

NAME

ctladm - CAM Target Layer control utility

SYNOPSIS

```

ctladm <command> [lun] [generic args] [command args]
ctladm tur <lun> [general options]
ctladm inquiry <lun> [general options]
ctladm reqsense <lun> [general options]
ctladm reportluns <lun> [general options]
ctladm read <lun> [general options] <-l lba> <-d datalen> <-f file/-> <-b blocksize_bytes> [-c cdbsize]
    [-N]
ctladm write <lun> [general options] <-l lba> <-d datalen> <-f file/-> <-b blocksize_bytes> [-c cdbsize]
    [-N]
ctladm readcap <lun> [general options] [-c cdbsize]
ctladm modesense <lun> <-m page | -l> [-P pc] [-d] [-S subpage] [-c size]
ctladm start <lun> [general options] [-i] [-o]
ctladm stop <lun> [general options] [-i] [-o]
ctladm synccache <lun> [general options] [-l lba] [-b blockcount] [-r] [-i] [-c cdbsize]
ctladm lunlist
ctladm delay <lun> <-l datamove/done> <-t secs> [-T oneshot/cont]
ctladm inject <-i action> <-p pattern> [-r lba,len] [-s len fmt [args]] [-c] [-d delete_id]
ctladm create <-b backend> [-B blocksize] [-d device_id] [-l lun_id] [-o name=value] [-s size_bytes]
    [-S serial_num] [-t device_type]
ctladm remove <-b backend> <-l lun_id> [-o name=value]
ctladm modify <-b backend> <-l lun_id> [-o name=value] <-s size_bytes>
ctladm devlist [-b backend] [-v] [-x]
ctladm port [-c] [-o on/off] [-w wwpn] [-W wwnn] [-O pp/vp] [-p targ_port] [-r targ_port] [-t fe_type]
ctladm portlist [-f frontend] [-i] [-l] [-p targ_port] [-q] [-v] [-x]
ctladm lunmap <-p targ_port> [-l pLUN] [-L cLUN]
ctladm dumppoa
ctladm dumpstructs
ctladm islist [-v] [-x]
ctladm islogout <-a | -c connection-id | -i name | -p portal>
ctladm isterminate <-a | -c connection-id | -i name | -p portal>
ctladm help

```

DESCRIPTION

The **ctladm** utility is designed to provide a way to access and control the CAM Target Layer (CTL). It provides a way to send SCSI commands to the CTL layer, and also provides some meta-commands that utilize SCSI commands. (For instance, the **lunlist** command is implemented using the SCSI REPORT

LUNS and INQUIRY commands.)

The **ctladm** utility has a number of primary functions, many of which require a device identifier. The device identifier takes the following form:

lun Specify the LUN number to operate on.

Many of the primary functions of the **ctladm** utility take the following optional arguments:

-C retries Specify the number of times to retry a command in the event of failure.

-D device Specify the device to open. This allows opening a device other than the default device, */dev/cam/ctl*, to be opened for sending commands.

-I id Specify the initiator number to use. By default, **ctladm** will use 7 as the initiator number.

Primary commands:

tur Send the SCSI TEST UNIT READY command to the device and report whether or not it is ready.

inquiry Send the SCSI INQUIRY command to the device and display some of the returned inquiry data.

reqsense Send the SCSI REQUEST SENSE command to the device and display the returned sense information.

reportluns Send the SCSI REPORT LUNS command to the device and display supported LUNs.

read Send a SCSI READ command to the device, and write the requested data to a file or stdout.

-l lba Specify the starting Logical Block Address for the READ. This can be specified in decimal, octal (starting with 0), hexadecimal (starting with 0x) or any other base supported by `strtoull(3)`.

-d datalen Specify the length, in 512 byte blocks, of the READ request.

-f file Specify the destination for the data read by the READ command. Either a filename or '-' for stdout may be specified.

- c *cdbsize*** Specify the minimum SCSI CDB (Command Data Block) size to be used for the READ request. Allowable values are 6, 10, 12 and 16. Depending upon the LBA and amount of data requested, a larger CDB size may be used to satisfy the request. (e.g., for LBAs above 0xffffffff, READ(16) must be used to satisfy the request.)
- b *blocksize*** Specify the blocksize of the underlying SCSI device, so the transfer length can be calculated accurately. The blocksize can be obtained via the SCSI READ CAPACITY command.
- N** Do not copy data to **ctladm** from the kernel when doing a read, just execute the command without copying data. This is to be used for performance testing.

write

Read data from a file or stdin, and write the data to the device using the SCSI WRITE command.

- l *lba*** Specify the starting Logical Block Address for the WRITE. This can be specified in decimal, octal (starting with 0), hexadecimal (starting with 0x) or any other base supported by strtoull(3).
- d *atalen*** Specify the length, in 512 byte blocks, of the WRITE request.
- f *file*** Specify the source for the data to be written by the WRITE command. Either a filename or '-' for stdin may be specified.
- c *cdbsize*** Specify the minimum SCSI CDB (Command Data Block) size to be used for the READ request. Allowable values are 6, 10, 12 and 16. Depending upon the LBA and amount of data requested, a larger CDB size may be used to satisfy the request. (e.g., for LBAs above 0xffffffff, READ(16) must be used to satisfy the request.)
- b *blocksize*** Specify the blocksize of the underlying SCSI device, so the transfer length can be calculated accurately. The blocksize can be obtained via the SCSI READ CAPACITY command.
- N** Do not copy data to **ctladm** to the kernel when doing a write, just execute the command without copying data. This is to be used for performance testing.

- readcap** Send the SCSI READ CAPACITY command to the device and display the device size and device block size. By default, READ CAPACITY(10) is used. If the device returns a maximum LBA of 0xffffffff, however, **ctladm** will automatically issue a READ CAPACITY(16), which is implemented as a service action of the SERVICE ACTION IN(16) opcode. The user can specify the minimum CDB size with the **-c** argument. Valid values for the **-c** option are 10 and 16. If a 10 byte CDB is specified, the request will be automatically reissued with a 16 byte CDB if the maximum LBA returned is 0xffffffff.
- modesense** Send a SCSI MODE SENSE command to the device, and display the requested mode page(s) or page list.
- m page** Specify the mode page to display. This option and the **-l** option are mutually exclusive. One of the two must be specified, though. Mode page numbers may be specified in decimal or hexadecimal.
 - l** Request that the list of mode pages supported by the device be returned. This option and the **-m** option are mutually exclusive. One of the two must be specified, though.
 - P pc** Specify the mode page control value. Possible values are:
 - 0 Current values.
 - 1 Changeable value bitmask.
 - 2 Default values.
 - 3 Saved values.
 - d** Disable block descriptors when sending the mode sense request.
 - S subpage** Specify the subpage used with the mode sense request.
 - c cdbsize** Specify the CDB size used for the mode sense request. Supported values are 6 and 10.
- start** Send the SCSI START STOP UNIT command to the specified LUN with the start bit set.
- i** Set the immediate bit in the CDB. Note that CTL does not support the immediate bit, so this is primarily useful for making sure that CTL returns the proper error.
- stop** Send the SCSI START STOP UNIT command to the specified LUN with the start bit cleared. We use an ordered tag to stop the LUN, so we can guarantee that all pending I/O

executes before it is stopped. (CTL guarantees this anyway, but **ctladm** sends an ordered tag for completeness.)

- i** Set the immediate bit in the CDB. Note that CTL does not support the immediate bit, so this is primarily useful for making sure that CTL returns the proper error.

synccache Send the SCSI SYNCHRONIZE CACHE command to the device. By default, SYNCHRONIZE CACHE(10) is used. If the specified starting LBA is greater than 0xffffffff or the length is greater than 0xffff, though, SYNCHRONIZE CACHE(16) will be used. The 16 byte command will also be used if the user specifies a 16 byte CDB with the **-c** argument.

- l lba** Specify the starting LBA of the cache region to synchronize. This option is a no-op for CTL. If you send a SYNCHRONIZE CACHE command, it will sync the cache for the entire LUN.
- b blockcount** Specify the length of the cache region to synchronize. This option is a no-op for CTL. If you send a SYNCHRONIZE CACHE command, it will sync the cache for the entire LUN.
- r** Specify relative addressing for the starting LBA. CTL does not support relative addressing, since it only works for linked commands, and CTL does not support linked commands.
- i** Tell the target to return status immediately after issuing the SYNCHRONIZE CACHE command rather than waiting for the cache to finish syncing. CTL does not support this bit.
- c cdbsize** Specify the minimum CDB size. Valid values are 10 and 16 bytes.

lunlist List all LUNs registered with CTL. Because this command uses the ioctl port, it will only work when the FETDs (Front End Target Drivers) are enabled. This command is the equivalent of doing a REPORT LUNS on one LUN and then an INQUIRY on each LUN in the system.

delay Delay commands at the given location. There are two places where commands may be delayed currently: before data is transferred ("datamove") and just prior to sending status to the host ("done"). One of the two must be supplied as an argument to the **-l** option. The **-t** option must also be specified.

- l *delayloc*** Delay command(s) at the specified location. This can either be at the data movement stage (*datamove*) or prior to command completion (*done*).
- t *delaytime*** Delay command(s) for the specified number of seconds. This must be specified. If set to 0, it will clear out any previously set delay for this particular location (*datamove* or *done*).
- T *delaytype*** Specify the delay type. By default, the **delay** option will delay the next command sent to the given LUN. With the **-T *cont*** option, every command will be delayed by the specified period of time. With the **-T *oneshot*** the next command sent to the given LUN will be delayed and all subsequent commands will be completed normally. This is the default.

inject

Inject the specified type of error for the LUN specified, when a command that matches the given pattern is seen. The sense data returned is in either fixed or descriptor format, depending upon the status of the D_SENSE bit in the control mode page (page 0xa) for the LUN.

Errors are only injected for commands that have not already failed for other reasons. By default, only the first command matching the pattern specified is returned with the supplied error.

If the **-c** flag is specified, all commands matching the pattern will be returned with the specified error until the error injection command is deleted with **-d** flag.

- i *action*** Specify the error to return:
 - aborted** Return the next matching command on the specified LUN with the sense key ABORTED COMMAND (0x0b), and the ASC/ASCQ 0x45,0x00 ("Select or reselect failure").
 - mediumerr** Return the next matching command on the specified LUN with the sense key MEDIUM ERROR (0x03) and the ASC/ASCQ 0x11,0x00 ("Unrecovered read error") for reads, or ASC/ASCQ 0x0c,0x02 ("Write error - auto reallocation failed") for write errors.
 - ua** Return the next matching command on the specified

LUN with the sense key UNIT ATTENTION (0x06) and the ASC/ASCQ 0x29,0x00 ("POWER ON, RESET, OR BUS DEVICE RESET OCCURRED").

custom Return the next matching command on the specified LUN with the supplied sense data. The **-s** argument must be specified.

-p pattern

Specify which commands should be returned with the given error.

read The error should apply to READ(6), READ(10), READ(12), READ(16), etc.

write The error should apply to WRITE(6), WRITE(10), WRITE(12), WRITE(16), WRITE AND VERIFY(10), etc.

rw The error should apply to both read and write type commands.

readcap The error should apply to READ CAPACITY(10) and READ CAPACITY(16) commands.

tur The error should apply to TEST UNIT READY commands.

any The error should apply to any command.

-r lba,len

Specify the starting lba and length of the range of LBAs which should trigger an error. This option is only applies when read and/or write patterns are specified. If used with other command types, the error will never be triggered.

-s len fmt [args]

Specify the sense data that is to be returned for custom actions. If the format is '-', len bytes of sense data will be read from standard input and written to the sense buffer. If len is longer than 252 bytes (the maximum allowable SCSI sense data length), it will be truncated to that length. The sense data format is described in cam_cdbparse(3).

-c The error injection should be persistent, instead of happening once. Persistent errors must be deleted with the **-d** argument.

-d *delete_id* Delete the specified error injection serial number. The serial number is returned when the error is injected.

port

Perform one of several CTL frontend port operations. Either get a list of frontend ports (**-l**), turn one or more frontends on or off (**-o** *on/off*), or set the World Wide Node Name (**-w** *wwnn*) or World Wide Port Name (**-W** *wwpn*) for a given port. One of **-l**, **-o**, or **-w** or **-W** must be specified. The WWNN and WWPn may both be specified at the same time, but cannot be combined with enabling/disabling or listing ports.

-c Create new frontend port using free *pp* and *vp*=0.

-o *on/off* Turn the specified CTL frontend ports on or off. If no port number or port type is specified, all ports are turned on or off.

-O *pp/vp* Specify generic options on the *ioctl* frontend port. At present, only *pp* and *vp* port numbers can be set.

-p *targ_port* Specify the frontend port number. The port numbers can be found in the frontend port list.

-r Remove port specified with (**-p** *targ_port*).

-t *fe_type* Specify the frontend type. Currently defined port types are "fc" (Fibre Channel), "scsi" (Parallel SCSI), "ioctl" (CTL *ioctl* interface), and "internal" (CTL CAM SIM).

-w *wwnn* Set the World Wide Node Name for the given port. The **-n** argument must be specified, since this is only possible to implement on a single port. As a general rule, the WWNN should be the same across all ports on the system.

-W *wwpn* Set the World Wide Port Name for the given port. The **-n** argument must be specified, since this is only possible to implement on a single port. As a general rule, the WWPn must be different for every port in the system.

portlist

List CTL frontend ports.

- f** *frontend* Specify the frontend type.
- i** Report target and connected initiators addresses.
- l** Report LUN mapping.
- p** *targ_port* Specify the frontend port number.
- q** Omit the header in the port list output.
- v** Enable verbose output (report all port options).
- x** Output the port list in XML format.

lunmap Change LUN mapping for specified port. If both *pLUN* and *cLUN* are specified -- LUN will be mapped. If *pLUN* is specified, but *cLUN* is not -- LUN will be unmapped. If neither *pLUN* nor *cLUN* are specified -- LUN mapping will be disabled, exposing all CTL LUNs.

- p** *targ_port* Specify the frontend port number.
- l** *pLUN* LUN number visible by specified port.
- L** *cLUN* CTL LUN number.

dumpooa Dump the OOA (Order Of Arrival) queue for each LUN registered with CTL.

dumpstructs Dump the CTL structures to the console.

create Create a new LUN. The backend must be specified, and depending upon the backend requested, some of the other options may be required. If the LUN is created successfully, the LUN configuration will be displayed. If LUN creation fails, a message will be displayed describing the failure.

- b** *backend* The **-b** flag is required. This specifies the name backend to use when creating the LUN. Examples are "ramdisk" and "block".
- B** *blocksize* Specify the blocksize of the backend in bytes.
- d** *device_id* Specify the LUN-associated string to use in the SCSI INQUIRY VPD

page 0x83 data.

- l *lun_id*** Request that a particular LUN number be assigned. If the requested LUN number is not available, the request will fail.
- o *name=value*** Specify a backend-specific name/value pair. Multiple **-o** arguments may be specified. Refer to the backend documentation for arguments that may be used.
- s *size_bytes*** Specify the size of the LUN in bytes. Some backends may allow setting the size (e.g. the ramdisk backend) and for others the size may be implicit (e.g. the block backend).
- S *serial_num*** Specify the serial number to be used in the SCSI INQUIRY VPD page 0x80 data.
- t *device_type*** Specify the numeric SCSI device type to use when creating the LUN. If this flag is not used, the type of LUN created is backend-specific. Not all LUN types are supported. Currently CTL supports Direct Access (type 0), Processor (type 3) and CD/DVD (type 5) LUNs. The backend requested may or may not support all of the LUN types that CTL supports.

remove Remove a LUN. The backend must be specified, and the LUN number must also be specified. Backend-specific options may also be specified with the **-o** flag.

- b *backend*** Specify the backend that owns the LUN to be removed. Examples are "ramdisk" and "block".
- l *lun_id*** Specify the LUN number to remove.
- o *name=value*** Specify a backend-specific name/value pair. Multiple **-o** arguments may be specified. Refer to the backend documentation for arguments that may be used.

modify Modify a LUN size. The backend, the LUN number, and the size must be specified.

- b *backend*** Specify the backend that owns the LUN to be modified. Examples are "ramdisk" and "block".

- l *lun_id*** Specify the LUN number to modify.
- o *name=value*** Specify a backend-specific name/value pair. Multiple **-o** arguments may be specified. Refer to the backend documentation for arguments that may be used.
- s *size_bytes*** Specify the size of the LUN in bytes. For the "block" backend, an "auto" keyword may be passed instead; this will make CTL use the size of backing file or device.

devlist Get a list of all configured LUNs. This also includes the LUN size and blocksize, serial number and device ID.

- b *backend*** Specify the backend. This restricts the LUN list to the named backend. Examples are "ramdisk" and "block".
- v** Be verbose. This will also display any backend-specific LUN attributes in addition to the standard per-LUN information.
- x** Dump the raw XML. The LUN list information from the kernel comes in XML format, and this option allows the display of the raw XML data. This option and the **-v** and **-b** options are mutually exclusive. If you specify **-x**, the entire LUN database is displayed in XML format.

islist Get a list of currently running iSCSI sessions. This includes initiator and target names and the unique connection IDs.

- v** Verbose mode.
- x** Dump the raw XML. The sessions list information from the kernel comes in XML format, and this option allows the display of the raw XML data.

islogout Ask the initiator to log out iSCSI sessions matching criteria.

- a** Log out all sessions.
- c** Specify connection ID.
- i** Specify initiator name.

- p** Specify initiator portal (hostname or IP address).
- isterninate** Forcibly terminate iSCSI sessions matching criteria.
- a** Terminate all sessions.
- c** Specify connection ID.
- i** Specify initiator name.
- p** Specify initiator portal (hostname or IP address).
- help** Display **ctladm** usage information.

OPTIONS

Number of additional configuration options may be specified for LUNs. Some options are global, others are backend-specific.

Global options:

- vendor* Specifies LUN vendor string up to 8 chars.
- product* Specifies LUN product string up to 16 chars.
- revision* Specifies LUN revision string up to 4 chars.
- scsiname* Specifies LUN SCSI name string.
- eui* Specifies LUN EUI-64 identifier.
- naa* Specifies LUN NAA identifier.
- uuid* Specifies LUN locally assigned RFC 4122 UUID identifier. EUI, NAA or UUID identifier should be set to UNIQUE value to allow EXTENDED COPY command access the LUN. Non-unique LUN identifiers may lead to data corruption. Some initiators may not support later introduced UUID identifiers.
- ident_info* Specified LUN identification information (string or 0x + hex).
- text_ident_info*

Specified LUN text identification information (UTF-8 string).

<i>ha_role</i>	Setting to "primary" or "secondary" overrides default role of the node in HA cluster, set by kern.camctl.ha_role sysctl.
<i>insecure_tpc</i>	Setting to "on" allows EXTENDED COPY command sent to this LUN access other LUNs on this host, not accessible otherwise. This allows to offload copying between different iSCSI targets residing on the same host in trusted environments.
<i>readcache</i>	Set to "off", disables read caching for the LUN, if supported by the backend.
<i>readonly</i>	Set to "on", blocks all media write operations to the LUN, reporting it as write protected.
<i>removable</i>	Set to "on", makes LUN removable.
<i>reordering</i>	Set to "unrestricted", allows target to process commands with SIMPLE task attribute in arbitrary order. Any data integrity exposures related to command sequence order shall be explicitly handled by the application client through the selection of appropriate commands and task attributes. The default value is "restricted". It improves data integrity, but may introduce some additional delays.
<i>serseq</i>	Set to "on" to fully serialize consecutive reads/writes. Set to "read" to fully serialize consecutive reads. Set to "soft" to slightly serialize consecutive reads. Set to "off" to allow them be issued in parallel. Parallel issue of consecutive operations may confuse logic of the backing file system, hurting performance; but it may improve performance of backing stores without prefetch/write-back.
<i>pblocksize</i>	
<i>pblockoffset</i>	Specify physical block size and offset of the device.
<i>ublocksize</i>	
<i>ublockoffset</i>	Specify UNMAP block size and offset of the device.
<i>rpm</i>	Specifies medium rotation rate of the device: 0 -- not reported, 1 -- non-rotating (SSD), >1024 -- value in revolutions per minute.
<i>formfactor</i>	Specifies nominal form factor of the device: 0 -- not reported, 1 -- 5.25", 2 -- 3.5", 3 -- 2.5", 4 -- 1.8", 5 -- less than 1.8".

temperature

reftemperature Specify current and reference (maximum) temperatures of the device.

provisioning_type

When UNMAP support is enabled, this option specifies provisioning type: "resource", "thin" or "unknown". Default value is "thin". Logical units without UNMAP support are reported as fully provisioned.

unmap

Setting to "on" or "off" controls UNMAP support for the logical unit. Default value is "on" if supported by the backend.

unmap_max_lba

unmap_max_descr

Specify maximum allowed number of LBAs and block descriptors per UNMAP command to report in Block Limits VPD page.

write_same_max_lba

Specify maximum allowed number of LBAs per WRITE SAME command to report in Block Limits VPD page.

avail-threshold

used-threshold

pool-avail-threshold

pool-used-threshold

Set per-LUN/-pool thin provisioning soft thresholds. LUN will establish UNIT ATTENTION condition if its or pool available space get below configured avail values, or its or pool used space get above configured used values. Pool thresholds are working only for ZVOL-backed LUNs.

writecache

Set to "off", disables write caching for the LUN, if supported by the backend.

Options specific for block backend:

file

Specifies file or device name to use for backing store.

num_threads Specifies number of backend threads to use for this LUN.

Options specific for ramdisk backend:

capacity Specifies capacity of backing store (maximum RAM for data). The default value is zero, that disables backing store completely, making all writes go to nowhere, while all reads return zeroes.

EXAMPLES

Send a SCSI TEST UNIT READY command to LUN 1.

```
ctladm tur 1
```

Display the list of mode pages supported by LUN 1.

```
ctladm modesense 1 -l
```

Display the saved version of the Control mode page (page 10) on LUN 0. Disable fetching block descriptors, and use a 10 byte MODE SENSE command instead of the default 6 byte command.

```
ctladm modesense 0 -m 10 -P 3 -d -c 10
```

Read the first 512 byte block from LUN 2 and dump it to the file

```
ctladm read 2 -l 0 -d 1 -b 512 -f - > foo
```

Read 10240 bytes from the file */tmp/bar* and write it to LUN 3. starting at LBA 0xff432140.

```
ctladm write 3 -l 0xff432140 -d 20 -b 512 -f /tmp/bar
```

Create a LUN with the "fake" ramdisk as a backing store. The LUN will claim to have a size of approximately 10 terabytes, while having no real data store (all written data are lost).

```
ctladm create -b ramdisk -s 10485760000000000
```

Create a thin provisioned LUN with a ramdisk as a backing store. The LUN will have maximal backing store capacity of 10 gigabytes, while reporting size of 10 terabytes,

```
ctladm create -b ramdisk -s 10T -o capacity=10G
```

Create a LUN using the block backend, specify the ZFS volume *tank/example* as the backing store, and specify the SCSI VPD page 0x80 and 0x83 serial number (**-S**) and device ID (**-d**). The size of the LUN will be derived from the size of the ZVOL.

```
ctladm create -b block -o file=/dev/zvol/tank/example -S MYSERIAL321 -d MYDEVID123
```

Use to specify generic options on ioctl frontend port, now it is only possible to set pp and/or vp port number.

```
ctladm port -c -O pp=11 -O vp=12
```

Remove specified targ_port.

```
ctladm port -r -p 4
```

Remove LUN 12, which is handled by the block backend, from the system.

```
ctladm remove -b block -l 12
```

List configured LUNs in the system, along with their backend and serial number. This works when the Front End Target Drivers are enabled or disabled.

```
ctladm devlist
```

List all LUNs in the system, along with their inquiry data and device type. This only works when the FETDs are enabled, since the commands go through the ioctl port.

```
ctladm lunlist
```

Inject a medium error on LUN 6 for every read that covers the first 512 blocks of the LUN.

```
ctladm inject 6 -i mediumerr -p read -r 0,512 -c
```

Inject a custom error on LUN 6 for the next TEST UNIT READY command only. This will result in a sense key of NOT READY (0x02), and an ASC/ASCQ of 0x04,0x02 ("Logical unit not ready, initializing command required").

```
ctladm inject 6 -i custom -p tur -s 18 "f0 0 02 s12 04 02"
```

SEE ALSO

cam(3), cam_cdbparse(3), cam(4), ctl(4), xpt(4), camcontrol(8), ctld(8), ctlstat(8)

HISTORY

The **ctladm** utility was originally written during the Winter/Spring of 2003 as an interface to CTL.

AUTHORS

Ken Merry <*ken@FreeBSD.org*>