon status and detailed connection status for all IPsec connections.

This option cannot be specified with the **-l**, **-c**, or **--status** options.
- F** A delimiter separated list of attribute names for the desired attribute values to be displayed. If no attribute names are specified, then values for all of the attributes will be displayed.

When this option is specified, only attribute values will be displayed. No attribute names will be displayed. The attribute values displayed will be separated by the delimiter which was specified with this option.

This option is useful when only attribute values are desired to be displayed, or when the values of only selected attributes are desired to be displayed.
- header** Display a header record, which is a delimiter separated list of attribute names for the attribute values that will be displayed. This header record will be the first record displayed. This option is only valid when used with the **-F** option.
- help** Display the help text for this command and exit.

EXAMPLES

List all existing IPsec connection names:

```
lsipsec -l
```

List the configuration of IPsec connection **sys1**:

```
lsipsec -c sys1
```

List the status of IPsec connection **j21**:

lsipsec --status j21

List detailed status information for the IPsec daemon and all IPsec connections:

lsipsec --statusall

ENVIRONMENT

None

BUGS

None

AUTHOR

IBM Austin

SEE ALSO

chipsec

NAME

lsled - list LEDs

SYNOPSIS

```
lsled -r sa -t {phys | virtuallpar | virtualsys}
-m managed-system [--filter "filter-data"]
-F [attribute-names] [--header] [--help]
```

DESCRIPTION

lsled lists LED information for the *managed-system*.

OPTIONS

- r** The type of LED resources to list. The only valid value is **sa** for System Attention (SA) LEDs.
- t** The type of System Attention (SA) LEDs to list. Valid values are **phys** for the physical SA LED for the *managed-system*, **virtuallpar** for virtual partition SA LEDs, or **virtualsys** for the virtual SA LED for the *managed-system*.
- m** The name of the managed system which has the LEDs to list. The name may either be the user-defined name for the managed system, or be in the form *tttt-mmm*sssssss*, where *tttt* is the machine type, *mmm* is the model, and *sssssss* is the serial number of the managed system. The *tttt-mmm*sssssss* form must be used if there are multiple managed systems with the same user-defined name.
- filter** The filter to apply to the virtual partition SA LEDs to be listed. A filter is used to select which virtual partition SA LEDs are to be listed. If no filter is used, then all of the virtual partition SA LEDs for the *managed-system* will be listed.

The filter data consists of filter name/value pairs, which are in comma separated value (CSV) format. The filter data must be enclosed in double quotes.

The format of the filter data is as follows:

```
"filter-name=value,filter-name=value,..."
```

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

```
""filter-name=value,value,...",..."
```

When a list of values is specified, the filter name/value pair must be enclosed in double quotes. Depending on the shell being used, nested double quote characters may need to be preceded by an escape character, which is usually a `'\'` character.

Valid filter names for this command are **lpar_names** or **lpar_ids**. Only one of these filters may be specified. Multiple partitions can be specified with the filter.

This option is not valid when listing the physical SA LED or the virtual SA LED for the *managed-system*.

- F** A delimiter separated list of attribute names for the desired attribute values to be displayed for each LED. If no attribute names are specified, then values for all of the attributes for each LED will be displayed.

When this option is specified, only attribute values will be displayed. No attribute names will be displayed. The attribute values displayed will be separated by the delimiter which was specified with this option.

This option is useful when only attribute values are desired to be displayed, or when the values of only selected attributes are desired to be displayed.

--header

Display a header record, which is a delimiter separated list of attribute names for the attribute values that will be displayed. This header record will be the first record displayed. This option is only valid when used with the **-F** option.

--help Display the help text for this command and exit.

EXAMPLES

Display the physical SA LED for the system:

```
lsled -m 9406-579*100103C -r sa -t phys
```

Display all of the virtual partition SA LEDs:

```
lsled -m system1 -r sa -t virtualpar
```

Display the virtual partition SA LEDs for partitions **lpar1** and **lpar2**:

```
lsled -m system1 -r sa -t virtualpar --filter  
"'lpar_names=lpar1,lpar2'"
```

ENVIRONMENT

None

BUGS

None

AUTHOR

IBM Austin

SEE ALSO

chled

NAME

lslic - list Licensed Internal Code levels

SYNOPSIS

```
lslic {-m managed-system | -e managed-frame [--allsystems] | -w | -c}
[-t {sys | power | syspower | powerfru | sriov}]
[-r {ibmwebsite | ftp | sftp | dvd | disk | mountpoint | usb}]
[-h host-name] [-u user-ID] [--passwd password] [-k SSH-private-key]
[-d directory]
[-F [attribute-names] [--header]] [--help]
```

DESCRIPTION

lslic lists Licensed Internal Code (LIC) levels installed, activated, and accepted. The LIC levels that are available to be retrieved from a repository can also be listed.

lslic can also list the LIC levels currently in the management console hard disk repository.

OPTIONS

- m** The name of the managed system for which LIC levels are to be listed. The name may either be the user-defined name for the managed system, or be in the form *tttt-mmm*sssssss*, where *tttt* is the machine type, *mmm* is the model, and *sssssss* is the serial number of the managed system. The *tttt-mmm*sssssss* form must be used if there are multiple managed systems with the same user-defined name.
- e** The name of the managed frame for which LIC levels are to be listed. The name may either be the user-defined name for the managed frame, or be in the form *tttt-mmm*sssssss*, where *tttt* is the machine type, *mmm* is the model, and *sssssss* is the serial number of the managed frame. The *tttt-mmm*sssssss* form must be used if there are multiple managed frames with the same user-defined name.
- allsystems**
List Licensed Internal Code levels on all managed systems in the managed frame specified with the **-e** option.
- w** List Licensed Internal Code levels for all managed frames which contain High Performance Switches.
- c** List Licensed Internal Code levels currently in the management console hard disk repository.
- t** The type of LIC levels to list. Valid values are **sys** for Managed System LIC levels only, **power** for Power LIC levels only, **syspower** for both Managed System and Power LIC levels, **powerfru** for Power FRU levels, or **sriov** for Single Root I/O Virtualization (SR-IOV) adapter levels.

If this option is omitted, it will default to **syspower** for the **-m** and **-e --allsystems** options, and to **power** for the **-e** and **-w** options. This option is not allowed with the **-c** option.
- r** The repository to query for LIC levels that are available to be retrieved. Valid values are **ibmwebsite** for the IBM service website, **ftp** for a remote FTP site, **sftp** for a remote secure FTP (SFTP) site, **dvd** for the DVD drive on the management console, **disk** for the internal hard disk drive on the management console, **mountpoint** for the specified mount point on the internal hard disk drive on the management console, or **usb** for a USB flash memory device.
- h** The host name or IP address of the remote FTP or SFTP server.

This option is required when displaying LIC levels that are available to be retrieved from a remote FTP or SFTP site. This option is not valid otherwise.
- u** The user ID to use to log in to the remote FTP or SFTP site.

This option is required when displaying LIC levels that are available to be retrieved from a remote FTP or SFTP site. This option is not valid otherwise.

--passwd

The password to use to log in to the remote FTP or SFTP site.

If this option is omitted when displaying LIC levels that are available to be retrieved from a remote FTP site, you will be prompted to enter the password. If both this option and the **-k** option are omitted when displaying LIC levels that are available to be retrieved from a remote SFTP site, you will be prompted to enter the password. The **--passwd** and **-k** options are mutually exclusive.

This option is only valid when displaying LIC levels that are available to be retrieved from a remote FTP or SFTP site.

- k** The name of the file that contains the SSH private key. If the file name is not fully qualified, the file must exist in the user's home directory on the management console.

Use the **ssh-keygen** command to generate the public and private SSH key pair. The **ssh-keygen** command is not allowed to write to the **.ssh** directory in the user's home directory on the management console, so when you run the command on the management console, you must specify both the directory and the file name for the private key. If you generate a key with a passphrase, you will be prompted to enter the passphrase when you run any management console command that uses the key.

If both this option and the **--passwd** option are omitted when displaying LIC levels that are available to be retrieved from a remote SFTP site, you will be prompted to enter the password. The **-k** and **--passwd** options are mutually exclusive.

This option is only valid when displaying LIC levels that are available to be retrieved from a remote SFTP site.

- d** The directory to use on the remote FTP or SFTP site, or the mount point location on the internal management console hard disk drive. If this option is not specified for FTP or SFTP, the **/opt/ccfw/data** directory will be used.

This option is only valid when displaying LIC levels that are available to be retrieved from a mount point or a remote FTP or SFTP site.

- F** A delimiter separated list of attribute names for the desired attribute values to be displayed for each LIC level. If no attribute names are specified, then values for all of the attributes for each LIC level will be displayed.

When this option is specified, only attribute values will be displayed. No attribute names will be displayed. The attribute values displayed will be separated by the delimiter which was specified with this option.

This option is useful when only attribute values are desired to be displayed, or when the values of only selected attributes are desired to be displayed.

--header

Display a header record, which is a delimiter separated list of attribute names for the attribute values that will be displayed. This header record will be the first record displayed. This option is only valid when used with the **-F** option.

- help** Display the help text for this command and exit.

EXAMPLES

Display Managed System LIC levels:

lslic -t sys -m mysystem

Display Power LIC levels:

lslic -t power -m 9406-570*1020304

Display Managed System and Power LIC levels and display levels available on the IBM service website:

lslic -t syspower -m mysystem -r ibmwebsite

Display Managed System LIC levels and display levels available from a remote FTP server:

**lslic -t sys -m mysystem -r ftp -h ftpHost.domainname
-u myuser --passwd mypasswd**

Display Managed System LIC levels and display levels available from a remote SFTP server using SSH keys for authentication:

**lslic -t sys -m system1 -r sftp -h sftpServer -u user1
-k /home/hmcuser1/keys/id_rsa -d liclevels/840**

Display Managed System and Power LIC levels for all managed systems in the managed frame:

lslic -e myframe --allsystems -t syspower

Display LIC levels for all managed frames which contain High Performance Switches:

lslic -w

Display LIC levels for all managed frames which contain High Performance Switches and display levels available on the IBM service website:

lslic -w -r ibmwebsite

Display Power FRU levels for a managed system:

lslic -m mysystem -t powerfru

Display Power FRU levels for a managed frame:

lslic -e 9458-100*02AB33C -t powerfru

Display Power FRU levels for all managed frames which contain High Performance Switches:

lslic -w -t powerfru

Display the firmware levels for all SR-IOV adapters in a managed system:

lslic -t sriov -m mysystem

Display the LIC levels currently in the management console hard disk repository:

lslic -c

ENVIRONMENT

None

BUGS

None

AUTHOR

IBM Austin

SEE ALSO

updlic

NAME

lslock - list lock information

SYNOPSIS

lslock -e *managed-frame*
[-F [*attribute-names*] [--header]] [--help]

DESCRIPTION

lslock lists which Hardware Management Console (HMC) owns the lock on the *managed-frame*. If the *managed-frame* is not locked, then "No results were found." will be displayed.

OPTIONS

-e The name of the managed frame for which to list lock information. The name may either be the user-defined name for the managed frame, or be in the form *ttt-mmm*sssssss*, where *ttt* is the type, *mmm* is the model, and *sssssss* is the serial number of the managed frame. The *ttt-mmm*sssssss* form must be used if there are multiple managed frames with the same user-defined name.

-F A delimiter separated list of attribute names for the desired attribute values to be displayed. If no attribute names are specified, then values for all of the attributes will be displayed.

When this option is specified, only attribute values will be displayed. No attribute names will be displayed. The attribute values displayed will be separated by the delimiter which was specified with this option.

This option is useful when only attribute values are desired to be displayed, or when the values of only selected attributes are desired to be displayed.

--header

Display a header record, which is a delimiter separated list of attribute names for the attribute values that will be displayed. This header record will be the first record displayed. This option is only valid when used with the **-F** option.

--help Display the help text for this command and exit.

EXAMPLES

Display which HMC owns the lock on the managed frame **myFrame**:

```
lslock -e myFrame
```

ENVIRONMENT

None

BUGS

None

AUTHOR

IBM Austin

SEE ALSO

rmlock

NAME

lslogon - list logon information

SYNOPSIS

lslogon -r { **webui** | **ssh** } { **-t** | **-u** } [**-F** *attribute-names*] [**--header**]
[**--help**]

DESCRIPTION

lslogon lists the logged on users or the tasks they are running on the Hardware Management Console (HMC).

OPTIONS

-r The type of logon information to list. Valid values are **webui** for Web user interface users or tasks, or **ssh** for SSH users or tasks.

-t Specify this option to list information about the tasks that the users logged onto the HMC are running.

Either this option or the **-u** option is required. The **-t** and **-u** options are mutually exclusive.

-u Specify this option to list information about the users logged onto the HMC.

Either this option or the **-t** option is required. The **-u** and **-t** options are mutually exclusive.

-F A delimiter separated list of attribute names for the desired attribute values to be displayed. If no attribute names are specified, then values for all of the attributes will be displayed.

When this option is specified, only attribute values will be displayed. No attribute names will be displayed. The attribute values displayed will be separated by the delimiter which was specified with this option.

This option is useful when only attribute values are desired to be displayed, or when the values of only selected attributes are desired to be displayed.

--header

Display a header record, which is a delimiter separated list of attribute names for the attribute values that will be displayed. This header record will be the first record displayed. This option is only valid when used with the **-F** option.

--help Display the help text for this command and exit.

EXAMPLES

List all of the tasks that users logged onto the HMC Web user interface are running:

```
lslogon -r webui -t
```

List all of the users remotely logged onto the HMC via SSH:

```
lslogon -r ssh -u
```

ENVIRONMENT

None

BUGS

None

AUTHOR

IBM Austin

SEE ALSO

termtask

NAME

lslparmigr - list partition migration information

SYNOPSIS

```
lslparmigr -r {lpar | mc | mempool | msp | procpool | sys |
             virtualio}
[-m managed-system]
[-t target-managed-system [--ip IP-address [-u user-ID]]]
[--redundantvios {0 | 1 | 2}] [--filter "filter-data"]
[-F [attribute-names] [--header]] [--help]
```

DESCRIPTION

lslparmigr lists partition migration information.

OPTIONS

-r The type of resources for which to list partition migration information.

Specify **lpar** to list partition migration information for all of the partitions in *managed-system*.

Specify **mc** to list partition mobility capabilities for the Hardware Management Console (HMC).

Specify **mempool** to list shared memory pool information for the destination managed system. Shared memory pool information is only listed if the shared memory pool has enough available memory for the partition specified with the **lpar_names** or **lpar_ids** filter, if there is an available paging space device for the partition, and if the paging VIOS redundancy requirement of the partition can be satisfied. The **redundant** attribute that is listed indicates whether or not the partition would be configured to use redundant paging VIOS partitions on the destination managed system (a value of **0** indicates the partition would be configured to not use redundant paging VIOS partitions, and a value of **1** indicates the partition would be configured to use redundant paging VIOS partitions).

Specify **msp** to list possible source and destination mover service partition (MSP) pairs for active partition migration of the partitions specified with the **lpar_names** or **lpar_ids** filter.

Specify **procpool** to list possible shared processor pools in the destination managed system that have enough available processing resources for the partition specified with the **lpar_names** or **lpar_ids** filter.

Specify **sys** to list the partition mobility capabilities for *managed-system*.

Specify **virtualio** to list possible and suggested mappings of the virtual fibre channel adapters and virtual SCSI adapters in the partition specified with the **lpar_names** or **lpar_ids** filter to virtual I/O servers in the destination managed system.

-m The name of the managed system for which to list partition migration information. The name may either be the user-defined name for the managed system, or be in the form *ttt-mmm*ssssss*, where *ttt* is the machine type, *mmm* is the model, and *ssssss* is the serial number of the managed system. The *ttt-mmm*ssssss* form must be used if there are multiple managed systems with the same user-defined name.

When listing shared memory pool, mover service partition, shared processor pool, or virtual I/O information, specify the source managed system for the partition migration with this option.

This option is not valid when listing partition mobility capabilities for the HMC. This option is required when listing all other types of partition migration information.

- t The name of the target, or destination, managed system for the partition migration. The name may either be the user-defined name for the managed system, or be in the form *tttt-mmm*sssssss*, where *tttt* is the machine type, *mmm* is the model, and *sssssss* is the serial number of the managed system. The *tttt-mmm*sssssss* form must be used if there are multiple managed systems with the same user-defined name.

This option is required when listing shared memory pool, mover service partition, shared processor pool, or virtual I/O information. This option is not valid when listing any other type of partition migration information.

- ip If the destination managed system is not managed by the same management console that is managing the source managed system, then use this option to specify the IP address or host name of the management console that is managing the destination managed system.

To use this option, SSH must be enabled on both management consoles. Also, you must run the **mkauthkeys** command once to set up SSH key authentication.

- u If the destination managed system is not managed by the same management console that is managing the source managed system, then use this option to specify the user ID to use on the management console that is managing the destination managed system. If you do not specify the user ID, then the user ID of the user that issues this command will be used.

--redundantvios,--redundantpgvios

When listing shared memory pool information, use this option to specify the paging VIOS redundancy requirement for the partition on the destination managed system. Valid values are **0** if the partition is not to use redundant paging VIOS partitions, **1** if the partition is required to use redundant paging VIOS partitions, or **2** if the partition is to use redundant paging VIOS partitions if possible.

This option is only valid when listing shared memory pool information. If this option is not specified when listing shared memory pool information, then the partition is required to use the same paging VIOS redundancy configuration on the destination managed system that the partition is currently using on the source managed system.

The **--redundantvios** and **--redundantpgvios** options are equivalent. The **--redundantpgvios** option is deprecated.

- filter** The filter(s) to apply to the resources to be listed. Filters are used to select which resources are to be listed. If no filters are used, then all of the resources will be listed. For example, partition information can be listed by using a filter to specify the names or IDs of the partitions to list. Otherwise, if no filter is used, then information for all of the partitions in the managed system will be listed.

The filter data consists of filter name/value pairs, which are in comma separated value (CSV) format. The filter data must be enclosed in double quotes.

The format of the filter data is as follows:

```
"filter-name=value,filter-name=value,..."
```

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

```
""filter-name=value,value,...",..."
```

When a list of values is specified, the filter name/value pair must be enclosed in double quotes. Depending on the shell being used, nested double quote characters may need to be preceded by an escape character, which is usually a `'\'` character.

Valid filter names:

lpar_names | lpar_ids

One of these filters is required when listing shared memory pool, mover service partition, shared processor pool, or virtual I/O information. When listing shared memory pool, shared processor pool, or virtual I/O information, only one partition can be specified. These filters are optional when listing partition information, and if specified, multiple partitions can be specified.

dest_sys_names

The user-defined name or the machine type, model, and serial number (*ttt-mmm*sssss*) for one or more destination managed systems. This filter is only valid when listing partition information.

The **lpar_names**, **lpar_ids**, and **dest_sys_names** filters are mutually exclusive.

This option is not valid when listing HMC or managed system information.

-F A delimiter separated list of attribute names for the desired attribute values to be displayed. If no attribute names are specified, then values for all of the attributes will be displayed.

When this option is specified, only attribute values will be displayed. No attribute names will be displayed. The attribute values displayed will be separated by the delimiter which was specified with this option.

This option is useful when only attribute values are desired to be displayed, or when the values of only selected attributes are desired to be displayed.

--header

Display a header record, which is a delimiter separated list of attribute names for the attribute values that will be displayed. This header record will be the first record displayed. This option is only valid when used with the **-F** option.

--help Display the help text for this command and exit.

EXAMPLES

List partition migration information for all partitions in the managed system **system1**:

lslparmigr -r lpar -m system1

List partition migration information for all partitions in the managed system **system1** that are currently migrating or waiting to migrate to managed system **system2**:

lslparmigr -r lpar -m system1 --filter "dest_sys_names=system2"

List the shared memory pool information for managed system **system2** for partition **lpar1** if **lpar1** is to use redundant paging VIOS partitions on **system2** if possible:

**lslparmigr -r mempool -m system1 -t system2 --filter "lpar_names=lpar1"
--redundantvios 2**

List possible source and destination MSP pairs for migrating partition **lpar1** from managed system **system1** to managed system **system2**:

```
lslparmigr -r msp -m system1 -t system2 --filter "lpar_names=lpar1"
```

List possible shared processor pools in managed system **system2** that partition **lpar1** can use when it is migrated from managed system **system1** to **system2**:

```
lslparmigr -r procpool -m system1 -t system2 --filter  
"lpar_names=lpar1"
```

List the partition mobility capabilities for managed system **system1**:

```
lslparmigr -r sys -m system1
```

List possible and suggested mappings of the virtual fibre channel and virtual SCSI adapters in the partition with ID **1** to virtual I/O servers in managed system **system2**:

```
lslparmigr -r virtualio -m system1 -t system2 --filter "lpar_ids=1"
```

List possible and suggested mappings of the virtual fibre channel and virtual SCSI adapters in partition **aix1** to virtual I/O servers in managed system **system2**, when **system2** is managed by the HMC with host name **hmc2**:

```
mkauthkeys --ip hmc2 -u hmc2user
```

```
lslparmigr -r virtualio -m system1 -t system2 --ip hmc2 -u hmc2user  
--filter "lpar_names=aix1"
```

ENVIRONMENT

None

BUGS

None

AUTHOR

IBM Austin

SEE ALSO

migr_lpar, **mkauthkeys**

NAME

lsiparutil - list utilization data

SYNOPSIS

To list utilization data collected for a managed system:

```
lsiparutil -r {hmc | lpar | pool | procpool | mempool |  
  sys | all}  
  -m managed-system  
  [-d number-of-days] [-h number-of-hours]  
  [--minutes number-of-minutes]  
  [--startyear year] [--startmonth month]  
  [--startday day] [--starthour hour]  
  [--startminute minute] [--endyear year]  
  [--endmonth month] [--endday day]  
  [--endhour hour] [--endminute minute]  
  [-n number-of-events] [-s sample-rate]  
  [--filter "filter-data"]  
  [-F [attribute-names] [--header]] [--help]
```

To list HMC settings for utilization data collection:

```
lsiparutil -r config [-m managed-system]  
  [-F [attribute-names] [--header]] [--help]
```

DESCRIPTION

lsiparutil lists utilization data collected for a *managed-system*. This command also lists the Hardware Management Console (HMC) settings for utilization data collection.

The HMC collects the following types of utilization data: sampling events, state change events, configuration change events, and Utility Capacity on Demand (CoD) processor usage events.

Sampling events are collected for the managed system, for each partition in the managed system, and for the physical processor pool, each shared processor pool, and the shared memory pool in the managed system. Sampling events are collected at the configured sample rate. They are also collected hourly, daily at midnight, and monthly at midnight on the first day of each month. Sampling events are also collected immediately following the collection of a Utility CoD processor usage event. Sampling events contain information about memory and processor utilization.

State change events are collected for the managed system, for each partition in the managed system, and for the HMC. State change events are collected when a state change occurs for the managed system, a partition, or when the HMC is started or shut down.

Configuration change events are collected for the managed system, for each partition in the managed system, for each shared processor pool in the managed system, for the shared memory pool in the managed system, and for the HMC. Configuration change events are collected when a configuration change affecting memory or processor resources occurs for the managed system, a partition, a shared processor pool, or the shared memory pool. Configuration change events are also collected when the local time is changed on the HMC.

Utility CoD processor usage events are collected for the managed system when a Utility CoD processor minute is used. All Utility CoD processor minutes used during a single minute are grouped into one event.

Hourly sampling events and all state change, configuration change, and Utility CoD processor usage events collected are only saved on the HMC for about 2 months. Daily sampling events are saved for about 2 years, and monthly sampling events are saved for about 10 years.

Sampling events are also saved in a snapshot file. Sampling events that are collected more frequently than

hourly are only saved in the snapshot file. All state change, configuration change, and Utility CoD processor usage events collected are also saved in the snapshot file, in addition to being saved with the hourly sampling events. The snapshot file is purged daily at midnight. Events more than 24 hours old are purged. Therefore, the snapshot file will only contain events that were collected during the last 24 to 48 hours.

Utilization data collection for managed systems is not automatically enabled. The **chlparutil** command can be used to enable utilization data collection and configure the sample rate.

OPTIONS

-r The type of system resources for which events are to be listed. Valid values are **hmc** for HMC, **lpar** for partitions, **pool** for physical processor pool, **procpool** for shared processor pools, **mem-pool** for shared memory pool, **sys** for managed system, **all** for HMC, partitions, physical processor pool, shared processor pools, shared memory pool, and managed system, and **config** for HMC configuration settings for utilization data collection.

-m The name of the managed system for which the collected utilization data or the HMC configuration settings is to be listed. The name may either be the user-defined name for the managed system, or be in the form *tttt-mmm*sssssss*, where *tttt* is the machine type, *mmm* is the model, and *sssssss* is the serial number of the managed system. The *tttt-mmm*sssssss* form must be used if there are multiple managed systems with the same user-defined name, or if the HMC does not currently have a connection to the managed system.

If this option is specified when listing HMC configuration settings, then the HMC configuration settings for the *managed-system* will be listed. Otherwise, the HMC configuration settings for all of the systems currently managed by this HMC, and for all of the systems for which utilization data has been previously collected by this HMC will be listed.

-d The number of days prior to today for which events will be listed. Events that were collected today will be listed, along with events that were collected during the past *number-of-days* days.

This option cannot be specified with the **-h**, **--minutes**, **--startyear**, **--startmonth**, **--startday**, **--starthour**, **--startminute**, **--endyear**, **--endmonth**, **--endday**, **--endhour**, or **--endminute** options.

-h The number of hours prior to the current time for which events will be listed.

This option cannot be specified with the **-d**, **--minutes**, **--startyear**, **--startmonth**, **--startday**, **--starthour**, **--startminute**, **--endyear**, **--endmonth**, **--endday**, **--endhour**, or **--endminute** options.

--minutes

The number of minutes prior to the current time for which events will be listed.

This option cannot be specified with the **-d**, **-h**, **--startyear**, **--startmonth**, **--startday**, **--starthour**, **--startminute**, **--endyear**, **--endmonth**, **--endday**, **--endhour**, or **--endminute** options.

--startyear

The starting year for which events will be listed. The default value for this option is **1970**.

This option cannot be specified with the **-d**, **-h**, or **--minutes** options.

--startmonth

The starting month for which events will be listed. Valid values are **1** for January through **12** for December. The default value for this option is **1** (January).

This option cannot be specified with the **-d**, **-h**, or **--minutes** options.

--startday

The starting day for which events will be listed. Valid values are **1** through **31**. The default value for this option is **1**.

This option cannot be specified with the **-d**, **-h**, or **--minutes** options.

--starthour

The starting hour for which events will be listed. Valid values are **0** for midnight through **23** for 11:00 pm. The default value for this option is **0** (midnight).

This option cannot be specified with the **-d**, **-h**, or **--minutes** options.

--startminute

The starting minute for which events will be listed. Valid values are **0** through **59**. The default value for this option is **0**.

This option cannot be specified with the **-d**, **-h**, or **--minutes** options.

--endyear

The ending year for which events will be listed. The default value for this option is now.

This option cannot be specified with the **-d**, **-h**, or **--minutes** options.

--endmonth

The ending month for which events will be listed. Valid values are **1** for January through **12** for December. The default value for this option is now.

This option cannot be specified with the **-d**, **-h**, or **--minutes** options.

--endday

The ending day for which events will be listed. Valid values are **1** through **31**. The default value for this option is now.

This option cannot be specified with the **-d**, **-h**, or **--minutes** options.

--endhour

The ending hour for which events will be listed. Valid values are **0** for midnight through **23** for 11:00 pm. The default value for this option is now.

This option cannot be specified with the **-d**, **-h**, or **--minutes** options.

--endminute

The ending minute for which events will be listed. Valid values are **0** through **59**. The default value for this option is now.

This option cannot be specified with the **-d**, **-h**, or **--minutes** options.

-n

The maximum number of events to be listed, starting with the most recent event. The number specified must be greater than 0.

If this option is not specified, and neither are any of the **-d**, **-h**, **--minutes**, **--startyear**, **--startmonth**, **--startday**, **--starthour**, **--startminute**, **--endyear**, **--endmonth**, **--endday**, **--endhour**, or **--endminute** options, then only the most recent event will be listed.

-s

Use this option to specify the sample rate for which utilization data is to be listed. Valid values are **h** for hourly sampling events and all state change, configuration change, and Utility CoD processor usage events, **d** for daily sampling events, **m** for monthly sampling events, and **s** for the snapshot file.

If this option is not specified then events in the snapshot file will be listed.

--filter The filter(s) to apply to the events to be listed. Filters are used to select which events for the specified resource type are to be listed.

The filter data consists of filter name/value pairs, which are in comma separated value (CSV) format. The filter data must be enclosed in double quotes.

The format of the filter data is as follows:

```
"filter-name=value,filter-name=value,..."
```

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

```
""filter-name=value,value,...",..."
```

When a list of values is specified, the filter name/value pair must be enclosed in double quotes. Depending on the shell being used, nested double quote characters may need to be preceded by an escape character, which is usually a `'\'` character.

Multiple values can be specified for each filter.

Valid filter names for this command:

- lpar_ids**
Specify partition ID(s)
- lpar_names**
Specify partition user-defined name(s)
- event_types**
Specify one or more of the values
sample, state_change, config_change,
utility_cod_proc_usage
- pools**
Specify shared processor pool ID(s)
- pool_names**
Specify shared processor pool
user-defined name(s)

Valid filters with **-r hmc**:

event_types

Valid filters with **-r lpar**:

event_types, lpar_ids | lpar_names,
pools | pool_names

Valid filters with **-r pool**:

event_types

Valid filters with **-r procpool**:

event_types, pools | pool_names

Valid filters with **-r mempool**:

event_types

Valid filters with **-r sys**:

event_types

Valid filters with **-r all**:

event_types, lpar_ids | lpar_names,
pools | pool_names

The **lpar_ids** and **lpar_names** filter will only apply to partition events, and the **pools** and **pool_names** filter will only apply to partition events and shared processor pool events

- F** A delimiter separated list of attribute names for the desired attribute values to be displayed for each event. If no attribute names are specified, then values for all of the attributes for each event will be displayed.

When this option is specified, only attribute values will be displayed. No attribute names will be displayed. The attribute values displayed will be separated by the delimiter which was specified with this option.

This option is useful when only attribute values are desired to be displayed, or when the values of only selected attributes are desired to be displayed.

Command attributes:

borrowed_pool_proc_units

The sum of the processing units being borrowed from powered off partitions with dedicated processors and the processing units available to be assigned to partitions, rounded down to a whole processor.

capped_cycles

The number of capped processing cycles utilized by this partition since the managed system was started.

configurable_pool_proc_units

The number of processing units assigned to all shared processor partitions, rounded up to a whole processor.

configurable_sys_mem

The amount of configurable system memory (in megabytes).

configurable_sys_proc_units

The number of configurable system processing units.

curr_5250_cpw_percent

The 5250 CPW percent assigned to the partition.

curr_avail_5250_cpw_percent

The 5250 CPW percent available to be assigned to partitions.

curr_avail_pool_proc_units

The number of processing units available to be assigned to partitions.

curr_avail_sys_mem

The amount of memory (in megabytes) available to be assigned to partitions.

curr_avail_sys_proc_units

The number of processing units available to be assigned to partitions.

curr_io_entitled_mem

The amount of I/O entitled memory (in megabytes) assigned to the shared memory partition.

curr_max_pool_mem

The maximum size (in megabytes) of the shared memory pool.

curr_mem

The amount of memory (in megabytes) assigned to the partition. For shared memory partitions, this is the amount of logical memory assigned to the partition.

curr_mem_weight

The current relative memory priority for the shared memory partition. The smaller the value, the lower the priority. Possible values are 0 - 255.

curr_pool_mem

The size (in megabytes) of the shared memory pool.

curr_proc_mode

The processing mode for the partition. Possible values are **ded** or **shared**.

curr_proc_units

The number of processing units assigned to the partition.

curr_procs

The number of processors or virtual processors assigned to the partition.

curr_reserved_pool_proc_units

The number of processing units that are reserved for temporary use by the uncapped partitions in the shared processor pool.

curr_shared_proc_pool_id

The unique integer identifier for the shared processor pool that the partition is in.

curr_shared_proc_pool_name

The user-defined name of the shared processor pool, at the time the event was collected, that the partition is in.

curr_sharing_mode

The sharing mode of the partition. Possible values are **keep_idle_procs**, **share_idle_procs**, **share_idle_procs_active**, **share_idle_procs_always**, **cap**, or **uncap**.

curr_uncap_weight

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

dedup_cycles

The number of processing cycles spent deduplicating data since Active Memory Deduplication was enabled for the shared memory pool. The processing cycles were donated by the paging Virtual I/O Server partitions assigned to the pool.

dedup_mem

The amount of logical memory (in megabytes) assigned to the partition that has been deduplicated.

dedup_pool_mem

The amount of memory (in megabytes) in the shared memory pool that is being used for deduplicated data.

entitled_cycles

The number of processing cycles to which the partition has been entitled since the managed system was started. This value is based on the number of processing units assigned to the partition, and may be greater than or smaller than the number of cycles actually used.

event_type

The type of event. Possible values are **sample**, **state_change**, **config_change**, or **utility_cod_proc_usage**.

idle_cycles

The number of cycles that the partition reported as idle. If the partition does not have the ability to report idle cycles, this value is 0.

lpar_curr_io_entitled_mem

The total amount of I/O entitled memory (in megabytes) assigned to all of the partitions in the shared memory pool.

lpar_dedup_mem

The total amount of partition logical memory (in megabytes) that has been deduplicated.

lpar_id

The unique integer identifier for the partition.

lpar_mapped_io_entitled_mem

The total amount of I/O entitled memory (in megabytes) currently mapped by all of the partitions in the shared memory pool.

lpar_name

The user-defined name of the partition at the time the event was collected.

lpar_run_mem

The total amount of logical memory (in megabytes) assigned to all of the partitions in the shared memory pool.

mapped_io_entitled_mem

The amount of I/O entitled memory (in megabytes) currently mapped by the shared memory partition.

max_pool_proc_units

This number, minus the reserved

processing units in the shared processor pool, is the maximum number of processing units that the partitions in the shared processor pool can use.

mem_dedup

Indicates whether Active Memory Deduplication is enabled for the shared memory pool. Possible values are **0** for disabled or **1** for enabled.

mem_mode

The memory mode for the partition.

Possible values are **ded** or **shared**.

mem_oveage_cooperation

The difference between the shared memory partition's assigned memory overcommitment and its actual overcommitment. A positive value means the partition is using less memory than system firmware has requested it to use.

name

The user-defined name of the managed system.

page_faults

The total number of page faults that have occurred since the shared memory pool was created.

page_in_delay

The total page-in delay, in microseconds, spent waiting for page faults since the shared memory pool was created.

phys_run_mem

The runtime amount of physical memory (in megabytes) allocated to the shared memory partition.

prev_time

The time on the HMC when the HMC time was changed.

proc_cycles_per_second

Processing cycles per second on one physical processor. This value is static for a particular managed system.

On POWER6 and later servers, this value is set to a fixed constant of 512,000,000.

resource_type

The type of system resource for which the event was collected. Possible values are **hmc**, **lpar**, **pool**, **procpool**, **mempool**, or **sys**.

run_latch_cycles

The number of non-idle cycles used by the partition, while the run latch was set, since the managed system was started.

POWER6 and POWER7 servers only

run_latch_instructions

The number of non-idle instructions performed by the partition, while the run latch was set, since the

managed system was started.

POWER6 and POWER7 servers only

run_mem_weight

The runtime relative memory priority for the shared memory partition. The smaller the value, the lower the priority. Possible values are 0 - 255.

sample_rate

The rate, in seconds, at which samples are obtained. This rate can be changed with the **chlpurutil** command.

shared_cycles_while_active

The number of dedicated processing cycles shared by this partition while it has been active since the managed system was started.

shared_proc_pool_id

The unique integer identifier for the shared processor pool.

shared_proc_pool_name

The user-defined name of the shared processor pool at the time the event was collected.

state

For system events, this is the state of the managed system at the time the event was collected. For partition events, this is the state of the partition at the time the event was collected.

sys_firmware_mem

The amount of memory (in megabytes) on the managed system that is being used by system firmware.

sys_firmware_pool_mem

The amount of memory (in megabytes) in the shared memory pool that is being used by system firmware.

sys_time

The time on the managed system that the sample was taken.

time

The time on the HMC that the event was collected.

time_cycles

The number of time cycles since the managed system was started.

time_shut_down

The time the HMC was shut down or rebooted.

total_instructions

The number of instructions performed by the partition since the managed system was started.

POWER8 and later servers only

total_instructions_execution_time

The amount of time instruction counts were collected for the partition since the managed system was started.

POWER8 and later servers only

total_pool_cycles

The total number of processing cycles available in the physical processor pool or shared processor pool since the managed system was started.

type_model_serial_num

The machine type, model, and serial number of the managed system.

uncapped_cycles

The number of uncapped processing cycles utilized by this partition since the managed system was started.

unreported_proc_min

The total number of Utility CoD processor minutes that have not been reported.

used_proc_min

The number of Utility CoD processor minutes that were used during the last minute.

utilized_pool_cycles

The number of processing cycles in the physical processor pool or shared processor pool that have been utilized since the managed system was started.

--header

Display a header record, which is a delimiter separated list of attribute names for the attribute values that will be displayed. This header record will be the first record displayed. This option is only valid when used with the **-F** option.

--help Display the help text for this command and exit.

EXAMPLES

List all of the monthly sampling events that were collected for the managed system **sys1** for the entire year 2005:

```
lslparutil -r all -m sys1 --startyear 2005 --endyear 2005 --endmonth 12 --endday 31 --endhour 23 -s m
```

List all of the hourly managed system sampling events that have been collected for the managed system with type, model, and serial number **9406-520*1000101** since midnight today:

```
lslparutil -r sys -m 9406-520*1000101 -d 0 --filter "event_types=sample" -s h
```

List all of the events that have been collected for the managed system **sys1** since January 1, 2006, and list a maximum of 25 events:

```
lslparutil -r all -m sys1 --startyear 2006 -n 25 -s h
```

List the 10 most recent hourly sampling events that were collected for partition **p1**:

```
lslparutil -r lpar -m sys1 -n 10 -s h --filter "event_types=sample,lpar_names=p1"
```

List all of the state change and configuration change events that were collected for the managed system **sys1** for the month of June in 2005:

```
lslparutil -r all -m sys1 --startyear 2005 --startmonth 6 --endyear 2005 --endmonth 6 --endday 30 --endhour 23 -s h --filter ""event_types=state_change,config_change""
```

List all of the Utility CoD processor usage events that were collected for the managed system **sys1** for the month of May in 2007:

```
lslparutil -r all -m sys1 --startyear 2007 --startmonth 5 --endyear 2007 --endmonth 5 --endday 31 --endhour 23 -s h --filter "event_types=utility_cod_proc_usage"
```

List all of the HMC events that have been collected for managed system **sys1** since February 1, 2006:

```
lslparutil -r hmc -m sys1 --startyear 2006 --startmonth 2 -s h
```

List the configuration settings for utilization data collection on this HMC for all managed systems:

```
lslparutil -r config
```

To calculate the default shared processor pool utilization in percent over a twelve hour time period:

```
lslparutil -m sys1 -r procpool --startyear 2006 --startmonth 2 --startday 23 --starthour 0 --endyear 2006 --endmonth 2 --endday 23 --endhour 12 --filter "event_types=sample,pool_names=DefaultPool" -F time,total_pool_cycles,utilized_pool_cycles
```

```
02/23/2006 12:00:01,134967149091025,467439053292
02/23/2006 11:00:02,134963299532241,467428119008
02/23/2006 10:00:01,134959313365305,467419269942
02/23/2006 09:00:01,134954622214624,467403199531
02/23/2006 08:00:02,134942086330068,467368397739
02/23/2006 07:00:01,134929553859752,467333227651
02/23/2006 06:00:01,134917026289150,467295577359
02/23/2006 05:00:02,134904482088726,467258616569
02/23/2006 04:00:01,134891946956456,467223704573
02/23/2006 03:00:01,134879415157938,467188374373
02/23/2006 02:00:01,134866883128692,467152556956
02/23/2006 01:00:02,134854347365860,467116506907
02/23/2006 00:00:03,134841811733640,467081011935
```

```
Pool utilization =
  (utilized_pool_cycles / total_pool_cycles) * 100
Pool utilization = ((467439053292 - 467081011935) /
  (134967149091025 - 134841811733640)) * 100
Pool utilization = 0.29%
```

To calculate the processor utilization for the shared processor partition with ID 7 over the last minute (the sampling rate is 1 minute):

```
lslparutil -m sys1 -r lpar --filter "lpar_ids=7", -n 2 -F time,lpar_id, capped_cycles,uncapped_cycles,entitled_cycles,time_cycles
```

03/29/2012 18:09:30,7,55844419525044,0,106332165468039,
1234978811234159
03/29/2012 18:08:30,7,55839107527864,0,106316742434446,
1234947965168143

Processor utilization % =
 $((\text{capped_cycles} + \text{uncapped_cycles}) / \text{entitled_cycles}) * 100$
Processor utilization % =
 $((55844419525044 - 55839107527864) + (0 - 0)) /$
 $(106332165468039 - 106316742434446) * 100$
Processor utilization % = 34.44%

Processor units utilized =
 $(\text{capped_cycles} + \text{uncapped_cycles}) / \text{time_cycles}$
Processor units utilized =
 $((55844419525044 - 55839107527864) + (0 - 0)) /$
 $(1234978811234159 - 1234947965168143)$
Processor units utilized = 0.17

To calculate the processor utilization for the dedicated processor partition with ID 8 over the last minute (the sampling rate is 1 minute):

**lsparutil -m sys1 -r lpar --filter "lpar_ids=8", -n 2
-F time,lpar_id,capped_cycles,idle_cycles,time_cycles**

03/29/2012 18:09:30,8,10650723610566,27015925200832,1234978811234159
03/29/2012 18:08:30,8,10619877543196,26994400296144,1234947965168143

Processor utilization % =
 $((\text{capped_cycles} - \text{idle_cycles}) / \text{capped_cycles}) * 100$
Processor utilization % = $((10650723610566 - 10619877543196) -$
 $(27015925200832 - 26994400296144)) /$
 $(10650723610566 - 10619877543196) * 100$
Processor utilization % = 30.22%

Processor units utilized =
 $(\text{capped_cycles} - \text{idle_cycles}) / \text{time_cycles}$
Processor units utilized = $((10650723610566 - 10619877543196) -$
 $(27015925200832 - 26994400296144)) /$
 $(1234978811234159 - 1234947965168143)$
Processor units utilized = 0.30

ENVIRONMENT

None

BUGS

None

AUTHOR

IBM Austin

SEE ALSO

chlparutil, rmlparutil

NAME

lsmediadev - list storage media devices

SYNOPSIS

lsmediadev [-F *attribute-names*] [--header] [--help]

DESCRIPTION

lsmediadev lists the storage media devices that are available for use on the Hardware Management Console (HMC).

OPTIONS

-F A delimiter separated list of attribute names for the desired attribute values to be displayed. If no attribute names are specified, then values for all of the attributes will be displayed.

When this option is specified, only attribute values will be displayed. No attribute names will be displayed. The attribute values displayed will be separated by the delimiter which was specified with this option.

This option is useful when only attribute values are desired to be displayed, or when the values of only selected attributes are desired to be displayed.

--header

Display a header record, which is a delimiter separated list of attribute names for the attribute values that will be displayed. This header record will be the first record displayed. This option is only valid when used with the **-F** option.

--help Display the help text for this command and exit.

EXAMPLES

List all of the storage media devices that are available for use on the HMC:

lsmediadev

ENVIRONMENT

None

BUGS

None

AUTHOR

IBM Austin

NAME

lsmemdev - list memory devices

SYNOPSIS

```
lsmemdev -r avail -m managed-system  
{ -p partition-names | --id partition-IDs }  
[--min minimum-size] [--max maximum-size] [--filter "filter-data"]  
[-F [attribute-names] [--header] [--help]
```

DESCRIPTION

lsmemdev lists block storage devices that can be used as reserved storage devices for the *managed-system*. Reserved storage devices are used as paging space devices for partitions that use shared memory. Reserved storage devices are also used for suspended partitions.

Block storage devices can be listed for any Virtual I/O Server (VIOS) partition in the managed system. However, a device cannot be added to a pool unless the VIOS partition is assigned to that pool.

On a managed system that supports partition suspend and resume operations, block storage devices are added to the reserved storage device pool. On a managed system that supports Active Memory Sharing but does not support partition suspend and resume operations, block storage devices are added to the shared memory pool.

OPTIONS

- r** The type of resources to list. The only valid value is **avail** to list the block storage devices that are available to add to a reserved storage device pool or a shared memory pool. Any device that is already in a pool will not be listed.
- m** The name of the managed system which has the VIOS partition(s) for which to list block storage devices. The name may either be the user-defined name for the managed system, or be in the form *ttt-mmm*sssss*, where *ttt* is the machine type, *mmm* is the model, and *sssss* is the serial number of the managed system. The *ttt-mmm*sssss* form must be used if there are multiple managed systems with the same user-defined name.
- p** The name of one or two VIOS partitions for which to list block storage devices. The partition(s) must be in the running state and must have an RMC connection to the Hardware Management Console (HMC).

If two VIOS partitions are specified, then only those block storage devices that can be accessed by both VIOS partitions are listed. The two partition names must be comma separated.

You can either use this option to specify the name(s) of the VIOS partition(s), or use the **--id** option to specify the ID(s). The **-p** and the **--id** options are mutually exclusive.

- id** The ID of one or two VIOS partitions for which to list block storage devices. The partition(s) must be in the running state and must have an RMC connection to the HMC.

If two VIOS partitions are specified, then only those block storage devices that can be accessed by both VIOS partitions are listed. The two partition IDs must be comma separated.

You can either use this option to specify the ID(s) of the VIOS partition(s), or use the **-p** option to specify the name(s). The **--id** and the **-p** options are mutually exclusive.

- min** The minimum size, in megabytes, of the block storage devices to list. If this option is not specified, a default value of 0 is used.
- max** The maximum size, in megabytes, of the block storage devices to list. If this option is not specified, no maximum size is used.
- filter** The filter(s) to apply to the block storage devices to be listed. Filters are used to select which devices are to be listed. If no filters are used, then all of the devices will be listed. For example,

only physical devices can be listed by using the **types** filter.

The filter data consists of filter name/value pairs, which are in comma separated value (CSV) format. The filter data must be enclosed in double quotes.

The format of the filter data is as follows:

```
"filter-name=value,filter-name=value,..."
```

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

```
""filter-name=value,value,...",..."
```

When a list of values is specified, the filter name/value pair must be enclosed in double quotes. Depending on the shell being used, nested double quote characters may need to be preceded by an escape character, which is usually a ``\`` character.

Only one value can be specified for each filter.

Valid filter names for this command:

redundant

Valid values are:

0 - devices which could never be accessed by another VIOS partition

1 - devices which possibly could be accessed by another VIOS partition

types

Valid values are:

logical - logical devices

phys - physical devices

-F A delimiter separated list of attribute names for the desired attribute values to be displayed for each device. If no attribute names are specified, then values for all of the attributes for each device will be displayed.

When this option is specified, only attribute values will be displayed. No attribute names will be displayed. The attribute values displayed will be separated by the delimiter which was specified with this option.

This option is useful when only attribute values are desired to be displayed, or when the values of only selected attributes are desired to be displayed.

--header

Display a header record, which is a delimiter separated list of attribute names for the attribute values that will be displayed. This header record will be the first record displayed. This option is only valid when used with the **-F** option.

--help Display the help text for this command and exit.

EXAMPLES

List all of the available block storage devices that can be accessed by VIOS partition **vios1**:

```
lsmemdev -r avail -m sys1 -p vios1
```

List all of the available block storage devices that can be accessed by both VIOS partitions **vios1** and **vios2**:

```
lsmemdev -r avail -m sys1 -p vios1,vios2
```

List the names and sizes of all of the available block storage devices that can be accessed by VIOS partition **vios1**:

```
lsmemdev -r avail -m sys1 -p vios1 -F device_name,size
```

List the available block storage devices which are between 1 GB and 4 GB in size and that can be accessed by both the VIOS partitions with IDs **1** and **2**:

```
lsmemdev -r avail -m sys1 --id 1,2 --min 1024 --max 4096
```

List all of the available physical block storage devices that can be accessed by the VIOS partition with ID **1**:

```
lsmemdev -r avail -m 9117-MMA*1234567 --id 1 --filter "types=phys"
```

ENVIRONMENT

None

BUGS

None

AUTHOR

IBM Austin

SEE ALSO

lshwres

NAME

lsmemopt - list memory optimization information

SYNOPSIS

```
lsmemopt -m managed-system [-r {sys | lpar}]  
[-o {currscore | calcscore}]  
[-p partition-names | --id partition-IDs]  
[-x partition-names | --xid partition-IDs]  
[--filter "filter-data"]  
[-F [attribute-names] [--header]] [--help]
```

DESCRIPTION

lsmemopt lists information, including status and progress information, about the most recent memory optimization operation that was started on the *managed-system*.

lsmemopt also lists affinity scores for the *managed-system* and for partitions in the *managed-system*. An affinity score is a measure of how good the processor-memory affinity is on the managed system or for a partition. The score is a number between 0 and 100, with 0 representing the worst affinity and 100 representing perfect affinity. Depending on the system configuration, scores of 100 might not be attainable. A partition with no processor and memory resources does not have an affinity score, and **none** will be displayed for its score.

lsmemopt can list current affinity scores, and can calculate and list potential affinity scores which could be attained after running a Dynamic Platform Optimization operation. Note that potential affinity scores may differ slightly from the actual affinity scores that are attained after running the Dynamic Platform Optimization operation.

OPTIONS

-m The name of the managed system for which to list memory optimization information. The name may either be the user-defined name for the managed system, or be in the form *ttt-mmm*ssssss*, where *ttt* is the machine type, *mmm* is the model, and *ssssss* is the serial number of the managed system. The *ttt-mmm*ssssss* form must be used if there are multiple managed systems with the same user-defined name.

-r The type of resources for which to list affinity scores. Valid values are **sys** to list the affinity score for the managed system, and **lpar** to list affinity scores for partitions.

If this option is omitted when listing affinity scores, the affinity score for the managed system will be listed.

This option is only valid when listing affinity scores.

-o The operation to perform. Valid values are **currscore** to list current affinity scores, and **calcscore** to calculate and list potential affinity scores.

If this option is omitted, information about the most recent memory optimization operation that was started on the managed system is listed.

-p The name of one or more potential partitions to prioritize for optimization. This list of partitions is used to calculate the potential affinity scores that could be attained if a Dynamic Platform Optimization operation that prioritized this list of partitions were run.

Multiple partition names must be comma separated.

You can either use this option to specify the names of the potential partitions to prioritize for optimization, or use the **--id** option to specify the IDs of the partitions. The **-p** and the **--id** options are mutually exclusive.

If neither this option nor the **--id** option is used to specify a list of potential partitions to prioritize for optimization, the managed system will use its own prioritization. In either case, the list of potential partitions to optimize defaults to all partitions except those that are specified in the list of potential partitions to exclude from the operation.

This option is only valid when **-o calcscore** is specified.

--id The ID of one or more potential partitions to prioritize for optimization. This list of partitions is used to calculate the potential affinity scores that could be attained if a Dynamic Platform Optimization operation that prioritizes this list of partitions were run.

A range of partition IDs may be specified. A range is specified by specifying the beginning partition ID, a dash, and the ending partition ID. Both the beginning and ending partition IDs are included in the range. For example, the range 1-5 includes partitions IDs 1, 2, 3, 4, and 5.

The list of partition IDs specified with this option must be comma separated and can include one or more partition IDs, one or more ranges of partition IDs, or both.

You can either use this option to specify the IDs of the potential partitions to prioritize for optimization, or use the **-p** option to specify the names of the partitions. The **--id** and the **-p** options are mutually exclusive.

If neither this option nor the **-p** option is used to specify a list of potential partitions to prioritize for optimization, the managed system will use its own prioritization. In either case, the list of potential partitions to optimize defaults to all partitions except those that are specified in the list of potential partitions to exclude from the operation.

This option is only valid when **-o calcscore** is specified.

-x The name of one or more potential partitions to exclude and protect from an optimization operation. This list of partitions is used to calculate the potential affinity scores that could be attained if a Dynamic Platform Optimization operation were run with this list of partitions to exclude.

Multiple partition names must be comma separated.

You can either use this option to specify the names of the potential partitions to exclude, or use the **--xid** option to specify the IDs of the partitions. The **-x** and the **--xid** options are mutually exclusive.

If neither this option nor the **--xid** option is used to specify a list of potential partitions to exclude, the list of potential partitions to exclude defaults to an empty list.

This option is only valid when **-o calcscore** is specified.

--xid The ID of one or more potential partitions to exclude and protect from an optimization operation. This list of partitions is used to calculate the potential affinity scores that could be attained if a Dynamic Platform Optimization operation were run with this list of partitions to exclude.

A range of partition IDs may be specified. A range is specified by specifying the beginning partition ID, a dash, and the ending partition ID. Both the beginning and ending partition IDs are included in the range. For example, the range 8-11 includes partitions IDs 8, 9, 10, and 11.

The list of partition IDs specified with this option must be comma separated and can include one or more partition IDs, one or more ranges of partition IDs, or both.

You can either use this option to specify the IDs of the potential partitions to exclude, or use the **-x**

option to specify the names of the partitions. The **--xid** and the **-x** options are mutually exclusive.

If neither this option nor the **-x** option is used to specify a list of potential partitions to exclude, the list of potential partitions to exclude defaults to an empty list.

This option is only valid when **-o calcscore** is specified.

--filter The filter to apply to the partition affinity scores to list. The filter is used to select which partitions to list. If a filter is not specified, then all partitions will be listed.

The filter data consists of filter name/value pairs, which are in comma separated value (CSV) format. The filter data must be enclosed in double quotes.

The format of the filter data is as follows:

```
"filter-name=value"
```

Note that the filters accept a comma separated list of values, as follows:

```
""filter-name=value,value,...""
```

When a list of values is specified, the filter name/value pair must be enclosed in double quotes. Depending on the shell being used, nested double quote characters may need to be preceded by an escape character, which is usually a `'\'` character.

Multiple values can be specified for each filter.

Valid filter names:

lpar_names | lpar_ids

This option is only valid when listing partition affinity scores.

-F A delimiter separated list of attribute names for the desired attribute values to be displayed. If no attribute names are specified, then values for all of the attributes will be displayed.

When this option is specified, only attribute values will be displayed. No attribute names will be displayed. The attribute values displayed will be separated by the delimiter which was specified with this option.

This option is useful when only attribute values are desired to be displayed, or when the values of only selected attributes are desired to be displayed.

--header

Display a header record, which is a delimiter separated list of attribute names for the attribute values that will be displayed. This header record will be the first record displayed. This option is only valid when used with the **-F** option.

--help Display the help text for this command and exit.

EXAMPLES

List information about the last memory optimization operation that was started:

```
lsmemopt -m system1
```

List just the status of the last memory optimization operation that was started:

```
lsmemopt -m 8233-E8B*1234321 -F status
```

List the current system affinity score:

```
lsmemopt -m sys -o currscore
```

List the current affinity scores for all partitions in the managed system:

```
lsmemopt -m sys -o currscore -r lpar
```

List the current affinity scores for partitions **lp2** and **lp3** only:

```
lsmemopt -m sys -o currscore -r lpar --filter ""lpar_names=lp2,lp3""
```

Calculate and list the potential system affinity score that could be attained if a Dynamic Platform Optimization operation were run to optimize all partitions in the system:

```
lsmemopt -m sys -o calcscore
```

Calculate and list the potential partition affinity scores that could be attained if a Dynamic Platform Optimization operation that prioritizes partitions **lp1** and **lp3** were run:

```
lsmemopt -m sys -o calcscore -r lpar -p lp1,lp3
```

Calculate and list the potential system affinity score that could be attained if a Dynamic Platform Optimization operation that prioritizes the partitions with IDs **2**, **4**, **6** through **9**, and **11**, and excludes the partition with ID **1**, were run:

```
lsmemopt -m sys -o calcscore --id 2,4,6-9,11 --xid 1
```

ENVIRONMENT

None

BUGS

None

AUTHOR

IBM Austin

SEE ALSO

optmem

NAME

lsmigrdbg - list partition migration debug resources

SYNOPSIS

lsmigrdbg -r {**config** | **file**}
[-**F** [*attribute-names*] [**--header**]] [**--help**]

DESCRIPTION

lsmigrdbg lists partition migration debug resources for this management console.

OPTIONS

- r** The type of partition migration debug resources to list. Valid values are **config** to list partition migration failure debug data collection configuration, and **file** to list available partition migration failure debug data files.
- F** A delimiter separated list of attribute names for the desired attribute values to be displayed. If no attribute names are specified, then values for all of the attributes will be displayed.
- When this option is specified, only attribute values will be displayed. No attribute names will be displayed. The attribute values displayed will be separated by the delimiter which was specified with this option.
- This option is useful when only attribute values are desired to be displayed, or when the values of only selected attributes are desired to be displayed.
- header** Display a header record, which is a delimiter separated list of attribute names for the attribute values that will be displayed. This header record will be the first record displayed. This option is only valid when used with the **-F** option.
- help** Display the help text for this command and exit.

EXAMPLES

List the partition migration failure debug data collection configuration for this management console:

lsmigrdbg -r config

List all of the partition migration failure debug data packages on this management console:

lsmigrdbg -r file

ENVIRONMENT

None

BUGS

None

AUTHOR

IBM Austin

SEE ALSO

migrdbg

NAME

lsnportlogin - list N_Port login information

SYNOPSIS

lsnportlogin [-m *managed-system*] [-w *wait-time*] [-d *detail-level*]
--filter "*filter-data*" [-F [*attribute-names*] [--header]] [--help]

DESCRIPTION

lsnportlogin lists WWPN login status information for virtual fibre channel client adapters configured in partitions or partition profiles.

OPTIONS

-m The name of the managed system for which to list WWPN login status information. The name may either be the user-defined name for the managed system, or be in the form *tttt-mmm*sssssss*, where *tttt* is the machine type, *mmm* is the model, and *sssssss* is the serial number of the managed system. The *tttt-mmm*sssssss* form must be used if there are multiple managed systems with the same user-defined name.

-w The maximum time, in minutes, to wait for each Virtual I/O Server (VIOS) command issued by the HMC to complete. If a VIOS command does not complete within the time specified, the query operation for the virtual fibre channel client adapter will be stopped.

wait-time must be a whole number. If wait-time is 0, the operation will not be timed out.

If this option is not specified, a default value of 3 minutes is used.

-d The level of detail requested from VIOS commands issued by the HMC. Valid values are **0** (none) through **5** (highest).

If this option is not specified, a default value of 1 is used.

--filter The filters to apply to the WWPN login status information to be listed. Filters are used to select the partitions or partition profiles for which WWPN login status information is to be listed.

The filter data consists of filter name/value pairs, which are in comma separated value (CSV) format. The filter data must be enclosed in double quotes.

The format of the filter data is as follows:

```
"filter-name=value,filter-name=value,..."
```

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

```
""filter-name=value,value,...",..."
```

When a list of values is specified, the filter name/value pair must be enclosed in double quotes. Depending on the shell being used, nested double quote characters may need to be preceded by an escape character, which is usually a `'\'` character.

Multiple values can be specified for each filter.

Valid filter names:

lpar_names | lpar_ids, profile_names

-F A delimiter separated list of attribute names for the desired attribute values to be displayed. If no attribute names are specified, then values for all of the attributes will be displayed.

When this option is specified, only attribute values will be displayed. No attribute names will be displayed. The attribute values displayed will be separated by the delimiter which was specified

with this option.

This option is useful when only attribute values are desired to be displayed, or when the values of only selected attributes are desired to be displayed.

Descriptions of selected command attributes:

wwpn_status

The WWPN status. Possible values are:

0 - WWPN is not activated

1 - WWPN is activated

2 - WWPN status is unknown

--header

Display a header record, which is a delimiter separated list of attribute names for the attribute values that will be displayed. This header record will be the first record displayed. This option is only valid when used with the **-F** option.

--help Display the help text for this command and exit.

EXAMPLES

List the WWPN login status information for partition **clientPartition**:

```
lsnportlogin -m sys --filter "lpar_names=clientPartition"
```

List the WWPN login status information for profile **defaultProf** for the partition with ID **5**:

```
lsnportlogin -m 9117-MMB*1234567 --filter "profile_names=defaultProf, lpar_ids=5"
```

List only the partition name, WWPN, and WWPN status, and separate the output values with a colon:

```
lsnportlogin -m sys --filter "lpar_names=clientPartition" -F lpar_name:wwpn:wwpn_status
```

ENVIRONMENT

None

BUGS

None

AUTHOR

IBM Austin

SEE ALSO

chnportlogin

NAME

lsprofspace - list profile space information

SYNOPSIS

lsprofspace -m *managed-system*
[-F [*attribute-names*] [--header]] [--help]

DESCRIPTION

lsprofspace lists information about the amount of space used for partition profiles in the Hardware Management Console's storage area on *managed-system*. This command also lists information about the amount of space available in that storage area for additional partition profiles.

OPTIONS

- m** The name of the managed system for which to list partition profile space information. The name may either be the user-defined name for the managed system, or be in the form *ttt-mmm*ssssss*, where *ttt* is the machine type, *mmm* is the model, and *ssssss* is the serial number of the managed system. The *ttt-mmm*ssssss* form must be used if there are multiple managed systems with the same user-defined name.
- F** A delimiter separated list of attribute names for the desired attribute values to be displayed. If no attribute names are specified, then values for all of the attributes will be displayed.

When this option is specified, only attribute values will be displayed. No attribute names will be displayed. The attribute values displayed will be separated by the delimiter which was specified with this option.

This option is useful when only attribute values are desired to be displayed, or when the values of only selected attributes are desired to be displayed.

Command attributes:

avail_space_in_bytes

The amount of space, in bytes, currently available in the HMC's storage area on *managed-system*.

avail_space_in_client_profiles

An estimate of the number of additional client (AIX, Linux, or IBM i) partition profiles that can be created on *managed-system* based on the space currently available in the HMC's storage area on *managed-system*. This estimate uses the average size of all client partition profiles that already exist on *managed-system*.

If you plan to create any additional partitions, you must account for the additional last valid configuration profile that the HMC will automatically create for each partition.

client_profile_bytes_in_use

The amount of space, in bytes, currently used in the HMC's storage area on *managed-system* for all of the client (AIX, Linux, and IBM i) partition profiles.

client_profiles_in_use

The actual number of client (AIX, Linux, and IBM i) partition profiles configured for *managed-system*.

vios_profile_bytes_in_use

The amount of space, in bytes, currently used in the HMC's storage area on *managed-system* for all of the Virtual I/O Server (VIOS) partition profiles.

vios_profiles_in_use

The actual number of VIOS partition profiles configured for *managed-system*.

--header

Display a header record, which is a delimiter separated list of attribute names for the attribute values that will be displayed. This header record will be the first record displayed. This option is only valid when used with the **-F** option.

--help Display the help text for this command and exit.

EXAMPLES

List profile space information for a managed system:

lsprofspace -m mySys

ENVIRONMENT

None

BUGS

None

AUTHOR

IBM Austin

NAME

lsproxy - list proxy settings

SYNOPSIS

lsproxy -t sslproxy
[-F *attribute-names*] [--header] [--help]

DESCRIPTION

lsproxy lists the Secure Sockets Layer (SSL) proxy settings used by the Hardware Management Console when performing call-home functions.

OPTIONS

-t The type of proxy settings to list. The only valid value is **sslproxy** to list the SSL proxy settings used when performing call-home functions.

-F A delimiter separated list of attribute names for the desired attribute values to be displayed. If no attribute names are specified, then values for all of the attributes will be displayed.

When this option is specified, only attribute values will be displayed. No attribute names will be displayed. The attribute values displayed will be separated by the delimiter which was specified with this option.

This option is useful when only attribute values are desired to be displayed, or when the values of only selected attributes are desired to be displayed.

--header

Display a header record, which is a delimiter separated list of attribute names for the attribute values that will be displayed. This header record will be the first record displayed. This option is only valid when used with the **-F** option.

--help Display the help text for this command and exit.

EXAMPLES

List the SSL proxy settings used when performing call-home functions:

```
lsproxy -t sslproxy
```

ENVIRONMENT

None

BUGS

None

AUTHOR

IBM Austin

SEE ALSO

chproxy

NAME

lspsm - list Power systems management console settings

SYNOPSIS

lspsm [-F *attribute-names*] [--header] [--help]

DESCRIPTION

lspsm lists Power systems management console configuration settings.

OPTIONS

-F A delimiter separated list of attribute names for the desired attribute values to be displayed. If no attribute names are specified, then values for all of the attributes will be displayed.

When this option is specified, only attribute values will be displayed. No attribute names will be displayed. The attribute values displayed will be separated by the delimiter which was specified with this option.

This option is useful when only attribute values are desired to be displayed, or when the values of only selected attributes are desired to be displayed.

--header

Display a header record, which is a delimiter separated list of attribute names for the attribute values that will be displayed. This header record will be the first record displayed. This option is only valid when used with the **-F** option.

--help Display the help text for this command and exit.

EXAMPLES

List the Power systems management console settings:

lspsm

ENVIRONMENT

None

BUGS

None

AUTHOR

IBM Austin

SEE ALSO

chpsm

NAME

lspwdpolicy - list password policy information

SYNOPSIS

lspwdpolicy -t {p | s} [-F *attribute-names*] [--header] [--help]

DESCRIPTION

lspwdpolicy lists Hardware Management Console (HMC) password policy information.

OPTIONS

-t The type of password policy information to list. Valid values are **p** to list all of the password policies on the HMC, and **s** to list password policy status information. Password policy status information indicates whether a password policy is active, and if so, which password policy is the active policy.

-F A delimiter separated list of attribute names for the desired attribute values to be displayed. If no attribute names are specified, then values for all of the attributes will be displayed.

When this option is specified, only attribute values will be displayed. No attribute names will be displayed. The attribute values displayed will be separated by the delimiter which was specified with this option.

This option is useful when only attribute values are desired to be displayed, or when the values of only selected attributes are desired to be displayed.

--header

Display a header record, which is a delimiter separated list of attribute names for the attribute values that will be displayed. This header record will be the first record displayed. This option is only valid when used with the **-F** option.

--help Display the help text for this command and exit.

EXAMPLES

List all of the HMC password policies:

```
lspwdpolicy -t p
```

List just the names of all of the HMC password policies:

```
lspwdpolicy -t p -F name
```

List HMC password policy status information:

```
lspwdpolicy -t s
```

ENVIRONMENT

None

BUGS

None

AUTHOR

IBM Austin

SEE ALSO

chpwdpolicy, mkpwdpolicy, rmpwdpolicy

NAME

`lspwrmgmt` - list power management settings

SYNOPSIS

To list power saver mode settings:

```
lspwrmgmt [-m managed-system] -r sys  
[-F [attribute-names] [--header]] [--help]
```

To list dynamic or idle power saver parameters:

```
lspwrmgmt -m managed-system -r sys -t {dynamic_parms | idle_parms}  
[-F [attribute-names] [--header]] [--help]
```

DESCRIPTION

`lspwrmgmt` lists the power management settings for managed systems.

OPTIONS

-m The name of the managed system for which to list power management settings. The name may either be the user-defined name for the managed system, or be in the form *ttt-mmm*ssssss*, where *ttt* is the machine type, *mmm* is the model, and *ssssss* is the serial number of the managed system. The *ttt-mmm*ssssss* form must be used if there are multiple managed systems with the same user-defined name.

When listing power saver mode settings, if this option is omitted power saver mode settings for all of the systems managed by this Hardware Management Console (HMC) will be listed.

This option is required when listing dynamic or idle power saver parameters.

-r The type of resources for which to list power management settings. The only valid value is **sys** for managed systems.

-t The type of power saver parameters to list. Valid values are **dynamic_parms** to list dynamic power saver tuning parameters, and **idle_parms** to list idle power saver parameters.

-F A delimiter separated list of attribute names for the desired attribute values to be displayed. If no attribute names are specified, then values for all of the attributes will be displayed.

When this option is specified, only attribute values will be displayed. No attribute names will be displayed. The attribute values displayed will be separated by the delimiter which was specified with this option.

This option is useful when only attribute values are desired to be displayed, or when the values of only selected attributes are desired to be displayed.

--header

Display a header record, which is a delimiter separated list of attribute names for the attribute values that will be displayed. This header record will be the first record displayed. This option is only valid when used with the **-F** option.

--help Display the help text for this command and exit.

EXAMPLES

List the power saver mode settings for all systems managed by this HMC:

```
lspwrmgmt -r sys
```

List just the supported power saver mode types for managed system **system1**:

```
lspwrmgmt -r sys -m system1 -F supported_power_saver_mode_types
```

List the dynamic power saver tuning parameters for managed system **8286-42A*1234567**:

lspwrmgmt -r sys -m 8286-42A*1234567 -t dynamic_parms

ENVIRONMENT

None

BUGS

None

AUTHOR

IBM Austin

SEE ALSO

chpwrmgmt

NAME

lsrefcode - list reference codes

SYNOPSIS

```
lsrefcode -r {sys | lpar} -m managed-system  
[-s {p | s}] [-n number] [--filter "filter-data"]  
[-F [attribute-names] [--header]] [--help]
```

DESCRIPTION

lsrefcode lists reference codes for the *managed-system* or for partitions in the *managed-system*.

OPTIONS

-r The type of reference codes to list. Valid values are **sys** for managed system reference codes, and **lpar** for partition reference codes.

-m The name of the managed system which has the reference codes to list. The name may either be the user-defined name for the managed system, or be in the form *tttt-mmm*sssssss*, where *tttt* is the machine type, *mmm* is the model, and *sssssss* is the serial number of the managed system. The *tttt-mmm*sssssss* form must be used if there are multiple managed systems with the same user-defined name.

-s When listing managed system reference codes, use this option to specify the service processor for which reference codes are to be listed. Valid values are **p** for the primary service processor and **s** for the secondary service processor. If this option is not specified, reference codes for the primary service processor will be listed.

This option is only valid when listing managed system reference codes.

-n The *number* of reference codes to list, starting with the current reference code, for the *managed-system* or for each partition. Reference codes are listed in order, with the most recent (current) reference code first.

The *number* specified must be greater than 0. If there are fewer reference codes available than *number*, only the available reference codes will be listed.

If this option is omitted, only the current reference code will be listed.

--filter The filter to apply to the partition reference codes to be listed. A filter is used to select the partitions for which reference codes are to be listed. If no filter is specified, then reference codes for all partitions in the *managed-system* will be listed.

The filter data consists of filter name/value pairs, which are in comma separated value (CSV) format. The filter data must be enclosed in double quotes.

The format of the filter data is as follows:

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

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

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

When a list of values is specified, the filter name/value pair must be enclosed in double quotes. Depending on the shell being used, nested double quote characters may need to be preceded by an escape character, which is usually a `\` character.

Valid filter names for this command are **lpar_names** or **lpar_ids**. Only one of these filters may be specified. Multiple partitions can be specified with the filter.

This option is not valid when listing managed system reference codes.

- F** A delimiter separated list of attribute names for the desired attribute values to be displayed for each reference code. If no attribute names are specified, then values for all of the reference code attributes will be displayed.

When this option is specified, only attribute values will be displayed. No attribute names will be displayed. The attribute values displayed will be separated by the delimiter which was specified with this option.

This option is useful when only attribute values are desired to be displayed, or when the values of only selected attributes are desired to be displayed.

--header

Display a header record, which is a delimiter separated list of attribute names for the attribute values that will be displayed. This header record will be the first record displayed. This option is only valid when used with the **-F** option.

- help** Display the help text for this command and exit.

EXAMPLES

List the current reference code for the managed system:

```
lsrefcode -r sys -m system1
```

List the last 10 reference codes for the managed system, and only display attribute values for each reference code, following a header of attribute names:

```
lsrefcode -r sys -m 9406-570*1543901A -n 10 -F --header
```

List the last 5 reference codes for the secondary service processor on the managed system:

```
lsrefcode -r sys -m system1 -s s -n 5
```

List the current reference code for each partition in the managed system, and only list the partition name, time stamp, and reference code value for each reference code, and separate the output values with a comma:

```
lsrefcode -r lpar -m system1 -F lpar_name,time_stamp,  
refcode
```

List the last 25 reference codes for partitions **p1** and **p2**:

```
lsrefcode -r lpar -m 9406-520*12345678 -n 25 --filter  
""lpar_names=p1,p2""
```

ENVIRONMENT

None

BUGS

None

AUTHOR

IBM Austin

NAME

lsrrstartlpar - list partition remote restart information

SYNOPSIS

```
lsrrstartlpar -r {lpar | mc | sys} [-m managed-system]  
[-filter "filter-data"] [-F [attribute-names] [--header]] [--help]
```

DESCRIPTION

lsrrstartlpar lists partition remote restart information.

OPTIONS

-r The type of resources for which to list partition remote restart information.

Specify **lpar** to list partition remote restart information for all of the partitions in *managed-system*.

Specify **mc** to list partition remote restart capabilities for the Hardware Management Console (HMC).

Specify **sys** to list partition remote restart capabilities for *managed-system*.

-m The name of the managed system for which to list partition remote restart information. The name may either be the user-defined name for the managed system, or be in the form *ttt-mmm*ssssss*, where *ttt* is the machine type, *mmm* is the model, and *ssssss* is the serial number of the managed system. The *ttt-mmm*ssssss* form must be used if there are multiple managed systems with the same user-defined name.

This option is not valid when listing partition remote restart information for the HMC. This option is required when listing all other types of partition remote restart information.

--filter The filter(s) to apply to the resources to be listed. Filters are used to select which resources are to be listed. If no filters are used, then all of the resources will be listed. For example, partition information can be listed by using a filter to specify the names or IDs of the partitions to list. Otherwise, if no filter is used, then information for all of the partitions in the managed system will be listed.

The filter data consists of filter name/value pairs, which are in comma separated value (CSV) format. The filter data must be enclosed in double quotes.

The format of the filter data is as follows:

```
"filter-name=value"
```

Note that the filters accept a comma separated list of values, as follows:

```
""filter-name=value,value,...""
```

When a list of values is specified, the filter name/value pair must be enclosed in double quotes. Depending on the shell being used, nested double quote characters may need to be preceded by an escape character, which is usually a `'\'` character.

Multiple values can be specified for each filter.

Valid filter names:

```
lpar_names | lpar_ids
```

This option is only valid when listing partition information.

-F A delimiter separated list of attribute names for the desired attribute values to be displayed. If no attribute names are specified, then values for all of the attributes will be displayed.

When this option is specified, only attribute values will be displayed. No attribute names will be displayed. The attribute values displayed will be separated by the delimiter which was specified with this option.

This option is useful when only attribute values are desired to be displayed, or when the values of only selected attributes are desired to be displayed.

--header

Display a header record, which is a delimiter separated list of attribute names for the attribute values that will be displayed. This header record will be the first record displayed. This option is only valid when used with the **-F** option.

--help Display the help text for this command and exit.

EXAMPLES

List partition remote restart information for all partitions in the managed system **system1**:

```
lsrrstartlpar -r lpar -m system1
```

List partition remote restart information for partition **lp1** only:

```
lsrrstartlpar -r lpar -m system1 --filter "lpar_names=lp1"
```

List the partition remote restart capabilities for the HMC:

```
lsrrstartlpar -r mc
```

List the partition remote restart capabilities for managed system **system1**:

```
lsrrstartlpar -r sys -m system1
```

ENVIRONMENT

None

BUGS

None

AUTHOR

IBM Austin

SEE ALSO

rrstartlpar

NAME

lsrsdevsize - list reserved storage device size

SYNOPSIS

```
lsrsdevsize -m managed-system
[-p partition-name | --id partition-ID]
[-f input-data-file | -i "input-data"]
[-F [attribute-names] [--header]] [--help]
```

DESCRIPTION

lsrsdevsize lists the minimum required reserved storage device size for a suspend capable or remote restart capable partition. The device size listed is in megabytes.

The device size can be listed for an existing partition, or for a partition that has not yet been created.

OPTIONS

- m** The name of the managed system which has, or will have, the partition for which to list the device size. The name may either be the user-defined name for the managed system, or be in the form *ttt-mmm*ssssss*, where *ttt* is the machine type, *mmm* is the model, and *ssssss* is the serial number of the managed system. The *ttt-mmm*ssssss* form must be used if there are multiple managed systems with the same user-defined name.
- p** The name of the partition for which to list the device size. The configuration defined in the last activated profile for the partition will be used to determine the device size for the partition. If the partition has never been activated, then the configuration defined in the default profile will be used. If the partition has not yet been created, then you cannot use this option.

If the partition for which to list the device size already exists, you should either use this option to specify the name of the partition, or use the **--id** option to specify the partition's ID. The **-p** and the **--id** options are mutually exclusive.

- id** The ID of the partition for which to list the device size. The configuration defined in the last activated profile for the partition will be used to determine the device size for the partition. If the partition has never been activated, then the configuration defined in the default profile will be used. If the partition has not yet been created, then you cannot use this option.

If the partition for which to list the device size already exists, you should either use this option to specify the ID of the partition, or use the **-p** option to specify the partition's name. The **--id** and the **-p** options are mutually exclusive.

- f** The name of the file containing the input data for this command. The input data consists of attribute name/value pairs, which are in comma separated value (CSV) format.

The format of the input data is as follows:

```
attribute-name=value,attribute-name=value,...
```

If the partition for which to list the device size already exists and it is specified with the **-p** or **--id** option, then all of the following attributes are optional. Any of the attributes that are specified will override the value in the last activated profile or default profile for the partition (neither profile will be modified). If the partition for which to list the device size does not exist, then all of the following attributes except **hpt_ratio** must be specified.

Valid attribute names:

lpar_env

Valid values are:

aixlinux - the partition will be an AIX or Linux partition

os400 - the partition will be an IBM i partition

max_mem

The maximum amount of memory for the partition, in megabytes.

max_procs

The maximum number of dedicated processors or the maximum number of virtual processors for the partition.

max_virtual_slots

The maximum number of virtual I/O adapter slots for the partition.

mem_mode

Valid values are:

ded - the partition will use dedicated memory

shared - the partition will use shared memory

hpt_ratio

The ratio of the hardware page table size to the maximum memory for the partition. Valid values are displayed by the **lshwres -r mem -m managed-system --level sys -F possible_hpt_ratios** command.

Input data for this command can be specified with this option or the **-i** option. The **-f** and the **-i** options are mutually exclusive.

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

Input data for this command can be specified with this option or the **-f** option. The **-i** and the **-f** options are mutually exclusive.

- F** A delimiter separated list of attribute names for the desired attribute values to be displayed. If no attribute names are specified, then values for all of the attributes will be displayed.

When this option is specified, only attribute values will be displayed. No attribute names will be displayed. The attribute values displayed will be separated by the delimiter which was specified with this option.

This option is useful when only attribute values are desired to be displayed, or when the values of only selected attributes are desired to be displayed.

--header

Display a header record, which is a delimiter separated list of attribute names for the attribute values that will be displayed. This header record will be the first record displayed. This option is only valid when used with the **-F** option.

- help** Display the help text for this command and exit.

EXAMPLES

List the minimum required reserved storage device size for a partition that has not yet been created:

```
lsrsdevsize -m system1 -i "lpar_env=aixlinux,max_mem=8192,max_procs=1,
max_virtual_slots=25,mem_mode=ded"
```

List the minimum required reserved storage device size for partition **p1** using the configuration defined in the last activated or default profile for the partition:

lssrsdevsize -m 9117-MMB*1000234 -p p1

List the minimum required reserved storage device size for the partition with ID **5** using the configuration defined in the last activated or default profile for the partition and a maximum memory of **4096** MB:

lssrsdevsize -m system1 --id 5 -i "max_mem=4096"

ENVIRONMENT

None

BUGS

None

AUTHOR

IBM Austin

SEE ALSO

lssyscfg

NAME

lssacfg - list Service Agent configuration information

SYNOPSIS

```
lssacfg -t { authuser | callhome | callhomeserver | custinfo | email |  
          ftpfirewall | ftpoffload | snmp | snmpopt | snmptrapnames }  
[-F [attribute-names]] [--header] [--help]
```

DESCRIPTION

lssacfg lists Service Agent configuration information.

OPTIONS

-t The type of Service Agent configuration information to list. Valid values are **authuser** for the users that are authorized to access information collected by Electronic Service Agent for the systems managed by this Hardware Management Console (HMC), **callhome** for the call-home states for the systems managed by this HMC, **callhomeserver** for the call-home server settings for this HMC, **custinfo** for customer information settings, **email** for customer email notification settings, **ftpfirewall** for FTP firewall settings, **ftpoffload** for FTP offload server settings, **snmp** for SNMP trap notification settings, **snmpopt** for SNMP options, and **snmptrapnames** to list all defined SNMP traps.

-F A delimiter separated list of attribute names representing the desired attribute values to display. If this option is specified without any attribute names, then all of the attributes will be displayed.

When this option is specified, only attribute values will be displayed. No attribute names will be displayed. The attribute values displayed will be separated by the delimiter which was specified with this option.

This option is useful when only attribute values are desired to be displayed, or when the values of only selected attributes are desired to be displayed.

--header

Display a header record, which is a delimiter separated list of attribute names for the attribute values that will be displayed. This header record will be the first record displayed. This option is only valid when used with the **-F** option.

--help Display the help text for this command and exit.

EXAMPLES

Display the users that are authorized to access information collected by Electronic Service Agent for the systems managed by this HMC:

```
lssacfg -t authuser
```

Display the call-home states for the systems managed by this HMC:

```
lssacfg -t callhome
```

Display the call-home server settings for this HMC:

```
lssacfg -t callhomeserver
```

Display the customer information settings for this HMC:

```
lssacfg -t custinfo
```

Display the Service Agent customer email notification settings:

```
lssacfg -t email
```

Display the FTP firewall settings used for offloading service information:

lssacfg -t ftpfirewall

Display the FTP server host name and directory used for offloading service information:

lssacfg -t ftpoffload -F host,directory

Display the Service Agent configuration for emitting SNMP trap notifications:

lssacfg -t snmp

Display the Service Agent SNMP options:

lssacfg -t snmpopt

Display all defined Service Agent SNMP traps:

lssacfg -t snmptrapnames

ENVIRONMENT

None

BUGS

None

AUTHOR

IBM Austin

SEE ALSO

chsacfg

NAME

lsstat - list statistics

SYNOPSIS

lsstat -r { **samplerate** | **vlan** } [**-m** *managed-system*]
[**-F** [*attribute-names*] [**--header**]] [**--help**]

DESCRIPTION

lsstat lists statistics collected for a *managed-system*. This command also lists the Hardware Management Console (HMC) settings, such as the sample rate, for statistics collection.

When listing statistics collected for a *managed-system*, only the last sample collected is listed.

OPTIONS

-r The type of resources to list. Valid values are **samplerate** for the sample rate for statistics collection, and **vlan** for Virtual Local Area Network (VLAN) statistics data.

-m The name of the managed system for which statistics data or the HMC settings for statistics collection are to be listed. The name may either be the user-defined name for the managed system, or be in the form *ttt-mmm*ssssss*, where *ttt* is the machine type, *mmm* is the model, and *ssssss* is the serial number of the managed system. The *ttt-mmm*ssssss* form must be used if there are multiple managed systems with the same user-defined name.

This option is required when listing statistics data. This option is not required when listing the HMC settings for statistics collection. If it is not specified, then the settings for all of the systems that support statistics collection and are currently managed by this HMC will be listed.

-F A delimiter separated list of attribute names for the desired attribute values to be displayed. If no attribute names are specified, then values for all of the attributes will be displayed.

When this option is specified, only attribute values will be displayed. No attribute names will be displayed. The attribute values displayed will be separated by the delimiter which was specified with this option.

This option is useful when only attribute values are desired to be displayed, or when the values of only selected attributes are desired to be displayed.

--header

Display a header record, which is a delimiter separated list of attribute names for the attribute values that will be displayed. This header record will be the first record displayed. This option is only valid when used with the **-F** option.

--help Display the help text for this command and exit.

EXAMPLES

List the VLAN statistics data for managed system **mySystem**:

```
lsstat -r vlan -m mySystem
```

List the sample rates for all managed systems that support VLAN statistics collection:

```
lsstat -r samplerate
```

ENVIRONMENT

None

BUGS

None

AUTHOR

IBM Austin

SEE ALSO

chstat

NAME

lssvc - list service settings

SYNOPSIS

lssvc -s snmp [-r {trapnames | trapconfig}]
[-F *attribute-names*] [--header] [--help]

DESCRIPTION

lssvc lists the settings for a service, such as SNMP, on the Hardware Management Console (HMC).

OPTIONS

- s** The service for which to list settings. The only valid value is **snmp** for Simple Network Management Protocol (SNMP) version 3.
- r** The type of service resources to list. Valid values are **trapnames** to list all of the traps supported by the SNMPv3 agent on the HMC, and **trapconfig** to list the SNMPv3 trap notification list. If this option is omitted, then general status information about the SNMPv3 agent on the HMC will be listed.
- F** A delimiter separated list of attribute names for the desired attribute values to be displayed. If no attribute names are specified, then values for all of the attributes will be displayed.

When this option is specified, only attribute values will be displayed. No attribute names will be displayed. The attribute values displayed will be separated by the delimiter which was specified with this option.

This option is useful when only attribute values are desired to be displayed, or when the values of only selected attributes are desired to be displayed.

--header

Display a header record, which is a delimiter separated list of attribute names for the attribute values that will be displayed. This header record will be the first record displayed. This option is only valid when used with the **-F** option.

- help** Display the help text for this command and exit.

EXAMPLES

List status information about the SNMPv3 agent on the HMC:

```
lssvc -s snmp
```

List all of the traps supported by the SNMPv3 agent on the HMC:

```
lssvc -s snmp -r trapnames
```

List the SNMPv3 trap notification list:

```
lssvc -s snmp -r trapconfig
```

ENVIRONMENT

None

BUGS

None

AUTHOR

IBM Austin

SEE ALSO

chsvc

NAME

lssvcevents - list console or serviceable events

SYNOPSIS

```
lssvcevents -t {console | hardware}  
[-d number-of-days | -i number-of-minutes]  
[-m managed-system]  
[--filter "filter-data"]  
[-F [attribute-names] [--header]] [--help]
```

DESCRIPTION

lssvcevents lists console events logged by the Hardware Management Console (HMC), or serviceable events.

OPTIONS

-t The type of events to list. Valid values are **console** for console events, or **hardware** for serviceable events.

-d The number of days prior to today for which events will be listed. Events that occurred today will be listed, along with any events that occurred during the past *number-of-days* days.

If this option is omitted when listing console events, console events that occurred within the past 7 days will be listed.

If this option is omitted when listing serviceable events, all serviceable events will be listed.

-i The number of minutes to go back and search for events. This search is based on the time that the event was initially created on the HMC, and is not affected by later updates to the event.

-m The name of the managed system for which serviceable events are to be listed. The name may either be the user-defined name for the managed system, or be in the form *tttt-mmm*sssssss*, where *tttt* is the machine type, *mmm* is the model, and *sssssss* is the serial number of the managed system. The *tttt-mmm*sssssss* form must be used if there are multiple managed systems with the same user-defined name.

If this option is omitted, all serviceable events will be listed.

This option is only valid when listing serviceable events.

--filter The filter to apply to the serviceable events to be listed. A filter is used to select which serviceable events are to be listed. For example, only open serviceable events can be listed by using a filter to specify the status (open) of the serviceable events to list. If a filter is not used, then all serviceable events will be listed.

The filter data consists of filter name/value pairs, which are in comma separated value (CSV) format. The filter data must be enclosed in double quotes.

The format of the filter data is as follows:

```
"filter-name=value,filter-name=value,..."
```

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

```
""filter-name=value,value,...""
```

When a list of values is specified, the filter name/value pair must be enclosed in double quotes. Depending on the shell being used, nested double quote characters may need to be preceded by an escape character, which is usually a ``\`` character.

Multiple values can be specified for each filter.

Valid filter names for serviceable events:

problem_nums

Specify event problem number(s)

status

Specify event status. Valid values are

open or **closed**.

approval_state

Specify event approval state. Valid values

are **approved** or **unapproved**.

This option is not valid when listing console events.

- F** A delimiter separated list of attribute names for the desired attribute values to be displayed for each event. If no attribute names are specified, then values for all of the attributes for each event will be displayed.

When this option is specified, only attribute values will be displayed. No attribute names will be displayed. The attribute values displayed will be separated by the delimiter which was specified with this option.

This option is useful when only attribute values are desired to be displayed, or when the values of only selected attributes are desired to be displayed.

--header

Display a header record, which is a delimiter separated list of attribute names for the attribute values that will be displayed. This header record will be the first record displayed. This option is only valid when used with the **-F** option.

- help** Display the help text for this command and exit.

EXAMPLES

List the serviceable events that occurred today:

```
lssvcevents -t hardware -d 0
```

List the console events that occurred within the past 3 days:

```
lssvcevents -t console -d 3
```

List all of the open serviceable events for the system **system1**:

```
lssvcevents -t hardware -m system1 --filter "status=open"
```

List only the problem numbers and status of all serviceable events for the system that occurred within the last 7 days, and separate the output values with a colon:

```
lssvcevents -t hardware -m 9406-570*101234A -d 7 -F  
problem_num:status
```

ENVIRONMENT

None

BUGS

None

NAME

lssyscfg - list system resources

SYNOPSIS

```
lssyscfg -r { lpar | prof | sys | sysprof | cage | frame }  
[-m managed-system | -e managed-frame]  
[--osrefresh] [--lastvalidcfg]  
[--filter "filter-data"]  
[-F [attribute-names] [--header]] [--help]
```

DESCRIPTION

lssyscfg lists the attributes of partitions, partition profiles, or system profiles for the *managed-system*. It can also list the attributes of the *managed-system*, and of all of the systems managed by this Hardware Management Console (HMC).

lssyscfg can also list the attributes of cages in the *managed-frame*, the attributes of the *managed-frame*, or the attributes of all of the frames managed by this HMC.

OPTIONS

-r The type of resources to list. Valid values are **lpar** for partitions, **prof** for partition profiles, **sys** for managed systems, **sysprof** for system profiles, **cage** for managed frame cages, and **frame** for managed frames.

-m The name of either the managed system to list, or the managed system which has the system resources to list. The name may either be the user-defined name for the managed system, or be in the form *ttt-**mmm*****ssssss***, where *ttt* is the machine type, *mmm* is the model, and *ssssss* is the serial number of the managed system. The *ttt-**mmm*****ssssss*** form must be used if there are multiple managed systems with the same user-defined name.

This option is required when listing partitions, partition profiles, or system profiles. This option is optional when listing managed systems, and if it is omitted, then all of the systems managed by this HMC will be listed. This option is not valid when listing managed frame cages or managed frames.

-e The name of either the managed frame to list, or the managed frame which contains the cages to list. The name may either be the user-defined name for the managed frame, or be in the form *ttt-**mmm*****ssssss***, where *ttt* is the type, *mmm* is the model, and *ssssss* is the serial number of the managed frame. The *ttt-**mmm*****ssssss*** form must be used if there are multiple managed frames with the same user-defined name.

This option is required when listing managed frame cages. This option is optional when listing managed frames, and if it is omitted, then all of the frames managed by this HMC will be listed. This option is not valid when listing partitions, partition profiles, system profiles, or managed systems.

--osrefresh

When listing partitions, specify this option to refresh the current operating system version information for the partitions first. If a partition does not have an active RMC connection to the management console, **Unknown** will be displayed for that partition's operating system version information.

Specifying this option may cause this command to take a long time to complete if many partitions are being listed or there are network issues.

This option is only valid when listing partitions.

--lastvalidcfg

When listing partition profiles, specify this option to include the last valid configuration partition profiles.

This option is only valid when listing partition profiles.

--filter The filter(s) to apply to the resources to be listed. Filters are used to select which resources of the specified resource type are to be listed. If no filters are used, then all of the resources of the specified resource type will be listed. For example, specific partitions can be listed by using a filter to specify the names or IDs of the partitions to list. Otherwise, if no filter is used, then all of the partitions in the managed system will be listed.

The filter data consists of filter name/value pairs, which are in comma separated value (CSV) format. The filter data must be enclosed in double quotes.

The format of the filter data is as follows:

```
"filter-name=value,filter-name=value,..."
```

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

```
""filter-name=value,value,...",..."
```

When a list of values is specified, the filter name/value pair must be enclosed in double quotes. Depending on the shell being used, nested double quote characters may need to be preceded by an escape character, which is usually a `'\'` character.

Multiple values can be specified for each filter.

Valid filter names for partitions:

lpar_names | lpar_ids | work_groups

Only one of these three filters may be specified.

Valid filter names for partition profiles:

lpar_names | lpar_ids, profile_names

Valid filter names for system profiles:

profile_names

This option is not valid when listing managed systems, managed frame cages, or managed frames.

-F A delimiter separated list of attribute names for the desired attribute values to be displayed for each resource. If no attribute names are specified, then values for all of the attributes for the resource will be displayed.

When this option is specified, only attribute values will be displayed. No attribute names will be displayed. The attribute values displayed will be separated by the delimiter which was specified with this option.

This option is useful when only attribute values are desired to be displayed, or when the values of only selected attributes are desired to be displayed.

To display the values of only selected properties of the **sriov_eth_logical_ports** or **vnic_adapters** partition profile attributes, use the following format:

```
attribute-name=property-name:property-name:...
```

If no property names are specified, then values for all of the properties for each attribute will be displayed.

Descriptions of selected command attributes:

capabilities

A comma separated list of managed system capabilities. Capabilities that are supported by the managed system are displayed in the list. Capabilities that are not supported are not displayed.

--header

Display a header record, which is a delimiter separated list of attribute names for the attribute values that will be displayed. This header record will be the first record displayed. This option is only valid when used with the **-F** option.

--help Display the help text for this command and exit.

EXAMPLES

List all systems managed by this HMC:

lssyscfg -r sys

List only the user-defined name, machine type and model, and serial number for all of the systems managed by this HMC, and separate the output values with a colon:

lssyscfg -r sys -F name:type_model:serial_num

List the managed system system1:

lssyscfg -r sys -m system1

List all of the capabilities supported by managed system system1:

lssyscfg -r sys -m system1 -F capabilities

List all partitions in the managed system, and only display attribute values for each partition, following a header of attribute names:

lssyscfg -r lpar -m 9406-570*1234567 -F --header

List the partitions lpar1, lpar2, and lpar3:

lssyscfg -r lpar -m system1 --filter "'lpar_names=lpar1, lpar2,lpar3'"

List only the names, IDs, and states of partitions lpar1, lpar2, and lpar3, and separate the output values with a comma:

lssyscfg -r lpar -m system1 --filter "'lpar_names=lpar1, lpar2,lpar3'" -F name,lpar_id,state

List the RMC connection state, operating system version, and DLPAR capabilities via RMC connection for partition aix1:

lssyscfg -r lpar -m system1 -F rmc_state,os_version,dlpar_mem_capable, dlpar_proc_capable,dlpar_io_capable --filter "lpar_names=aix1"

List the operating system version and remote OS shutdown capability for partition IBMi1:

```
lssyscfg -r lpar -m system1 -F os_version,remote_osshutdown_capable --filter "lpar_names=IBMi1"
```

Refresh the operating system version information for partition aix1, then list the partition:

```
lssyscfg -r lpar -m system1 --filter "lpar_names=aix1" --osrefresh
```

List all partition profiles in the managed system:

```
lssyscfg -r prof -m 9406-570*1234567
```

List all partition profiles in the managed system and include the last valid configuration partition profiles:

```
lssyscfg -r prof -m 9406-570*1234567 --lastvalidcfg
```

List all partition profiles defined for partition lpar2:

```
lssyscfg -r prof -m system1 --filter "lpar_names=lpar2"
```

List the partition profiles prof1 and prof2 defined for the partition that has an ID of 2:

```
lssyscfg -r prof -m system1 --filter "lpar_ids=2,  
"profile_names=prof1,prof2""
```

List the values of all properties for all SR-IOV ethernet logical ports in all partition profiles in the managed system:

```
lssyscfg -r prof -m system1 -F sriov_eth_logical_ports
```

List the values of specific properties for all SR-IOV ethernet logical ports in all partition profiles in the managed system:

```
lssyscfg -r prof -m system1 -F sriov_eth_logical_ports=config_id:  
adapter_id:logical_port_id:port_vlan_id
```

List all system profiles defined for the managed system:

```
lssyscfg -r sysprof -m 9406-520*100128A
```

List the system profile sysprof1:

```
lssyscfg -r sysprof -m system1 --filter "profile_names=  
sysprof1"
```

List all frames managed by this HMC:

```
lssyscfg -r frame
```

List the managed frame myFrame:

```
lssyscfg -r frame -e myFrame
```

List all cages in the managed frame:

```
lssyscfg -r cage -e 9119-59*000012C
```

ENVIRONMENT

None

BUGS

None

AUTHOR

IBM Austin

SEE ALSO

chsyscfg, mksyscfg, rmsyscfg, lshwres

NAME

lssysconn - list system connections

SYNOPSIS

lssysconn -r { **all** | **nondiscover** } [-F *attribute-names*] [--header]
[--help]

DESCRIPTION

lssysconn lists connection information for all of the systems and frames managed by this Hardware Management Console (HMC). Connection information for all systems and frames to which this HMC is connected or attempting to connect is listed.

lssysconn also lists IP addresses that cannot be automatically discovered by this HMC when using DHCP. If this HMC is set up as a DHCP server on a private network, whenever the Remove Connection task or the **rmsysconn** command is run to remove a managed system or a managed frame from the HMC, the HMC places the IP address(es) of that system or frame in a list of removed IP addresses. Any IP address in that list will not be rediscovered when reattached to the HMC. The **lssysconn -r nondiscover** command can be used to display the contents of that list of removed IP addresses.

OPTIONS

-r The type of resources for which to list connection information. Valid values are **all** for all managed systems and managed frames, and **nondiscover** for all IP addresses that cannot be automatically discovered by this HMC when using DHCP.

-F A delimiter separated list of attribute names for the desired attribute values to be displayed for each resource. If no attribute names are specified, then values for all of the attributes for each resource will be displayed.

When this option is specified, only attribute values will be displayed. No attribute names will be displayed. The attribute values displayed will be separated by the delimiter which was specified with this option.

This option is useful when only attribute values are desired to be displayed, or when the values of only selected attributes are desired to be displayed.

--header

Display a header record, which is a delimiter separated list of attribute names for the attribute values that will be displayed. This header record will be the first record displayed. This option is only valid when used with the **-F** option.

--help Display the help text for this command and exit.

EXAMPLES

List connection information for all systems and frames managed by this HMC:

```
lssysconn -r all
```

List only the IP address and connection state for all of the systems and frames managed by this HMC, and separate the output values with a colon:

```
lssysconn -r all -F ipaddr:state
```

List all IP addresses that cannot be automatically discovered by this HMC when using DHCP:

```
lssysconn -r nondiscover
```

ENVIRONMENT

None

BUGS

None

AUTHOR

IBM Austin

SEE ALSO

lssyscfg, mksysconn, rmsysconn

NAME

lssysplan - list system plans

SYNOPSIS

lssysplan [-F *[attribute-names]*] [--header] [--help]

DESCRIPTION

lssysplan lists the system plan files in the system plan file directory on the Hardware Management Console (HMC).

OPTIONS

-F A delimiter separated list of attribute names for the desired attribute values to be displayed for each file. If no attribute names are specified, then values for all of the attributes for each file will be displayed.

When this option is specified, only attribute values will be displayed. No attribute names will be displayed. The attribute values displayed will be separated by the delimiter which was specified with this option.

This option is useful when only attribute values are desired to be displayed, or when the values of only selected attributes are desired to be displayed.

The possible attributes that can be listed are: **name**, **description**, **source**, **version**, and **date**.

--header

Display a header record, which is a delimiter separated list of attribute names for the attribute values that will be displayed. The header record will be the first record displayed. This option is only valid when used with the **-F** option.

--help Display the help text for this command and exit.

EXAMPLES

List all of the system plan files on this HMC:

lssysplan

List the system plan files on this HMC, and only display attribute values for each file, following a header of attribute names:

lssysplan -F --header

List only the names, dates, and descriptions of the system plan files on this HMC:

lssysplan -F name,date,description

List only the names and descriptions of the system plan files on this HMC, following a header of attribute names:

lssysplan -F name,description --header

ENVIRONMENT

None

BUGS

None

AUTHOR

IBM Austin

SEE ALSO

deploysysplan, cpsysplan, mksysplan, rmsysplan

NAME

lstskey - list trusted system key information

SYNOPSIS

lstskey -m *managed-system* [-F [*attribute-names*]] [--header]] [--help]

DESCRIPTION

lstskey lists trusted system key information for the *managed-system*. The trusted system key is used for virtual Trusted Platform Module (vTPM) data encryption.

OPTIONS

-m The name of the managed system for which to list trusted system key information. The name may either be the user-defined name for the managed system, or be in the form *ttt-mmm*ssssss*, where *ttt* is the machine type, *mmm* is the model, and *ssssss* is the serial number of the managed system. The *ttt-mmm*ssssss* form must be used if there are multiple managed systems with the same user-defined name.

-F A delimiter separated list of attribute names for the desired attribute values to be displayed. If no attribute names are specified, then values for all of the attributes will be displayed.

When this option is specified, only attribute values will be displayed. No attribute names will be displayed. The attribute values displayed will be separated by the delimiter which was specified with this option.

This option is useful when only attribute values are desired to be displayed, or when the values of only selected attributes are desired to be displayed.

--header

Display a header record, which is a delimiter separated list of attribute names for the attribute values that will be displayed. This header record will be the first record displayed. This option is only valid when used with the **-F** option.

--help Display the help text for this command and exit.

EXAMPLES

List the trusted system key information for managed system **system1**:

```
lstskey -m system1
```

ENVIRONMENT

None

BUGS

None

AUTHOR

IBM Austin

SEE ALSO

chtskey

NAME

lsusrzca - Displays welcome text

SYNOPSIS

lsusrzca -t { **w** | **b** } [-**F** [*attribute-names*] [--**header**]] [--**help**]

DESCRIPTION

lsusrzca displays the Web user interface welcome text or the SSH banner text that is shown before users log onto the Hardware Management Console (HMC).

OPTIONS

-t The type of text to display. Valid values are **w** for the welcome text that is displayed on the Web user interface before users log onto the HMC, and **b** for the SSH banner text that is displayed before users remotely log into the HMC using SSH.

-F A delimiter separated list of attribute names for the desired attribute values to be displayed. If no attribute names are specified, then values for all of the attributes will be displayed.

When this option is specified, only attribute values will be displayed. No attribute names will be displayed. The attribute values displayed will be separated by the delimiter which was specified with this option.

This option is useful when only attribute values are desired to be displayed, or when the values of only selected attributes are desired to be displayed.

--header

Display a header record, which is a delimiter separated list of attribute names for the attribute values that will be displayed. This header record will be the first record displayed. This option is only valid when used with the **-F** option.

--help Display the help text for this command and exit.

EXAMPLES

Display the welcome text that is displayed on the Web user interface before a user logs in:

```
lsusrzca -t w
```

Display the SSH banner text that is displayed before a user logs in remotely using SSH:

```
lsusrzca -t b
```

ENVIRONMENT

None

BUGS

None

AUTHOR

IBM Austin

SEE ALSO

chusrzca

NAME

lsvet - list Capacity on Demand advanced functions activation information

SYNOPSIS

lsvet -t {code | hist} -m managed-system
[-F *attribute-names*] [--header] [--help]

DESCRIPTION

lsvet lists Capacity on Demand (CoD) advanced functions activation information for the *managed-system*. CoD advanced functions include PowerVM and Enterprise Enablement.

CoD advanced functions are sometimes referred to as Virtualization Engine systems technologies.

OPTIONS

- t** The type of information to list. Valid values are **code** for information used to generate CoD advanced functions activation codes, and **hist** for the CoD advanced functions activation history log.
- m** The name of the managed system for which information is to be listed. The name may either be the user-defined name for the managed system, or be in the form *ttt-mmm*sssssss*, where *ttt* is the machine type, *mmm* is the model, and *sssssss* is the serial number of the managed system. The *ttt-mmm*sssssss* form must be used if there are multiple managed systems with the same user-defined name.
- F** A delimiter separated list of attribute names representing the desired attribute values to display. If this option is specified without any attribute names, then all of the attributes will be displayed.

When this option is specified, only attribute values will be displayed. No attribute names will be displayed. The attribute values displayed will be separated by the delimiter which was specified with this option.

This option is useful when only attribute values are desired to be displayed, or when the values of only selected attributes are desired to be displayed.

--header

Display a header record, which is a delimiter separated list of attribute names for the attribute values that will be displayed. This header record will be the first record displayed. This option is only valid when used with the **-F** option.

- help** Display the help text for this command and exit.

EXAMPLES

Display activation code generation information:

```
lsvet -m sys1 -t code
```

Display the activation history log:

```
lsvet -m 9117-570*1001213 -t hist
```

ENVIRONMENT

None

BUGS

None

AUTHOR

IBM Austin

SEE ALSO

chvet

NAME

migrcfg - migrate configuration

SYNOPSIS

migrcfg -m *managed-system* **-t 1 -f** *file* [**--help**]

DESCRIPTION

migrcfg migrates partition configuration data to a *managed-system*.

OPTIONS

- m** The name of the managed system to which to migrate the partition configuration data. The name may either be the user-defined name for the managed system, or be in the form *ttt-mmm*sssssss*, where *ttt* is the machine type, *mmm* is the model, and *sssssss* is the serial number of the managed system. The *ttt-mmm*sssssss* form must be used if there are multiple managed systems with the same user-defined name.
- t** The migration data type. The only valid value is **1**.
- f** The name of the file on floppy diskette that contains the partition configuration data to migrate.
- help** Display the help text for this command and exit.

EXAMPLES

migrcfg -m mySystem **-t 1 -f** cfgFile

ENVIRONMENT

None

BUGS

None

AUTHOR

IBM Austin

SEE ALSO

lssyscfg

NAME

migrdbg - partition migration debug tools

SYNOPSIS

To enable or disable automatic partition migration failure debug data collection:

migrdbg -o {e | d} [--help]

To manually collect partition migration debug data:

**migrdbg -o c [-m *managed-system*] [-t *target-managed-system*]
[--help]**

To remove a partition migration debug data package:

migrdbg -o r {-f *file* | --all} [--help]

DESCRIPTION

migrdbg provides partition migration debug tools.

OPTIONS

-o The operation to perform.

Specify **e** to enable automatic collection of partition migration debug data on this management console. When enabled, debug data will be collected automatically for every partition migration that fails on this management console. If you are performing partition migrations to destination managed systems that are managed by other management consoles, you should enable automatic collection of partition migration debug data on those other management consoles also.

Specify **d** to disable automatic collection of partition migration debug data on this management console.

Specify **c** to manually collect partition migration debug data on this management console after a partition migration failure occurs when automatic collection of partition migration debug data on this management console is not enabled. It is recommended that you collect the debug data as close to the partition migration failure as possible. Performing another operation that affects the same source or destination managed system may cause the partition migration debug data for that managed system to be overwritten. If you perform multiple partition migrations with a single **migrpar** command, it is recommended that you use automatic collection of partition migration debug data, since only debug data for the most recent partition migration is likely to be available after the **migrpar** command completes. If a failed partition migration was performed using different source and destination management consoles, you should collect partition migration debug data on both management consoles.

Specify **r** to remove a partition migration debug data package from this management console.

-m The name of the source managed system for the partition migration for which you want to collect debug data. The name may either be the user-defined name for the managed system, or be in the form *tttt-mmm*ssssss*, where *tttt* is the machine type, *mmm* is the model, and *ssssss* is the serial number of the managed system. The *tttt-mmm*ssssss* form must be used if there are multiple managed systems with the same user-defined name.

If the source managed system is not managed by this management console, the debug data must be collected from the management console that manages the source managed system.

-t The name of the target, or destination, managed system for the partition migration for which you want to collect debug data. The name may either be the user-defined name for the managed system, or be in the form *tttt-mmm*ssssss*, where *tttt* is the machine type, *mmm* is the model, and *ssssss* is the serial number of the managed system. The *tttt-mmm*ssssss* form must be used if there are multiple managed systems with the same user-defined name.

If the destination managed system is not managed by this management console, the debug data must be collected from the management console that manages the destination managed system.

- f** The file name of the partition migration debug data package to remove from this management console.
- all** Specify this option to remove all partition migration debug data packages from this management console.
- help** Display the help text for this command and exit.

EXAMPLES

Enable automatic partition migration failure debug data collection on this management console:

migrdbg -o e

Manually collect partition migration debug data for the failed partition migration from managed system **system1** to managed system **system2**:

migrdbg -o c -m system1 -t system2

Manually collect partition migration debug data for the failed partition migration from managed system **S1** to a managed system managed by another management console:

migrdbg -o c -m S1

Manually collect partition migration debug data for the failed partition migration from a managed system managed by another management console to managed system **S2**:

migrdbg -o c -t S2

Remove a partition migration debug data package from this management console:

migrdbg -o r -f lpmfdc-140728152436.tgz

Remove all partition migration debug data packages from this management console:

migrdbg -o r --all

ENVIRONMENT

None

BUGS

None

AUTHOR

IBM Austin

SEE ALSO

lsmigrdbg

NAME

migrpar - perform a partition migration operation

SYNOPSIS

To perform a partition migration operation:

```
migrpar -o { m | r | s | v }  
  -m managed-system [-t target-managed-system]  
  [--ip IP-address [-u user-ID]]  
  { -p partition-name | --id partition-ID | --all }  
  [-n profile-name]  
  [--redundantvios { 0 | 1 | 2 } ] [--mpio { 1 | 2 } ]  
  [--vlanbridge { 1 | 2 } ] [--vsi { 1 | 2 } ]  
  [--uuid { 1 | 2 } ] [--protectstorage { 1 | 2 } ]  
  [--requireerr { 1 | 2 } ] [--usecurrdata ]  
  [ { -f input-data-file | -i "input-data" } ]  
  [-w wait-time] [-d detail-level] [--force]  
  [-v] [--nodetails]  
  [--help]
```

To set system attributes related to partition migration operations:

```
migrpar -o set [-m managed-system] -r sys  
  { -f input-data-file | -i "input-data" }  
  [--help]
```

To set mover service partition (MSP) attributes related to partition migration operations:

```
migrpar -o set -m managed-system -r lpar  
  { -p partition-name | --id partition-ID }  
  { -f input-data-file | -i "input-data" }  
  [--help]
```

DESCRIPTION

migrpar performs partition migration operations.

Partition migration operations can be performed for AIX, Linux, or IBM i partitions only.

OPTIONS

-o The operation to perform. Valid values are **m** to validate then migrate one or more partitions if validation succeeds, **r** to recover from a failed partition migration, **s** to stop a partition migration, **v** to validate one or more partition migrations, and **set** to set attributes related to partition migration operations.

A recover operation should be issued, if possible, on the management console that is managing the source managed system.

A stop operation must be issued on the management console that is managing the source managed system.

-m The name of the source managed system for the partition migration operation. The name may either be the user-defined name for the managed system, or be in the form *ttt-mmm*sssss*, where *ttt* is the machine type, *mmm* is the model, and *sssss* is the serial number of the managed system. The *ttt-mmm*sssss* form must be used if there are multiple managed systems with the same user-defined name.

If this option is not specified for a system set operation, then the set operation will be performed for all systems that are currently managed by this Hardware Management Console (HMC).

-t The name of the target, or destination, managed system for the partition migration operation. The name may either be the user-defined name for the managed system, or be in the form *tttt-*mmm***ssssss**, where *tttt* is the machine type, *mmm* is the model, and *ssssss* is the serial number of the managed system. The *tttt-*mmm***ssssss** form must be used if there are multiple managed systems with the same user-defined name.

This option is required when migrating partitions or validating partition migrations. This option is not valid when performing any other operation.

--ip If the destination managed system is not managed by the same management console that is managing the source managed system, then use this option to specify the IP address or host name of the management console that is managing the destination managed system.

SSH must be enabled on both management consoles. Also, you must run the **mkauthkeys** command once to set up SSH key authentication.

This option is required when migrating partitions, validating partition migrations, or recovering from a failed partition migration if the destination managed system is not managed by the same management console that is managing the source managed system. This option is not valid when performing any other operation.

-u If the destination managed system is not managed by the same management console that is managing the source managed system, then use this option to specify the user ID to use on the management console that is managing the destination managed system. If you do not specify the user ID, then the user ID of the user that issues this command will be used.

-p The name of the partition for which to perform the operation.

One or more partition names can be specified When migrating partitions or validating partition migrations. When performing any other operation, only a single partition name can be specified.

Multiple partition names must be comma separated. The maximum number of partition names that can be specified is displayed by the **lsiparmigr -r mc -F max_lpars_per_migrpar_cmd** command.

The **-p**, **--id**, and **--all** options are mutually exclusive.

--id The ID of the partition for which to perform the operation.

One or more partition IDs can be specified When migrating partitions or validating partition migrations. When performing any other operation, only a single partition ID can be specified.

Multiple partition IDs must be comma separated. The maximum number of partition IDs that can be specified is displayed by the **lsiparmigr -r mc -F max_lpars_per_migrpar_cmd** command.

The **--id**, **-p**, and **--all** options are mutually exclusive.

--all Specify this option to perform the operation for all AIX, Linux, and IBM i partitions.

This option is only valid when migrating partitions or stopping partition migration operations.

The **--all**, **-p**, and **--id** options are mutually exclusive.

-n The name of the partition profile to create for the migrated partition on the destination managed system. If this option is omitted when migrating a partition, then the last activated profile for the partition will be replaced with the current partition configuration on the destination managed system.

This option is only valid when migrating a single partition or validating a single partition migration. When migrating more than one partition or validating more than one partition migration, the **multiple_profile_names** attribute must be specified instead.

--redundantvios,--redundantpgvios

When validating or migrating partitions that use shared memory, use this option to specify whether the partitions are to be configured to use redundant paging VIOS partitions on the destination managed system.

When validating or migrating suspended partitions, use this option to specify whether the partitions are to be configured to use redundant VIOS partitions on the destination managed system for accessing the devices that contain the data needed to resume the suspended partitions.

When validating or migrating remote restart capable partitions, use this option to specify whether the partitions are to be configured to use redundant VIOS partitions on the destination managed system for accessing the devices that contain the data needed to perform remote restarts of the partitions.

Valid values are **0** if the partitions are to be configured to not use redundant VIOS partitions, **1** if the partitions are to be configured to use redundant VIOS partitions, or **2** if the partitions are to be configured to use redundant VIOS partitions if possible. If this option is not specified, then each partition will be configured to use the same VIOS redundancy configuration on the destination managed system that the partition is currently using on the source managed system.

This option is only valid when validating or migrating partitions that use shared memory, partitions that are suspended, or partitions that are remote restart capable. This option is not valid when performing any other operation.

The **--redundantvios** and **--redundantpgvios** options are equivalent. The **--redundantpgvios** option is deprecated.

--mpio When validating or migrating partitions, use this option to specify whether the management console is required to maintain an equivalent multipath I/O (MPIO) configuration of each partition's virtual SCSI and virtual fibre channel adapters on the destination managed system. Valid values are **1** if the management console is required to maintain an equivalent MPIO configuration, or **2** if the management console is not required to maintain an equivalent MPIO configuration, but should do so if possible. If this option is not specified, then this option defaults to the value **1**.

If this option has a value of **1** and the management console cannot guarantee that an equivalent MPIO configuration of all of a partition's virtual SCSI and virtual fibre channel adapters can be maintained on the destination managed system, then an error will occur and the operation will fail for that partition.

This option is only valid when validating or migrating partitions. This option is not valid when performing any other operation.

--vlanbridge

When validating or migrating partitions, use this option to specify whether each partition's virtual ethernet adapters are required to be configured so that they are bridged on the same VLAN to an external network on the destination managed system. Valid values are **1** if virtual ethernet adapters are required to be bridged, or **2** if virtual ethernet adapters are not required to be bridged, but should be bridged if possible. If this option is not specified, then this option defaults to the value **1**.

If this option has a value of **1** and the management console cannot guarantee that a partition's virtual ethernet adapters can be bridged, then an error will occur and the operation will fail for that

partition.

A value of **2** cannot be specified with this option when a value of **1** is specified with the **--vsi** option.

This option is only valid when validating or migrating partitions. This option is not valid when performing any other operation.

--vsi When validating or migrating partitions, use this option to specify whether each partition's virtual ethernet adapter Virtual Station Interface (VSI) profiles are required to be configured on the destination managed system. Valid values are **1** if VSI profiles are required to be configured, or **2** if VSI profiles are not required to be configured, but should be configured if possible. If this option is not specified, then this option defaults to the value of the **--vlanbridge** option.

If this option has a value of **1** and the management console cannot configure a partition's VSI profiles on the destination managed system, an error will occur and the operation will fail for that partition.

A value of **1** cannot be specified with this option when a value of **2** is specified with the **--vlanbridge** option.

This option is only valid when validating or migrating partitions. This option is not valid when performing any other operation.

--uuid Use this option to allow partitions to be migrated to a destination managed system that is managed by a management console at Version 7 Release 7.7 or earlier. These releases do not support partition universally unique identifiers (UUIDs).

Normally, a partition's UUID is transferred to the destination managed system during a migration operation. However, older management consoles do not support partition UUIDs and the migration operation will fail. Specifying a value of **2** with this option will allow the migration operation to proceed, and the UUIDs of all of the migrating partitions will be lost. Note that the loss of a partition's UUID can cause higher level management software to lose the ability to track the location of the partition.

Valid values for this option are **1** if partition UUIDs are required to be migrated, or **2** if partition UUIDs are not required to be migrated, but should be migrated if possible. If this option is not specified, it defaults to the value **1**.

This option is only valid when validating or migrating partitions. This option is not valid when performing any other operation.

--protectstorage

Use this option to allow suspended partitions to be migrated.

After a suspended partition has been migrated to another managed system, it is exposed to accidental reassignment of its virtual storage devices while it remains suspended. Since there is no way to prevent this exposure, it is recommended that a suspended partition be resumed before being migrated.

Valid values for this option are **1** if each suspended partition's virtual storage devices are required to be protected, or **2** if each suspended partition's virtual storage devices are not required to be protected. If this option is not specified, then this option defaults to the value **1**. To migrate suspended partitions, you must specify a value of **2** and assume responsibility for ensuring the integrity of each suspended partition's virtual storage devices.

This option is only valid when validating or migrating partitions. This option is not valid when performing any other operation.

--requirerr

When validating or migrating partitions, use this option to specify whether each partition is required to be capable of simplified remote restart on the destination managed system. Valid values are **1** if each partition is required to be capable of simplified remote restart on the destination managed system, or **2** if each partition is not required to be capable of simplified remote restart on the destination managed system but should be if possible.

If this option has a value of **1** and the simplified remote restart capability of a migrating partition cannot be enabled on the destination managed system, an error will occur and the operation will fail for that partition. If this option has a value of **2** and the simplified remote restart capability of a migrating partition cannot be enabled on the destination managed system, the operation will not fail for that partition, even if that partition was capable of simplified remote restart on the source managed system. If this option is not specified, each partition's simplified remote restart capability on the source managed system will be maintained on the destination managed system. If that is not possible for a partition, an error will occur and the operation will fail for that partition.

This option is only valid when validating or migrating partitions. This option is not valid when performing any other operation.

--usecrrdata

Use this option to allow partitions to be migrated when Virtual I/O Servers that are hosting their virtual SCSI or virtual fibre channel adapters are shutdown and the storage adapter data that has been previously collected for the partitions is stale. The partitions will be migrated using the stale data.

This option is only valid when validating or migrating partitions. This option is not valid when performing any other operation.

-f The name of the file containing the input data for this command. The input data consists of attribute name/value pairs, which are in comma separated value (CSV) format.

The format of the input data is as follows:

attribute-name=value,attribute-name=value,...

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

"attribute-name=value,value,...",...

When a list of values is specified, the attribute name/value pair must be enclosed in double quotes. Depending on the shell being used, nested double quote characters may need to be preceded by an escape character, which is usually a `'\'` character.

Attributes that can be specified when validating or migrating a single partition:

dest_lpar_id

The partition ID to use on the destination managed system

virtual_fc_mappings

Comma separated list of virtual fibre channel adapter mappings, with each mapping having the following format:

virtual-slot-number/vios-lpar-name/vios-lpar-ID

`[[vios-virtual-slot-number][vios-fc-port-name]]`

The first 2 '/' characters must be present. The third '/' character is optional, but it must be present if *vios-virtual-slot-number* or *vios-fc-port-name* is specified. The last '/' character is optional, but it must be present if *vios-fc-port-name* is specified.

Optional values may be omitted. Optional values are *vios-lpar-name* or *vios-lpar-ID* (one of those values is required, but not both), *vios-virtual-slot-number*, and *vios-fc-port-name*.

For example:

4//1/14/fcs0 specifies a mapping of the virtual fibre channel client adapter with slot number **4** to the virtual fibre channel server adapter with slot number **14** in the VIOS partition with ID **1** on the destination managed system. In addition, the mapping specifies to use physical fibre channel port **fcs0**.

virtual_scsi_mappings

Comma separated list of virtual SCSI adapter mappings, with each mapping having the following format:

`virtual-slot-number/vios-lpar-name/vios-lpar-ID
[vios-virtual-slot-number]`

The first 2 '/' characters must be present. The last '/' character is optional, but it must be present if *vios-virtual-slot-number* is specified. Optional values may be omitted. Optional values are *vios-lpar-name* or *vios-lpar-ID* (one of those values is required, but not both), and *vios-virtual-slot-number*.

For example:

12/vios1//16 specifies a mapping of the virtual SCSI adapter with slot number **12** to slot number **16** on the VIOS partition **vios1** on the destination managed system.

vnic_mappings

Comma separated list of virtual NIC mappings, with each mapping having the following format:

`virtual-slot-number/ded[vios-lpar-name]/
[vios-lpar-ID]/[sriov-adapter-ID]/
[sriov-physical-port-ID]/[capacity]`

All 6 '/' characters must be present. Optional values may be omitted.

For example:

5/ded/vios1//1/2/4.5 specifies a mapping of the

virtual NIC in slot number **5** to a virtual NIC on the destination managed system with a backing device hosted by VIOS partition **vios1**. In addition, the mapping specifies to use a logical port on physical port **2** of SR-IOV adapter **1** and a logical port capacity of **4.5%** for the virtual NIC backing device.

vswitch_mappings

Comma separated list of virtual switch mappings, with each mapping having the following format:

vlan-ID/source-vswitch-name/dest-vswitch-name

For example:

6/srcSwitch/dstSwitch specifies a mapping of VLAN ID **6** on virtual switch **srcSwitch** to the virtual switch **dstSwitch** on the destination managed system.

source_msp_name

source_msp_id

source_msp_ipaddr

dest_msp_name

dest_msp_id

dest_msp_ipaddr

shared_proc_pool_name | **shared_proc_pool_id**

primary_rs_vios_name | **primary_rs_vios_id**

primary_paging_vios_name | **primary_paging_vios_id**

These attributes are deprecated. Use the

primary_rs_vios_name | **primary_rs_vios_id**

attributes instead.

concurr_migration_perf_level

Valid values are **1 - 5** and **1r - 5r**. Specify a value of **1r - 5r** if the concurrency level is required and the migration operation should fail if the resources required for the concurrency level are not available.

inactive_prof_policy

Valid values are:

prof - use the last activated profile for the partition when validating or migrating an inactive partition

config - use the current configuration of the partition when validating or migrating an inactive partition

Attributes that can be specified when validating or migrating one or more partitions if the **--all** option is not specified. When the **--all** option is specified, only the **source_msps** and **dest_msps** attributes are allowed. For each of the following attributes, *lpar-name* and *lpar-ID* is the name or ID of the migrating partition. If the **-p** option is used to specify the migrating partitions, then *lpar-name* is required and *lpar-ID* can be omitted. If the **--id** option is used to specify the migrating partitions, then *lpar-ID* is required and *lpar-name* can be omitted:

multiple_dest_lpar_ids

Comma separated list of partition IDs to use on the destination managed system. The format is:

lpar-name/lpar-ID/dest-lpar-ID

For example:

lp16//16 specifies that migrating partition **lp16** is to use partition ID **16** on the destination managed system.

multiple_virtual_fc_mappings

Comma separated list of virtual fibre channel adapter mappings, with each mapping having the following format:

```
lpar-namelpar-ID/virtual-slot-number/vios-lpar-name  
vios-lpar-ID/vios-virtual-slot-number[/  
vios-fc-port-name]
```

See the **virtual_fc_mappings** attribute format for more information.

multiple_virtual_scsi_mappings

Comma separated list of virtual SCSI adapter mappings, with each mapping having the following format:

```
lpar-namelpar-ID/virtual-slot-number/vios-lpar-name  
vios-lpar-ID/vios-virtual-slot-number
```

See the **virtual_scsi_mappings** attribute format for more information.

multiple_vnic_mappings

Comma separated list of virtual NIC mappings, with each mapping having the following format:

```
lpar-namelpar-ID/virtual-slot-number/ded/  
vios-lpar-name[/vios-lpar-ID]/  
sriov-adapter-ID[/sriov-physical-port-ID]/  
capacity
```

See the **vnic_mappings** attribute format for more information.

multiple_vswitch_mappings

Comma separated list of virtual switch mappings, with each mapping having the following format:

```
lpar-namelpar-ID/vlan-ID/source-vswitch-name  
dest-vswitch-name
```

See the **vswitch_mappings** attribute format for more information.

multiple_shared_proc_pool_names

Comma separated list of the names of the shared processor pools to use on the destination managed system. The format is:

```
lpar-namelpar-ID/shared-proc-pool-name
```

For example:

lp3//sp6 specifies that the migrating partition **lp3** is to use shared processor pool **sp6** on the destination managed system.

multiple_shared_proc_pool_ids

Comma separated list of the IDs of the shared processor pools to use on the destination managed system. The format is:

lpar-name||lpar-ID/shared-proc-pool-ID

For example:

/3/6 specifies that the migrating partition with ID **3** is to use the shared processor pool with ID **6** on the destination managed system.

multiple_primary_rs_vios_names

Comma separated list of the names of the VIOS partitions to use on the destination managed system as the primary VIOS partitions for accessing the reserved storage devices of the migrating partitions. The format is:

lpar-name||lpar-ID/vios-lpar-name

multiple_primary_rs_vios_ids

Comma separated list of the IDs of the VIOS partitions to use on the destination managed system as the primary VIOS partitions for accessing the reserved storage devices of the migrating partitions. The format is:

lpar-name||lpar-ID/vios-lpar-ID

multiple_profile_names

Comma separated list of names of the partition profiles to create for the migrating partitions on the destination managed system. The format is:

lpar-name||lpar-ID/profile-name

For example:

lpar1//migprof specifies to create the partition profile **migprof** for migrating partition **lpar1** on the destination managed system.

multiple_concurr_migration_perf_levels

Comma separated list of the concurrency levels to use. The format is:

lpar-name||lpar-ID/concurr-migration-perf-level

See the **concurr_migration_perf_level** attribute for more information.

multiple_inactive_prof_policies

Comma separated list of the profile policies to use for inactive partitions. The format is:

lpar-name||lpar-ID/inactive-prof-policy

See the **inactive_prof_policy** attribute for more information.

source_msps

Comma separated list of source MSPs to use for the operation. The format is:

MSP-name/MSP-ID/[MSP-IP-address]

Both '/' characters must be present, but optional values may be omitted. Optional values are *MSP-name* or *MSP-ID* (one of those values is required, but not both), and *MSP-IP-address*.

For example:

msp1//9.1.2.3 specifies to use IP address **9.1.2.3** on source MSP **msp1** for the operation.

dest_msps

Comma separated list of destination MSPs to use for the operation. The format is:

MSP-name/MSP-ID/[MSP-IP-address]

Both '/' characters must be present, but optional values may be omitted. Optional values are *MSP-name* or *MSP-ID* (one of those values is required, but not both), and *MSP-IP-address*.

For example:

msp1// specifies to use any IP address on destination MSP **msp1** for the operation.

Attributes that can be specified for a system set operation:

inactive_prof_policy

Sets the default profile policy to use when validating or migrating inactive partitions when the system is the source managed system for the operation. You can override this value for a partition by specifying it on the validation or migration operation.

Valid values are:

prof - use the last activated profile for partitions when validating or migrating inactive partitions

config - use the current configuration of partitions when validating or migrating inactive partitions

allow_inactive_source_storage_vios

Valid values are:

0 - disable virtual SCSI and virtual fibre channel adapter data collection for all AIX, Linux, and IBM i partitions in the managed system.

1 - enable virtual SCSI and virtual fibre channel adapter data collection for all AIX, Linux, and IBM i partitions in the managed system.

The data is collected from the VIOS partitions hosting the client adapters to allow the

partitions to be migrated when VIOS partitions on the source managed system are shutdown.

Attributes that can be specified for an MSP set operation:

num_active_migrations_configured

concurr_migration_perf_level

Sets the default concurrency level to use on the MSP. You can override this value for a partition validation or migration operation by specifying it on the operation.

Valid values are **1 - 5**

Input data for this command can be specified with this option or the **-i** option. The **-f** and the **-i** options are mutually exclusive.

Input data can be specified when migrating partitions, validating partition migrations, or performing a set operation. This option is not valid when performing any other operation.

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

Input data for this command can be specified with this option or the **-f** option. The **-i** and the **-f** options are mutually exclusive.

Input data can be specified when migrating partitions, validating partition migrations, or performing a set operation. This option is not valid when performing any other operation.

- r** The type of resource for which to set attributes related to partition migration operations. Valid values are **sys** for managed system, and **lpar** for mover service partition.
- w** The maximum time, in minutes, to wait for operating system commands issued by the management console to the partitions to be migrated to complete. If an operating system command does not complete within the time specified, the partition migration operation for that partition will be stopped.

wait-time must be a whole number. If wait-time is 0, the partition migration operation will not be timed out.

If this option is not specified, a default value of 3 minutes is used.

This option is valid when migrating partitions or validating partition migrations. This option is not valid when performing any other operation.

- d** The level of detail requested from operating system commands issued by the management console to all partitions participating in migrations. Valid values are **1** (lowest) through **5** (highest).

If this option is not specified, a default value of 1 is used.

This option is valid when migrating partitions or validating partition migrations. This option is not valid when performing any other operation.

- force** This option allows you to force a recover operation to proceed when errors are encountered.

This option is only valid when recovering from a failed partition migration.

- v** Specify this option to enable verbose mode for the partition migration operation. When verbose mode is enabled, detail messages and warning messages are displayed for successful partition

migrations. Detail messages and warning messages are displayed for partition migrations that fail, regardless of whether this option is specified.

When the **--nodetails** option is also specified, no detail messages are displayed.

--nodetails

Specify this option to suppress the display of all detail messages.

--help Display the help text for this command and exit.

EXIT STATUS

This command has the following return codes:

0	Success
100	Partial success

When an operation is performed for multiple partitions, this value is returned when the operation for at least one partition succeeds, and the operation for at least one partition fails.

Any other value means the command failed.

EXAMPLES

In the following examples nested double quote characters are preceded by an escape character (`'\'`). The escape characters are required to run the commands locally on an HMC.

Validate the operation to migrate partition **mylpar** from managed system **system1** to managed system **system2**:

```
migrpar -o v -m system1 -t system2 -p mylpar
-i "\"virtual_scsi_mappings=12/vios1,13/vios2/\",dest_msp_name=vios1,
source_msp_name=vios"
```

Validate the operation to migrate the partitions with IDs **3**, **6**, and **7** from managed system **system1** to managed system **system2**:

```
migrpar -o v -m system1 -t system2 --id 3,6,7
-i "\"multiple_virtual_scsi_mappings=/3/12/vios1//32,/3/13/vios2//33,
/7/12/vios1//72\",source_msp=vios1//9.3.4.5,dest_msp=vios1//"
```

Migrate the partition with ID **5** from managed system **system1** to managed system **system2**:

```
migrpar -o m -m system1 -t system2 --id 5 -n migprof
-i "virtual_fc_mappings=5//1/18/fcs0,dest_msp_id=1,source_msp_id=3"
```

Migrate the partition **smp** from managed system **system1** to managed system **system2**, configure the partition to use redundant paging VIOS partitions on **system2**, and use **vios1** as the primary paging VIOS for the partition:

```
migrpar -o m -m system1 -t system2 -p smp --redundantvios 1
-i "primary_rs_vios_name=vios1"
```

Migrate the partitions **lp4** and **lp5** from managed system **system1** to managed system **system2**:

```
migrpar -o m -m system1 -t system2 -p lp4,lp5 --mpio 2
-i "\"multiple_dest_lpar_ids=lp4//21,lp5//22\",
multiple_shared_proc_pool_names=lp5//sp1"
```

Migrate all AIX, Linux, and IBM i partitions from managed system **m11** to managed system **m13**:

```
migrpar -o m -m m11 -t m13 --all -i  
"\"source_msps=m11-vios1//,m11-vios2/\"  
\"dest_msps=m13-vios1//,m13-vios2/\""
```

Stop the partition migration that is just starting for partition **mylpar**:

```
migrpar -o s -m system1 -p mylpar
```

Recover the failed partition migration of partition **mylpar**:

```
migrpar -o r -m system1 -p mylpar
```

Migrate the partition **aix1** from managed system **system1** to managed system **system2**, when **system2** is managed by the HMC with the host name **hmc2**:

```
mkauthkeys --ip hmc2 -u hmc2user
```

```
migrpar -o m -m system1 -t system2 -p aix1 --ip hmc2 -u hmc2user
```

Set the inactive profile migration policy for all systems currently managed by this HMC:

```
migrpar -o set -r sys -i "inactive_prof_policy=prof"
```

Set partition migration attributes for MSP **vios2**:

```
migrpar -o set -r lpar -m sys1 -p vios2  
-i "num_active_migrations_configured=10,  
concurr_migration_perf_level=3"
```

ENVIRONMENT

None

BUGS

None

AUTHOR

IBM Austin

SEE ALSO

lsiparmigr, mkauthkeys

NAME

mkaccfg - create access control object

SYNOPSIS

```
mkaccfg -t {resourcerole | taskrole}
{-f configuration-file | -i "configuration-data"}
[--help]
```

DESCRIPTION

mkaccfg creates a new access control role.

OPTIONS

- t** The type of access control role to create. Valid values are **resourcerole** for managed resource role and **taskrole** for task role.
- f** The name of the file containing the configuration data needed to create the access control role. The configuration data consists of attribute name/value pairs, which are in comma separated value (CSV) format. These attribute name/value pairs form a configuration record. A line feed marks the end of a configuration record. There can only be one configuration record in the file.

The format of a configuration record is as follows:

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

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

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

When a list of values is specified, the attribute name/value pair must be enclosed in double quotes. Depending on the shell being used, nested double quote characters may need to be preceded by an escape character, which is usually a `'\'` character.

Attribute names for managed resource roles (all attributes are required):

- name**
name of the managed resource role to create
- resources**
comma separated list of managed resource objects

Attribute names for task roles (all attributes are required):

- name**
name of the task role to create
- resources**
comma separated list of tasks
- parent**
name of the parent task role on which to base this task role. Valid values are **hmcsuperadmin**, **hmcoperator**, **hmcviewer**, **hmcpe**, **hmcservicerep**, or a user-defined task role.

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

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

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

--help Display the help text for this command and exit.

EXAMPLES

Create a managed resource role using the configuration data in the file **/tmp/mr1file**:

```
mkaccfg -t resourcerole -f /tmp/mr1file
```

Create a task role named **tr1**:

```
mkaccfg -t taskrole -i "name=tr1,parent=hmcsuperadmin,  
"resources=cec:ChangeCoD+ListCoDInformation+  
ListCECProperty,lpar:ChangeLPARProperty+ListLPARProperty+  
CreateProfile""
```

ENVIRONMENT

None

BUGS

None

AUTHOR

IBM Austin

SEE ALSO

chaccfg, **lsaccfg**, **rmaccfg**

NAME

mkauthkeys - manage SSH authentication keys

SYNOPSIS

To add an SSH key as an authorized key:

```
mkauthkeys {-a | --add} "string" [--help]
```

To remove an SSH key from the authorized keys:

```
mkauthkeys {-r | --remove} {"string" | -u user-ID} [--help]
```

To set up SSH key authentication for partition mobility and partition remote restart operations between this Hardware Management Console (HMC) and another management console:

```
mkauthkeys [-g] --ip IP-address -u user-ID [--passwd password]  
[-t {rsa | dsa}] [--help]
```

To test SSH key authentication for partition mobility and partition remote restart operations from this HMC to another management console:

```
mkauthkeys --test --ip IP-address [-u user-ID] [--help]
```

DESCRIPTION

mkauthkeys manages SSH authentication keys.

mkauthkeys can be used to add or remove SSH keys as authorized keys for the user which issues this command.

mkauthkeys can also be used to set up SSH key authentication for partition mobility and partition remote restart operations between this HMC and another management console.

OPTIONS

-a,--add

Adds the specified SSH key *string* as an authorized key for the user which issues this command.

The *string* is added to the user's **\$HOME/.ssh/authorized_keys2** file.

-r,--remove

Removes the specified SSH key *string* from the authorized keys for the user which issues this command, or removes all of the SSH keys generated for the user specified in *string* from the authorized keys for the user which issues this command.

If this option and the **-u** option is specified, then all of the authorized keys for *user-ID* will be removed.

The keys are removed from the user's **\$HOME/.ssh/authorized_keys2** file.

string To add or remove an SSH key as an authorized key for the user which issues this command, specify the SSH key to add or remove with this option.

To remove all of the SSH keys generated for a user from the authorized keys for the user which issues this command, specify the user for which the SSH keys to remove were generated.

--ip

The IP address or host name of the remote management console with which to set up partition mobility and partition remote restart operations.

Public and private SSH keys will be generated for the ccfw user on this HMC. The keys will be stored on this HMC in the **\$HOME/.ssh/ccfw** directory for the user which issues this command. Then, if SSH keys have not already been set up on the remote management console for the remote user, then the saved public key will be sent to the remote management console and stored in the remote user's **\$HOME/.ssh/authorized_keys2** file.

-u When removing all authorized SSH keys for a user, use this option to specify the user ID.

When setting up SSH key authentication for partition mobility and partition remote restart operations between this HMC and another management console, use this option to specify the user ID to use on the remote management console.

--passwd

When setting up SSH key authentication for partition mobility and partition remote restart operations between this HMC and another management console, use this option to specify the password for the user ID to use on the remote management console. If this option is omitted, you will be prompted to enter the password.

-t The type of SSH keys to generate when setting up SSH key authentication for partition mobility and partition remote restart operations between this HMC and another management console. Valid values are **rsa** for RSA or **dsa** for DSA. If this option is not specified, then RSA keys will be generated.

-g Specify this option when setting up SSH key authentication for partition mobility and partition remote restart operations between this HMC and another management console to allow partition mobility and partition remote restart operations to be performed from the remote management console to this HMC. If this option is not specified, then partition mobility and partition remote restart operations can only be performed from this HMC to the remote management console.

When this option is specified, public and private SSH keys will also be generated for the ccfw user on the remote management console. The keys will be stored on the remote management console in the remote user's **\$HOME/.ssh/ccfw** directory. Then, if SSH keys have not already been set up on this HMC for the user that issues this command, then the saved public key will be sent to this HMC and stored in this user's **\$HOME/.ssh/authorized_keys2** file.

--test Tests SSH key authentication for partition mobility and partition remote restart operations from this HMC to the management console specified with the **--ip** option.

If the return code from this command is 0, then SSH key authentication is working properly.

--help Display the help text for this command and exit.

EXAMPLES

To add the SSH key **ssh-rsa AAAB3NzaC1yc2EAAA joe@somehost**, which was generated for the user **joe@somehost**, as an authorized key for your user ID:

```
mkauthkeys -a "ssh-rsa AAAB3NzaC1yc2EAAA joe@somehost"
```

To remove the SSH key **ssh-rsa AAAB3NzaC1yc2EAAA joe@somehost**, which was generated for user **joe@somehost**, from the authorized keys for your user ID:

```
mkauthkeys -r "ssh-rsa AAAB3NzaC1yc2EAAA joe@somehost"
```

To remove all of the SSH keys generated for the user **joe@somehost** from the authorized keys for your user ID:

```
mkauthkeys -r joe@somehost
```

To remove all of the authorized SSH keys for user **hmcUser1**:

```
mkauthkeys -r -u hmcUser1
```

To set up partition mobility and partition remote restart operations from your user ID on this HMC to the HMC with host name **hmc2** for the user **hmc2user** on HMC **hmc2** (you will be prompted for **hmc2user**'s

password):

```
mkauthkeys --ip hmc2 -u hmc2user
```

Check if partition mobility and partition remote restart operations from your user ID on this HMC to the HMC with host name **hmc2** for the user **hmc2user** on HMC **hmc2** have been set up:

```
mkauthkeys --ip hmc2 -u hmc2user --test
```

ENVIRONMENT

None

BUGS

None

AUTHOR

IBM Austin

NAME

mkcodpool - create a CoD Power enterprise pool

SYNOPSIS

```
mkcodpool -p pool-name -f file  
[-a "attributes"] [-v] [--help]
```

DESCRIPTION

mkcodpool creates a Capacity on Demand (CoD) Power enterprise pool.

OPTIONS

- p** A name for the Power enterprise pool.
- f** The name of the XML file that contains the Power enterprise pool configuration data. If *file* is not fully qualified, *file* must exist in the user's home directory on the Hardware Management Console (HMC).

If *file* exists on removable media, the media must be present in the removable media device and the device must be mounted with the **mount** command before this command is issued. The **lsmediadev** command can be used to display all of the removable media devices on the HMC.

- a** The Power enterprise pool attributes. The attribute data consists of attribute name/value pairs, which are in comma separated value (CSV) format.

The format of the attribute data is as follows:

```
"attribute-name=value,value,...",...
```

When a list of values is specified, the attribute name/value pair must be enclosed in double quotes. Depending on the shell being used, nested double quote characters may need to be preceded by an escape character, which is usually a `'\'` character.

Valid attribute names:

mcs

Comma separated list of managing HMCs for the pool, with each managing HMC having the following format:

```
host-name-or-IP-address/user-ID/[password]
```

Both `'/'` characters must be present, but *password* may be omitted. If *password* is omitted, you will be prompted to enter the password.

Note: do not include the HMC which is running this command in the list.

For example:

hmc1/user1/ prompts the user for **user1**'s password then uses that password and the user ID **user1** to log into the HMC with host name **hmc1** to add that HMC to the pool.

- v** Specify this option to enable verbose mode. When verbose mode is enabled, warning messages and informational messages are displayed for successful create pool operations. Warning messages and informational messages are displayed for create pool operations that partially succeed or fail, regardless of whether this option is specified.
- help** Display the help text for this command and exit.

EXIT STATUS

This command has the following return codes:

0 Success
100 Partial success

This value is returned when at least one operation in the configuration file succeeded, and at least one failed.

Any other value means the command failed.

EXAMPLES

Create the Power enterprise pool **pool1** using the configuration file **pool0195.xml** on a USB flash memory device (the USB flash memory device must already be connected to the HMC) and add HMC **hmc1** to the pool:

```
lsmediadev (to obtain mount points)
```

```
mount /media/sdb1
```

```
mkcodpool -p pool1 -f /media/sdb1/pool0195.xml  
-a "mcs=hmc1/myuserID/mypassword" -v
```

Create the Power enterprise pool **myPool** using the configuration file **pool.xml** in the user's home directory on the HMC and add the HMCs **hmc1** and **hmc2** to the pool (you will be prompted to enter the password for each HMC):

```
mkcodpool -p myPool -f pool.xml -a ""mcs=hmc1/myuserID/,  
hmc2/myuserID/"" -v
```

ENVIRONMENT

None

BUGS

None

AUTHOR

IBM Austin

SEE ALSO

chcodpool, **lscodpool**, **lsmediadev**

NAME

mkhmcusr - create a Hardware Management Console user

SYNOPSIS

```
mkhmcusr -u user-name -a task-role [-d description]  
[--passwd password] [-M number-of-days]  
[--auth {local | kerberos | ldap}]  
[--remoteuser remote-user-name] [--help]
```

or

```
mkhmcusr {-f input-data-file | -i "input-data"} [--help]
```

DESCRIPTION

mkhmcusr creates a Hardware Management Console (HMC) user.

OPTIONS

-u The user name of the HMC user to create. The user name cannot be longer than 32 characters, and it must begin with a letter.

You can either use this option, or use the **name** attribute with the **-f** or **-i** option, to specify the user name. The **-u**, **-f**, and **-i** options are mutually exclusive.

-a The access control task role for this user. Valid values are **hmcsuperadmin**, **hmcoperator**, **hmcviewer**, **hmcpe**, **hmcservicerep**, or a user-defined task role.

You can either use this option, or use the **taskrole** attribute with the **-f** or **-i** option, to specify the access control task role. The **-a**, **-f**, and **-i** options are mutually exclusive.

-d The description for this user. *description* can be any string.

You can either use this option, or use the **description** attribute with the **-f** or **-i** option, to specify the description. If a description is not specified, then the description will be set to "HMC User". The **-d**, **-f**, and **-i** options are mutually exclusive.

--passwd

The password for this user. The password must be at least 7 characters in length.

A password cannot be specified when creating a remotely authenticated Kerberos or LDAP user.

You can either use this option, or use the **passwd** attribute with the **-f** or **-i** option, to specify the password. If this option is omitted or the **-f** or **-i** option is specified and the **passwd** attribute is omitted, you will be prompted to enter the password. The **--passwd**, **-f**, and **-i** options are mutually exclusive.

-M The number of days until the password for this user expires.

A password expiration cannot be specified when creating a remotely authenticated Kerberos or LDAP user.

You can either use this option, or use the **pwage** attribute with the **-f** or **-i** option, to specify the password expiration. If a password expiration is not specified, then this user's password will never expire. The **-M**, **-f**, and **-i** options are mutually exclusive.

--auth The authentication type for this user. Valid values are **local** for local authentication, **kerberos** for remote Kerberos authentication, and **ldap** for remote LDAP authentication.

You can either use this option, or use the **authentication_type** attribute with the **-f** or **-i** option, to specify the authentication type. If an authentication type is not specified, then the authentication

type for this user will be set to local authentication. The **--auth**, **-f**, and **-i** options are mutually exclusive.

--remoteuser

The remote user ID used for remote Kerberos authentication for this user. This is the user's Kerberos principal. The format of a typical Kerberos principal is *primary/instance@REALM*.

You can either use this option, or use the **remote_user_name** attribute with the **-f** or **-i** option, to specify the remote user ID. A remote user ID must be specified when creating a remotely authenticated Kerberos user. The **--remoteuser**, **-f**, and **-i** options are mutually exclusive.

-f The name of the file containing the input data for this command. The input data consists of attribute name/value pairs, which are in comma separated value (CSV) format.

The format of the input data is as follows:

attribute-name=value,attribute-name=value,...

Valid attribute names for this command:

name

taskrole

Valid values are **hmcsuperadmin**, **hmcoperator**, **hmcviewer**, **hmcpe**, **hmcserVICerEP**, or a user-defined task role

[resourcerole]

[description]

[passwd]

Local users only

[pwage]

Local users only
number of days

[min_pwage]

Local users only
number of days

[authentication_type]

Valid values are:

local - local authentication

kerberos - remote Kerberos authentication

ldap - remote LDAP authentication

[session_timeout]

number of minutes

[verify_timeout]

number of minutes

[idle_timeout]

number of minutes

[inactivity_expiration]

number of days

[remote_webui_access]

Valid values are:

0 - do not allow this user to log in remotely to the
HMC Web user interface

1 - allow this user to log in remotely to the
HMC Web user interface

[remote_ssh_access]

Valid values are:

0 - do not allow this user to log in remotely to the

HMC using SSH
1 - allow this user to log in remotely to the
HMC using SSH
[remote_user_name]
Kerberos users only

Input data for this command can be specified with this option, the **-i** option, or any of the other command options. The **-f** and the **-i** options are mutually exclusive, and they cannot be specified if any of the other command options are specified.

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

Input data for this command can be specified with this option, the **-f** option, or any of the other command options. The **-i** and the **-f** options are mutually exclusive, and they cannot be specified if any of the other command options are specified.

- help** Display the help text for this command and exit.

EXAMPLES

Create the user **sysadmin**:

```
mkhmcusr -u sysadmin -a hmcsuperadmin --passwd aielkw3j  
-M 180 -d "System Administrator"
```

or

```
mkhmcusr -i "name=sysadmin,taskrole=hmcsuperadmin,passwd=  
aielkw3j,pwage=180,description=System Administrator"
```

Create the user **myhmcuser** (the user's password must be entered when prompted):

```
mkhmcusr -u myhmcuser -a hmcviewer
```

or

```
mkhmcusr -i "name=myhmcuser,taskrole=hmcviewer"
```

Create a remotely authenticated Kerberos user **krbuser**:

```
mkhmcusr -u krbuser -a hmcoperator --auth kerberos  
--remoteuser krbuser/hmcoperator@EXAMPLE.COM
```

or

```
mkhmcusr -i "name=krbuser,taskrole=hmcoperator,  
authentication_type=kerberos,remote_user_name=  
krbuser/hmcoperator@EXAMPLE.COM"
```

Create a remotely authenticated LDAP user **ldapuser**:

```
mkhmcusr -u ldapuser -a hmcsuperadmin --auth ldap
```

or

```
mkhmcusr -i "name=ldapuser,taskrole=hmcsuperadmin,  
authentication_type=ldap"
```

ENVIRONMENT

None

BUGS

None

AUTHOR

IBM Austin

SEE ALSO

chhmcusr, lshmcusr, rmhmcusr, lsaccfg

NAME

mkprofdata - create profile data

SYNOPSIS

To recreate profile data:

```
mkprofdata -r sys -o recreate -m managed-system -s sys [-v]  
 [--help]
```

To convert profile data to XML:

```
mkprofdata -r sys -o createxml  
 {-m managed-system | -b binary-file}  
 -x XML-file [--help]
```

DESCRIPTION

mkprofdata recreates the profile data for *managed-system* based on the current configuration of *managed-system*. **This operation will overwrite the current profile data for the managed system so it should only be performed if the profile data for the managed system is corrupted.** This operation will create a single profile for each partition in the managed system. This profile will match the current configuration of the partition on the managed system. Any profile attributes not stored on the managed system, including the profile name, will be set to default values. Since no existing partition profiles can be recovered, when this operation completes only the newly created profile will exist for each partition along with the last valid configuration profile that the HMC automatically creates for every partition.

mkprofdata can also be used to convert the profile data for a managed system to XML format and save it to a file.

OPTIONS

- r** The type of resource for which to recreate or convert profile data. The only valid value is **sys** for managed system.
- o** The operation to perform. Valid values are **recreate** to recreate the profile data for a managed system based on the current configuration of the managed system, and **createxml** to convert the profile data for a managed system to XML format and save it to a file.
- m** The name of the managed system for which to perform the profile data operation. The name may either be the user-defined name for the managed system, or be in the form *tttt-mmm*sssssss*, where *tttt* is the machine type, *mmm* is the model, and *sssssss* is the serial number of the managed system. The *tttt-mmm*sssssss* form must be used if there are multiple managed systems with the same user-defined name.

If this option is specified when converting profile data to XML format, the current profile data on the managed system will be used. The **-m** and the **-b** options are mutually exclusive.

- s** The source for the profile data recreate operation. The only valid value is **sys** for managed system. The profile data for the managed system will be recreated based on the current configuration of the managed system.
- b** The name of the files containing the binary profile data to be converted to XML format. *binary-file* must be a fully qualified file name (it must include the directory name). Both files containing the binary profile data, *binary-file* and *binary-file.dir*, must be present for the operation to be performed.

If the files are on removable media, the media must be present in the removable media device and the device must be mounted with the **mount** command before this command is issued. The **lsmedev** command can be used to display all of the removable media devices on the management console.

The **-b** and the **-m** options are mutually exclusive.

-x The name of the files to which the profile data is to be written in XML format. If *XML-file* is not fully qualified, *XML-file* will be written to the user's home directory. Two files will be written: *XML-file.xml* and *XML-file_dir.xml*.

If the files are to be written to removable media, the media must be present in the removable media device and the device must be mounted with the **mount** command before this command is issued. The **lsmediadev** command can be used to display all of the removable media devices on the management console.

-v Specify this option to enable verbose mode for the profile data recreate operation. When verbose mode is enabled, progress messages are displayed.

--help Display the help text for this command and exit.

EXAMPLES

Recreate the profile data for the managed system **mySys**:

```
mkprofdata -r sys -o recreate -m mySys -s sys -v
```

Convert the current profile data on managed system **mySys** to XML format and save it in the files **mySysXML.xml** and **mySysXML_dir.xml** in the user's home directory:

```
mkprofdata -r sys -o createxml -m mySys -x mySysXML
```

Convert the profile data in the files **/home/hscpe/sysA** and **/home/hscpe/sysA.dir** to XML format and save it in the files **sysAXML.xml** and **sysAXML_dir.xml** in the user's home directory:

```
mkprofdata -r sys -o createxml -b /home/hscpe/sysA -x sysAXML
```

ENVIRONMENT

None

BUGS

None

AUTHOR

IBM Austin

NAME

mkpwdpolicy - create a password policy

SYNOPSIS

mkpwdpolicy {-f *input-data-file* | -i "*input-data*"} [--help]

DESCRIPTION

mkpwdpolicy creates a password policy for the Hardware Management Console (HMC).

mkpwdpolicy only creates a password policy. The **chpwdpolicy** command must be used to activate the password policy.

OPTIONS

-f The name of the file containing the data to create the password policy. The data consists of attribute name/value pairs, which are in comma separated value (CSV) format.

The format of the data is as follows:

attribute-name=value,attribute-name=value,...

Valid attribute names for this command:

name

The name of the password policy to create.

[description]

[min_pwage]

The number of days that must elapse before a password can be changed. Defaults to 1 day.

[pwage]

The number of days that can elapse before a password expires and must be changed. A value of **99999** indicates no password expiration. Defaults to 180 days.

[warn_pwage]

The number of days prior to password expiration when a warning message will begin to be displayed. Defaults to 7 days.

[min_length]

The minimum password length. Defaults to 8 characters.

[hist_size]

The number of times a password must be changed before a password can be reused. This value cannot exceed 50, and defaults to 10.

[min_digits]

The minimum number of digits that a password must contain. Defaults to 0.

[min_uppercase_chars]

The minimum number of uppercase characters that a password must contain. Defaults to 1.

[min_lowercase_chars]

The minimum number of lowercase characters that a password must contain. Defaults to 6.

[min_special_chars]

The minimum number of special characters that a password must contain. Special characters include symbols, punctuation, and white space characters.

Defaults to 0.

Brackets around an attribute name indicate that the attribute is optional.

Comments are allowed in the input file. Comments must begin with the '#' character.

The data to create the password policy is required to be specified with this option or the **-i** option. The **-f** and the **-i** options are mutually exclusive.

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

The data to create the password policy is required to be specified with this option or the **-f** option. The **-i** and the **-f** options are mutually exclusive.

- help** Display the help text for this command and exit.

EXAMPLES

Create a password policy:

```
mkpwdpolicy -i "name=xyzPolicy,description=Company xyz policy,  
pwage=90,min_digits=2,min_uppercase_chars=0,min_lowercase_chars=0"
```

ENVIRONMENT

None

BUGS

None

AUTHOR

IBM Austin

SEE ALSO

chpwdpolicy, **lspwdpolicy**, **mpwdpolicy**

NAME

mksyscfg - create system resources

SYNOPSIS

```
mksyscfg -r {lpar | prof | sysprof} -m managed-system
[-f configuration-file | -i "configuration-data"]
[-o save {-p partition-name | --id partition-ID}
-n profile-name [--force]]
[--help]
```

DESCRIPTION

mksyscfg creates partitions, partition profiles, or system profiles for the *managed-system*.

mksyscfg can also be used to save the current configuration of a partition to a partition profile.

USAGE IN A POWERVM MANAGEMENT MASTER MODE ENVIRONMENT

When the *managed-system* is in PowerVM management master mode, the **mksyscfg -r lpar** command may be used to create a partition, however the following attributes are not allowed: **profile_name** and **sync_curr_profile**. The partition will be configured immediately with the specified attributes. The Hardware Management Console (HMC) must be the current PowerVM management master for the *managed-system* when running the **mksyscfg -r lpar** command.

The **mksyscfg -r prof** and **mksyscfg -r sysprof** commands are not allowed.

OPTIONS

-r The type of system resources to create. Valid values are **lpar** for partitions, **prof** for partition profiles, and **sysprof** for system profiles.

When a partition is created, the default profile for the partition is also created.

-m The name of the managed system for which the system resources are to be created. The name may either be the user-defined name for the managed system, or be in the form *tttt-mmm*sssssss*, where *tttt* is the machine type, *mmm* is the model, and *sssssss* is the serial number of the managed system. The *tttt-mmm*sssssss* form must be used if there are multiple managed systems with the same user-defined name.

-f The name of the file containing the configuration data needed to create the system resources. The configuration data consists of attribute name/value pairs, which are in comma separated value (CSV) format. These attribute name/value pairs form a configuration record. A line feed marks the end of a configuration record. The file must contain one configuration record for each resource to be created, and each configuration record must be for the same resource type.

The format of a configuration record is as follows:

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

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

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

When a list of values is specified, the attribute name/value pair must be enclosed in double quotes. Depending on the shell being used, nested double quote characters may need to be preceded by an escape character, which is usually a ``\`` character.

Attribute names for partitions (see below for attribute names that are common to both partitions and partition profiles):

name

name of the partition to create

[lpar_id]**profile_name**

name of the default profile to create

lpar_env

Valid values are **aixlinux**, **os400**, or **vioserver**

[allow_perf_collection]

Valid values are:

0 - do not allow

1 - allow

This attribute replaces the **shared_proc_pool_util_auth** attribute. Setting this attribute also sets the **shared_proc_pool_util_auth** attribute to the same value.

[shared_proc_pool_util_auth]

Valid values are:

0 - do not allow authority

1 - allow authority

This attribute has been deprecated. Use the **allow_perf_collection** attribute instead. Setting this attribute also sets the **allow_perf_collection** attribute to the same value.

[lpar_avail_priority]

Valid values are **0** - **255**.

[msp]

Virtual I/O server only

Valid values are:

0 - the partition is not a mover service partition

1 - the partition is a mover service partition

[time_ref]

Valid values are:

0 - the partition is not a time reference partition

1 - the partition is a time reference partition

[suspend_capable]

AIX, Linux, and IBM i only

Valid values are:

0 - the partition cannot be suspended

1 - the partition can be suspended

[simplified_remote_restart_capable]

AIX, Linux, and IBM i only

Valid values are:

0 - remote restart without using a reserved storage device cannot be performed for the partition

1 - remote restart without using a reserved storage device can be performed for the partition

[remote_restart_capable]

AIX, Linux, and IBM i only

Valid values are:

- 0** - remote restart using a reserved storage device cannot be performed for the partition
 - 1** - remote restart using a reserved storage device can be performed for the partition
- [primary_rs_vios_name | primary_rs_vios_id]**
Remote restart capable partitions only
- [secondary_rs_vios_name | secondary_rs_vios_id]**
Remote restart capable partitions only
To set no secondary reserved storage VIOS specify **none** for **secondary_rs_vios_id**
- [rs_device_name]**
Remote restart capable partitions only
Name of the device on the primary reserved storage VIOS partition
- [vtpm_enabled]**
AIX, Linux, and VIOS only
Valid values are:
 - 0** - virtual Trusted Platform Module (vTPM) is not enabled for this partition
 - 1** - vTPM is enabled for this partition
- [os400_restricted_io_mode]**
IBM i only
Valid values are:
 - 0** - disable IBM i restricted I/O mode
 - 1** - enable IBM i restricted I/O mode
- [sync_curr_profile]**
Specifies whether this partition's current configuration is to be synchronized with its current active profile. Valid values are:
 - 0** - disable synchronization
 - 1** - enable synchronization
 - 2** - suspend synchronization until the next time a profile is activated or applied
- [hardware_mem_encryption]**
Valid values are:
 - 0** - disable hardware-accelerated encryption
 - 1** - enable hardware-accelerated encryption
- [hardware_mem_expansion]**
Valid values are:
 - 0** - disable hardware-accelerated Active Memory Expansion
 - 1** - enable hardware-accelerated Active Memory Expansion
- [powervm_mgmt_capable]**
Linux only
Valid values are:
 - 0** - do not allow this partition to provide PowerVM management functions
 - 1** - enable this partition to provide PowerVM management functions
- [migration_disabled]**
AIX, Linux, and IBM i only
Valid values are:
 - 0** - do not disable partition migration for this

partition
1 - disable partition migration for this
 partition

Attribute names for partition profiles (see below for attribute names that are common to both partition profiles and partitions):

name
 name of the partition profile to create
lpar_name | lpar_id
 name or ID of the partition for which
 to create the profile

Attribute names for both partitions and partition profiles:

[all_resources]
 Valid values are:
0 - do not use all the managed system
 resources
1 - use all the managed system resources
 (this option is not valid for IBM i
 partitions on IBM System p5 or
 eServer p5 servers)
min_mem
 megabytes
desired_mem
 megabytes
max_mem
 megabytes
[min_num_huge_pages]
 AIX and Linux only
[desired_num_huge_pages]
 AIX and Linux only
[max_num_huge_pages]
 AIX and Linux only
[mem_mode]
 Valid values are:
ded - dedicated memory
shared - shared memory
[desired_io_entitled_mem]
 Specify the number of megabytes or specify **auto** for
 automatic I/O entitled memory management mode
[mem_weight]
[primary_paging_vios_name | primary_paging_vios_id]
[secondary_paging_vios_name | secondary_paging_vios_id]
 To set no secondary paging VIOS specify **none** for
secondary_paging_vios_id
[mem_expansion]
 Valid values are:
0 - disable Active Memory Expansion
1.00-10.00 - expansion factor
[hpt_ratio]
 Ratio of hardware page table size to the maximum
 memory for the partition. Valid values are displayed
 by the **lshwres -r mem -m managed-system**
--level sys -F possible_hpt_ratios command.

[bsr_arrays]**[proc_mode]**

Valid values are:

ded - dedicated processors

shared - shared processors

[min_procs]**[desired_procs]****[max_procs]****[min_proc_units]****[desired_proc_units]****[max_proc_units]****[sharing_mode]**

Valid values for partitions using dedicated processors are:

keep_idle_procs - never share processors

share_idle_procs - share processors only when partition is inactive

share_idle_procs_active - share processors only when partition is active

share_idle_procs_always - always share processors

Valid values for partitions using shared processors are:

cap - capped

uncap - uncapped

[uncap_weight]**[shared_proc_pool_name | shared_proc_pool_id]****[affinity_group_id]**

Memory and processor affinity group in which the partition will participate. Valid values are

none and **1 - 255**.

[io_slots]

Comma separated list of I/O slots, with each I/O slot having the following format:

slot-DRC-index[slot-IO-pool-ID]/is-required

Both '/' characters must be present, but optional values may be omitted. Optional values are *slot-IO-pool-ID*.

Valid values for *is-required*:

0 - no

1 - yes

For example:

21030002/3/1 specifies an I/O slot with a DRC index of **21030002**, it is assigned to I/O pool **3**, and it is a required slot.

[lpar_io_pool_ids]

comma separated

[load_source_slot]

IBM i only

DRC index of I/O slot, DRC index of SR-IOV logical port, or virtual slot number

[alt_restart_device_slot]

IBM i only

DRC index of I/O slot, DRC index of HEA or SR-IOV logical port, or virtual slot number

[console_slot]

IBM i only

DRC index of I/O slot, DRC index of HEA or SR-IOV logical port, virtual slot number, or the value **hmc**

[alt_console_slot]

IBM i only

DRC index of I/O slot

[op_console_slot]

IBM i only

DRC index of I/O slot

[auto_start]

Valid values are:

0 - off

1 - on

[boot_mode]

AIX, Linux, and virtual I/O server only

Valid values are:

norm - normal

dd - diagnostic with default boot list

ds - diagnostic with stored boot list

of - Open Firmware OK prompt

sms - System Management Services

[power_ctrl_lpar_ids | power_ctrl_lpar_names]

comma separated

[conn_monitoring]

Valid values are:

0 - off

1 - on

[hsl_pool_id]

IBM i only

Valid values are:

0 - HSL OptiConnect is disabled

1 - HSL OptiConnect is enabled

[virtual_opti_pool_id]

IBM i only

Valid values are:

0 - virtual OptiConnect is disabled

1 - virtual OptiConnect is enabled

[max_virtual_slots]

[virtual_eth_adapters]

Comma separated list of virtual ethernet adapters, with each adapter having the following format:

*virtual-slot-number[/is-IEEE/port-vlan-ID/
[additional-vlan-IDs]/[trunk-priority]/*

is-required[[*virtual-switch*][[*MAC-address*]/
[*allowed-OS-MAC-addresses*][*QoS-priority*]]]

The first 5 *'/* characters must be present. The 6th *'/* character is optional, but it must be present if *virtual-switch* or any of the values following *virtual-switch* are specified. The last 3 *'/* characters are optional, but all 3 must be present if *MAC-address*, *allowed-OS-MAC-addresses*, or *QoS-priority* is specified.

Optional values may be omitted. Optional values are *additional-vlan-IDs*, *trunk-priority*, *virtual-switch*, *MAC-address*, *allowed-OS-MAC-addresses*, and *QoS-priority*.

Valid values for *is-IEEE* and *is-required*:

- 0** - no
- 1** - yes

If values are specified for *additional-vlan-IDs*, they must be comma separated.

Valid values for *trunk-priority*:

- 0** - this adapter is not a trunk adapter (default value)
- 1 - 15** - this adapter is a trunk adapter with the specified priority

If *MAC-address* is not specified, a unique MAC address will be automatically generated for the adapter. If it is specified, it must be specified as 12 hexadecimal characters. It is highly recommended that you do not specify *MAC-address* so that the MAC address will be automatically generated.

allowed-OS-MAC-addresses can be a comma separated list of 1 to 4 MAC addresses, where each MAC address is specified as 12 hexadecimal characters, or it can be one of the following values:

- all** - all OS defined MAC addresses are allowed (default value)
- none** - no OS defined MAC addresses are allowed

Valid values for *QoS-priority*:

- none** - do not use Quality of Service (QoS) (default value)
- 0 - 7** - the QoS priority level to use

For example:

13/1/5/6,7/2/1

specifies a virtual ethernet adapter with a virtual slot number of **13**, it is IEEE 802.1Q compatible, it has a port virtual LAN ID of **5**, additional virtual LAN

IDs of **6** and **7**, it is a trunk adapter with a trunk priority of **2**, and it is required. It also uses the default virtual switch, it uses an automatically generated MAC address, it allows all OS defined MAC addresses, and it does not use QoS.

4/1/3//0/1///fecd537a910b,fecd537a910c/2

specifies a virtual ethernet adapter with a virtual slot number of **4**, it is IEEE 802.1Q compatible, it has a port virtual LAN ID of **3**, no additional virtual LAN IDs, it is not a trunk adapter, it is required, it uses the default virtual switch, it uses an automatically generated MAC address, it allows the OS defined MAC addresses FE:CD:53:7A:91:0B and FE:CD:53:7A:91:0C, and it uses a QoS priority level of **2**.

[virtual_eth_vsi_profiles]

Comma separated list of virtual ethernet adapter Virtual Station Interface (VSI) profile attributes:

virtual-slot-number/VSI-manager-ID/VSI-type-ID/VSI-type-version

All 3 '/' characters must be present, and all values are required.

For example:

15/2/1193040/1

specifies a VSI profile for the virtual ethernet adapter in slot **15**. The VSI profile has a VSI manager ID of **2**, a VSI type ID of **1193040**, and a VSI type version of **1**.

[virtual_fc_adapters]

Comma separated list of virtual fibre channel adapters, with each adapter having the following format:

*virtual-slot-number/client-or-server/
[remote-lpar-ID]/[remote-lpar-name]/
remote-slot-number/[wwpns]/is-required*

All 6 '/' characters must be present, but optional values may be omitted. Optional values are *remote-lpar-ID* or *remote-lpar-name* (one of those values is required, but not both).

wwpns is optional for a client adapter, and is not allowed for a server adapter. If *wwpns* is not specified, WWPNs will be automatically generated for a client adapter. It is highly recommended that you do not specify *wwpns* so that WWPNs will be automatically generated.

Valid values for *client-or-server*:

client
server

Valid values for *is-required*:

0 - no

1 - yes

For example:

4/client/vios_p1/16/1

specifies a virtual fibre channel **client** adapter with a virtual slot number of **4**, a remote (server) partition name of **vios_p1**, a remote (server) slot number of **16**, and it is required. WWPNs will be automatically generated for this client adapter.

[virtual_scsi_adapters]

Comma separated list of virtual SCSI adapters, with each adapter having the following format:

```
virtual-slot-number/client-or-server/
[remote-lpar-ID]/[remote-lpar-name]/
[remote-slot-number]/is-required
```

All 5 '/' characters must be present, but optional values may be omitted. Optional values for server adapters are *remote-lpar-ID*, *remote-lpar-name*, and *remote-slot-number*. Optional values for client adapters are *remote-lpar-ID* or *remote-lpar-name* (one of those values is required, but not both).

Valid values for *client-or-server*:

client

server

Valid values for *is-required*:

0 - no

1 - yes

For example:

14/client/2//13/0

specifies a virtual SCSI **client** adapter with a virtual slot number of **14**, a remote (server) partition ID of **2**, a remote (server) slot number of **13**, and it is not required.

[virtual_serial_adapters]

Comma separated list of virtual serial adapters, with each adapter having the following format:

```
virtual-slot-number/client-or-server/
[supports-HMC]/[remote-lpar-ID]/
```

*[remote-lpar-name]/[remote-slot-number]/
is-required*

All 6 '/' characters must be present, but optional values may be omitted. Optional values for server adapters are *supports-HMC*, *remote-lpar-ID*, *remote-lpar-name*, and *remote-slot-number*. Optional values for client adapters are *remote-lpar-ID* or *remote-lpar-name* (one of those values is required, but not both), and the *supports-HMC* value is not allowed.

Valid values for *client-or-server*:

client

server

Valid values for *supports-HMC*:

0 - no

Valid values for *is-required*:

0 - no

1 - yes

For example:

14/server/0///0

specifies a virtual serial **server** adapter with a virtual slot number of **14**, it does not support an HMC connection, any client adapter is allowed to connect to it, and it is not required.

[vnic_adapters]

AIX, Linux, and IBM i only

Comma separated list of virtual NICs, with each virtual NIC consisting of a list of colon separated property name/value pairs, which can be specified in any order:

property-name=value:property-name=value:...

Property names:

slot_num

Required

port_vlan_id

Port VLAN ID or **0** to disable VLAN tag insertions for untagged frames (default value)

pvid_priority

Valid values:

An integer between **0** and **7**, inclusive.

Default value is **0**.

allowed_vlan_ids

Comma separated list of VLAN IDs, or one of the

following values:

all - all VLAN IDs are allowed (default value)

none - no VLAN IDs are allowed

mac_addr

12 hexadecimal characters

If not specified, a unique MAC address will be automatically generated for the virtual NIC.

It is highly recommended that you use an automatically generated MAC address.

allowed_os_mac_addr

Comma separated list of MAC addresses, each specified as 12 hexadecimal characters. Other valid values:

all - all OS defined MAC addresses are allowed (default value)

none - no OS defined MAC addresses are allowed

is_required

Valid values:

0 - no

1 - yes (default value)

auto_priority_failover

Valid values:

0 - automatic priority failover is disabled

1 - automatic priority failover is enabled (default value)

backing_devices

Required

Comma separated list of virtual NIC backing devices, with each backing device having the following format:

```
sriov/vios-lpar-name/vios-lpar-ID/  
sriov-adapter-ID/sriov-physical-port-ID/  
[capacity][/failover-priority]
```

The first 5 '/' characters must be present. The 6th '/' character is optional, but it must be present if *failover-priority* is specified.

Optional values may be omitted. Optional values are *capacity*, either *vios-lpar-name* or *vios-lpar-ID*, and *failover-priority*.

Specify *capacity* as a percentage with up to 2 decimal places.

Specify *failover-priority* as a value between **1** and **100** inclusive, with **1** being the highest priority. It defaults to **50** if not specified.

For example:

slot_num=5:backing_devices=sriov/vios1//1/0//1,sriov/vios2//2/0//10 specifies a virtual NIC with a virtual slot number of **5** and 2 backing devices. The

first backing device is on VIOS **vios1** and physical port **0** of SR-IOV adapter **1** and has a failover priority of **1**. The second backing device is on VIOS **vios2** and physical port **0** of SR-IOV adapter **2** and has a failover priority of **10**. All the other virtual NIC properties are configured with default values.

[hca_adapters]

AIX, Linux, and virtual I/O server only
Comma separated list of Host Channel adapters (HCA), with each adapter having the following format:

adapter-ID/GUID/capability

All 3 values must be specified for each adapter.

Valid values for *capability*:

- 1** - low
- 2** - medium
- 3** - high
- 4** - dedicated

For example:

23000cff/255000000609/3 specifies an HCA with an adapter ID of **23000cff**, a Globally Unique ID (GUID) of 2:55:00:00:00:06:09, and a capability setting of high.

[lhea_logical_ports]

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

*adapter-ID/port-group/physical-port-ID/
logical-port-ID/[allowed-VLAN-IDs]/
allowed-OS-MAC-addresses]*

The first 4 '/' characters must be present. The last '/' character is optional, but it must be present if *allowed-OS-MAC-addresses* is specified. Optional values may be omitted. Optional values are *allowed-VLAN-IDs* and *allowed-OS-MAC-addresses*.

If values are specified for *allowed-VLAN-IDs*, they must be comma separated, or one of the following values can be specified:

- all** - allowed to participate in all VLANs
- none** - not allowed to participate in any VLAN (default value)

allowed-OS-MAC-addresses can be a comma separated list of 1 to 4 MAC addresses, where each MAC address is specified as 12 hexadecimal characters, or it can be

one of the following values:

all - all OS defined MAC addresses are allowed
(default value)

none - no OS defined MAC addresses are allowed

For example:

23000000/2/0/1/3 specifies logical port **1** for physical port **0** belonging to port group **2** of the Host Ethernet adapter (HEA) with an adapter ID of **23000000**. This logical port is allowed to participate in the VLAN with an ID of **3** only. It also allows all OS defined MAC addresses.

23000000/2/1/5/all/fecd537a910b,fecd537a910c specifies logical port **5** for physical port **1** belonging to port group **2** of the HEA with an adapter ID of **23000000**. This logical port is allowed to participate in all VLANs, and it allows the OS defined MAC addresses FE:CD:53:7A:91:0B and FE:CD:53:7A:91:0C.

[lhea_capabilities]

Comma separated list of LHEA capabilities, with each capability having one of the following formats:

adapter-ID/capability

or

adapter-ID/5/ieq/nieq/qp/cq/mr

where *ieq* (interruptible event queues), *nieq* (non-interruptible event queues), *qp* (queue pairs), *cq* (completion queues), and *mr* (memory regions) each specify the resource amount in addition to the base minimum.

Valid values for *capability*:

0 - base minimum

1 - low

2 - medium

3 - high

4 - dedicated

For example:

23000000/3 sets the LHEA capability for the HEA with an adapter ID of **23000000** to high.

[sni_device_ids]

AIX, Linux, and virtual I/O server only
Comma separated list of Switch Network Interface (SNI) adapter device IDs

[work_group_id]

Specify **none** if you are not using

NAME

mksysconn - create system connection

SYNOPSIS

To establish a connection and add a system or frame to the Hardware Management Console (HMC):

```
mksysconn --ip IP-address [-r {sys | frame}]  
[--passwd password] [--force] [--help]
```

To enable all systems and frames to be automatically discovered by the HMC when using DHCP:

```
mksysconn -o auto [--help]
```

DESCRIPTION

mksysconn establishes a connection from the Hardware Management Console (HMC) to a system in the network and adds the system to the systems managed by the HMC. If a connection to the system cannot be established because the network is down, the service processor for the system is down, or too many other HMCs have already established a connection to the system, the system will be added as a managed system which is in the No Connection state.

mksysconn can also establish a connection from the HMC to a frame in the network and add the frame to the frames managed by the HMC. If a connection to the frame cannot be established because the network is down, the bulk power assembly (BPA) for the frame is down, or too many other HMCs have already established a connection to the frame, the frame will be added as a managed frame which is in the No Connection state.

If your HMC is set up as a DHCP server on a private network, **mksysconn -o auto** can enable all systems and frames to be automatically discovered by the HMC in the event that the Remove Connection task or the **rmsysconn** command was previously run to remove a system or frame from the HMC.

If your HMC is set up as a DHCP server on a private network, do not use this command to establish HMC connections to managed systems and frames that are DHCP clients. This command is intended for use on a public network only, where the systems are set to use static IP addresses. Using this command in a DHCP environment establishes a temporary connection that will not continue to work properly over network configuration changes, system power loss, and service repair actions.

OPTIONS

--ip To connect to a system and add the system to the systems managed by the HMC, specify the IP address or host name of the service processor for the system.

To connect to a frame and add the frame to the frames managed by the HMC, specify the IP address or host name of one side of the bulk power assembly (BPA) for the frame. Note that to properly connect to a frame, it is recommended that you run this command twice, once for each side of the BPA, in order to connect to both sides of the BPA for the frame.

-r The type of resource to which to connect and add to the HMC. Valid values are **sys** for system and **frame** for frame. If this option is omitted, the HMC determines the resource type by determining if the specified IP address belongs to a system or a frame. If the HMC is unable to determine if the specified IP address belongs to a system or a frame, this command will fail unless the **--force** option is specified.

--passwd

The HMC Access password for the system or the frame to which to connect. If this option is omitted, you will be prompted to enter the password.

--force If the HMC is unable to determine if the specified IP address belongs to a system or a frame, use this option to force the HMC to attempt to establish a connection. The HMC will use the value specified with the **-r** option to determine the type of connection to establish. If the **-r** option is not specified, the HMC will attempt to establish a connection to a system.

- o** If your HMC is set up as a DHCP server on a private network, specify **auto** with this option to enable the HMC to automatically discover all systems and frames. You may need to do this if the Remove Connection task or the **rmsysconn** command was previously run to remove a system or frame from the HMC, and you want the HMC to be able to manage a system or frame with the same IP address again. Whenever a managed system or a managed frame is removed from the HMC when using DHCP, the HMC places the IP address(es) of that system or frame in a list of removed IP addresses. Any IP address in that list will not be automatically rediscovered by the HMC. The **mksysconn -o auto** command removes all IP addresses from that list. To remove a specific IP address from that list, you can run the **rmsysconn -o rediscover** command instead.

auto is the only value that can be specified with this option. When this option is specified, all other options are ignored.

- help** Display the help text for this command and exit.

EXAMPLES

Connect to and add the system with the host name **sys1.company.com**:

```
mksysconn --ip sys1.company.com --passwd sys1pw
```

Connect to and add the system with the IP address 9.3.152.145 (the HMC Access password for the system must be entered when prompted):

```
mksysconn --ip 9.3.152.145
```

Connect to and add the frame with the host names **frame1_A.company.com** and **frame1_B.company.com**:

```
mksysconn --ip frame1_A.company.com -r frame --passwd  
frame1pw  
mksysconn --ip frame1_B.company.com -r frame --passwd  
frame1pw
```

To enable all systems and frames to be automatically discovered by the HMC when using DHCP:

```
mksysconn -o auto
```

ENVIRONMENT

None

BUGS

None

AUTHOR

IBM Austin

SEE ALSO

lssysconn, **rmsysconn**, **lssyscfg**

NAME

mksysplan - make system plan file

SYNOPSIS

```
mksysplan -f file-name -m managed-system  
[--check] [-d "description"]  
[--noprobe] [--novios] [--nohwdisc] [--noinvscout]  
[-v] [--help]
```

DESCRIPTION

mksysplan creates a system plan file that represents the information known about the hardware, partitions, profiles, and partition provisioning for the *managed-system*.

By default, mksysplan will perform additional inventory probes of active partitions, attempt hardware discovery for inactive partitions or unallocated hardware, and gather additional information from Virtual I/O Server (VIOS) partitions. Additional information from VIOS partitions is only gathered for POWER5 and POWER6 servers.

OPTIONS

-f Specifies the file name that will contain the system plan that this command creates. If the file does not exist, the command will create it in the system plan file directory on the Hardware Management Console (HMC). If the file exists, the command will overwrite the contents of the file unless the **--check** option is specified.

The file name can only consist of the alphanumeric characters (upper and lower case) and a set of special characters (comma, period, hyphen, underscore, and space). The file name cannot begin with a period or hyphen.

The file name must end with the **.sysplan** suffix.

-m Specifies the name of the managed system that the command will use to get the information to create the system plan file.

The name can be either the user-defined name for the managed system, or it can be in the form *ttt-
mmm*sssssss*, where *ttt* is the machine type, *mmm* is the model, and *sssssss* is the serial number of the managed system. The *ttt-
mmm*sssssss* form must be used if there are multiple managed systems with the same user-defined name.

--check

Specifies that this command will check the system plan directory to see if the file name specified with the **-f** option already exists. If the file does exist, the command will fail with an error message.

-d Specifies a description that the command will add to the created file.

--noprobe

Limit the inventory gathering to obtain only the PCI slot devices without any further inventory probes of active partitions, without any refresh of inactive partition or unallocated hardware information, and without gathering additional information from VIOS partitions.

--novios

Do not gather additional information from VIOS partitions.

This option is always in effect for POWER7 and later servers.

--nohwdisc

Do not perform hardware discovery for inactive partitions or unallocated hardware.

--noinvscout

Do not perform additional inventory probes of active partitions.

-v Displays verbose output during command processing, in addition to the default messages.

--help Display the help text for this command and exit.

EXAMPLES

Make a system plan file **sysplan.sysplan** that represents the configuration of the managed system **mySystem**:

```
mksysplan -f sysplan.sysplan -m mySystem
```

Make a system plan file **sysplan.sysplan** that represents the configuration of the managed system **9406-570*3413556**, and only survey PCI devices in slots:

```
mksysplan -f sysplan.sysplan -m 9406-570*3413556 -d "New 570 system plan for 5/11/2007" --noprobe
```

ENVIRONMENT

None

BUGS

None

AUTHOR

IBM Austin

SEE ALSO

cpsysplan, deploysysplan, lssysplan, rmsysplan

NAME

mkvterm - open a virtual terminal session

SYNOPSIS

mkvterm -m *managed-system*
{**-p** *partition-name* | **--id** *partition-ID*} [**--help**]

DESCRIPTION

mkvterm opens a virtual terminal session for an AIX, Linux, or virtual I/O server partition.

After establishing a virtual terminal session, the `~.` character sequence can be entered in the terminal window to terminate it, or the **rmvterm** command can be used to force the session to be closed.

A partition can only have one open virtual terminal session at a time.

OPTIONS

-m The name of the managed system which has the partition for which to open the virtual terminal session. The name may either be the user-defined name for the managed system, or be in the form *ttt-mmm*sssssss*, where *ttt* is the machine type, *mmm* is the model, and *sssssss* is the serial number of the managed system. The *ttt-mmm*sssssss* form must be used if there are multiple managed systems with the same user-defined name.

-p The name of the partition for which to open the virtual terminal session.

You must either use this option to specify the name of the partition, or use the **--id** option to specify the partition's ID. The **-p** and the **--id** options are mutually exclusive.

--id The ID of the partition for which to open the virtual terminal session.

You must either use this option to specify the ID of the partition, or use the **-p** option to specify the partition's name. The **--id** and the **-p** options are mutually exclusive.

--help Display the help text for this command and exit.

EXAMPLES

Open a virtual terminal session for partition **p1**:

```
mkvterm -m mySystem -p p1
```

Open a virtual terminal session for the partition with an ID of **1**:

```
mkvterm -m 9406-570*12345678 --id 1
```

ENVIRONMENT

None

BUGS

None

AUTHOR

IBM Austin

SEE ALSO

rmvterm

NAME

monhmc - monitor HMC subsystems and system resources

SYNOPSIS

To monitor Hardware Management Console (HMC) subsystems:

```
monhmc -s {hmcsvr | rmc} [-n <interval>]  
 [--help]
```

To monitor HMC system resources:

```
monhmc -r {disk | proc | mem | swap} [-n <interval>]  
 [--help]
```

DESCRIPTION

monhmc provides a dynamic real-time view of HMC related subsystems and system resources. The **top**, **watch**, and **df** commands are used to implement this command. Refer to documentation for those commands for additional information.

OPTIONS

-s The HMC subsystem to monitor. Valid values are **hmcsvr** for HMC server processes, and **rmc** for Resource Monitoring and Control (RMC) processes.

Either this option or the **-r** option is required. The **-s** and the **-r** options are mutually exclusive.

-r The HMC system resource to monitor. Valid values are **disk** for filesystem disk space usage, **proc** for processor (CPU) usage, **mem** for memory usage, and **swap** for swap space usage.

Either this option or the **-s** option is required. The **-r** and the **-s** options are mutually exclusive.

-n The interval between updates in seconds. The default value for this option is 4 seconds.

If an interval of 0 is specified, then statistics are displayed only once, and this command exits. If any other interval is specified, or if this option is omitted, then statistics are updated every *interval* seconds, and this command runs until interrupted with Ctrl-c.

--help Display the help text for this command and exit.

EXAMPLES

Monitor HMC processor usage and update the statistics every 10 seconds:

```
monhmc -r proc -n 10
```

Display RMC subsystem statistics once:

```
monhmc -s rmc -n 0
```

ENVIRONMENT

None

BUGS

None

AUTHOR

IBM Austin

NAME

optmem - optimize memory

SYNOPSIS

To start a mirrored memory defragmentation operation on a managed system:

```
optmem -m managed-system -o start -t mirror  
[-q quantity] [--minutes time-limit]
```

To start a Dynamic Platform Optimization operation on a managed system:

```
optmem -m managed-system -o start -t affinity  
[-p partition-names | --id partition-IDs]  
[-x partition-names | --xid partition-IDs]
```

To stop a memory optimization operation on a managed system:

```
optmem -m managed-system -o stop [--optid ID]
```

DESCRIPTION

optmem performs memory optimization operations on the *managed-system*. The memory optimization operations that can be performed with this command are mirrored memory defragmentation operations and Dynamic Platform Optimization operations.

Dynamic Platform Optimization operations dynamically optimize the placement of partitions to maximize the processor-memory affinity of the partitions, which improves system performance. You can specify a list of partitions to prioritize for optimization, and a list of partitions to protect from the optimization. Partitions which are not specified in either list could be either positively or negatively impacted by the optimization. The **lsmemopt** command can be used to evaluate the current system affinity and the potential system affinity that could be attained by running a Dynamic Platform Optimization operation.

When a memory optimization operation is in progress, another memory optimization operation cannot be started. In addition, operations that change the memory or processor configuration in the system are not allowed.

System performance will degrade during a memory optimization operation. A memory optimization operation may take a long time to complete.

OPTIONS

-m The name of the managed system on which to perform the memory optimization operation. The name may either be the user-defined name for the managed system, or be in the form *ttt-
mmm*ssssss*, where *ttt* is the machine type, *mmm* is the model, and *ssssss* is the serial number of the managed system. The *ttt-
mmm*ssssss* form must be used if there are multiple managed systems with the same user-defined name.

-o The operation to perform. Valid values are **start** to start a memory optimization operation, and **stop** to stop the memory optimization operation that is currently in progress. The managed system may not be able to stop the operation immediately, so it may continue to run for several more seconds.

Dynamic Platform Optimization operations should not be stopped. Stopping a Dynamic Platform Optimization operation before it has completed could leave the system in an affinity state that is much worse than before the operation started.

-t The type of memory optimization to perform. Valid values are **mirror** to perform a mirrored memory defragmentation operation, and **affinity** to perform a Dynamic Platform Optimization operation.

-q The amount of mirrored memory to be made available by the defragmentation operation. Specify only the additional amount of mirrored memory to be made available; do not include the amount of mirrored memory that is currently available. *quantity* must be in megabytes, it must be a

multiple of the memory region size for the *managed-system*, and it must be greater than 0.

As much additional mirrored memory as possible, up to *quantity*, will be made available by the defragmentation operation. If this option is not specified, then the maximum amount of mirrored memory will be made available.

--minutes

The time limit, in minutes, for the mirrored memory defragmentation operation. If the operation does not complete within the time limit specified, the operation will be stopped.

The managed system may not be able to stop the operation immediately when the time limit expires, so it may continue to run for several more seconds.

If this option is not specified, the mirrored memory defragmentation operation will not be timed out.

-p The name of one or more partitions to prioritize for optimization in a Dynamic Platform Optimization operation.

Multiple partition names must be comma separated.

You can either use this option to specify the names of the partitions to prioritize for optimization, or use the **--id** option to specify the IDs of the partitions. The **-p** and the **--id** options are mutually exclusive.

If neither this option nor the **--id** option is used to specify a list of partitions to prioritize, the managed system will optimize partitions based on its own prioritization. In either case, the managed system will attempt to optimize all partitions except those that are specified in the list of partitions to exclude from the operation.

--id The ID of one or more partitions to prioritize for optimization in a Dynamic Platform Optimization operation.

A range of partition IDs may be specified. A range is specified by specifying the beginning partition ID, a dash, and the ending partition ID. Both the beginning and ending partition IDs are included in the range. For example, the range 1-5 includes partitions IDs 1, 2, 3, 4, and 5. The list of partition IDs specified with this option must be comma separated and can include one or more partition IDs, one or more ranges of partition IDs, or both.

You can either use this option to specify the IDs of the partitions to prioritize for optimization, or use the **-p** option to specify the names of the partitions. The **--id** and the **-p** options are mutually exclusive.

If neither this option nor the **-p** option is used to specify a list of partitions to prioritize, the managed system will optimize partitions based on its own prioritization. In either case, the managed system will attempt to optimize all partitions except those that are specified in the list of partitions to exclude from the operation.

-x The name of one or more partitions to exclude and protect from the Dynamic Platform Optimization operation. The processor-memory affinity for the partitions will not be impacted by the operation.

Multiple partition names must be comma separated.

You can either use this option to specify the names of the partitions to exclude, or use the **--xid** option to specify the IDs of the partitions. The **-x** and the **--xid** options are mutually exclusive.

If neither this option nor the **--xid** option is used to specify a list of partitions to exclude, no partitions will be excluded from the operation.

--xid The ID of one or more partitions to exclude and protect from the Dynamic Platform Optimization operation. The processor-memory affinity for the partitions will not be impacted by the operation.

A range of partition IDs may be specified. A range is specified by specifying the beginning partition ID, a dash, and the ending partition ID. Both the beginning and ending partition IDs are included in the range. For example, the range 8-11 includes partitions IDs 8, 9, 10, and 11.

The list of partition IDs specified with this option must be comma separated and can include one or more partition IDs, one or more ranges of partition IDs, or both.

You can either use this option to specify the IDs of the partitions to exclude, or use the **-x** option to specify the names of the partitions. The **--xid** and the **-x** options are mutually exclusive.

If neither this option nor the **-x** option is used to specify a list of partitions to exclude, no partitions will be excluded from the operation.

--optid The ID of the memory optimization operation to stop. If this option is omitted, then the memory optimization operation currently in progress is stopped. The ID is displayed by the **lsmemopt -m managed-system** command.

Although this option is not required to stop a memory optimization operation, it is highly recommended that it be specified to prevent the accidental cancellation of the wrong operation.

--help Display the help text for this command and exit.

EXAMPLES

Start a mirrored memory defragmentation operation with no time limit to make an additional **256** MB of mirrored memory available:

```
optmem -m system1 -o start -t mirror -q 256
```

Start a mirrored memory defragmentation operation with a time limit of 1 hour to make as much mirrored memory available as possible:

```
optmem -m 8233-E8B*1234321 -o start -t mirror --minutes 60
```

Stop the mirrored memory defragmentation operation currently in progress:

```
optmem -m system1 -o stop
```

Start a Dynamic Platform Optimization operation to optimize all partitions in the system:

```
optmem -m sys -o start -t affinity
```

Start a Dynamic Platform Optimization operation that prioritizes partitions **lp1** and **lp3** for optimization:

```
optmem -m sys -o start -t affinity -p lp1,lp3
```

Start a Dynamic Platform Optimization operation that prioritizes the partitions with IDs **2**, **4**, **6** through **9**, and **11**, for optimization, and excludes the partition with ID **1**:

```
optmem -m sys -o start -t affinity --id 2,4,6-9,11 --xid 1
```

ENVIRONMENT

None

BUGS

None

AUTHOR

IBM Austin

SEE ALSO

lsmemopt

NAME

pedbg - Product Engineering debug tools

SYNOPSIS

pedbg {**-d** {**on** | **off**} | **-j** {**on** | **off**} | **-l** {**rmc** | **se**} | **-c** | **-s** | **-q** {**n**} | **-r**} [**--help**]

DESCRIPTION

pedbg provides debug tools for Product Engineer/Support Personnel. This command requires PE authority to run, i.e, only hscpe user can access this command.

OPTIONS

- d** Turn on or off various debug tracing. This option will restart subsystems to start/stop internal tracing.
- j** Turn on or off Just in Time compiler. Turn off Just in Time compiler will result in loss of performance.
- l** List internal subsystem information. If rmc is specified, the list of RMC daemons will be displayed, along with its state. If se is specified, the IBM.ServiceEvent resource entries are displayed.
- c** Collect various logs and javacore. This option can copy the data collected onto a USB flash memory device or leave a zip file in the /dump directory.
- s** Collect various managed system dumps that were sent to the HMC. This option can copy the data collected onto a USB flash memory device or leave a zip file in the /dump directory.
- q** Collect logs or managed system dumps in quiet mode. Not prompted. Used with either the -c or the -s option.

Options for log collection:

- 1 = network info only
- 2 = network info + base logs
- 3 = network info + base logs + extended logs
- 4 = all logs - network info + base logs + extended logs + archives
- 5 = collect files in /home/hscpe/ibmsupt only
- 7 = collect RMC ctsnap only
- 9 = run prompt to copy files to media

Options for managed system dumps:

- 1 collect all system dump types
- 2 collect FSPDUMP
- 3 collect SYSDUMP
- 4 collect LOGDUMP
- 5 collect SMADUMP
- 6 collect PWRDUMP
- 8 delete all system dumps
- 9 = run prompt to copy files to media
- 11 collect LPADUMP

- r** Remove the log files collected that are put in /dump.
- help** Display the help text for this command and exit.

EXAMPLES

The following command turns on debug:

```
pedbg -d on
```

The following command disables debug:

pedbg -d off

ENVIRONMENT

None

BUGS

None

AUTHOR

IBM Austin

NAME

`pesh` - provides PE shell access

SYNOPSIS

To obtain full shell access to a Hardware Management Console (HMC):

```
pesh serial-number-of-HMC-machine
```

To obtain full shell access to a virtual HMC:

```
pesh unique-ID-of-virtual-HMC
```

DESCRIPTION

pesh provides full shell access to the HMC for product engineering and support personnel. **pesh** takes the serial number of the HMC machine or unique ID of the virtual HMC where full shell access is requested, then prompts the user for a one day password obtained from the support organization. If the password is valid, the user is granted full shell access. Only the `hscpe` user can run this command.

EXAMPLES

On an HMC machine, run the **lshmc -v** command to list the Vital Product Data (VPD) information:

```
lshmc -v
```

```
"vpd=*FC ????????  
*VC 20.0  
*N2 Wed Aug 26 08:13:14 CDT 2015  
*FC ????????  
*DS Hardware Management Console  
*TM 7042-CR7  
*SE 10DA45C  
*MN IBM  
*PN 42D01T  
*SZ 8363966464  
*OS Embedded Operating Systems  
*NA 9.3.242.97  
*FC ????????  
*DS Platform Firmware  
*RM V8R8.4.0.0  
"
```

Pass the serial number following the SE tag to the **pesh** command:

```
pesh 10DA45C
```

You will be prompted for a password.

On a virtual HMC, run the **lshmc -v** command to list the Vital Product Data (VPD) information:

```
lshmc -v
```

```
"vpd=*FC ????????  
*VC 20.0  
*N2 Wed Aug 26 13:26:59 CDT 2015  
*FC ????????  
*DS Hardware Management Console  
*TM V4c7-f45
```

*SE 9c4bef5
*MN IBM
*PN N/A
*SZ 4152320000
*OS Embedded Operating Systems
*NA 9.53.178.75
*FC ???????
*DS Platform Firmware
*RM V8R8.4.0.0
*UVMID 4c7f:459c:4980:bef5
"

Pass the unique ID following the UVMID tag with or without the :
separators to the **pesh** command:

pesh 4c7f:459c:4980:bef5

or

pesh 4c7f459c4980bef5

You will be prompted for a password.

ENVIRONMENT

None

BUGS

None

AUTHOR

IBM Austin

NAME

refdev - refresh device

SYNOPSIS

```
refdev -m managed-system { -p partition-name | --id partition-ID }  
[-w wait-time] [-d detail-level] [-v] [--help]
```

DESCRIPTION

refdev refreshes the partition and profile data used for remote restart of the specified partition. The partition and profile data, referred to as remote restart data, is stored on the Hardware Management Console (HMC) hard disk for partitions that are capable of simplified remote restart. For partitions that are capable of remote restart but not simplified remote restart, the remote restart data is stored on the reserved storage device assigned to the partition.

refdev also refreshes the virtual SCSI and virtual fibre channel adapter data collected from the Virtual I/O Server (VIOS) partitions hosting the specified partition's client adapters. This data is used during partition migration when one or more of the VIOS partitions hosting the migrating partition's client adapters on the source managed system are shutdown. This data is stored on the HMC hard disk and is collected for all AIX, Linux, and IBM i partitions in the managed system when data collection is enabled for the managed system.

The HMC always attempts to automatically keep the remote restart data and the storage adapter data for a partition up to date. However, there may be times when the HMC is unable to update the data for a partition. If that happens, the user can run this command to force the data to be refreshed.

OPTIONS

-m The name of the managed system on which to perform the operation. The name may either be the user-defined name for the managed system, or be in the form *ttt-mmm*ssssss*, where *ttt* is the machine type, *mmm* is the model, and *ssssss* is the serial number of the managed system. The *ttt-mmm*ssssss* form must be used if there are multiple managed systems with the same user-defined name.

-p The name of the partition for which to perform the operation.

You must either use this option to specify the name of the partition, or use the **--id** option to specify the partition's ID. The **-p** and the **--id** options are mutually exclusive.

--id The ID of the partition for which to perform the operation.

You must either use this option to specify the ID of the partition, or use the **-p** option to specify the partition's name. The **--id** and the **-p** options are mutually exclusive.

-w The maximum time, in minutes, to wait for VIOS commands issued by the HMC to complete. If a VIOS command does not complete within the time specified, the refresh operation will be stopped.

wait-time must be a whole number. If *wait-time* is 0, the refresh operation will not be timed out.

If this option is not specified, a default value of 3 minutes is used.

-d The level of detail requested from VIOS commands issued by the HMC. Valid values are **0** (none) through **5** (highest).

If this option is not specified, a default value of 1 is used.

-v Specify this option to enable verbose mode for the refresh operation. When verbose mode is enabled, detail messages and warning messages are displayed for a successful operation. Detail messages and warning messages are always displayed for a refresh operation that fails, regardless of whether this option is specified.

--help Display the help text for this command and exit.

EXAMPLES

Refresh the remote restart data for partition **remoteRestartPartition**:

```
refdev -m sys1 -p remoteRestartPartition
```

ENVIRONMENT

None

BUGS

None

AUTHOR

IBM Austin

NAME

rmaccfg - remove access control object

SYNOPSIS

rmaccfg -t { **resource** | **resourcerole** | **taskrole** }
-n *name* [**--help**]

DESCRIPTION

rmaccfg removes an access control role from the Hardware Management Console (HMC), or it removes inactive managed resource objects assigned to an HMC user.

OPTIONS

-t The type of access control objects to remove. Valid values are **resource** for inactive managed resource objects, **resourcerole** for a managed resource role, and **taskrole** for a task role.

-n When removing inactive managed resource objects assigned to an HMC user, use this option to specify the user name of the HMC user from which to remove the inactive objects.

When removing an access control role, use this option to specify the name of the role to remove. Note that the predefined task roles **hmcsuperadmin**, **hmcoperator**, **hmcviewer**, **hmcpe**, and **hmcservicerep** cannot be removed.

If the specified access control role is currently assigned to one or more HMC users, you will be prompted to confirm that you want to remove the role.

--help Display the help text for this command and exit.

EXAMPLES

Remove the managed resource role **lpar_role**:

```
rmaccfg -t resourcerole -n lpar_role
```

Purge all inactive managed resource objects from the user **superuser**:

```
rmaccfg -t resource -n superuser
```

ENVIRONMENT

None

BUGS

None

AUTHOR

IBM Austin

SEE ALSO

chaccfg, **lsaccfg**, **mkaccfg**, **lshmcusr**

NAME

`rmdump` - remove a dump

SYNOPSIS

`rmdump -f file [--force] [--help]`

DESCRIPTION

`rmdump` removes a managed system dump or a managed frame dump from the Hardware Management Console (HMC).

OPTIONS

-f The file name of the dump to be removed.

The **`lsdump -h -F name`** command can be used to list the file names of the dumps on the HMC.

--force Specify this option to force an in progress dump file to be removed. An in progress dump file should only be removed if the offload of the dump was abnormally terminated and the in progress dump file was not cleaned up. Removing an in progress dump file for a dump that is currently being offloaded will cause the offload process to start over.

--help Display the help text for this command and exit.

EXAMPLES

Remove a system dump from the HMC:

```
rmdump -f SYSDUMP.1000D6P.00000001.20100806170032.gz
```

ENVIRONMENT

None

BUGS

None

AUTHOR

IBM Austin

SEE ALSO

`lsdump`

NAME

rmfile - remove file

SYNOPSIS

rmfile -t {krbkeyfile | ldapcert} [--help]

DESCRIPTION

rmfile removes a file of the specified type from the Hardware Management Console (HMC).

OPTIONS

-t The type of file to remove. Valid values are **krbkeyfile** for the Kerberos service key (keytab) file (/etc/krb5.keytab), and **ldapcert** for the LDAP Certificate Authority (CA) certificate file.

After the successful completion of this command to remove the Kerberos service key file, the HMC must be rebooted for the change to take effect.

--help Display the help text for this command and exit.

EXAMPLES

Remove the Kerberos service key file from the HMC:

```
rmfile -t krbkeyfile
```

Remove the LDAP CA certificate file from the HMC:

```
rmfile -t ldapcert
```

ENVIRONMENT

None

BUGS

None

AUTHOR

IBM Austin

SEE ALSO

getfile

NAME

rmhmcusr - remove Hardware Management Console users

SYNOPSIS

```
rmhmcusr {-u user-name |  
-t {all | local | kerberos | ldap | automanage}}  
[--help]
```

DESCRIPTION

rmhmcusr removes Hardware Management Console (HMC) users.

OPTIONS

- u** The user name of the HMC user to remove. The HMC users **root** and **hscroot** cannot be removed.
- Either this option or the **-t** option must be specified. The **-u** and the **-t** options are mutually exclusive.
- t** The type of HMC users to be removed. Valid values are **all** for all users (except **root**, **hscroot**, and **hscpe**), **local** for all users with an authentication type of local (except **root**, **hscroot**, and **hscpe**), **kerberos** for all users with an authentication type of Kerberos, **ldap** for all users with an authentication type of LDAP, and **automanage** for all users with an authentication type of LDAP auto managed.
- Either this option or the **-u** option must be specified. The **-t** and the **-u** options are mutually exclusive.
- help** Display the help text for this command and exit.

EXAMPLES

Remove the user **tester**:

```
rmhmcusr -u tester
```

Remove all users with an authentication type of Kerberos:

```
rmhmcusr -t kerberos
```

ENVIRONMENT

None

BUGS

None

AUTHOR

IBM Austin

SEE ALSO

chhmcusr, **lshmcusr**, **mkhmcusr**

NAME

rmlock - remove lock

SYNOPSIS

rmlock -e *managed-frame* [--help]

DESCRIPTION

rmlock forces a Hardware Management Console (HMC) lock on the *managed-frame* to be released.

OPTIONS

- e** The name of the managed frame to unlock. The name may either be the user-defined name for the managed frame, or be in the form *ttt-mmm*sssssss*, where *ttt* is the type, *mmm* is the model, and *sssssss* is the serial number of the managed frame. The *ttt-mmm*sssssss* form must be used if there are multiple managed frames with the same user-defined name.
- help** Display the help text for this command and exit.

EXAMPLES

Unlock the managed frame **myFrame**:

```
rmlock -e myFrame
```

ENVIRONMENT

None

BUGS

None

AUTHOR

IBM Austin

SEE ALSO

lslock

NAME

rmlparutil - remove utilization data

SYNOPSIS

rmlparutil -m *managed-system* [--help]

DESCRIPTION

rmlparutil removes the utilization data collected for a managed system from the Hardware Management Console (HMC).

OPTIONS

-m The name of the managed system for which collected utilization data is to be removed. The name may either be the user-defined name for the managed system, or be in the form *ttt-mmm*sssssss*, where *ttt* is the machine type, *mmm* is the model, and *sssssss* is the serial number of the managed system. The *ttt-mmm*sssssss* form must be used if there are multiple managed systems with the same user-defined name, or if the HMC does not currently have a connection to the managed system.

Removing the utilization data for a managed system does not disable the collection of utilization data for that managed system. Use the **chlp** command to disable the collection of utilization data.

--help Display the help text for this command and exit.

EXAMPLES

Remove the utilization data collected for the managed system **mySystem**:

```
rmlparutil -m mySystem
```

Remove the utilization data collected for the managed system with the type, model, and serial number **9406-520*98765432**:

```
rmlparutil -m 9406-520*98765432
```

ENVIRONMENT

None

BUGS

None

AUTHOR

IBM Austin

SEE ALSO

chlp, **lslparutil**

NAME

rmprofdata - remove profile data

SYNOPSIS

rmprofdata -m *managed-system* **-f** *file* [**--help**]

DESCRIPTION

rmprofdata removes a profile data backup file for the *managed-system*.

OPTIONS

-m The name of the managed system for which to remove the profile data backup file. The name may either be the user-defined name for the managed system, or be in the form *ttt-mmm*sssssss*, where *ttt* is the machine type, *mmm* is the model, and *sssssss* is the serial number of the managed system. The *ttt-mmm*sssssss* form must be used if there are multiple managed systems with the same user-defined name, or if the HMC does not currently have a connection to the managed system.

-f The name of the profile data backup file to be removed. If *file* is not fully qualified, *file* will be removed from the `/var/hsc/profiles/serial-number` directory on the HMC (*serial-number* is the serial number of the managed system).

To remove profile data from removable media, the media must be present in the removable media device and the device must be mounted with the **mount** command before this command is issued. The **lsmediadev** command can be used to display all of the removable media devices on the HMC.

--help Display the help text for this command and exit.

EXAMPLES

Remove the profile data backup file called **backup1** (the file will be removed from the `/var/hsc/profiles/3413444` directory on the HMC):

```
rmprofdata -m 9406-570*3413444 -f backup1
```

Remove the profile data backup file called **myFile** on a floppy diskette (a floppy diskette must have already been inserted into the diskette drive):

```
mount /media/floppy
```

```
rmprofdata -m mySystem -f /media/floppy/myFile
```

ENVIRONMENT

None

BUGS

None

AUTHOR

IBM Austin

SEE ALSO

bkprofdata, **lsmediadev**, **rstprofdata**

NAME

rmpwdpolicy - remove a password policy

SYNOPSIS

rmpwdpolicy -n *policy-name* [**--help**]

DESCRIPTION

rmpwdpolicy removes a password policy from the Hardware Management Console (HMC).

Only user-defined password policies can be removed. The active password policy cannot be removed.

OPTIONS

-n The name of the password policy to remove.

--help Display the help text for this command and exit.

EXAMPLES

Remove the password policy **xyzPolicy**:

```
rmpwdpolicy -n xyzPolicy
```

ENVIRONMENT

None

BUGS

None

AUTHOR

IBM Austin

SEE ALSO

chpwdpolicy, **lspwdpolicy**, **mkpwdpolicy**

NAME

rmsyscfg - remove a system resource

SYNOPSIS

```
rmsyscfg -r {lpar | prof | sysprof} -m managed-system  
[-n resource-name] [-p partition-name]  
[--id partition-ID] [--help]
```

DESCRIPTION

rmsyscfg removes a partition, a partition profile, or a system profile from the *managed-system*.

OPTIONS

-r The type of system resource to remove. Valid values are **lpar** for a partition, **prof** for a partition profile, and **sysprof** for a system profile.

When a partition is removed, all of the partition profiles that are defined for that partition are also removed.

When a partition profile is removed, any system profiles that contain just that one partition profile are also removed.

-m The name of the managed system from which the system resource is to be removed. The name may either be the user-defined name for the managed system, or be in the form *tttt-mmm*sssssss*, where *tttt* is the machine type, *mmm* is the model, and *sssssss* is the serial number of the managed system. The *tttt-mmm*sssssss* form must be used if there are multiple managed systems with the same user-defined name.

-n The name of the system resource to remove.

To remove a partition, you must either use this option to specify the name of the partition to remove, or use the **--id** option to specify the partition's ID. The **-n** and the **--id** options are mutually exclusive when removing a partition.

To remove a partition profile or a system profile, you must use this option to specify the name of the profile to remove.

-p The name of the partition which has the partition profile to remove. This option is only valid when removing a partition profile.

To remove a partition profile, you must either use this option to specify the name of the partition which has the partition profile to remove, or use the **--id** option to specify the partition's ID. The **-p** and the **--id** options are mutually exclusive.

--id The partition's ID.

To remove a partition, you must either use this option to specify the ID of the partition to remove, or use the **-n** option to specify the partition's name. The **--id** and the **-n** options are mutually exclusive when removing a partition.

To remove a partition profile, you must either use this option to specify the ID of the partition that has the profile to remove, or use the **-p** option to specify the partition's name. The **--id** and the **-p** options are mutually exclusive when removing a partition profile.

This option is not valid when removing a system profile.

--help Display the help text for this command and exit.

EXAMPLES

Remove the partition **partition5**:

rmsyscfg -r lpar -m system1 -n partition5

Remove the partition with ID 5:

rmsyscfg -r lpar -m system1 --id 5

Remove the partition profile **prof1** for partition **lpar3**:

rmsyscfg -r prof -m system1 -n prof1 -p lpar3

Remove the system profile **sysprof1**:

rmsyscfg -r sysprof -m 9406-520*34134441 -n sysprof1

ENVIRONMENT

None

BUGS

None

AUTHOR

IBM Austin

SEE ALSO

chsyscfg, lssyscfg, mksyscfg

NAME

rmsysconn - remove system connection

SYNOPSIS

```
rmsysconn -o { remove | reset | rediscover }  
  { -m managed-system | -e managed-frame | --ip IP-address }  
  [--help]
```

DESCRIPTION

rmsysconn removes or resets a connection from the Hardware Management Console (HMC) to a managed system or a managed frame.

If your HMC is set up as a DHCP server on a private network, you should not use this command to remove HMC connections to managed systems and managed frames. The remove option of this command is intended for use on a public network only, where the HMC is not set up as a DHCP server. Using the remove option of this command prevents the HMC from managing any system or frame at that IP address, even though that address may still be assigned through DHCP.

If your HMC is set up as a DHCP server on a private network and you do use this command to remove a connection, the HMC places the IP address(es) of the managed system or managed frame being removed, or places the IP address being removed, in a list of removed IP addresses. Any IP address in that list will not be automatically rediscovered by the HMC. If you want to remove an IP address from that list so that the IP address can be automatically discovered and the HMC can manage a system or frame at that IP address in the future, use the rediscover option of this command. If you want to remove all IP addresses from that list, you can use the **mksysconn -o auto** command.

OPTIONS

-o The operation to perform. Valid values are **remove**, **reset**, or **rediscover**.

When **remove** is specified, the HMC disconnects from the specified managed system or the specified managed frame. If all connections to the managed system or managed frame are removed, then the managed system or managed frame is removed from the HMC.

When **reset** is specified, the HMC disconnects from the specified managed system or the specified managed frame then attempts to reconnect. The **reset** operation is useful for retrying to establish a connection to a managed system or a managed frame that is in the "No Connection" state.

When **rediscover** is specified, the specified IP address is removed from the HMC's list of removed IP addresses. This will allow the HMC to automatically reconnect to that IP address when using DHCP.

-m The name of the managed system to remove or to reset the connection to. The name may either be the user-defined name for the managed system, or be in the form *tttt-mmm*sssssss*, where *tttt* is the machine type, *mmm* is the model, and *sssssss* is the serial number of the managed system. The *tttt-mmm*sssssss* form must be used if there are multiple managed systems with the same user-defined name.

If the managed system has a single service processor, the connection to the service processor is removed or reset. When the connection is removed, the managed system is removed from the HMC.

If the managed system has two service processors, the connections to both service processors for the managed system are removed or reset. When the connections to both service processors for the managed system are removed, the managed system is removed from the HMC. To remove or reset the connection to just one service processor for the managed system, use the **--ip** option.

Either this option, the **--ip** option, or the **-e** option is required for a remove or a reset operation.

This option is not valid for a rediscover operation.

The **-m**, **--ip**, and the **-e** options are mutually exclusive.

- e** The name of the managed frame to remove or to reset the connection to. The name may either be the user-defined name for the managed frame, or be in the form *tttt-mmm*sssssss*, where *tttt* is the type, *mmm* is the model, and *sssssss* is the serial number of the managed frame. The *tttt-
mmm*sssssss* form must be used if there are multiple managed frames with the same user-defined name.

When this option is used to specify the managed frame to remove or to reset the connection to, the connections to both sides of the bulk power assembly (BPA) for the managed frame are removed or reset. When the connections to both sides of the BPA for the managed frame are removed, the managed frame is removed from the HMC.

To remove or reset the connection to just one side of the BPA for the managed frame, use the **--ip** option.

Either this option, the **--ip** option, or the **-m** option is required for a remove or reset operation. This option is not valid for a rediscover operation.

The **-e**, **--ip**, and the **-m** options are mutually exclusive.

- ip** To remove or reset the connection to a managed system with a single service processor, you can use this option to specify the IP address or host name of the service processor for the managed system to remove or to reset the connection to. Alternatively, you can use the **-m** option to specify the name of the managed system to remove or to reset the connection to.

To remove or reset the connection to just one service processor for a managed system with two service processors, you can use this option to specify the IP address or host name of the service processor to remove or to reset the connection to. To remove or reset the connection to both service processors for a managed system, use the **-m** option. Note that a managed system with two service processors will not be removed from the HMC until the connections to both service processors for the managed system are removed.

To remove or reset the connection to just one side of the BPA for a managed frame, you can use this option to specify the IP address or host name of the BPA side to remove or to reset the connection to. To remove or reset the connection to both sides of the BPA for a managed frame, use the **-e** option. Note that a managed frame will not be removed from the HMC until the connections to both sides of the BPA for the managed frame are removed.

To allow the HMC to automatically rediscover an IP address when using DHCP, use this option to specify the IP address to be automatically rediscovered.

Either this option, the **-m** option, or the **-e** option is required for a remove or reset operation. This option is required for a rediscover operation.

The **--ip**, **-m**, and the **-e** options are mutually exclusive.

- help** Display the help text for this command and exit.

EXAMPLES

Disconnect from the managed system `mySys` and remove it from the HMC:

```
rmsysconn -o remove -m mySys
```

Reset the connection to the managed system service processor with IP address **9.3.152.145**:

```
rmsysconn -o reset --ip 9.3.152.145
```

Disconnect from the managed frame myFrame and remove it from the HMC:

```
rmsysconn -o remove -e myFrame
```

Reset the connection to one side of the BPA (with IP address **9.3.152.29**) for a managed frame:

```
rmsysconn -o reset --ip 9.3.152.29
```

Allow the IP address **10.0.0.127** to be automatically rediscovered when using DHCP:

```
rmsysconn -o rediscover --ip 10.0.0.127
```

ENVIRONMENT

None

BUGS

None

AUTHOR

IBM Austin

SEE ALSO

mksysconn, lssyscfg, lssysconn

NAME

rmsysplan - remove system plan

SYNOPSIS

rmsysplan -f *file-name* [--help]

DESCRIPTION

rmsysplan removes a system plan file from the system plan file directory on the Hardware Management Console (HMC).

OPTIONS

-f Specifies the name of the file that will be removed. Only files in the system plan file directory on the HMC will be removed.

--help Display the help text for this command and exit.

EXAMPLES

Remove the system plan file **sysplan.sysplan**:

rmsysplan -f sysplan.sysplan

ENVIRONMENT

None

BUGS

None

AUTHOR

IBM Austin

SEE ALSO

deploysysplan, cpsysplan, lssysplan, mksysplan

NAME

rmvterm - close a virtual terminal session

SYNOPSIS

rmvterm -m *managed-system*
{**-p** *partition-name* | **--id** *partition-ID*} [**--help**]

DESCRIPTION

rmvterm forces the closure of a virtual terminal session for an AIX, Linux, or virtual I/O server partition.

To close the virtual terminal session normally, enter the `~.` character sequence in the terminal window.

OPTIONS

- m** The name of the managed system which has the partition for which to close the virtual terminal session. The name may either be the user-defined name for the managed system, or be in the form *ttt-mmm*sssssss*, where *ttt* is the machine type, *mmm* is the model, and *sssssss* is the serial number of the managed system. The *ttt-mmm*sssssss* form must be used if there are multiple managed systems with the same user-defined name.
- p** The name of the partition for which to close the virtual terminal session.

You must either use this option to specify the name of the partition, or use the **--id** option to specify the partition's ID. The **-p** and the **--id** options are mutually exclusive.
- id** The ID of the partition for which to close the virtual terminal session.

You must either use this option to specify the ID of the partition, or use the **-p** option to specify the partition's name. The **--id** and the **-p** options are mutually exclusive.
- help** Display the help text for this command and exit.

EXAMPLES

Close a virtual terminal session for partition **p1**:

```
rmvterm -m mySystem -p p1
```

Close a virtual terminal session for the partition with an ID of **1**:

```
rmvterm -m 9406-570*12345678 --id 1
```

ENVIRONMENT

None

BUGS

None

AUTHOR

IBM Austin

SEE ALSO

mkvterm

NAME

`rnvi` - HMC restricted-nvi text editor

SYNOPSIS

`rnvi -f file [--help]`

DESCRIPTION

`rnvi` enables users to edit a text file in a restricted mode. The `rnvi` command invokes the `nvi` command in a **chroot** environment. The `rnvi` command must be issued from the user's home directory and only one file may be specified on the command line. The user may edit a file in a subdirectory located in their home directory by specifying the relative path name. When this command is issued for the first time, a temporary directory, called `.rnvi_tmp`, is created in the user's home directory to store the editor's temporary files. Files are not recoverable from a crashed editing session.

Note: when the editor starts, the message "stderr: No such file or directory" may be displayed. This message can be safely ignored.

OPTIONS

-f The name of the text file to edit. The file must be located within the user's home directory.

--help Display the help text for this command and exit.

EXAMPLES

Edit the file named `example.txt` in the user's home directory:

```
rnvi -f example.txt
```

Edit the file named `example.txt` in the subdirectory called `myfiles` in the user's home directory:

```
rnvi -f myfiles/example.txt
```

ENVIRONMENT

None

BUGS

None

AUTHOR

IBM Austin

SEE ALSO

`nvi`, `mkdir`, `rm`

NAME

rrstartlpar - perform a partition remote restart operation

SYNOPSIS

To perform a partition remote restart operation:

```
rrstartlpar -o {restart | validate | cancel | cleanup | recover}  
  -m managed-system [-t target-managed-system]  
  [-ip IP-address [-u user-ID]]  
  {-p partition-name | --id partition-ID}  
  [--redundantvios {0 | 1 | 2}] [--mpio {1 | 2}]  
  [--vlanbridge {1 | 2}] [--retaindev] [--usecurrdata]  
  [--noconnection]  
  [{-f input-data-file | -i "input-data"}]  
  [-w wait-time] [-d detail-level] [--force]  
  [-v] [--nodetails]  
  [--help]
```

To set management console attributes related to partition remote restart operations:

```
rrstartlpar -o set -r mc  
  {-f input-data-file | -i "input-data"}  
  [--help]
```

DESCRIPTION

rrstartlpar performs partition remote restart operations.

Remote restart operations can be performed for AIX, Linux, or IBM i partitions only.

A partition can be remotely restarted only when the source managed system is in the Power Off In Progress, Power Off, Error, Error - Dump in progress, Initializing, or No Connection state. The destination managed system must be in the Operating state.

OPTIONS

-o The operation to perform. Valid values are **restart** to validate then remotely restart the partition if validation succeeds, **validate** to validate the operation to remotely restart the partition, **cancel** to stop the remote restart of the partition, **cleanup** to clean up after the successful remote restart of the partition, **recover** to recover from the failed remote restart of the partition, and **set** to set attributes related to partition remote restart operations.

After the successful remote restart of a partition, you should perform a clean up operation on the source managed system once the source managed system is in Operating or Standby state.

-m The name of the source managed system for the partition remote restart operation. The name may either be the user-defined name for the managed system, or be in the form *ttt-mmm*ssssss*, where *ttt* is the machine type, *mmm* is the model, and *ssssss* is the serial number of the managed system. The *ttt-mmm*ssssss* form must be used if there are multiple managed systems with the same user-defined name.

-t The name of the target, or destination, managed system for the partition remote restart operation. The name may either be the user-defined name for the managed system, or be in the form *ttt-mmm*ssssss*, where *ttt* is the machine type, *mmm* is the model, and *ssssss* is the serial number of the managed system. The *ttt-mmm*ssssss* form must be used if there are multiple managed systems with the same user-defined name.

This option is required for all operations except a clean up operation.

--ip If the destination managed system is not managed by the same management console that is managing the source managed system, use this option to specify the IP address or host name of the management console that is managing the destination managed system.

SSH must be enabled on both management consoles. Also, you must run the **mkauthkeys** command once to set up SSH key authentication.

This option is valid only for partitions that support simplified remote restart operations. This option is required when remotely restarting partitions, validating partition remote restart operations, or recovering from a failed partition remote restart if the destination managed system is not managed by the same management console that is managing the source managed system. This option is not valid when performing any other operation.

- u** If the destination managed system is not managed by the same management console that is managing the source managed system, use this option to specify the user ID to use on the management console that is managing the destination managed system. If you do not specify the user ID, then the user ID of the user that issues this command will be used.
- p** The name of the partition for which to perform the remote restart operation.

You must either use this option to specify the name of the partition, or use the **--id** option to specify the partition's ID. This option must be used when performing a cancel or recover operation.

The **-p** and the **--id** options are mutually exclusive.

- id** The ID of the partition for which to perform the remote restart operation.

You must either use this option to specify the ID of the partition, or use the **-p** option to specify the partition's name. The **-p** option must be used when performing a cancel or recover operation.

The **--id** and the **-p** options are mutually exclusive.

--redundantvios

When validating or performing the remote restart of a partition using a reserved storage device or a partition that uses shared memory, use this option to specify whether the partition is to be configured to use redundant VIOS partitions on the destination managed system for accessing the reserved storage device or the paging device assigned to the partition.

Valid values are **0** if the partition is to be configured to not use redundant VIOS partitions, **1** if the partition is to be configured to use redundant VIOS partitions, or **2** if the partition is to be configured to use redundant VIOS partitions if possible. If this option is not specified, then the partition will be configured to use the same VIOS redundancy configuration on the destination managed system that the partition is currently using on the source managed system.

This option is only valid when performing a validate or remote restart operation.

- mpio** When validating or remotely restarting a partition, use this option to specify whether the management console is required to maintain an equivalent multipath I/O (MPIO) configuration of each of the partition's virtual SCSI and virtual fibre channel adapters on the destination managed system. Valid values are **1** if the management console is required to maintain an equivalent MPIO configuration, or **2** if the management console is not required to maintain an equivalent MPIO configuration, but should do so if possible. If this option is not specified, then this option defaults to the value **1**.

If this option has a value of **1** and the management console cannot guarantee that an equivalent MPIO configuration of all of the partition's virtual SCSI and virtual fibre channel adapters can be maintained on the destination managed system, then an error will occur and the operation will fail.

This option is only valid when performing a validate or remote restart operation.

--vlanbridge

When validating or remotely restarting a partition, use this option to specify whether each of the partition's virtual ethernet adapters are required to be configured so that they are bridged on the same VLAN to an external network on the destination managed system. Valid values are **1** if virtual ethernet adapters are required to be bridged, or **2** if virtual ethernet adapters are not required to be bridged, but should be bridged if possible. If this option is not specified, then this option defaults to the value **1**.

If this option has a value of **1** and the management console cannot guarantee that the partition's virtual ethernet adapters can be bridged, then an error will occur and the operation will fail.

This option is only valid when performing a validate or remote restart operation.

--usecurrdata

Specify this option to allow a partition to be remotely restarted when the partition configuration data saved in the remote restart data for the partition does not match the current configuration of the partition. The partition will be remotely restarted with the saved configuration data.

This option is only valid when performing a remote restart operation.

--retaindev

When performing a clean up operation after a successful remote restart of a partition using a reserved storage device, specify this option to prevent the reserved storage device assigned to the partition from being removed from the reserved storage device pool on the source managed system.

This option is only valid when performing a clean up operation.

--noconnection

Specify this option to allow a partition to be remotely restarted when the source managed system is in the No Connection state.

This option is only valid when performing a validate or simplified remote restart operation.

-f The name of the file containing the input data for this command. The input data consists of attribute name/value pairs, which are in comma separated value (CSV) format.

The format of the input data is as follows:

attribute-name=value,attribute-name=value,...

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

"attribute-name=value,value,...",...

When a list of values is specified, the attribute name/value pair must be enclosed in double quotes. Depending on the shell being used, nested double quote characters may need to be preceded by an escape character, which is usually a `'\'` character.

Attributes that can be specified when validating or remotely restarting a partition:

virtual_fc_mappings

Comma separated list of virtual fibre channel adapter mappings, with each mapping having the following format:

*virtual-slot-number/vios-lpar-name/vios-lpar-ID/
[vios-virtual-slot-number]/[vios-fc-port-name]*

All four '/' characters must be present. Optional values may be omitted. Optional values are *vios-lpar-name* or *vios-lpar-ID* (one of those values is required, but not both), *vios-virtual-slot-number*, and *vios-fc-port-name*.

For example:

4//1/14/fcs0 specifies a mapping of the virtual fibre channel client adapter with slot number **4** to the virtual fibre channel server adapter with slot number **14** in the VIOS partition with ID **1** on the destination managed system. In addition, the mapping specifies to use physical fibre channel port **fcs0**.
shared_proc_pool_name | shared_proc_pool_id

Attributes that can be specified for a set operation:

auto_cleanup_enabled

Automatic clean up of remotely restarted partitions can be performed on the source managed system once the source managed system is in Operating state and all of the VIOS partitions that were hosting client virtual SCSI and virtual fibre channel adapters for the remotely restarted partitions are running with RMC connections to the HMC.

Valid values are:

- 0** - this HMC is not to perform automatic clean up
- 1** - this HMC is to perform automatic clean up (default)

Input data for this command can be specified with this option or the **-i** option. The **-f** and the **-i** options are mutually exclusive.

Input data can be specified when remotely restarting partitions, validating partition remote restart operations, or performing a set operation. This option is not valid when performing any other operation.

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

Input data for this command can be specified with this option or the **-f** option. The **-i** and the **-f** options are mutually exclusive.

Input data can be specified when remotely restarting partitions, validating partition remote restart operations, or performing a set operation. This option is not valid when performing any other operation.

- r** The type of resource for which to set attributes related to partition remote restart operations. Valid values are **mc** for management console.
- w** The maximum time, in minutes, to wait for operating system commands issued by the management console to the partition being remotely restarted to complete. If an operating system command does not complete within the time specified, the partition remote restart operation will be stopped.

wait-time must be a whole number. If *wait-time* is 0, the partition remote restart operation will not

be timed out.

If this option is not specified, a default value of 3 minutes is used.

This option is only valid when performing a validate or remote restart operation.

- d** The level of detail requested from operating system commands issued by the management console to all partitions participating in the remote restart operation. Valid values are **1** (lowest) through **5** (highest).

If this option is not specified, a default value of 1 is used.

This option is only valid when performing a validate or remote restart operation.

- force** When performing a clean up or recover operation, specify this option to force the operation to proceed when errors are encountered.

When performing a simplified remote restart operation for a suspended partition, you must specify this option. A suspended partition will be restarted (resumed) on the destination managed system therefore it will not remain suspended.

- v** Specify this option to enable verbose mode for the partition remote restart operation. When verbose mode is enabled, detail messages and warning messages are displayed for a successful partition remote restart operation. Detail messages and warning messages are displayed for partition remote restart operation that fails, regardless of whether this option is specified.

When the **--nodetails** option is also specified, no detail messages are displayed.

--nodetails

Specify this option to suppress the display of all detail messages.

- help** Display the help text for this command and exit.

EXAMPLES

Validate the operation to remotely restart partition **mylpar** from managed system **system1** on managed system **system2**:

```
rrstartlpar -o validate -m system1 -t system2 -p mylpar
```

Remotely restart the partition with ID **5** from managed system **system1** on managed system **system2**:

```
rrstartlpar -o restart -m system1 -t system2 --id 5
```

Remotely restart the partition **lp1** from managed system **system1** on managed system **system2**:

```
rrstartlpar -o restart -m system1 -t system2 -p lp1  
-i ""virtual_fc_mappings=4/vios1//14/fcs0,5/vios2//14/fcs1""
```

Remotely restart the partition **lp1** from managed system **system1** on managed system **system2** when **system1** is in No Connection state:

```
rrstartlpar -o restart -m system1 -t system2 -p lp1 --noconnection
```

Clean up on the source managed system after the successful remote restart of the partition with ID **5**:

```
rrstartlpar -o cleanup -m system1 --id 5
```

Stop the remote restart of partition **mylpar**:

```
rrstartlpar -o cancel -m system1 -t system2 -p mylpar
```

Recover the failed remote restart of partition **mylpar**:

```
rrstartlpar -o recover -m system1 -t system2 -p mylpar
```

Remotely restart the partition **aix1** from managed system **system1** to managed system **system2**, when **system2** is managed by the HMC with the host name **hmc2**:

```
mkauthkeys --ip hmc2 -u hmc2user
```

```
rrstartlpar -o restart -m system1 -t system2 -p aix1 --ip hmc2  
-u hmc2user
```

Disable automatic clean up for this HMC:

```
rrstartlpar -o set -r mc -i "auto_cleanup_enabled=0"
```

ENVIRONMENT

None

BUGS

None

AUTHOR

IBM Austin

SEE ALSO

lsrrstartlpar, lssyscfg, mkauthkeys

NAME

`rsthwres` - restore hardware resources

SYNOPSIS

To restore memory or processing resources:

```
rsthwres -r {mem | proc} -m managed-system  
[-p partition-name | --id partition-ID]
```

To restore physical I/O slots:

```
rsthwres -r io -m managed-system  
[-p partition-name | --id partition-ID]  
[-l slot-DRC-index]
```

To restore Host Ethernet Adapter (HEA) resources:

```
rsthwres -r hea -m managed-system  
[-p partition-name | --id partition-ID]  
[-l HEA-adapter-ID]  
[-g port-group --logport logical-port-ID]
```

To restore Single Root I/O Virtualization (SR-IOV) resources:

```
rsthwres -r sriov -m managed-system  
[-p partition-name | --id partition-ID]  
[-l adapter-ID] [--logport logical-port-ID]
```

DESCRIPTION

`rsthwres` restores the hardware resource configuration of partitions in the *managed-system*. This operation may need to be performed after a dynamic logical partitioning (DLPAR) operation fails.

OPTIONS

-r The type of hardware resources to restore. Valid values are **mem** for memory, **proc** for processing resources, **io** for physical I/O slots, **hea** for Host Ethernet Adapter (HEA) resources, and **sriov** for Single Root I/O Virtualization (SR-IOV) resources.

-m The name of the managed system which has the partitions for which to restore the hardware resources. The name may either be the user-defined name for the managed system, or be in the form *ttt-mmm*sssss*, where *ttt* is the machine type, *mmm* is the model, and *sssss* is the serial number of the managed system. The *ttt-mmm*sssss* form must be used if there are multiple managed systems with the same user-defined name.

-p The name of the partition for which to restore the hardware resources.

To restore hardware resources for a single partition, you must either use this option to specify the name of the partition, or use the **--id** option to specify the partition's ID. Otherwise, hardware resources for all partitions in the *managed-system* will be restored.

The **-p** and the **--id** options are mutually exclusive.

--id The ID of the partition for which to restore the hardware resources.

To restore hardware resources for a single partition, you must either use this option to specify the ID of the partition, or use the **-p** option to specify the partition's name. Otherwise, hardware resources for all partitions in the *managed-system* will be restored.

The **--id** and the **-p** options are mutually exclusive.

-l When restoring physical I/O slots, specify the DRC index of the physical I/O slot to restore. If this option is omitted and a partition is specified with the **-p** or **--id** option, then all physical I/O slots assigned to the specified partition will be restored. If this option is omitted and a partition is not specified, then all physical I/O slots in the *managed-system* will be restored.

When restoring HEA resources, specify the adapter ID of the HEA to restore. If this option is omitted and a partition is specified with the **-p** or **--id** option, then all HEA resources, including Logical Host Ethernet Adapters and logical ports, assigned to the specified partition will be restored. If this option is omitted and a partition is not specified, then all HEA resources in the *managed-system* will be restored. This option is required when the **-g** and **--logport** options are specified to restore a specific logical port.

When restoring SR-IOV resources, specify the adapter ID of the SR-IOV adapter to restore. If this option is omitted, then all the SR-IOV resources for the specified partitions on the managed system will be included in the restore.

-g The port group containing the HEA logical port to restore. This option is required when the **--logport** option is specified to restore a specific logical port.

--logport

The ID of the HEA logical port or the ID of the SR-IOV logical port to restore.

When restoring SR-IOV resources, if this option is omitted, then all the SR-IOV logical ports for the specified SR-IOV adapters and partitions will be included in the restore.

--help Display the help text for this command and exit.

EXAMPLES

Restore the physical I/O slots for all partitions:

```
rsthwres -r io -m mySystem
```

Restore the physical I/O slot with DRC index **21010003**:

```
rsthwres -r io -m 9406-570*12345678 -l 21010003
```

Restore memory resources for partition **p1**:

```
rsthwres -r mem -m 9406-570*12345678 -p p1
```

Restore processing resources for the partition with ID **1**:

```
rsthwres -r proc -m mySystem --id 1
```

Restore all HEA resources in the managed system:

```
rsthwres -r hea -m mySystem
```

Restore all HEA resources for partition **p1**:

```
rsthwres -r hea -m mySystem -p p1
```

Restore the logical port with ID **3** in port group **2** of the HEA with an adapter ID of **23000010**:

```
rsthwres -r hea -m 9117-MMA*1112223 --logport 3 -g 2 -l 23000010
```

Restore all SR-IOV resources in the managed system:

```
rsthwres -r sriov -m mySystem
```

Restore all SR-IOV resources for partition **p1**:

```
rsthwres -r sriov -m mySystem -p p1
```

Restore the SR-IOV logical port with ID **27004001** of the SR-IOV adapter with adapter ID **1**:

```
rsthwres -r sriov -m mySystem -l 1 --logport 27004001
```

ENVIRONMENT

None

BUGS

None

AUTHOR

IBM Austin

SEE ALSO

chhwres, **lshwres**

NAME

rstprofdata - restore profile data

SYNOPSIS

rstprofdata -m *managed-system* **-l** *restore-type* [**-f** *file*]
[**--help**]

DESCRIPTION

rstprofdata restores profile data for the *managed-system* from a backup file. **rstprofdata** can also be used to initialize the profile data for the *managed-system*.

The **bkprofdata** command can be used to back up profile data for the *managed-system*.

OPTIONS

- m** The name of the managed system for which to restore or initialize profile data. The name may either be the user-defined name for the managed system, or be in the form *tttt-mmm*sssssss*, where *tttt* is the machine type, *mmm* is the model, and *sssssss* is the serial number of the managed system. The *tttt-mmm*sssssss* form must be used if there are multiple managed systems with the same user-defined name.
- l** The type of restore to perform. Valid values are:
- 1** - full restore from the backup file
 - 2** - merge the current profile data and profile data from the backup file, giving priority to the backup data. If the profile data conflicts, the backup data is restored over the current data.
 - 3** - merge the current profile data and profile data from the backup file, giving priority to the current data. If the profile data conflicts, the backup data is not restored over the current data.
 - 4** - initialize the profile data. All partitions, partition profiles, and system profiles for the managed system will be deleted.
- f** The name of the backup file that contains the profile data to be restored. If *file* is not fully qualified, *file* must exist in the */var/hsc/profiles/serial-number* directory on the HMC (*serial-number* is the serial number of the managed system).

To restore the profile data from removable media, the media must be present in the removable media device and the device must be mounted with the **mount** command before this command is issued. The **lsmediadev** command can be used to display all of the removable media devices on the HMC.

This option is required for all restore types except **4** (initialize).

- help** Display the help text for this command and exit.

EXAMPLES

Do a full restore of the profile data from the file **myFile** on a USB flash memory device (a USB flash memory device must already be connected to the HMC):

lsmediadev (to obtain mount points)

mount /media/sdb1

rstprofdata -m mySystem -l 1 -f /media/sdb1/myFile

Restore the profile data from the file **backup1**, and give priority to the current profile data (the file **backup1** exists in the **/var/hsc/profiles/3413444** directory on the HMC):

rstprofdata -m 9406-570*3413444 -l 3 -f backup1

Initialize the profile data:

rstprofdata -m 9406-570*3413444 -l 4

ENVIRONMENT

None

BUGS

None

AUTHOR

IBM Austin

SEE ALSO

bkprofdata, lsmediadev

NAME

rstupgdata - restore upgrade data

SYNOPSIS

rstupgdata -r {dvd | usb} [--help]

DESCRIPTION

rstupgdata copies the current partition profile data, Hardware Management Console (HMC) user data, and HMC configuration data previously saved on removable media by the HMC Save Upgrade Data task onto the HMC hard drive. After running this command, the HMC must be restarted to recover the data.

Note: Time on the HMC must be set after the data has been recovered.

OPTIONS

-r If **dvd** is specified, the upgrade data will be read from DVD-RAM media.

If **usb** is specified, the upgrade data will be read from a USB flash memory device.

The DVD-RAM media or USB flash memory device must be present and the device must be mounted with the **mount** command before running this command. The **lsmediadev** command can be used to display the mount points for the removable media devices on the HMC.

--help Display the help text for this command and exit.

EXAMPLES

Restore upgrade data from DVD-RAM to the HMC hard drive:

lsmediadev (to obtain mount points)

mount /media/cdrom

rstupgdata -r dvd

Restore upgrade data from a USB flash memory device to the HMC hard drive:

lsmediadev (to obtain mount points)

mount /media/sdb1

rstupgdata -r usb

ENVIRONMENT

None

BUGS

None

AUTHOR

IBM Austin

SEE ALSO

lsmediadev, **saveupgdata**

NAME

saveupgdata - save upgrade data

SYNOPSIS

saveupgdata -r {disk | diskusb} [--help]

DESCRIPTION

saveupgdata stores the current partition profile data, Hardware Management Console (HMC) user data, and HMC configuration data in a portion of the HMC hard drive that the HMC upgrade process cannot erase. The HMC automatically recovers the data after the upgrade process completes.

This command should be run immediately prior to performing an HMC upgrade. If the HMC is rebooted without performing an upgrade after running this command, this command must be run again.

OPTIONS

-r If **disk** is specified, upgrade data will only be saved on the HMC hard drive.

If **diskusb** is specified, then in addition to storing the upgrade data on the HMC hard drive, this command will copy the same upgrade data onto a USB data storage device. The USB device must be inserted into a USB port on the HMC before running this command.

--help Display the help text for this command and exit.

EXAMPLES

Save upgrade data to the HMC hard drive:

```
saveupgdata -r disk
```

Save upgrade data to the HMC hard drive and to a USB flash drive:

```
saveupgdata -r diskusb
```

ENVIRONMENT

None

BUGS

None

AUTHOR

IBM Austin

NAME

sendfile - transfer a file from the HMC to a remote system

SYNOPSIS

```
sendfile -f file -h host-name -u user-ID  
[--passwd password] -d remote-directory  
[-n remote-file-name] [-s] [-k key-file] [--help]
```

DESCRIPTION

sendfile transfers a file from the Hardware Management Console (HMC) to a remote system using File Transfer Protocol (FTP) or secure FTP (SFTP).

OPTIONS

- f** The name of the file to transfer.
- h** The host name or IP address of the remote system to which to transfer the file.
- u** The user ID to use to log in to the remote system.
- passwd**
The password to use to log in to the remote system. If this option is omitted, you will be prompted to enter the password.
- d** The directory on the remote system to which to transfer the file.
- n** The file name to give the file on the remote system. If this option is omitted, the file is given the same name that it has on the HMC.
- s** Use secure FTP to transfer the file. SSH will be used to transfer the file, therefore the remote system must have SSH installed and running.
- k** The name of the identity key file for public key authentication. This is the key file generated by the **ssh-keygen** command. If the matching public key file resides on the remote system and the passphrase is empty, you will not be required to enter a password with this command.

This option is only valid when using secure FTP to transfer the file.
- help** Display the help text for this command and exit.

EXAMPLES

Transfer a file using normal FTP. The file will be written to **/home/myid/myfile** on the remote system:

```
sendfile -f /home/joe/myfile -h myhost@company.com  
-d /home/myid -u myid
```

Transfer a file using secure FTP. The file will be written to **/tmp/xfile** on the remote system:

```
sendfile -f /home/joe/myfile -h myhost@company.com -d /tmp  
-n xfile -u myid -s
```

Generate a public key on the HMC then copy to the remote system for sendfile to use public key authentication with SFTP:

```
ssh-keygen -t rsa -f mykey
```

```
scp mykey me@myhost:/home/me/.ssh/authorized_keys2
```

```
sendfile -f /home/me/myfile -h myhost@company.com -d /tmp  
-s -k mykey
```

ENVIRONMENT

None

BUGS

None

AUTHOR

IBM Austin

NAME

startdump - start dump

SYNOPSIS

```
startdump { -m managed-system | -e managed-frame }  
-t { pss | sp | sys | resource } [-s { a | b | p | s }]  
[-l location-code] [-r resource-selector] [--help]
```

DESCRIPTION

startdump initiates a dump on the *managed-system* or the *managed-frame*.

The **lsdump** command can be used to determine when the dump has completed and is available. The **get-dump** command can then be used to offload the dump from the *managed-system* or the *managed-frame* to the Hardware Management Console (HMC).

OPTIONS

-m The name of the managed system on which to initiate the dump. The name may either be the user-defined name for the managed system, or be in the form *tttt-mmm*sssssss*, where *tttt* is the machine type, *mmm* is the model, and *sssssss* is the serial number of the managed system. The *tttt-mmm*sssssss* form must be used if there are multiple managed systems with the same user-defined name.

This option is required when initiating a service processor dump, a system dump, or a resource dump. This option is not valid otherwise.

-e The name of the managed frame on which to initiate the dump. The name may either be the user-defined name for the managed frame, or be in the form *tttt-mmm*sssssss*, where *tttt* is the type, *mmm* is the model, and *sssssss* is the serial number of the managed frame. The *tttt-mmm*sssssss* form must be used if there are multiple managed frames with the same user-defined name.

This option is required when initiating a power subsystem dump. This option is not valid otherwise.

-t The dump to initiate. Valid values are **pss** for a power subsystem dump, **sp** for a service processor dump, **sys** for a system dump, or **resource** for a resource dump.

When a system dump is started, the *managed-system* will be halted. When the dump has completed, the *managed-system* will be restarted. Service processor dumps, resource dumps, and power subsystem dumps do not cause any disruption to the managed system.

Resource dumps are only supported on POWER6 and later servers.

-s When initiating a power subsystem dump, use this option to specify the side of the managed frame's bulk power assembly (BPA) on which to initiate the dump. Valid values are **a** for side A and **b** for side B.

When initiating a service processor dump of the primary or secondary service processor, use this option to specify the service processor on which to initiate the dump. Valid values are **p** for the primary service processor and **s** for the secondary service processor. If this option is not specified and the **-l** option is not specified, the service processor dump will be initiated on the primary service processor.

When initiating a service processor dump, this option and the **-l** option are mutually exclusive.

-l When initiating a service processor dump of a node service processor, use this option to specify the location code of the node service processor on which to initiate the dump.

This option is only valid when initiating a service processor dump. This option and the **-s** option are mutually exclusive.

-r When initiating a resource dump, use this option to specify the resource selector that you have obtained from your hardware service representative. If this option is not specified, no filter is applied to the resource dump. An unfiltered resource dump can take an extended period of time to complete and may prevent other dumps from processing.

This option is only valid when initiating a resource dump.

--help Display the help text for this command and exit.

EXAMPLES

Initiate a system dump on the managed system **9406-570*12345678** (confirmation will be required):

```
startdump -m 9406-570*12345678 -t sys
```

Initiate a dump of the primary service processor on the managed system **sys1**:

```
startdump -m sys1 -t sp -s p
```

Initiate a dump of the node service processor with location code **U78A2.001.DQDGARC-P9-C42**:

```
startdump -m sys1 -t sp -l U78A2.001.DQDGARC-P9-C42
```

Initiate a power subsystem dump on side A of the BPA for the managed frame **myFrame**:

```
startdump -e myFrame -t pss -s a
```

ENVIRONMENT

None

BUGS

None

AUTHOR

IBM Austin

SEE ALSO

getdump, lsdump, lsfru

NAME

termtask - terminate a task

SYNOPSIS

termtask -r {webui | ssh} -s session -t task [--help]

DESCRIPTION

termtask terminates a user's task that is running on the Hardware Management Console (HMC).

The **lslogon** command can be used to list information about the users logged onto the HMC and the tasks they are running.

OPTIONS

- r** The type of task to terminate. Valid values are **webui** for a Web user interface task, or **ssh** for an SSH task.
- s** When terminating a Web user interface task, use this option to specify the session ID of the user that is running the task to terminate.

When terminating an SSH task, use this option to specify the TTY ID of the user that is running the task to terminate.
- t** When terminating a Web user interface task, use this option to specify the task ID of the task to terminate.

When terminating an SSH task, use this option to specify the process ID (PID) of the task to terminate.

You can specify **all** with this option to terminate all tasks that the user associated with the *session* is running, and then log off the user.
- help** Display the help text for this command and exit.

EXAMPLES

Terminate all of the HMC Web user interface tasks that the user associated with session **4** is running, and log off the user:

```
termtask -r webui -s 4 -t all
```

Terminate the task with process ID **11644** that the user remotely logged into the HMC via the SSH session **pts/5** is running:

```
termtask -r ssh -s pts/5 -t 11644
```

ENVIRONMENT

None

BUGS

None

AUTHOR

IBM Austin

SEE ALSO

lslogon

NAME

updhmc - update the Hardware Management Console

SYNOPSIS

```
updhmc -t { disk | dvd | ftp | sftp | nfs | usb }  
[-h host-name] [-u user-ID] [--passwd password] [-k SSH-private-key]  
[-l mount-location] [-o mount-command-options]  
[-f file] [-r] [-c] [-n device-name] [--help]
```

DESCRIPTION

updhmc updates the Hardware Management Console (HMC) by installing a corrective service package.

OPTIONS

-t The type of location which contains the corrective service ISO image. Valid values are **disk** for the HMC hard disk, **dvd** for a DVD drive, **ftp** for an FTP site, **sftp** for a secure FTP (SFTP) site, **nfs** for an NFS file system, or **usb** for a USB data storage device.

To install a corrective service package located on a USB data storage device, the device must be connected to the HMC and must be mounted with the **mount** command before this command is issued. The **lsmediadev** command can be used to display the mount point for the USB device.

-h The host name or IP address of the remote server where the corrective service ISO image is located.

This option is required when the ISO image is located on a remote FTP, SFTP, or NFS server. Otherwise, this option is not valid.

-u The user ID to use to log in to the remote FTP or SFTP server.

This option is required when the ISO image is located on a remote FTP or SFTP server. Otherwise, this option is not valid.

--passwd

The password to use to log in to the remote FTP or SFTP server.

If this option is omitted and the corrective service ISO image is located on a remote FTP server, you will be prompted to enter the password. If both this option and the **-k** option are omitted and the corrective service ISO image is located on a remote SFTP server, you will be prompted to enter the password. The **--passwd** and **-k** options are mutually exclusive.

This option is only valid when the ISO image is located on a remote FTP or SFTP server.

-k The name of the file that contains the SSH private key. If the file name is not fully qualified, the file must exist in the user's home directory on the HMC.

Use the **ssh-keygen** command to generate the public and private SSH key pair. The **ssh-keygen** command is not allowed to write to the **.ssh** directory in the user's home directory on the HMC, so when you run the command on the HMC, you must specify both the directory and the file name for the private key. If you generate a key with a passphrase, you will be prompted to enter the passphrase when you run any HMC command that uses the key.

If both this option and the **--passwd** option are omitted and the corrective service ISO image is located on a remote SFTP server, you will be prompted to enter the password. The **-k** and **--passwd** options are mutually exclusive.

This option is only valid when the ISO image is located on a remote SFTP server.

-l The mount location defined on the NFS server where the corrective service ISO image is located.

This option is required when the ISO image is located on a remote NFS server. Otherwise, this option is not valid.

- o The options to be passed to the **mount** command used to mount the NFS file system where the corrective service ISO image is located.

This command uses NFS version 4 by default.

This option is only valid when the ISO image is located on a remote NFS server.

- f The name of the corrective service ISO image file.

This option is required when the ISO image is located on the HMC hard disk, a remote FTP, SFTP, or NFS server, or on a USB data storage device. Otherwise, this option is not valid.

- r Specify this option to restart the HMC after applying the corrective service package.
- c Specify this option to remove the corrective service ISO image file from the HMC hard disk after applying the corrective service package.

This option is only valid when the ISO image is located on the HMC hard disk.

- n The name of the remote virtual media device where the corrective service ISO image is located. The **lsmediadev** command can be used to display the device name of a remote virtual media device.

This option is only valid when **-t dvd** is specified. This option is not required when the ISO image is located on a DVD in the local DVD drive.

- help Display the help text for this command and exit.

EXAMPLES

Install the corrective service package located on a DVD in the local DVD drive then reboot the HMC:

```
updhmc -t dvd -r
```

Install the corrective service package located on remote virtual media:

lsmediadev (to obtain the remote virtual media device name)

```
updhmc -t dvd -n /dev/sr1
```

Install a corrective service package located on an FTP server (you will be prompted to enter the user's password):

```
updhmc -t ftp -h ftpServer -u user1 -f /tmp/Update1.iso
```

Install a corrective service package located on an SFTP server and use SSH keys for authentication:

```
updhmc -t sftp -h sftpServer -u user1 -k /home/hmcuser/keys/id_rsa  
-f /tmp/Update1.iso
```

Install a corrective service package located in an NFS file system:

```
updhmc -t nfs -h 9.3.121.61 -l /hmc/images -f 830/update.iso
```

Install a corrective service package located on a USB data storage device (the USB device must already be connected to the HMC):

lsmediadev (to obtain mount points)

mount /media/sdb1

updhmc -t usb -f /media/sdb1/update2.iso

ENVIRONMENT

None

BUGS

None

AUTHOR

IBM Austin

SEE ALSO

lsmediadev

NAME

updlic - update Licensed Internal Code (LIC)

SYNOPSIS

To retrieve, install, and activate LIC updates:

```
updlic -o a { -m managed-system | -e managed-frame [--allsystems] |  
  -w }  
  [-t {sys | power | syspower | io | all}]  
-r {ibmwebsite | ftp | sftp | dvd | disk | mountpoint |  
  usb }  
-l {latest | latestconcurrent | sss | ccc,ppp |  
  release1_level1,release2_level2,... }  
  [-h host-name] [-u user-ID]  
  [--passwd password] [-d directory]  
  [--allowpoweroff] [--nopreaccept] [-q] [-v]
```

To retrieve and install, but not activate LIC updates:

```
updlic -o i { -m managed-system | -e managed-frame [--allsystems] |  
  -w }  
  [-t {sys | power | syspower | io | all}]  
-r {ibmwebsite | ftp | sftp | dvd | disk | mountpoint |  
  usb }  
-l {latest | latestconcurrent | sss | ccc,ppp |  
  release1_level1,release2_level2,... }  
  [-h host-name] [-u user-ID]  
  [--passwd password] [-d directory] [-q] [-v]
```

To remove the most recently installed LIC updates and activate the previously accepted level:

```
updlic -o r { -m managed-system | -e managed-frame [--allsystems] |  
  -w }  
  [-t {sys | power | syspower}]  
  [--allowpoweroff] [-q] [-v]
```

To change a managed system's LIC update control to management console:

```
updlic -o h -m managed-system [-v]
```

To change a managed system's LIC update control to Operating System:

```
updlic -o o -m managed-system [-v]
```

To disruptively activate LIC updates:

```
updlic -o d { -m managed-system | -e managed-frame [--allsystems] |  
  -w }  
  [-t {sys | power | syspower}]  
  [--allowpoweroff] [-q] [-v]
```

To accept currently activated LIC updates:

```
updlic -o c { -m managed-system | -e managed-frame [--allsystems] |  
  -w }  
  [-t {sys | power | syspower}] [-v]
```

To reject installed LIC updates:

```
updlic -o j { -m managed-system | -e managed-frame [--allsystems] |  
  -w }  
  [-t {sys | power | syspower}] [-v]
```

To upgrade Licensed Internal Code to a new release:

```
updlc -o u {-m managed-system | -e managed-frame [--allsystems] |  
  -w }  
  [-t {sys | power | syspower}]  
-r {ibmwebsite | ftp | sftp | dvd | disk | mountpoint |  
  usb }  
  [-l {latest | release1_level1,release2_level2,...}]  
  [-h host-name] [-u user-ID]  
  [--passwd password] [-d directory]  
  [--allowpoweroff] [-q] [-v]
```

To check system readiness:

```
updlc -o k {-m managed-system | -e managed-frame [--allsystems] |  
  -w }  
  [-v]
```

To synchronize LIC:

```
updlc -o s {-m managed-system | -e managed-frame [--allsystems] |  
  -w }  
  [-t {sys | power | syspower}] [-v]
```

To reboot the secondary service processor:

```
updlc -o b -m managed-system [-v]
```

To perform hypervisor and partition firmware activation:

```
updlc -o v -m managed-system [-v]
```

To activate Single Root I/O Virtualization (SR-IOV) adapter firmware updates:

```
updlc -o f -m managed-system -t sriov  
  --subtype {adapterdriver | adapterdriver,adapter }  
  -s SRIOV-adapter-ID [-v]
```

To purge firmware update files from the management console hard disk repository:

```
updlc -o p --ecnumber EC-number
```

DESCRIPTION

updlc updates Licensed Internal Code (LIC) on the *managed-system*, *managed-frame*, or on all managed frames which contain High Performance Switches.

OPTIONS

-o The operation to perform. Valid values are **a**, **i**, **r**, **h**, **o**, **d**, **c**, **j**, **u**, **k**, **s**, **b**, **v**, **f** and **p**.

Use the **a** operation to retrieve, install, and activate LIC updates. Previously activated updates will be automatically accepted unless the **--noproaccept** option is specified.

Use the **i** operation to retrieve and install, but not activate, LIC updates.

Use the **r** operation to remove the most recently installed LIC updates and activate the previously accepted level.

Use the **h** operation to change LIC update control to management console. This option is only valid for a managed system.

Use the **o** operation to change LIC update control to Operating System. This option is only valid for a managed system.

Use the **d** operation to disruptively activate LIC updates.

Use the **c** operation to accept currently activated LIC updates (copy T to P).

Use the **j** operation to reject installed LIC updates (copy P to T).

Use the **u** operation to upgrade Licensed Internal Code to a new release.

Use the **k** operation to check system readiness for LIC operations.

Use the **s** operation to synchronize LIC on redundant components. LIC will be synchronized from the primary Service Processor to the secondary Service Processor, and from the Side-A Bulk Power Controller to the Side-B Bulk Power Controller.

Use the **b** operation to reboot the secondary service processor. This option is only valid for a managed system.

Use the **v** operation to perform hypervisor and partition firmware activation. Managed System LIC is not activated by this option. This option is only valid for a managed system.

Use the **f** operation to activate the firmware updates for an SR-IOV adapter. The SR-IOV adapter must be configured in shared mode. The updates that are activated are the updates that were installed with Managed System LIC updates.

Use the **p** operation to purge firmware update files from the management console hard disk repository.

-m The name of the managed system on which to update LIC. The name may either be the user-defined name for the managed system, or be in the form *tttt-mmm*ssssss*, where *tttt* is the machine type, *mmm* is the model, and *ssssss* is the serial number of the managed system. The *tttt-mmm*ssssss* form must be used if there are multiple managed systems with the same user-defined name.

-e The name of the managed frame on which to update LIC. The name may either be the user-defined name for the managed frame, or be in the form *tttt-mmm*ssssss*, where *tttt* is the machine type, *mmm* is the model, and *ssssss* is the serial number of the managed frame. The *tttt-mmm*ssssss* form must be used if there are multiple managed frames with the same user-defined name.

-w Update LIC on all managed frames which contain High Performance Switches.

--allsystems

Update LIC on all managed systems in the managed frame specified with the **-e** option.

-t The type of LIC updates for which to perform the operation. Valid values are **sys** for Managed System LIC updates only, **power** for Power LIC updates only, **syspower** for both Managed System and Power LIC updates, **io** for I/O LIC updates only, **all** for Managed System, Power, and I/O LIC updates, or **sriov** for SR-IOV adapter updates.

If this option is omitted, it will default to **syspower** for the **-m** and **-e --allsystems** options, and will default to **power** for the **-e** and **-w** options.

--subtype

The type of SR-IOV adapter updates to activate. Valid values are **adapterdriver** for adapter driver firmware updates only, or **adapterdriver,adapter** for adapter driver firmware updates and adapter firmware updates.

-s The adapter ID of the SR-IOV adapter to update.

One or more SR-IOV adapter IDs can be specified. Multiple adapter IDs must be comma separated. The adapters are updated sequentially.

A temporary I/O outage occurs for several minutes for each SR-IOV adapter while it is updated. No progress messages are displayed while an adapter is updated, but a message is displayed when the update of each adapter completes if **-v** is specified.

- l** The LIC level to retrieve. Valid values are **latest**, **latestconcurrent**, *sss*, *ccc,ppp* or *release1_level1,release2_level2,...*

This option is required for **a** and **i** operations. It is optional for **u** operation; if omitted it will default to **latest**. This option is not valid otherwise.

Specify **latest** to retrieve the latest LIC updates, even if disruptive.

Specify **latestconcurrent** to retrieve the latest concurrent LIC updates.

Specify *sss* to retrieve a specific level of Managed System or Power LIC updates, even if disruptive. *sss* is the three character identifier of the specific level to retrieve. This is only valid when the LIC type is either Managed System only or Power only.

Specify *ccc,ppp* to retrieve a specific level of Managed System and Power LIC updates, even if disruptive. *ccc* is the three character identifier of the specific level of Managed System LIC updates to retrieve. *ppp* is the three character identifier of the specific level of Power LIC updates to retrieve. This is only valid when the LIC type is both Managed System and Power.

Specify *release1_level1,release2_level2,...* to retrieve specific levels of LIC updates, even if disruptive. The level specified in each entry indicates the desired level for all components which are running the release specified in the entry.

- r** The repository from which to retrieve the LIC updates. Valid values are **ibmwebsite** for the IBM service website, **ftp** for a remote FTP site, **sftp** for a remote secure FTP (SFTP) site, **dvd** for the DVD drive on the management console, **disk** for the internal hard disk drive on the management console, **mountpoint** for the specified mount point on the internal hard disk drive on the management console, or **usb** for a USB flash memory device.

This option is required for **a**, **i**, and **u** operations. This option is not valid otherwise.

- h** The host name or IP address of the remote FTP or SFTP server.

This option is required when retrieving LIC updates from a remote FTP or SFTP site. This option is not valid otherwise.

- u** The user ID to use to log in to the remote FTP or SFTP site.

This option is required when retrieving LIC updates from a remote FTP or SFTP site. This option is not valid otherwise.

--passwd

The password to use to log in to the remote FTP or SFTP site.

If this option is omitted when retrieving LIC updates from a remote FTP site, you will be prompted to enter the password. If both this option and the **-k** option are omitted when retrieving LIC updates from a remote SFTP site, you will be prompted to enter the password. The **--passwd** and **-k** options are mutually exclusive.

This option is only valid when retrieving LIC updates from a remote FTP or SFTP site.

- k The name of the file that contains the SSH private key. If the file name is not fully qualified, the file must exist in the user's home directory on the management console.

Use the **ssh-keygen** command to generate the public and private SSH key pair. The **ssh-keygen** command is not allowed to write to the **.ssh** directory in the user's home directory on the management console, so when you run the command on the management console, you must specify both the directory and the file name for the private key. If you generate a key with a passphrase, you will be prompted to enter the passphrase when you run any management console command that uses the key.

If both this option and the **--passwd** option are omitted when retrieving LIC updates from a remote SFTP site, you will be prompted to enter the password. The **-k** and **--passwd** options are mutually exclusive.

This option is only valid when retrieving LIC updates from a remote SFTP site.

- d The directory to use on the remote FTP or SFTP site, or the mount point location on the internal management console hard disk drive. If this option is not specified for FTP or SFTP, the **/opt/ccfw/data** directory will be used.

This option is only valid when retrieving LIC updates from a mount point or a remote FTP or SFTP site.

--allowpoweroff

Use this option to allow the management console to power off all managed systems in the same managed frame during a disruptive update of the power subsystem, if necessary. **Use this option carefully. It allows the management console to power off managed systems which were not explicitly targeted by this command.**

- q Use this option to query the concurrency status of the specified update. The update is not performed when this option is specified.

This option allows you to determine if the specified update is disruptive. If so, you can shut down operating systems or configure I/O devices offline prior to performing the specified update.

When this option is used with **-m**, one of the following return codes will be returned:

- 0 - no updates are available
- 1 - all updates are concurrent
- 2 - Managed System and Power updates are concurrent.
I/O updates are disruptive.
- 3 - Managed System and Power updates are disruptive.
I/O updates are concurrent.
- 4 - all updates are disruptive

Any other return code value indicates that an error occurred.

When this option is used with **-w**, one of the following return codes will be returned:

- 0 - no updates are available
- 1 - all updates are concurrent
- 4 - all updates are disruptive

Any other return code value indicates that an error occurred.

When this option is not used, a return code value of zero indicates success. Any other value indicates that an error occurred.

--ecnumber

The engineering change (EC) number associated with the firmware update files to purge from the management console hard disk repository.

Specify **all** to purge all firmware update files in the management console hard disk repository.

Specify *pppprrr* to purge all firmware update files in the management console hard disk repository for the specified platform *pppp* and release *rrr*.

Specify *pppprrr_III* to purge all firmware update files in the management console hard disk repository for the specified platform *pppp*, release *rrr*, and level *III*.

--nopraccept

Specify this option if you do not want the previously activated LIC updates to be automatically accepted prior to starting the current operation. If the previously activated LIC updates were previously accepted or you are installing downlevel LIC updates, this option will have no effect.

-v Specify this option to enable verbose mode for the operation. When verbose mode is enabled, progress messages are displayed. Note that some operations may not have any progress messages.

--help Display the help text for this command and exit.

EXAMPLES

To retrieve, install, and activate the latest available Managed System LIC and Power LIC updates for all managed systems in the managed frame from the IBM service website:

```
updlc -e myframe --allsystems -o a -t syspower -l latest -r ibmwebsite
```

To retrieve and install specific Managed System and Power LIC levels from a remote FTP server:

```
updlc -m mysystem -o i -t syspower -l 22,21 -r ftp  
-h ftphost -u ftpuser --passwd userpw
```

To retrieve, install, and activate the latest available Managed System LIC updates from a remote SFTP server using SSH keys for authentication:

```
updlc -m system1 -o a -t sys -l latest -r sftp -h sftpServer  
-u sftpuser1 -k /home/hmcuser1/keys/id_rsa
```

To retrieve, install and activate specific Managed System and Power LIC levels from the IBM service website:

```
updlc -m mysystem -o a -t syspower  
-l 01EH330_100,02EB330_098,02BP240_219 -r ibmwebsite
```

To retrieve, install and activate specific Managed System and Power LIC levels for all managed systems in the managed frame from the IBM service website:

```
updlc -e myframe --allsystems -o a -t syspower  
-l 01ES340_061,02EP340_052 -r ibmwebsite
```

To remove the most recently installed Managed System and Power LIC levels and activate the previous levels:

```
updlc -m 9406-570*101234A -o r -t syspower
```

To query whether the latest LIC updates from the IBM service website are concurrent or disruptive:

**updlc -m 9406-570*101234A -o a -t all -l latest
-r ibmwebsite -q**

To change LIC update control to management console:

updlc -m mysystem -o h

To change LIC update control to Operating System:

updlc -m mysystem -o o

To disruptively activate LIC updates:

updlc -m mysystem -o d

To accept the currently activated LIC updates:

updlc -m mysystem -o c

To accept the currently activated LIC updates for all managed systems in the managed frame:

updlc -e myframe --allsystems -o c

To reject installed LIC updates:

updlc -m mysystem -o j

To check system readiness on a managed system:

updlc -m mysystem -o k

To retrieve and install the latest concurrent LIC updates for all managed frames which contain High Performance Switches from a remote FTP server:

**updlc -w -o i -l latestconcurrent -r ftp -h ftphost -u ftpuser
--passwd userpw**

To activate the adapter driver firmware and adapter firmware updates for the SR-IOV adapter with adapter ID 1:

updlc -m mysystem -o f -t sriov --subtype adapterdriver,adapter -s 1

To activate only the adapter driver firmware updates for the SR-IOV adapters with adapter IDs 1 and 2:

updlc -m mysystem -o f -t sriov --subtype adapterdriver -s 1,2

To purge all firmware update files from the management console hard disk repository for the 01AM platform and the 780 release:

updlc -o p --ecnumber 01AM780

ENVIRONMENT

None

BUGS

None

AUTHOR

IBM Austin

SEE ALSO

lslic

NAME

`viosvrcmd` - run Virtual I/O Server command

SYNOPSIS

```
viosvrcmd -m managed-system {-p partition-name | --id partition-ID}  
-c "command" [--admin] [--help]
```

DESCRIPTION

viosvrcmd runs an I/O server command line interface (`ioscli`) command on a Virtual I/O Server (VIOS).

The `ioscli` commands are passed from the Hardware Management Console (HMC) to the VIOS over an RMC session. RMC does not allow interactive execution of `ioscli` commands.

OPTIONS

-m The name of the managed system which has the VIOS on which to run the command. The name may either be the user-defined name for the managed system, or be in the form `ttt-mmm*ssssss`, where `ttt` is the machine type, `mmm` is the model, and `ssssss` is the serial number of the managed system. The `ttt-mmm*ssssss` form must be used if there are multiple managed systems with the same user-defined name.

-p The name of the VIOS partition on which to run the command.

You must either use this option to specify the name of the partition, or use the **--id** option to specify the partition's ID. The **-p** and the **--id** options are mutually exclusive.

--id The ID of the VIOS partition on which to run the command.

You must either use this option to specify the ID of the partition, or use the **-p** option to specify the partition's name. The **--id** and the **-p** options are mutually exclusive.

-c The I/O server command line interface (`ioscli`) command to run on the VIOS. *command* must be enclosed in double quotes.

--admin

Specify this option to run the command as root on the VIOS. The HMC will run `oem_setup_env` immediately before running *command*.

Only HMC users with the **ViosAdminOp** task in their task roles are authorized to specify the **--admin** option. Either the `mkacfg` or the `chacfg` command must be used to add the **ViosAdminOp** task to a task role. The **ViosAdminOp** task can only be added to custom task roles that are based on the `hmcoperator` task role. Neither the `hmcoperator` task role nor any of the other pre-defined task roles contain the **ViosAdminOp** task.

--help Display the help text for this command and exit.

EXAMPLES

Run the `ioslevel` command on VIOS `vios_lpar`:

```
viosvrcmd -m mySystem -p vios_lpar -c "ioslevel"
```

Run the `lsdev -virtual` command on the VIOS that has a partition ID of 4:

```
viosvrcmd -m 9406-570*A0001234 --id 4 -c "lsdev -virtual"
```

ENVIRONMENT

None

BUGS

None

