



# LIBVIRT ADMIN API

A DIFFERENT KIND OF MANAGEMENT  
FOR LIBVIRT

Erik Skultety - Associate Software Engineer  
[eskultet@redhat.com](mailto:eskultet@redhat.com)

# WHAT IS LIBVIRT ?

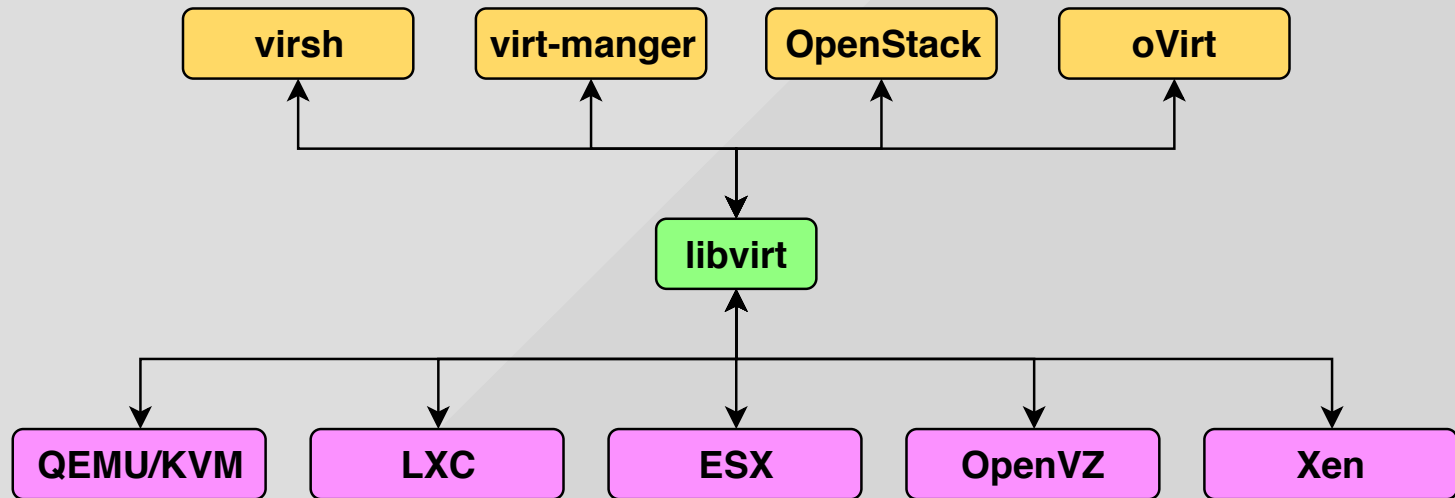
## VIRTUALIZATION TOOLSET

**TOOLSET TO UTILIZE VIRTUALIZATION CAPABILITIES OF VARIOUS HYPERVISORS**

- C library
- libvirtd daemon (backend)
- virsh CLI (frontend)

# LIBVIRT

## MIDDLEWARE LAYER



# LIBVIRT

## MIDDLEWARE LAYER

- abstract layer dealing with specifics of the underlying hypervisor
- secure, remote, management of
  - virtual machines
  - virtual networks
  - storage

# LIBVIRT

## WHY DO WE EVEN NEED IT?

- means for remote management of a VM
- store VM's state in a safe manner

## CAN IT MANAGE ITSELF?

- originally no support for runtime configuration
- *libvirtd.conf* → persistent configuration only
- restart is needed every time daemon's configuration is changed

# LIBVIRT-ADMIN

## WHAT IS IT ?

- new set of APIs
- independent library
- virt-admin binary

## WHY SHOULD I CARE ?

- live reconfiguration of daemon settings
- monitoring (threadpool monitoring, client monitoring, ...)

# LIBVIRT-ADMIN

## WHAT CAN IT DO ?

- reconfiguration of logging parameters
- threadpool management
- client management
  - listing all connected clients
  - client details
  - force disconnect of a client
  - modification of client limits

# LIBVIRT-ADMIN

## THREAD POOL MANAGEMENT

**SERVER LOAD MIGHT CAUSE ALL WORKER THREADS TO BE OCCUPIED...**

- next request to arrive would block → delay in service delivery
- identify which one of libvirtd's servers is hitting the limit on active workers

```
# virt-admin srv-list
```

```
Id      Name
```

```
-----
```

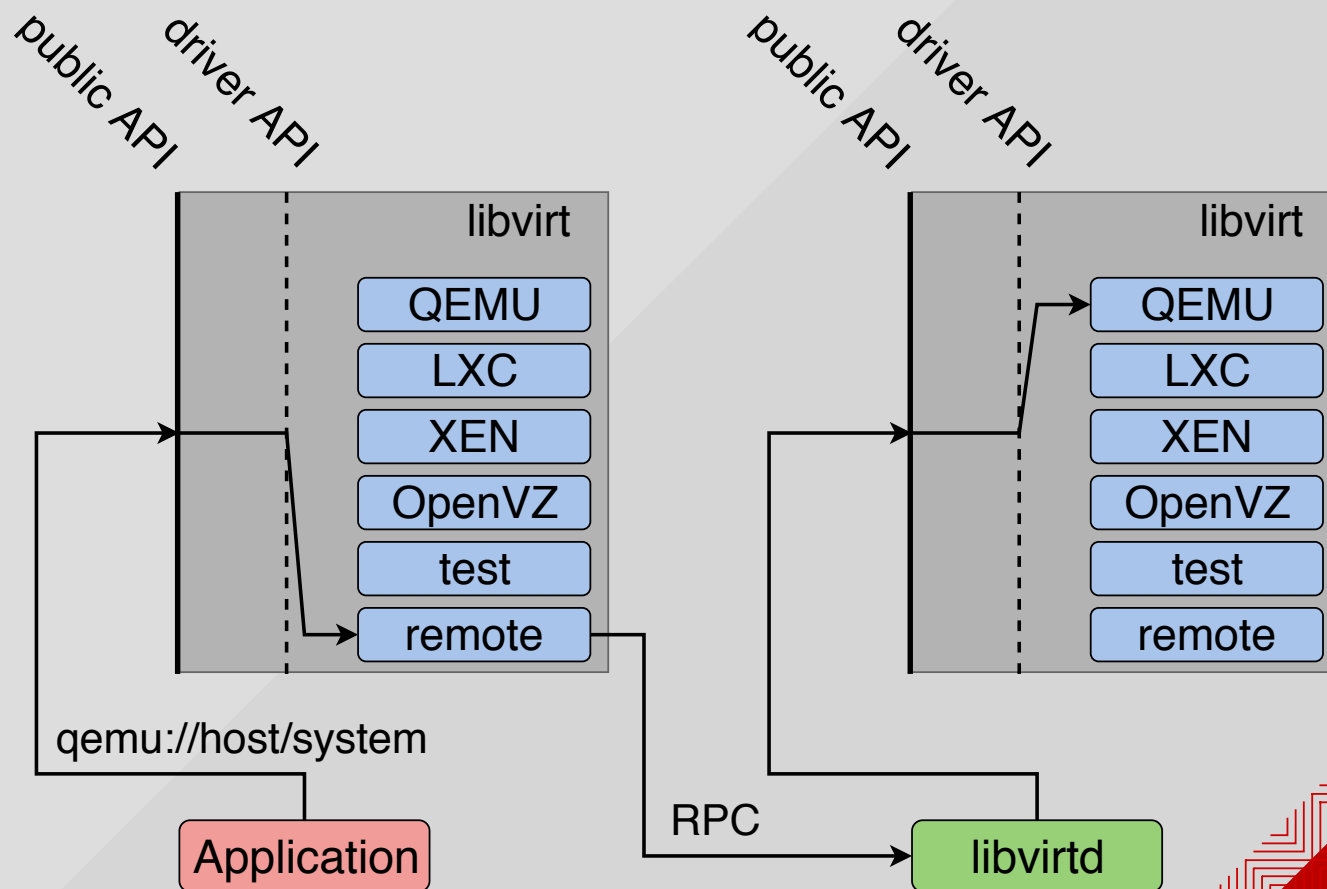
```
0       admin
```

```
1       libvirtd
```

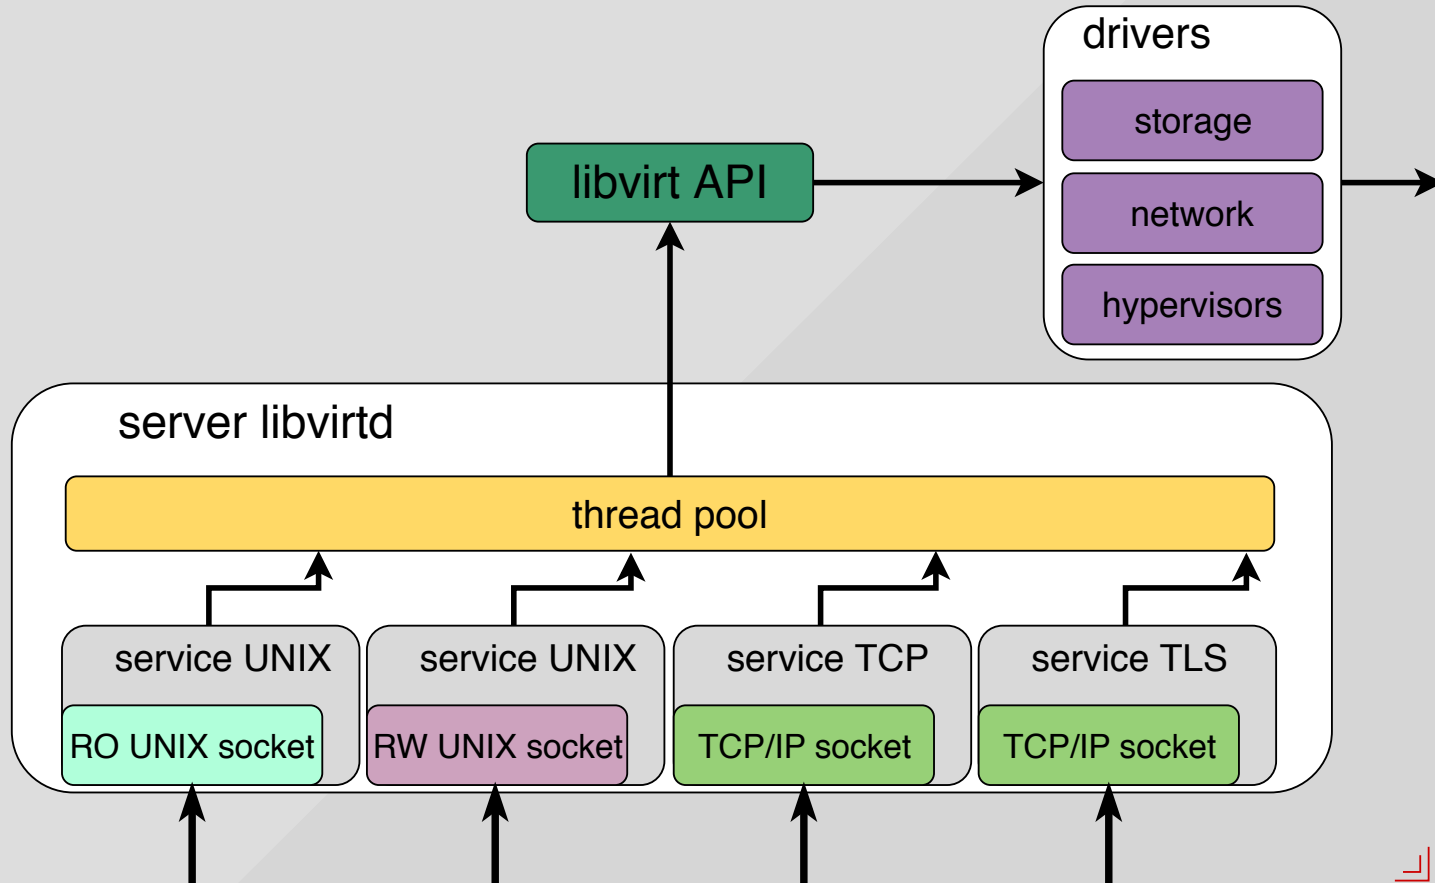


# LIBVIRT INTERNALS

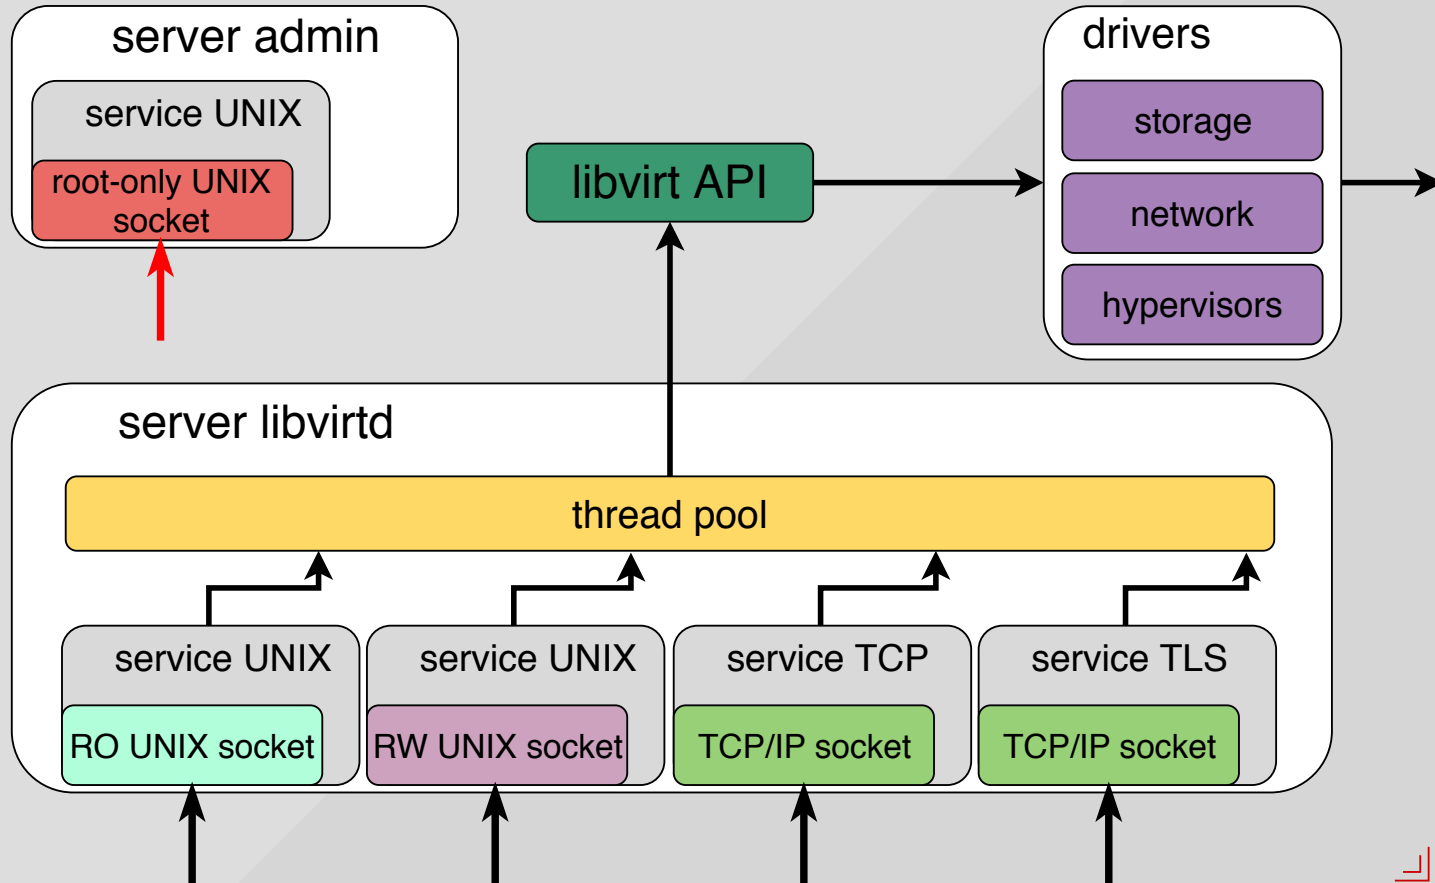
## DETOUR



# LIBVIRT TD UNDERNEATH



# LIBVIRT D UNDERNEATH



# LIBVIRTD UNDERNEATH

## JOBS - WHEN A REQUEST ARRIVES...

- ① event loop extracts raw data from a file descriptor
- ② deserializes the header of the message and creates a job for worker threads (decoded header + message payload)
- ③ worker thread deserializes the message payload to extract the procedure arguments
- ④ dispatch the registered method associated with the RPC procedure
- ⑤ serialize the results and send a reply back to the client

# LIBVIRT-ADMIN

## THREAD POOL MANAGEMENT CONT.

**SO WE ADDED MORE RESOURCES, HOW CAN WE MAKE USE OF THEM...**

- retrieve thread pool runtime info
  - standard vs priority worker thread
- modify thread pool's parameters to allow spawning more threads or kill existing ones

```
# virt-admin srv-threadpool-info libvirtd  
# virt-admin srv-threadpool-set libvirtd
```

# LIBVIRT-ADMIN

## CLIENT MANAGEMENT

### CONNECTIONS, CONNECTIONS, CONNECTIONS...

- how many clients will be connected
- setting the limit too low vs too high
- runtime support for tuning the max number of allowed clients

```
# virt-admin srv-clients-info libvirtd  
# virt-admin srv-clients-set libvirtd --max-clients --  
max-unauth-clients
```

# LIBVIRT-ADMIN

## CLIENT MANAGEMENT CONT.

- policy enforcement by 3rd party management applications
- provide detailed information about connected clients
  - IP (PID, GID, UID), timestamp, authentication, readonly connection
- list all connections
- force disconnect of a client

```
# virt-admin srv-clients-list libvirtd  
# virt-admin client-info libvirtd <ID>  
# virt-admin client-disconnect libvirtd <ID>
```

# LIBVIRT-ADMIN

## LOGGING

### DAEMON OR DOMAINS ENTER AN INVALID STATE, BUT...

- debugging is not enabled
- changing the config requires daemon restart
- we want to inspect current state
- modify the logging settings on the fly



```
2016-08-11 15:03:18.854+0000: 13787: info: virObjectRef:296: OBJECT_REF: obj=0x7f71bc0009c0
2016-08-11 15:03:18.854+0000: 13787: info: virObjectRef:296: OBJECT_REF: obj=0x7f71bc0009c0
2016-08-11 15:03:18.854+0000: 13787: info: virEventRunDefaultImpl:311: running default event implementation
2016-08-11 15:03:18.854+0000: 13787: info: virThreadJobSet:96: Thread 13787 (virNetServerHandleJob) is now running job remoteDispatchDomainBlockRebase
2016-08-11 15:03:18.854+0000: 13787: debug: virEventPollCleanupTimeouts:525: Cleanup 3
2016-08-11 15:03:18.854+0000: 13787: debug: remoteDispatchDomainBlockRebaseHelper:3823: server=0x560454ae47c0 client=0x560454b1ca0 msg=0x560454b1960 rerr=0x7f71c89e3ac0 args=0x7f71b4002000 ret=0x7f71b4002050
2016-08-11 15:03:18.854+0000: 13784: debug: virEventPollCleanupHandles:574: Cleanup 15
2016-08-11 15:03:18.854+0000: 13784: debug: virEventPollMakePollFDs:401: Prepare n=1 w=1, f=10 e=1 d=0
2016-08-11 15:03:18.854+0000: 13784: debug: virEventPollMakePollFDs:401: Prepare n=2, f=10 e=1 d=0
2016-08-11 15:03:18.854+0000: 13784: debug: virEventPollMakePollFDs:401: Prepare n=3, f=13 e=1 d=0
2016-08-11 15:03:18.854+0000: 13784: debug: virEventPollMakePollFDs:401: Prepare n=4, f=14 e=1 d=0
2016-08-11 15:03:18.854+0000: 13787: info: virObjectNew:205: OBJECT_UNREF: obj=0x7f71b40020d0 massname=virDomain
2016-08-11 15:03:18.854+0000: 13784: debug: virEventPollMakePollFDs:401: Prepare n=4 w=5, f=15 e=1 d=0
2016-08-11 15:03:18.854+0000: 13787: info: virObjectRef:296: OBJECT_REF: obj=0x7f71b4000c50
2016-08-11 15:03:18.854+0000: 13784: debug: virEventPollMakePollFDs:401: Prepare n=5 w=6, f=7 e=1 d=0
2016-08-11 15:03:18.854+0000: 13787: debug: virDomainBlockRebase:190: dom=0x7f71b40020d0, (VM: name=dummy, uuid=6640aced-8958-4edd-9480-5f7d5284707d), disk=vda, base=/tmp/ajmg, bandwidth=0, flags=a
```

# LIBVIRT-ADMIN

```
LOGGING
2016-08-11 15:03:18.854+0000: 13784: debug: virEventPollMakePollFDs:401: Prepare n=6 w=7, f=17 e=1 d=0
2016-08-11 15:03:18.854+0000: 13787: info: virObjectRef:296: OBJECT_REF: obj=0x7f7174376fa0
2016-08-11 15:03:18.854+0000: 13784: debug: virEventPollMakePollFDs:401: Prepare n=7 w=8, f=16 e=0 d=0
2016-08-11 15:03:18.854+0000: 13787: info: virObjectRef:296: OBJECT_REF: obj=0x560454ae40b0
2016-08-11 15:03:18.854+0000: 13784: debug: virEventPollMakePollFDs:401: Prepare n=8 w=9, f=16 e=1 d=0
2016-08-11 15:03:18.854+0000: 13787: debug: virAccessManagerCheckDomains:234: manager=0x560454ae40b0(name=stack) driver=QEMU domain=0x7f717476610 perm=22
2016-08-11 15:03:18.854+0000: 13784: debug: virEventPollMakePollFDs:401: Prepare n=9 w=10, f=21 e=1 d=0
2016-08-11 15:03:18.854+0000: 13787: debug: virAccessManagerCheckDomains:234: manager=0x560454ae4200(name=none) driver=QEMU domain=0x7f717476610 perm=22
2016-08-11 15:03:18.854+0000: 13784: debug: virEventPollMakePollFDs:401: Prepare n=10 w=38, f=23 e=25 d=0
2016-08-11 15:03:18.854+0000: 13787: info: virObjectUnref:259: OBJECT_UNREF: obj=0x560454ae40b0
2016-08-11 15:03:18.854+0000: 13784: debug: virEventPollMakePollFDs:401: Prepare n=11 w=39, f=24 e=25 d=0
2016-08-11 15:03:18.854+0000: 13787: info: virObjectRef:296: OBJECT_REF: obj=0x7f71742b2ba0
2016-08-11 15:03:18.854+0000: 13784: debug: virEventPollMakePollFDs:401: Prepare n=12 w=40, f=25 e=25 d=0
2016-08-11 15:03:18.854+0000: 13787: info: virObjectRef:296: OBJECT_REF: obj=0x7f71742b2ba0
2016-08-11 15:03:18.854+0000: 13784: debug: virEventPollMakePollFDs:401: Prepare n=13 w=41, f=26 e=1 d=0
2016-08-11 15:03:18.854+0000: 13787: debug: qemuDomainObjBeginJobInternal:2788: Starting job: modify (vm=0x7f7174376fa0 name=dummy, current job=none async=none)
2016-08-11 15:03:18.854+0000: 13784: debug: virEventPollMakePollFDs:401: Prepare n=14 w=42, f=27 e=1 d=0
2016-08-11 15:03:18.854+0000: 13787: debug: qemuDomainObjBeginJobInternal:2829: Started job: modify (async=none vm=0x7f7174376fa0 name=dummy)
2016-08-11 15:03:18.854+0000: 13784: debug: virEventPollCalculateTimeout:338: Calculate expiry of 3 timers
2016-08-11 15:03:18.854+0000: 13787: info: virObjectUnref:259: OBJECT_UNREF: obj=0x7f71742b2ba0
2016-08-11 15:03:18.854+0000: 13784: debug: virEventPollCalculateTimeout:346: Got a timeout scheduled for 1470927803854
2016-08-11 15:03:18.854+0000: 13787: debug: qemuDomainMaybeBreakMagically:16374: Maybe performing magic on disk=0x7f7174126e50, path=/tmp/ajmg, flags=2
2016-08-11 15:03:18.854+0000: 13784: debug: virEventPollCalculateTimeout:359: Schedule timeout then=1470927803854 now=1470927798854
2016-08-11 15:03:18.855+0000: 13784: debug: virEventPollCalculateTimeout:369: Timeout at 1470927803854 due in 5000 ms
2016-08-11 15:03:18.855+0000: 13787: info: virObjectUnref:259: OBJECT_UNREF: obj=0x7f7174376fa0
2016-08-11 15:03:18.855+0000: 13784: debug: qemuDomainObjEndJob:2983: Stopping job: modify (async=none vm=0x7f7174376fa0 name=dummy)
2016-08-11 15:03:18.855+0000: 13787: info: virObjectUnref:259: OBJECT_UNREF: obj=0x7f71742b2ba0
2016-08-11 15:03:18.855+0000: 13787: info: virObjectUnref:259: OBJECT_UNREF: obj=0x7f7174376fa0
2016-08-11 15:03:18.855+0000: 13787: info: virObjectUnref:261: OBJECT_UNREF: obj=0x7f71bc0020d0
2016-08-11 15:03:18.855+0000: 13787: info: virObjectUnref:261: OBJECT_UNREF: obj=0x7f71bc0020d0
2016-08-11 15:03:18.855+0000: 13787: debug: virDomainDispose:316: release domain 0x7f71b40020d0 dummy 6640aced-8958-4edd-9480-5f7d5284707d
2016-08-11 15:03:18.855+0000: 13787: info: virObjectUnref:259: OBJECT_UNREF: obj=0x7f71b4000c50
2016-08-11 15:03:18.855+0000: 13787: debug: virThreadJobClear:121: Thread 13787 (virNetServerHandleJob) finished job remoteDispatchDomainBlockRebase with ret=-1
2016-08-11 15:03:18.855+0000: 13787: info: virObjectUnref:259: OBJECT_UNREF: obj=0x7f71bc0009c0
2016-08-11 15:03:18.855+0000: 13787: info: virNetServerProgramSendError:153: prog=536903814 ver=1 proc=266 type=1 serial=12 msg=0x560454b1960 rerr=0x7f71c89e3ac0
2016-08-11 15:03:18.855+0000: 13787: debug: virNetMessageEncodePayload:386: Encode length as 216
2016-08-11 15:03:18.855+0000: 13787: debug: virNetServerClientSendMessageLocked:1480: msg=0x560454b1960 proc=266 len=216 offset=0
2016-08-11 15:03:18.855+0000: 13787: info: virNetServerClientSendMessageLocked:1488: RPC_SERVER_CLIENT_MSG_TX_QUEUE: client=0x560454b1ca0 len=216 prog=536903814 vers=1 proc=266 type=1 status=1 serial=12
2016-08-11 15:03:18.855+0000: 13787: debug: virNetServerClientCalculateHandleMode:166: tfs=(nil) hls=-1, rx=0x560454b1960 tx=0x560454b1960
2016-08-11 15:03:18.855+0000: 13787: debug: virNetServerClientCalculateHandleMode:201: mode=3
2016-08-11 15:03:18.855+0000: 13787: info: virEventPollUpdateHandle:152: EVENT_POLL_UPDATE_HANDLE: watch=42 events=3
2016-08-11 15:03:18.855+0000: 13787: debug: virEventPollInterruptLocked:726: Interrupting
2016-08-11 15:03:18.855+0000: 13787: info: virObjectUnref:259: OBJECT_UNREF: obj=0x7f71bc0009c0
```

**qemuDomainMaybeBreakMagically:16374 : Maybe performing magic**

**block copy still active: disk 'vda' already in active block job**

# LIBVIRT-ADMIN

## LOGGING CONT.

- switch on debugging by modifying the global logging level
  - *debug, info, warning, error*
- provide a custom set of filters (extremely verbose modules)
- define a custom set of logging outputs
  - *file, stderr, syslog, journald*

```
# virt-admin dmn-log-info libvirtd
# virt-admin dmn-log-define libvirtd --level --filters
--outputs
```

# FUTURE IMPROVEMENTS

- client discover events
- support for remote access
- logging client's activity

# HOW TO REACH US

## HELP

- <libvirt-users@redhat.com>
- API documentation
- IRC #virt (irc.oftc.net)
- <http://wiki.libvirt.org>

## DEVELOPMENT

- <libvir-list@redhat.com>
- contributor guidelines
- <http://libvirt.org>

# Q & A

**THANK YOU**