



# LDSTATE

## Server Storage Visualization

Michael Stumpf  
Curtis Quintal

2013-08-01  
PowerEdge C Engineering

# ldstate

---

- ◉ No way to see all storage on a Linux server
- ◉ Many storage tools are borderline unusable
  - inconsistent
  - to get all info, 100+ commands—then sift and sort
- ◉ **ldstate**, with 1 command:
  - shows all storage, from all adapters
  - presents a dashboard
  - monitors health
  - offers simple, *usable* scripting that *anyone* can use

## example: storage tool output

[illegible][illegible]

# 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	sd1	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*



# simple data export

```
[192.168.9.63 root@dcs-linpe35 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 ease of use

# 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
  ldsize : 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
  ldsize : 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
  ldsize : 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
  ldsize : 3809270
  ldstate : optimal
  ldtype : RAID10
  toolname : arcconf
  toolver : 1.2 (B20425)
```

# supported devices

---

- ◉ LSI MegaRAID
- ◉ MegaRAID Foreign Arrays
- ◉ LSI IR/IT (1, 2, soon 3)
- ◉ Adaptec (Series 6/7/8)
- ◉ on-board drives (PCH)
- ◉ Linux software raid (mdraid)
- ◉ USB storage devices
- ◉ LSI/Adaptec BBU and Supercaps





# acquiring USB info

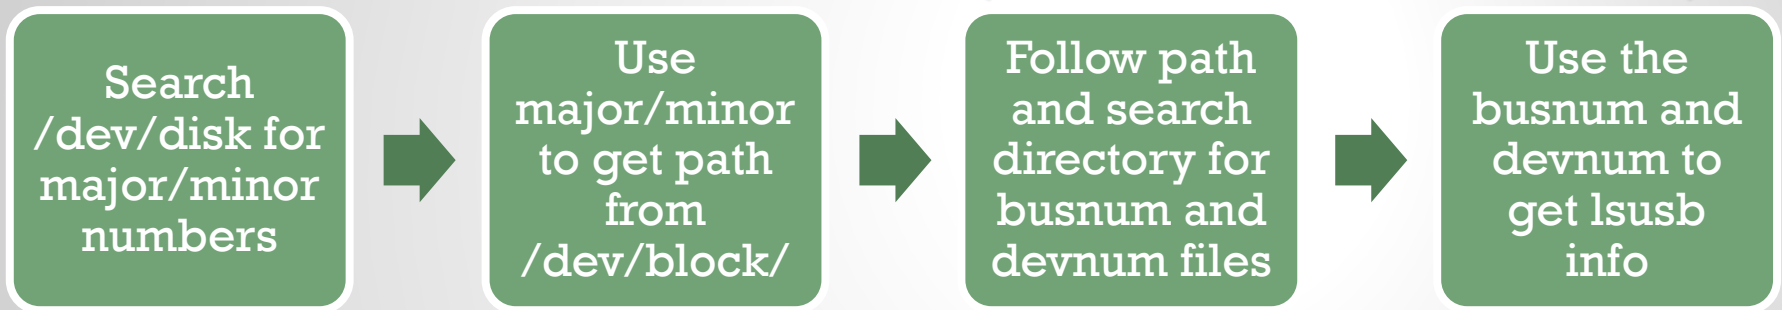
---

- 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
  - so `ldstate` cross-references serial number and product id to find
  - 8+ steps per device
- `ldstate` does this for you

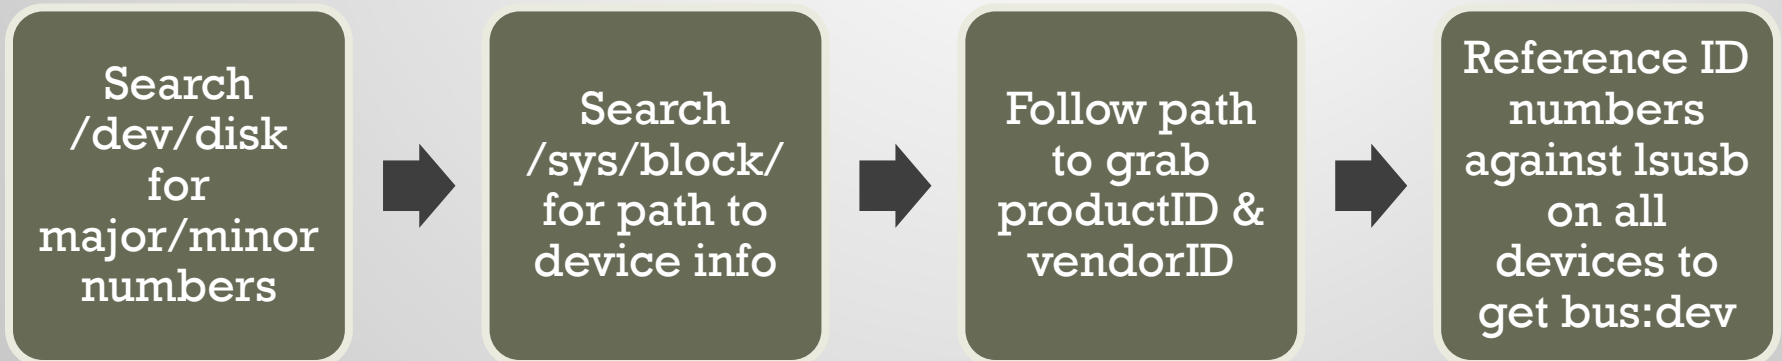


# process to get USB info

## Newer Linux OS (EL 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

# quick RAID create and destroy

---

- 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)
- ldstaate offers **ONE** common, predictive, minimal syntax

*Do what the user means!*

**ldstaate create\_raid 5 all** – make a big array on every drive found

**ldstaate create\_raid 1 adaptec each** – make as many mirror pairs as possible



# alarm silence

- ◉ Adapter alarms are extremely annoying
- ◉ and hard to pin down...
- ◉ but ldstate already has all info about all adapters, so...



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



# predictive failure

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

ldstate leverages smartctl to detect drive failures...  
this may provide a more conservative approach.

# command forking


---



- ◎ ldstate forks to run commands in parallel
- ◎ Run times dramatically reduced
  - From 24s to 6s in one example
- ◎ Crucial for lots-of-storage configs

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

# time saved via fork



---

Thank you for your time.

**Questions?**