



LDSTATE

Server Storage Visualization

Michael Stumpf
Curtis Quintal

2013-08-01
PowerEdge C Engineering

ldstate

- No clear way to see all storage on a Linux server
- Third-party tools can be challenging
 - 100+ commands and 10+ of minutes
- **ldstate**, with 1 command:
 - shows all storage, from all adapters
 - presents a dashboard
 - monitors health
 - offers simple, *usable* scripting that *anyone* can use

[illegible]

```

Linux CTR Encryption : No
Disable L0BMM : No
Support Breakerrior : No
Power Savings : No

Supported PD Operations
=====

Force Online : Yes
Force Offline : Yes
Force Rebuild : Yes
Any force failed : No
Any force Goodbad : No
Any Missing Replace : No
Any Clear : No
Any Locate : No
Support Temperature : No
Disable Cache : No
Disable Cooldown on SMART : No
Disable Cooldown on SSD on SMART Error : No
Disable SSD Patrol Read : No
PCI Correctly configured : Yes
Disable Spin Down of Disks : No
Disable Spin Down of Disks : No
Spin Down State : No
Counters : No

Memory Correctable Errors : 0
Memory Uncorrectable Errors : 0

Cluster Information
=====

Cluster Permitted
Cluster Active

Fault Settings
=====

Parity/Split : 0
Background Rate : 1%
Write Size : 64KB
Flush Time : 6 seconds
Eject Policy : Disabled
Head Load : Adaptive
Head Load Bad : Disabled
SMART IO : Disabled
SMART Mode : Mode 0
Alarm Disable : Yes
Garcion Drive : None
RAID Config : Unknown
Verify LED Show Drive Activity : No
HIDS Continue on Error : No

```

Cache Mem 8Kb	Disabled
Cached I/O	Yes
Cache Mem 6	Yes
Alarm Disable	Yes
Corruption Mode	Normal
Lock Config	Yes
Dirty LED Shows Drive Activity	No
RTOS Control on Error	Yes
Spin Down Mode	Normal
RTOS Service Type	ATA
Allow Mix in Enclosure	Yes
Allow 500/500/ATA Mix in V0	Yes
Allow 550/545/ATA Mix in V0	Yes
Allow 600/550 Mix in V0	No
Allow ATA Mix in Cluster	Yes
Non Chained Enclosures	16
Unusable Cyls	16
Disable Web BIOS	Yes
Force P0 Mapping	Yes
BIOS Enumerate V0s	Yes
Restore Hot Spare on Insertion	No
Remove Enclosure Devices	Yes
Maintain Po fail history	Yes
Disable Partitioning	Yes
Zero Based Enclosure Enumeration	No
Nonbootable old members	Yes
LED Show Drive Activity	Yes
Cluster Disable	Yes
ACS SSI Only	Yes
Auto Detect BackPlane Enable	SGPIO/12c SEP
Use ACS Only	Yes
Disable LED Header	No
Use LED during POST	Yes
Rebuild/draining	No
Disable Online Controller Reset	No
Rebuild/Time	No
Non-Certified Hard Disk Drives	Allow
Track Single Span K16 in K16	Yes
Max LD per array	16
Non-RAID Controller	Yes
Max power savings option is met	allows for LRS, only T18 power conditions
Default spin down time in minutes:	38
Minim 3800	No
T18 & 6 in Flash	No
Auto Enhanced Input	No
Auto Enhanced With Support	No
Disable J16 Mirror	No
Unusable Serial SAs	No
Time taken to detect OME	68s

ldstate Solves this Problem

```
[192.168.9.63 root@dcs-linpe35 ldstate]# ./ldstate info
```

Adapters:

```
onboard.0    onboard
sas2008.0    SAS9202-16e  (fw: 13.00.01.00  driver: 14.00.00.00)
sas2008.1    SAS9202-16e  (fw: 13.00.01.00  driver: 14.00.00.00)
```

Logical drives:

```
onboard.0    ld0          ok          2 drv    1000gb RAID0    [sda,sdb]
```

Physical drives:

onboard.0		ok	27C	491gb	SATA	sda	ST9500620NS	(AA07)	[9XF0VGXT]
onboard.0		ok	26C	491gb	SATA	sdb	ST9500620NS	(AA07)	[9XF0VH58]
sas2008.1	3:0	rdy	42C	600gb	SAS	sdc	ST3600057SS	(ES65)	[6SL4GRXN]
sas2008.1	3:1	rdy	41C	600gb	SAS	sdm	ST3600057SS	(ES65)	[6SL4F5HL]
sas2008.1	3:2	rdy	43C	600gb	SAS	sdl	ST3600057SS	(ES65)	[6SL4D0P9]
sas2008.1	3:3	rdy	39C	600gb	SAS	sdk	ST3600057SS	(ES65)	[6SL4EFCX]
sas2008.1	3:4	rdy	40C	600gb	SAS	sdj	ST3600057SS	(ES65)	[6SL4DNCL]
sas2008.1	3:5	rdy	37C	600gb	SAS	sdi	ST3600057SS	(ES65)	[6SL4J5BD]
sas2008.1	3:6	rdy	35C	600gb	SAS	sdh	ST3600057SS	(ES65)	[6SL4EK9W]
sas2008.1	3:7	rdy	36C	600gb	SAS	sdg	ST3600057SS	(ES65)	[6SL4F8F0]
sas2008.1	3:8	rdy	31C	600gb	SAS	sdf	ST3600057SS	(ES65)	[6SL4F5BJ]
sas2008.1	3:9	rdy	30C	600gb	SAS	sde	ST3600057SS	(ES65)	[6SL4DT4H]
sas2008.1	3:10	FAIL		0.00gb	SATA		()	[]	
sas2008.1	3:11	rdy	26C	600gb	SAS	sdd	ST3600057SS	(ES65)	[6SL4F3L8]

Full storage state: ***25.5 pages condensed to 1 page***

Easy Data Export

```
[192.168.9.63 root@dcslinpe35 ldstate]# ./ldstate info --csv
adp,onboard,0,onboard,,
adp,sas2008,0,SAS9202-16e,13.00.01.00,14.00.00.00
adp,sas2008,1,SAS9202-16e,13.00.01.00,14.00.00.00
pd,onboard,0,,,rdy,28C,491gb,SATA,,,,sda,ST9500620NS (AA07) [9XF0VGXT]
pd,onboard,0,,,rdy,28C,491gb,SATA,,,,sdb,ST9500620NS (AA07) [9XF0VH58]
pd,sas2008,1,3,0,rdy,43C,600gb,SAS,,,,sdc,ST3600057SS (ES65) [6SL4GRXN]
pd,sas2008,1,3,1,rdy,42C,600gb,SAS,,,,sdm,ST3600057SS (ES65) [6SL4F5HL]
pd,sas2008,1,3,2,rdy,44C,600gb,SAS,,,,sdl,ST3600057SS (ES65) [6SL4D0P9]
pd,sas2008,1,3,3,rdy,40C,600gb,SAS,,,,sdk,ST3600057SS (ES65) [6SL4EFCX]
pd,sas2008,1,3,4,rdy,41C,600gb,SAS,,,,sdj,ST3600057SS (ES65) [6SL4DNCL]
pd,sas2008,1,3,5,rdy,38C,600gb,SAS,,,,sdi,ST3600057SS (ES65) [6SL4J5BD]
pd,sas2008,1,3,6,rdy,36C,600gb,SAS,,,,sdh,ST3600057SS (ES65) [6SL4EK9W]
pd,sas2008,1,3,7,rdy,37C,600gb,SAS,,,,sdg,ST3600057SS (ES65) [6SL4F8F0]
pd,sas2008,1,3,8,rdy,32C,600gb,SAS,,,,sdf,ST3600057SS (ES65) [6SL4F5BJ]
pd,sas2008,1,3,9,rdy,32C,600gb,SAS,,,,sde,ST3600057SS (ES65) [6SL4DT4H]
pd,sas2008,1,3,10,FAIL,,0.00gb,SATA,,,,, () []
pd,sas2008,1,3,11,rdy,28C,600gb,SAS,,,,sdd,ST3600057SS (ES65) [6SL4F3L8]
```

CSV format, for import into Excel

Health Monitoring

```
NORMAL:    ld0:ok      pd-sda:optimal  
pd-sdb:optimal      pd-sdc:ready  
pd-sdd:ready      pd-sde:ready
```

```
WARNING:   ld0:ok      ld1:degraded  
pd-sdb:optimal      pd-sdc:ready  
pd-sde:ready      pd-sdd:optimal
```

```
CRITICAL:  ld0:ok      pd-sda:optimal  
pd-sdb:optimal      pd-sdc:ready  
pd-sdd:failed      pd-sde:ready
```

Succinct overview of storage health
Integrates with PEC agent & open source consoles

Key/Value View

Show all collected
key/value LD data

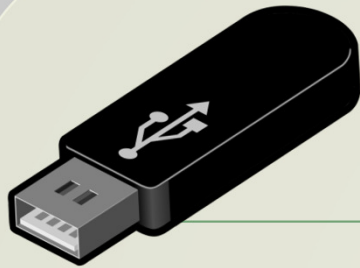
```
ldmembers : 3
  ldnum : 1
ldpdlist : 0:26,0:27,0:28
  ldsiz : 3809270
  ldstate : optimal
  ldtype : RAID5
toolname : arcconf
toolver : 1.2 (B20425)
```

```
[192.168.9.61 root@dcs-linpe35 ldstate]# ./ldstate keyval ld
LOGICAL DRIVES:
```

```
  adpnum : 1
  adptype : adaptec
  ldbootable : yes
  ldinprogress :
  ldinprogressspt :
  ldmembers : 2
  ldnum : 0
  ldpdlist : 0:24,0:25
  ldsiz : 1904630
  ldstate : optimal
  ldtype : RAID1
  toolname : arcconf
  toolver : 1.2 (B20425)
```

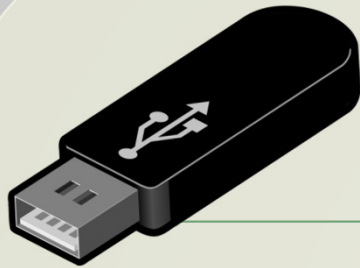
```
  adpnum : 1
  adptype : adaptec
  ldbootable : no
  ldinprogress :
  ldinprogressspt :
  ldmembers : 3
  ldnum : 1
  ldpdlist : 0:26,0:27,0:28
  ldsiz : 3809270
  ldstate : optimal
  ldtype : RAID5
  toolname : arcconf
  toolver : 1.2 (B20425)
```

```
  adpnum : 1
  adptype : adaptec
  ldbootable : no
  ldinprogress :
  ldinprogressspt :
  ldmembers : 4
  ldnum : 2
  ldpdlist : 0:29,0:30,0:31,0:32
  ldsiz : 3809270
  ldstate : optimal
  ldtype : RAID10
  toolname : arcconf
  toolver : 1.2 (B20425)
```



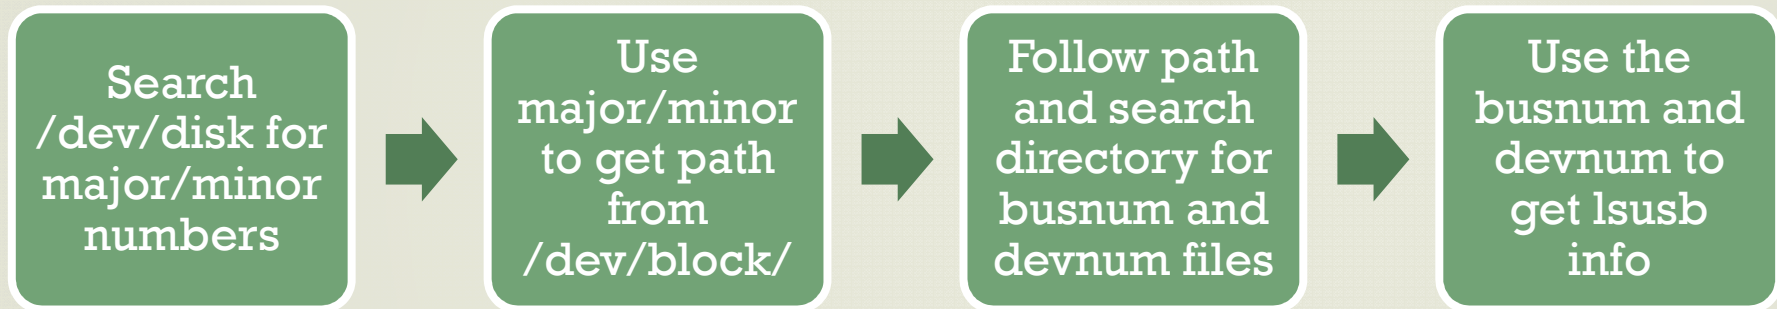
Acquiring USB Info

- Most info derived from **lsusb** (Linux kernel USB service)
- **lsusb** requires that you know bus and device numbers for the device
 - However, bus and device numbers are not clearly listed anywhere
 - Cross reference serial number, product id to find
- Overall, locating USB drives requires 8+ steps

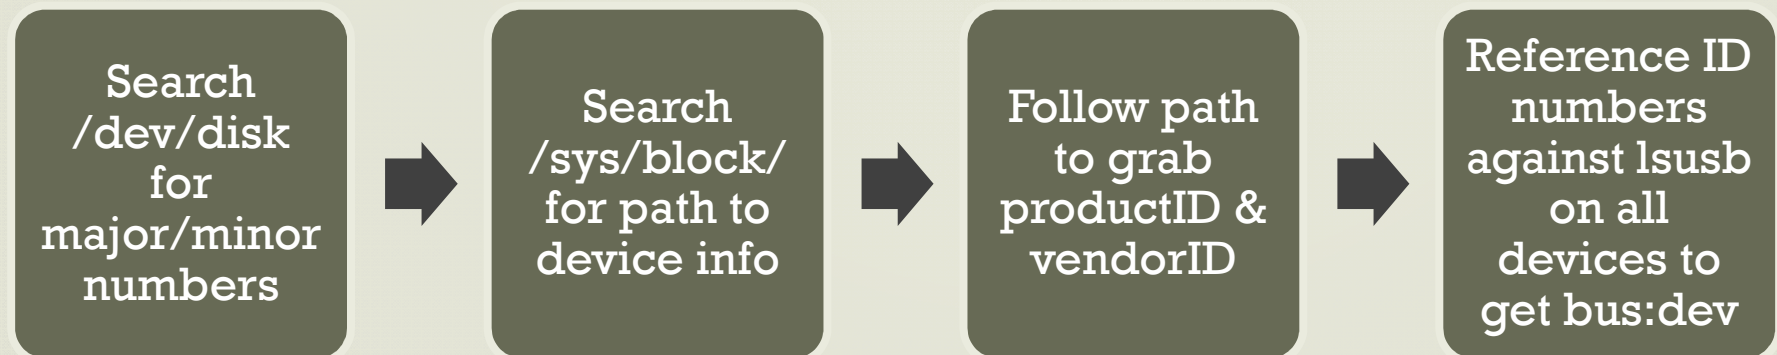


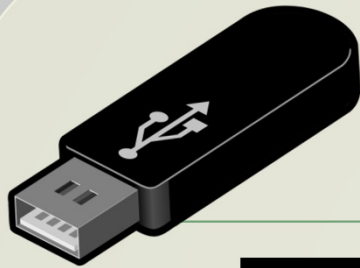
Process to get USB Info

Newer Linux OS (6.X and later)



Older Linux OS





Simple USB Info Summary

```
adpnum : 0
adptype : usb
pdmajor : 8
pdminor : 32
pdname : sdc
pdproduct : MicroVault Flash Drive
pdserial : 9B2001207030003676
pdsg : sg2
pdstate : optimal
pdsz : 4083
pdvendor : Sony Corp.
usbbusnum : 1
usbdevnum : 18
usbmaxpower : 200mA
usbtype : 2
```

A single `ldstate` command gives full USB device info

Create and Destroy RAID

- ⦿ A user trying to just get something up quickly faces **FIVE** completely different command sets:
 - Adaptec (arcconf)
 - MegaRaid (megaccli64) for RAID 0,1,5
 - MegaRaid (megaccli64) for RAID 10 / spanned
 - SAS2008 (sas2ircu)
 - Linux Software Raid (mdadm)
- ⦿ ldsta offers **ONE** common, predictive, minimal syntax

Do what the user means!

ldstate create_raid 5 all – make a big array on every drive found

ldstate create_raid 1 adaptec each – make as many mirror pairs as possible



Alarm Silence

- ◉ Adapter alarms are extremely annoying
- ◉ Locating the machine causing the noise can be very difficult
- ◉ ldstate aggregates all info about all adapters, so...



ldstate can silence all alarms *with **one** command*

Predictive Failure

```
smartctl  
pdsmtstateraw : FAILED! ←  
  pdstate : failed  
  pdstateraw : Ready (RDY) ← sas2ircu  
    pdsz : 953869  
    pdtemp : 25C  
  toolname : sas2ircu  
  toolver : 13.00.00.00
```

we use smartctl to detect drive failures...
that third-party tools cannot sense!

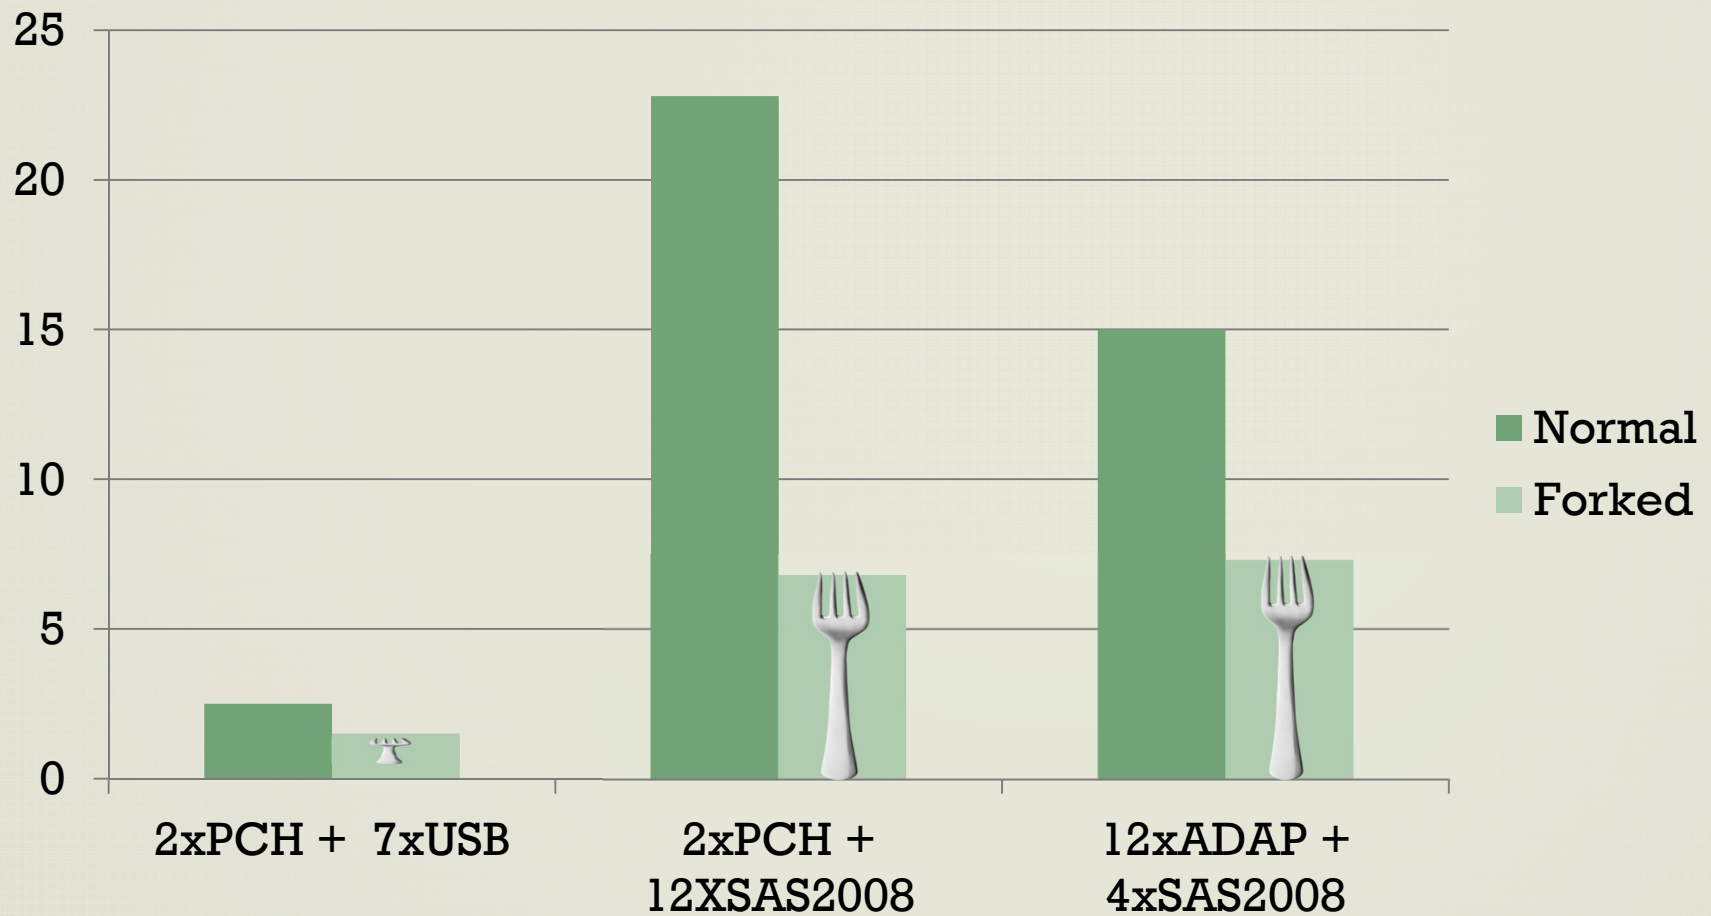
Command Forking



- Uses fork to create child process to run commands in parallel
- Reduces run times dramatically
 - -Reduced from 24s to 6s in one example.
- Allows for processing of all OS detected drives in reasonable time

-The fork is simple yet powerful tool, much like ldstate

Command Forking



Supported Devices

- ◉ LSI MegaRaid
- ◉ MegaRaid Foreign Arrays
- ◉ LSI IR/IT 2
- ◉ Adaptec (Series 6/7)
- ◉ PCH (onboard)
- ◉ Linux Software mdraid
- ◉ USB Devices
- ◉ LSI/Adaptec BBU and Supercaps
- ◉ LSI 1068e*

*Coming Soon

Thank you for your time.

Questions?