

mhVTL

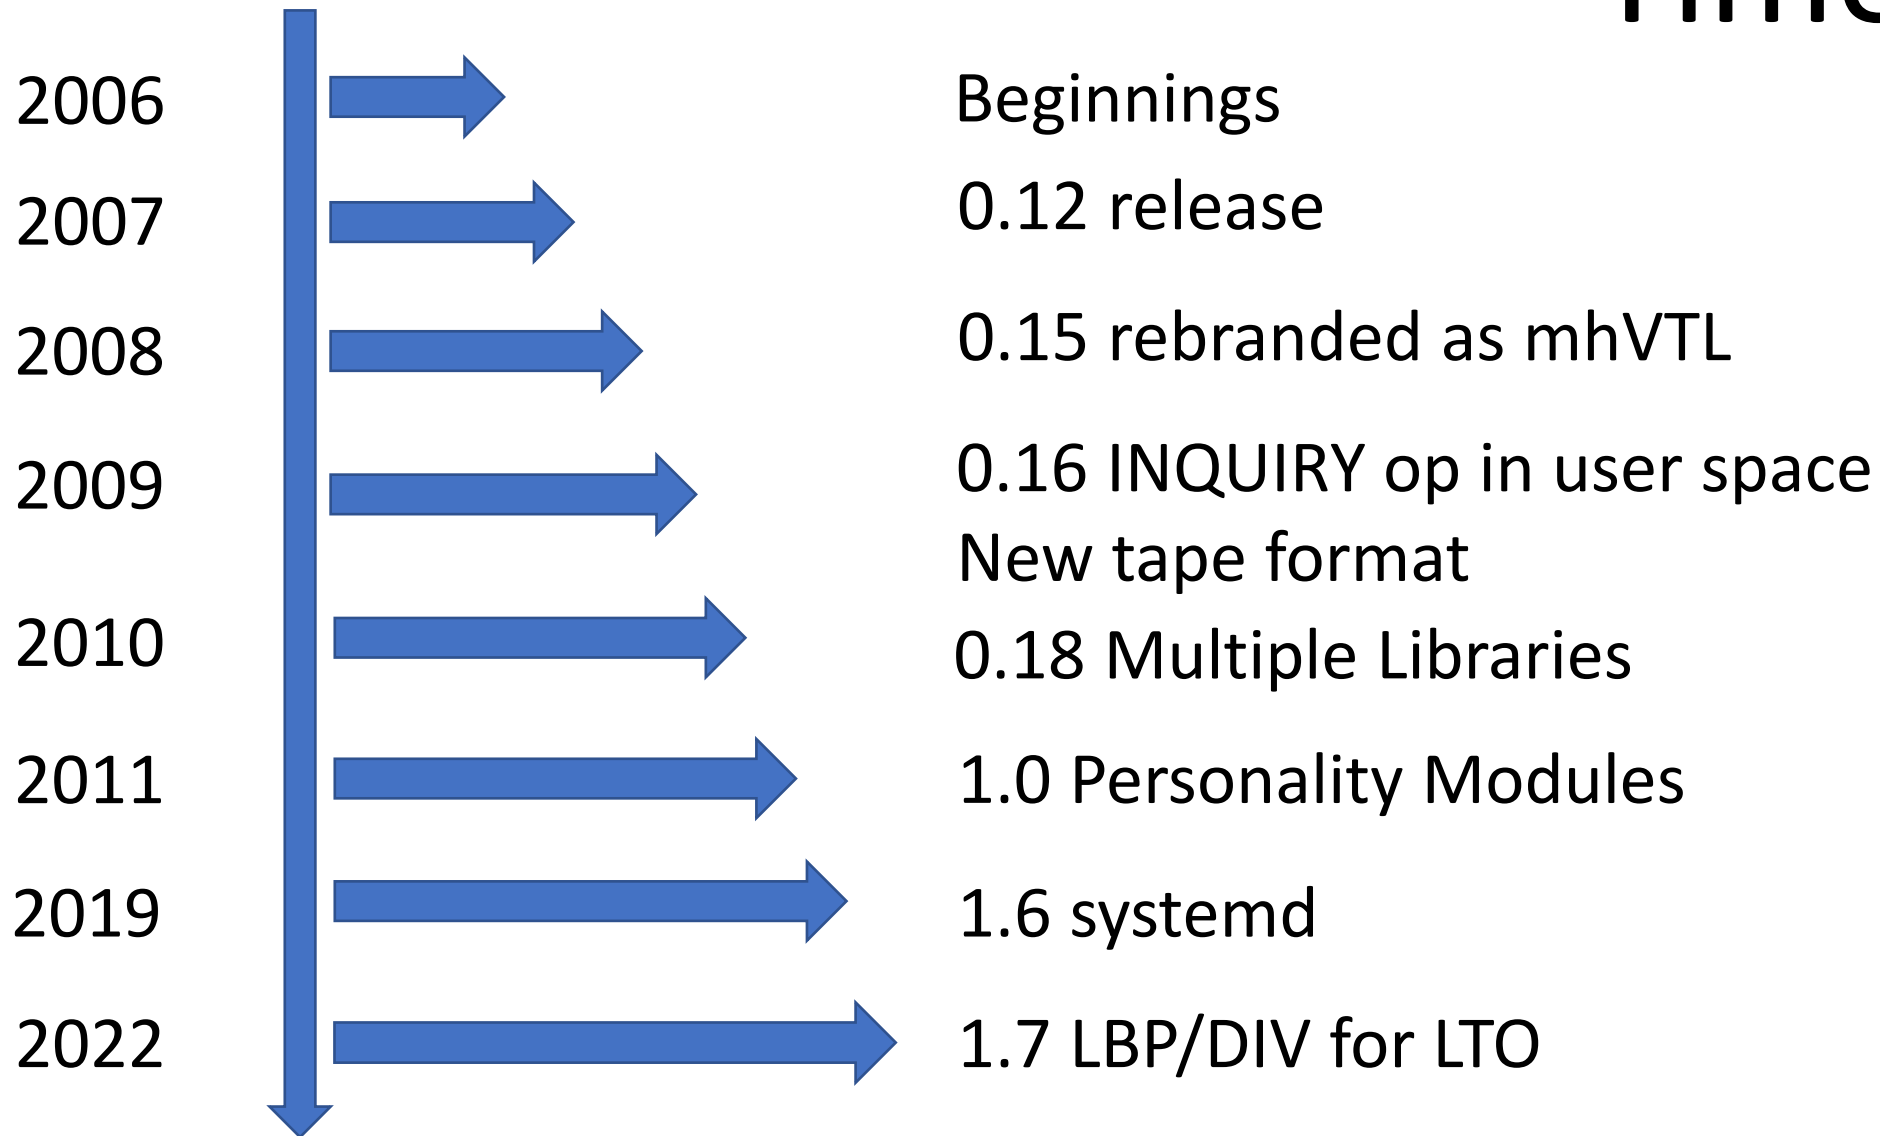
Mark Harvey

EOS 2023 workshop

inspiration



Timeline




Development lifecycle

Attempt to configure NBU



What failed ?



Implement SCSI OP code




Design Goals



Emulate real hardware

Flexibility



Target audience

Technical Support:

Education:

Sales:

Development / CI:



Major contributions


SGL:

- Multiple library support
- New virtual media format



SUSE:

- systemd



Future



Hurdles

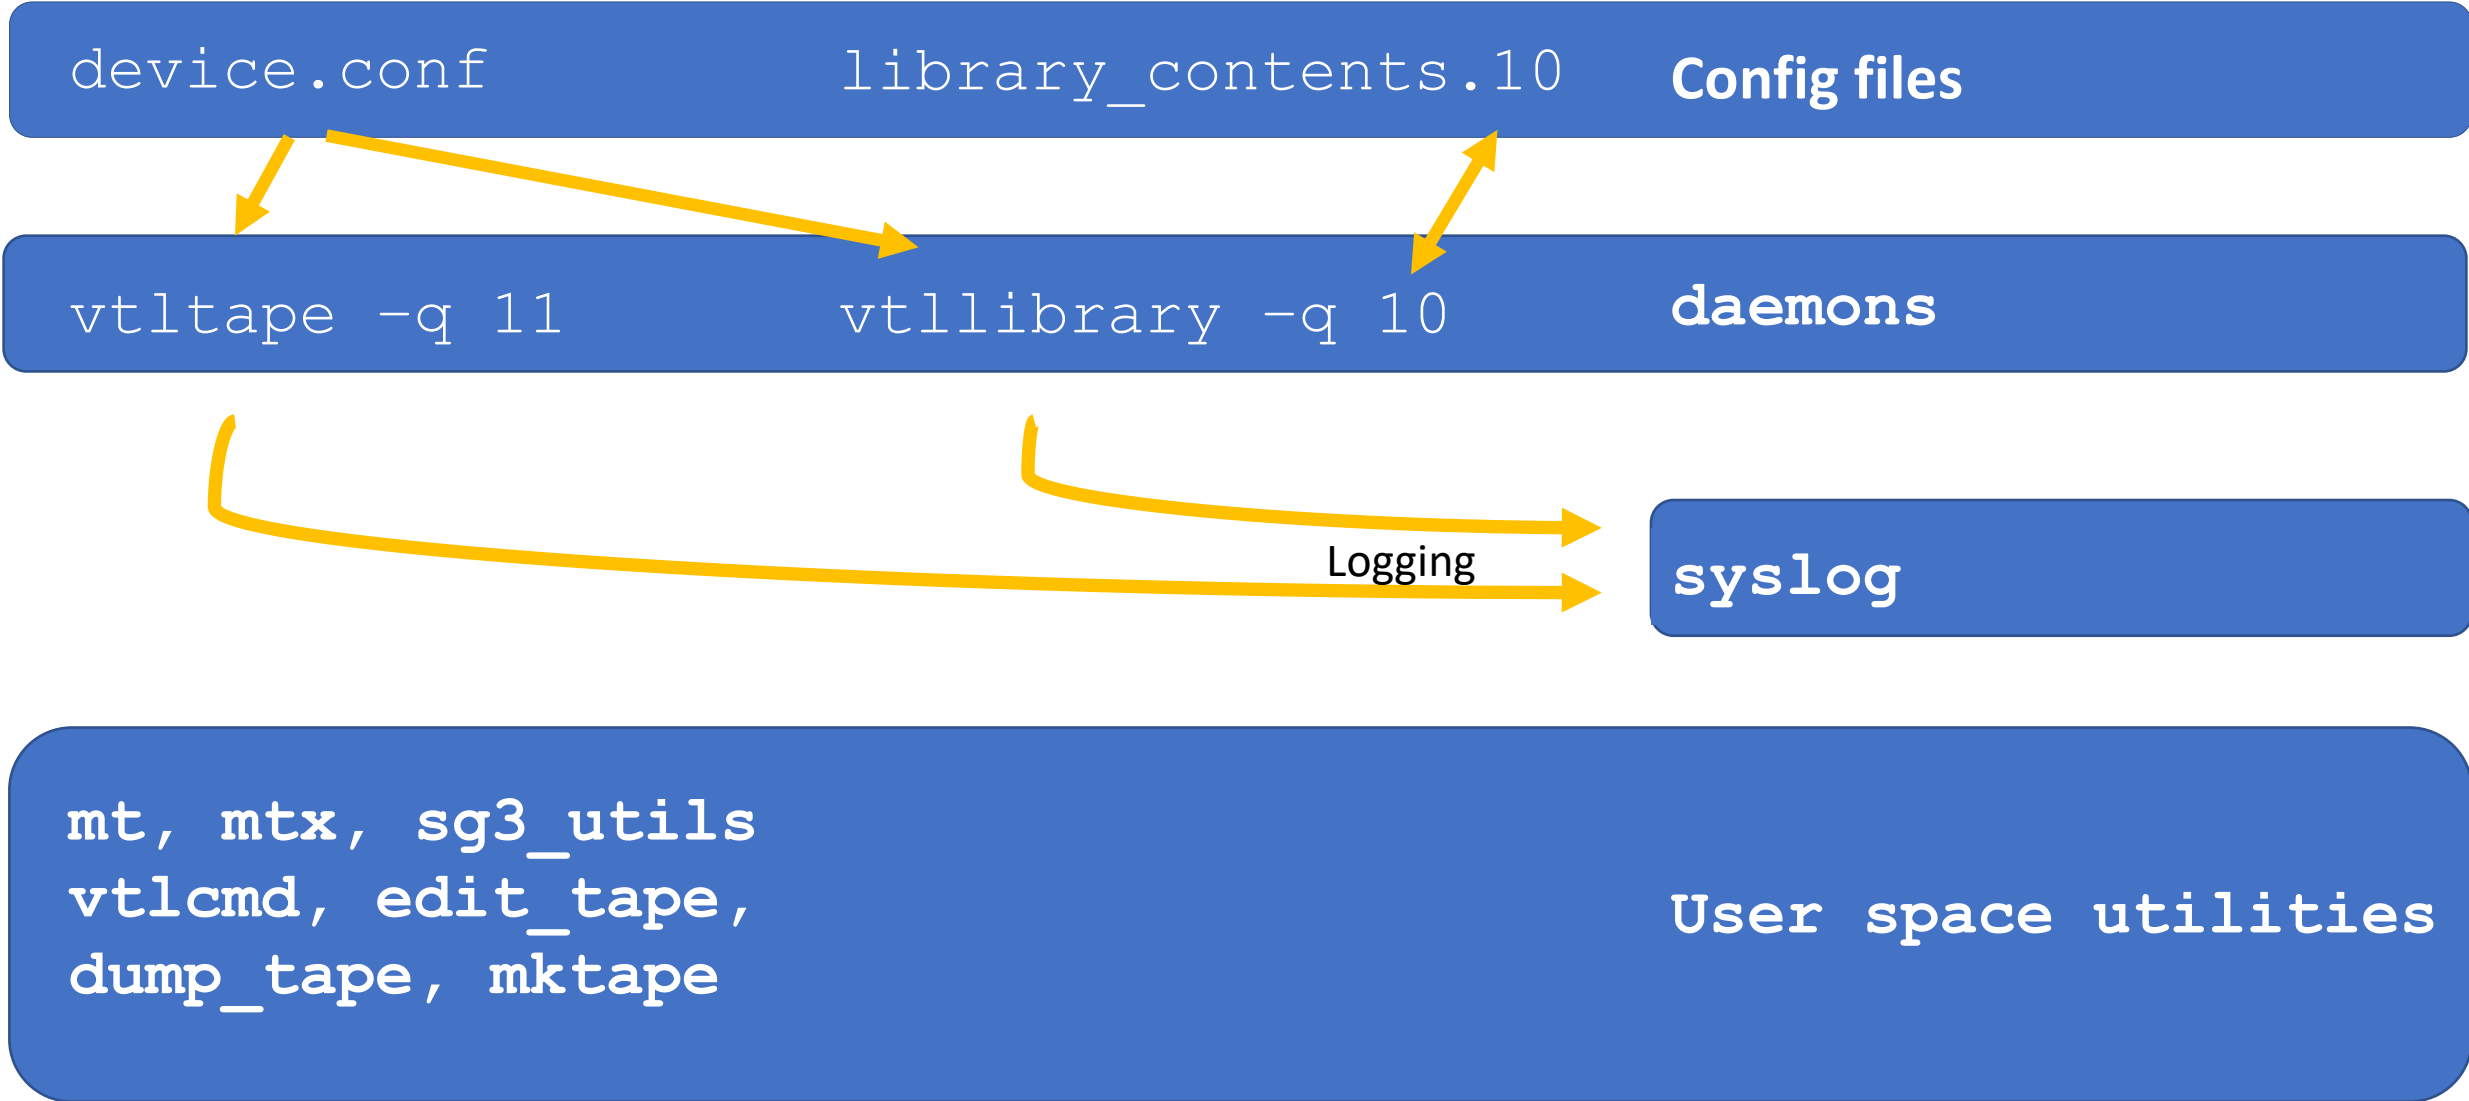




Switch Topics

Now switching from development / background
to a technical discussion / overview of the mhVTL

Architecture





/etc/mhvtl/device.conf

Device 'index' – Used as a unique identifier

```
Library: 10 CHANNEL: 00 TARGET: 00 LUN: 00
```

```
Vendor identification: STK  
Product identification: L700
```

```
Unit serial number: XYZZY_A
```

```
NAA: 10:22:33:44:ab:00:00:00
```

```
Home directory: /opt/mhvtl
```

```
PERSIST: False
```

```
Backoff: 400
```

```
# fifo: /var/tmp/mhvtl
```

```
Drive: 11 CHANNEL: 00 TARGET: 01 LUN: 00  
Library: 10 Slot: 01
```

```
Vendor identification: IBM  
Product identification: ULT3580-TD8
```

```
#Product revision level: 2160
```

```
Unit serial number: XYZZY_A1
```

```
NAA: 10:22:33:44:ab:00:01:00
```

```
Compression: factor 1 enabled 1
```

```
Compression type: lzo
```

```
Backoff: 400
```

```
# fifo: /var/tmp/mhvtl
```

Product Identifier: **Select correct 'personality module' (emulation)**

```
static struct tape_drives_table {  
    char *name;  
    void (*init)(struct lu_phy_attr *);  
} tape_drives[] = {  
    { "ULT3580-TD1",    init_ult3580_td1 },  
    { "ULT3580-TD2",    init_ult3580_td2 },  
    { "ULT3580-TD3",    init_ult3580_td3 },  
    { "ULT3580-TD4",    init_ult3580_td4 },  
    { "ULT3580-TD5",    init_ult3580_td5 },  
    { "ULT3580-TD6",    init_ult3580_td6 },  
    { "ULT3580-TD7",    init_ult3580_td7 },  
    { "ULT3580-TD8",    init_ult3580_td8 },  
    { "ULT3580-TD9",    init_ult3580_td9 },  
    { "ULTRIUM-TD1",    init_ult3580_td1 },  
    { "ULTRIUM-TD2",    init_ult3580_td2 },  
    { "ULTRIUM-HH2",    init_ult3580_td2 },  
    { "ULTRIUM-TD3",    init_ult3580_td3 },  
    { "ULTRIUM-TD4",    init_ult3580_td4 },  
    { "ULTRIUM-TD5",    init_ult3580_td5 },  
    { "ULTRIUM-TD6",    init_ult3580_td6 },  
    { "ULTRIUM-TD7",    init_ult3580_td7 },  
    { "ULTRIUM-TD8",    init_ult3580_td8 },  
    { "ULTRIUM-TD9",    init_ult3580_td9 },  
    { "ULTRIUM-TD10",   init_ult3580_td10 },  
    { "ULTRIUM-TD11",   init_ult3580_td11 },  
    { "ULTRIUM-TD12",   init_ult3580_td12 },  
    { "ULTRIUM-TD13",   init_ult3580_td13 },  
    { "ULTRIUM-TD14",   init_ult3580_td14 },  
    { "ULTRIUM-TD15",   init_ult3580_td15 },  
    { "ULTRIUM-TD16",   init_ult3580_td16 },  
    { "ULTRIUM-TD17",   init_ult3580_td17 },  
    { "ULTRIUM-TD18",   init_ult3580_td18 },  
    { "ULTRIUM-TD19",   init_ult3580_td19 },  
    { "ULTRIUM-TD20",   init_ult3580_td20 },  
    { "ULTRIUM-TD21",   init_ult3580_td21 },  
    { "ULTRIUM-TD22",   init_ult3580_td22 },  
    { "ULTRIUM-TD23",   init_ult3580_td23 },  
    { "ULTRIUM-TD24",   init_ult3580_td24 },  
    { "ULTRIUM-TD25",   init_ult3580_td25 },  
    { "ULTRIUM-TD26",   init_ult3580_td26 },  
    { "ULTRIUM-TD27",   init_ult3580_td27 },  
    { "ULTRIUM-TD28",   init_ult3580_td28 },  
    { "ULTRIUM-TD29",   init_ult3580_td29 },  
    { "ULTRIUM-TD30",   init_ult3580_td30 },  
    { "ULTRIUM-TD31",   init_ult3580_td31 },  
    { "ULTRIUM-TD32",   init_ult3580_td32 },  
    { "ULTRIUM-TD33",   init_ult3580_td33 },  
    { "ULTRIUM-TD34",   init_ult3580_td34 },  
    { "ULTRIUM-TD35",   init_ult3580_td35 },  
    { "ULTRIUM-TD36",   init_ult3580_td36 },  
    { "ULTRIUM-TD37",   init_ult3580_td37 },  
    { "ULTRIUM-TD38",   init_ult3580_td38 },  
    { "ULTRIUM-TD39",   init_ult3580_td39 },  
    { "ULTRIUM-TD40",   init_ult3580_td40 },  
    { "ULTRIUM-TD41",   init_ult3580_td41 },  
    { "ULTRIUM-TD42",   init_ult3580_td42 },  
    { "ULTRIUM-TD43",   init_ult3580_td43 },  
    { "ULTRIUM-TD44",   init_ult3580_td44 },  
    { "ULTRIUM-TD45",   init_ult3580_td45 },  
    { "ULTRIUM-TD46",   init_ult3580_td46 },  
    { "ULTRIUM-TD47",   init_ult3580_td47 },  
    { "ULTRIUM-TD48",   init_ult3580_td48 },  
    { "ULTRIUM-TD49",   init_ult3580_td49 },  
    { "ULTRIUM-TD50",   init_ult3580_td50 },  
    { "ULTRIUM-TD51",   init_ult3580_td51 },  
    { "ULTRIUM-TD52",   init_ult3580_td52 },  
    { "ULTRIUM-TD53",   init_ult3580_td53 },  
    { "ULTRIUM-TD54",   init_ult3580_td54 },  
    { "ULTRIUM-TD55",   init_ult3580_td55 },  
    { "ULTRIUM-TD56",   init_ult3580_td56 },  
    { "ULTRIUM-TD57",   init_ult3580_td57 },  
    { "ULTRIUM-TD58",   init_ult3580_td58 },  
    { "ULTRIUM-TD59",   init_ult3580_td59 },  
    { "ULTRIUM-TD60",   init_ult3580_td60 },  
    { "ULTRIUM-TD61",   init_ult3580_td61 },  
    { "ULTRIUM-TD62",   init_ult3580_td62 },  
    { "ULTRIUM-TD63",   init_ult3580_td63 },  
    { "ULTRIUM-TD64",   init_ult3580_td64 },  
    { "ULTRIUM-TD65",   init_ult3580_td65 },  
    { "ULTRIUM-TD66",   init_ult3580_td66 },  
    { "ULTRIUM-TD67",   init_ult3580_td67 },  
    { "ULTRIUM-TD68",   init_ult3580_td68 },  
    { "ULTRIUM-TD69",   init_ult3580_td69 },  
    { "ULTRIUM-TD70",   init_ult3580_td70 },  
    { "ULTRIUM-TD71",   init_ult3580_td71 },  
    { "ULTRIUM-TD72",   init_ult3580_td72 },  
    { "ULTRIUM-TD73",   init_ult3580_td73 },  
    { "ULTRIUM-TD74",   init_ult3580_td74 },  
    { "ULTRIUM-TD75",   init_ult3580_td75 },  
    { "ULTRIUM-TD76",   init_ult3580_td76 },  
    { "ULTRIUM-TD77",   init_ult3580_td77 },  
    { "ULTRIUM-TD78",   init_ult3580_td78 },  
    { "ULTRIUM-TD79",   init_ult3580_td79 },  
    { "ULTRIUM-TD80",   init_ult3580_td80 },  
    { "ULTRIUM-TD81",   init_ult3580_td81 },  
    { "ULTRIUM-TD82",   init_ult3580_td82 },  
    { "ULTRIUM-TD83",   init_ult3580_td83 },  
    { "ULTRIUM-TD84",   init_ult3580_td84 },  
    { "ULTRIUM-TD85",   init_ult3580_td85 },  
    { "ULTRIUM-TD86",   init_ult3580_td86 },  
    { "ULTRIUM-TD87",   init_ult3580_td87 },  
    { "ULTRIUM-TD88",   init_ult3580_td88 },  
    { "ULTRIUM-TD89",   init_ult3580_td89 },  
    { "ULTRIUM-TD90",   init_ult3580_td90 },  
    { "ULTRIUM-TD91",   init_ult3580_td91 },  
    { "ULTRIUM-TD92",   init_ult3580_td92 },  
    { "ULTRIUM-TD93",   init_ult3580_td93 },  
    { "ULTRIUM-TD94",   init_ult3580_td94 },  
    { "ULTRIUM-TD95",   init_ult3580_td95 },  
    { "ULTRIUM-TD96",   init_ult3580_td96 },  
    { "ULTRIUM-TD97",   init_ult3580_td97 },  
    { "ULTRIUM-TD98",   init_ult3580_td98 },  
    { "ULTRIUM-TD99",   init_ult3580_td99 },  
    { "ULTRIUM-TD100",  init_ult3580_td100 },  
};
```



mhVTL HBA (Host Bus Adapter): lsscsi -H

HBA number assigned at load time

```
[root@mhvtldev-localdomain mhvtl]# lsscsi -H
[0]   virtio_scsi
[1]   ahci
[2]   ata_piix
[3]   ata_piix
[4]   mptspi
[5]   mhvtl
```

mhVTL devices attached to HBA: lsscsi -g

Channel/Target/LUN

```
[root@mhvtldev-localdomain mhvtl]# lsscsi 5 -g
[5:0:0:0]   mediumx STK      L700      0107   /dev/sch0   /dev/sg9
[5:0:1:0]   tape   IBM      ULT3580-TD8 0107   /dev/st2    /dev/sg7
[5:0:2:0]   tape   IBM      ULT3580-TD8 0107   /dev/st1    /dev/sg6
[5:0:3:0]   tape   IBM      ULT3580-TD8 0107   /dev/st0    /dev/sg5
[5:0:4:0]   tape   IBM      ULT3580-TD8 0107   /dev/st3    /dev/sg8
```

mhVTL daemons

PID associated with each target (syslog)

```
[root@mhvtldev-localdomain mhvtl]# ps -ef | egrep "vltape|vllibrary"
root      268817   1  0 18:43 ?        00:00:00 /usr/bin/vltape -F -q12 -v1
root      268818   1  0 18:43 ?        00:00:00 /usr/bin/vltape -F -q13 -v1
root      268819   1  0 18:43 ?        00:00:00 /usr/bin/vllibrary -F -q10 -v1
root      268824   1  0 18:43 ?        00:00:00 /usr/bin/vllibrary -F -q10 -v1
root      268830   1  0 18:43 ?        00:00:00 /usr/bin/vllibrary -F -q10 -v1

markh@mhvtldev-localdomain ~$ grep 12 /etc/mhvtl/device.conf
Drive: 12 CHANNEL: 00 TARGET: 02 LUN: 00
```

Queue/index "-q <num>"



Device 'index'

```
markh@mhvtldev-localdomain ~$ sudo ipcs -q
----- Message Queues -----
key          msqid        owner         perms         used-bytes   messages
0x4d61726b  0            root          660           0             0
```

- The device index is unique
- Used as the IPC Message queue (0x4d61726b) offset.
 - Message queue used for inter-communicate between running daemon (either between each other or user->daemon):

```
# vtlcmd <index> <command>
```

 - TapeAlert, change daemon verbosity, MAP (I/E element) control etc
 - Inter-process hand-shake during library<->drive loading/unloading
- Accessing kernel driver:
 - Where the index is appended to /dev/mhvtl<>
/dev/mhvtl11 == device index 11 in device.conf

Library 10: load tape into drive index '11'

pid

Message queue ID

```
-localdomain /usr/bin/vtllibrary[365316]: CDB (216) (delay 485): a5 00 00 01 03 e8 01 f4 00 00 00 00
-localdomain /usr/bin/vtllibrary[365316]: smc_move_medium(): MOVE MEDIUM (216) **
-localdomain /usr/bin/vtllibrary[365316]: smc_move_medium(): moving from slot 1000 to slot 500 using transport 1, Invert media: no
-localdomain /usr/bin/vtllibrary[365316]: is_drive_empty(): 10: Sending "mount_state" to snd_id 11
-localdomain /usr/bin/vtltape[365321]: processMessageQ(): 11: Received message "mount_state" from snd_id 10
-localdomain /usr/bin/vtltape[365321]: send_msg_and_log(): 11: Replying to snd_id 10 with "Not occupied"
-localdomain /usr/bin/vtllibrary[365316]: is_drive_empty(): 10: Received "Not occupied" from snd_id 11
-localdomain /usr/bin/vtllibrary[365316]: move_slot2drive(): About to send cmd: 'lload E01001L8' to drive 1
-localdomain /usr/bin/vtltape[365321]: processMessageQ(): 11: Received message "lload E01001L8" from snd_id 10
-localdomain /usr/bin/vtltape[365321]: loadTape(): Media type 'LT08' loaded with S/No. : E01001L8_1678315011
-localdomain /usr/bin/vtltape[365321]: rewind_tape(): Media is writable
-localdomain /usr/bin/vtltape[365321]: return_sense(): [Key/ASC/ASCQ] [06 28 00]
-localdomain /usr/bin/vtltape[365321]: loadTape(): Media is writable
-localdomain /usr/bin/vtltape[365321]: loadTape(): Setting MediumDensityCode to LT08 (0x5e) Media type: 0x88
-localdomain /usr/bin/vtltape[365321]: processMessageQ(): pcl: "E01001L8", barcode: "E01001L8"
-localdomain /usr/bin/vtltape[365321]: send_msg_and_log(): 11: Replying to snd_id 10 with "Loaded OK: E01001L8"
-localdomain /usr/bin/vtllibrary[365316]: check_tape_load(): 10: Received "Loaded OK: E01001L8" from snd_id 11
-localdomain /usr/bin/vtllibrary[365316]: move_cart(): Setting src slot access true
```

Sense Code : Not ready to ready change, medium may have changed.

Thank you

Email : markh794@gmail.com

LinkedIn: [linkedin.com/in/markh794](https://www.linkedin.com/in/markh794)

Source : <https://github.com/markh794/mhvttl.git>

Primary web site : <https://sites.google.com/site/linuxvttl2>