



Vdsm

Dan
Kenigsberg

Vdsm is Free

What is Vdsm

Responsibilities

Why Vdsm

Architecture

API

Storage

Architecture

Storage API

Thin

Provisioning

Bright Future

Vdsm

Virtual Desktop and Server Management Daemon KVM Forum 2011, Vancouver

Red Hat

Dan Kenigsberg

August 15, 2011





Agenda

Vdsm

Dan
Kenigsberg

Vdsm is Free

What is Vdsm
Responsibilities
Why Vdsm
Architecture
API
Storage
Architecture
Storage API
Thin
Provisioning
Bright Future

- 1 **Vdsm is Free**
 - What is Vdsm
 - Responsibilities
 - Why Vdsm
 - Architecture
 - API
 - Storage Architecture
 - Storage API
 - Thin Provisioning
 - Bright Future



What is Vdsm?

Vdsm

Dan
Kenigsberg

Vdsm is Free

What is Vdsm
Responsibilities
Why Vdsm
Architecture
API
Storage
Architecture
Storage API
Thin
Provisioning
Bright Future

- Vdsm is a node management API.
- High level API for managing ahypervisor node.
- Abstracts low level details of underlying Linux environments

• [Vdsm: Provisioning and API for RHEV4](#)

• [Vdsm: Normal nodes and other Linux distros \(patches welcome\)](#)



What is Vdsm?

Vdsm

Dan
Kenigsberg

Vdsm is Free

What is Vdsm
Responsibilities
Why Vdsm
Architecture
API
Storage
Architecture
Storage API
Thin
Provisioning
Bright Future

- Vdsm is a node management API.
- High level API for managing a hypervisor node.
- Abstracts low level details of underlying Linux environments

• [Vdsm: A Node Management API](#)
(normal, vdsms and other Linux distros patches welcome)



What is Vdsm?

Vdsm

Dan
Kenigsberg

Vdsm is Free

What is Vdsm
Responsibilities
Why Vdsm
Architecture
API
Storage
Architecture
Storage API
Thin
Provisioning
Bright Future

- Vdsm is a node management API.
- High level API for managing a hypervisor node.
- Abstracts low level details of underlying Linux environments
 - Today: RHEL-5 RHEL6 RHEV-H
 - Tomorrow: Fedora and other Linux distros (patches welcome)



What is Vdsm?

Vdsm

Dan
Kenigsberg

Vdsm is Free

- What is Vdsm
- Responsibilities
- Why Vdsm
- Architecture
- API
- Storage
- Architecture
- Storage API
- Thin
- Provisioning
- Bright Future

- Vdsm is a node management API.
- High level API for managing a hypervisor node.
- Abstracts low level details of underlying Linux environments
 - Today: RHEL-5 RHEL6 RHEV-H
 - Tomorrow: Fedora and other Linux distros (patches welcome)



What is Vdsm?

Vdsm

Dan
Kenigsberg

Vdsm is Free

What is Vdsm
Responsibilities
Why Vdsm
Architecture
API
Storage
Architecture
Storage API
Thin
Provisioning
Bright Future

- Vdsm is a node management API.
- High level API for managing a hypervisor node.
- Abstracts low level details of underlying Linux environments
 - Today: RHEL-5 RHEL6 RHEV-H
 - Tomorrow: Fedora and other Linux distros (patches welcome)



Vdsm Responsibilities

Vdsm

Dan
Kenigsberg

Vdsm is Free

What is Vdsm
Responsibilities
Why Vdsm
Architecture
API
Storage
Architecture
Storage API
Thin
Provisioning
Bright Future

- VM lifecycle (libvirt + guest agent).
- Host bootstrap/registration.
- Monitoring host and VMs.
- Network configuration.
- Storage management.
- Policy management
 - scheduler
 - ksm
 - thin provisioning ("a running VM should keep running")
 - *storage/network/cpu QoS*
 - *ballooning.*



Vdsm Responsibilities

Vdsm

Dan
Kenigsberg

Vdsm is Free

What is Vdsm
Responsibilities
Why Vdsm
Architecture
API
Storage
Architecture
Storage API
Thin
Provisioning
Bright Future

- VM lifecycle (libvirt + guest agent).
- Host bootstrap/registration.
- Monitoring host and VMs.
- Network configuration.
- Storage management.
- Policy management
 - scheduler
 - ksm
 - thin provisioning ("a running VM should keep running")
 - *storage/network/cpu QoS*
 - *ballooning.*



Vdsm Responsibilities

Vdsm

Dan
Kenigsberg

Vdsm is Free

What is Vdsm
Responsibilities
Why Vdsm
Architecture
API
Storage
Architecture
Storage API
Thin
Provisioning
Bright Future

- VM lifecycle (libvirt + guest agent).
- Host bootstrap/registration.
- Monitoring host and VMs.
- Network configuration.
- Storage management.
- Policy management
 - scheduler
 - ksm
 - thin provisioning ("a running VM should keep running")
 - *storage/network/cpu QoS*
 - *ballooning.*



Vdsm Responsibilities

Vdsm

Dan
Kenigsberg

Vdsm is Free

What is Vdsm
Responsibilities
Why Vdsm
Architecture
API
Storage
Architecture
Storage API
Thin
Provisioning
Bright Future

- VM lifecycle (libvirt + guest agent).
- Host bootstrap/registration.
- Monitoring host and VMs.
- Network configuration.
- Storage management.
- Policy management
 - scheduler
 - ksm
 - thin provisioning ("a running VM should keep running")
 - *storage/network/cpu QoS*
 - *ballooning.*



Vdsm Responsibilities

Vdsm

Dan
Kenigsberg

Vdsm is Free

What is Vdsm
Responsibilities
Why Vdsm
Architecture
API
Storage
Architecture
Storage API
Thin
Provisioning
Bright Future

- VM lifecycle (libvirt + guest agent).
- Host bootstrap/registration.
- Monitoring host and VMs.
- Network configuration.
- Storage management.
- Policy management
 - scheduler
 - ksm
 - thin provisioning ("a running VM should keep running")
 - *storage/network/cpu QoS*
 - *ballooning.*



Vdsm Responsibilities

Vdsm

Dan
Kenigsberg

Vdsm is Free

What is Vdsm
Responsibilities
Why Vdsm
Architecture
API
Storage
Architecture
Storage API
Thin
Provisioning
Bright Future

- VM lifecycle (libvirt + guest agent).
- Host bootstrap/registration.
- Monitoring host and VMs.
- Network configuration.
- Storage management.
- Policy management
 - scheduler
 - ksm
 - thin provisioning (“a running VM should keep running”)
 - *storage/network/cpu QoS*
 - *ballooning*.



Why Vdsm

Vdsm

Dan
Kenigsberg

Vdsm is Free

What is Vdsm
Responsibilities

Why Vdsm

Architecture

API

Storage

Architecture

Storage API

Thin

Provisioning

Bright Future

- qemu is good:

\$ qemu-kvm & yay! we have a virtual machine (but see the fine print).

```
/usr/libexec/qemu-kvm -S -M rhel6.0.0 -cpu Conroe -enable-kvm -m 2048  
-smp 1,sockets=1,cores=1,threads=1 -name z-win7x86-1 -uuid e3e19b36-f6b7-4ab9-b604-1f8b5c471bda  
-smbios type=1,manufacturer=Red Hat,product=RHEL,version=6Server-6.1.0.2.e16_1,serial=50C1C6F0-  
-nodefconfig -nodefaults -chardev socket,id=charmonitor,path=/var/lib/libvirt/qemu/z-win7x86-1.  
-mon chardev=charmonitor,id=monitor,mode=control -rtc base=2011-08-04T06:17:36 -boot  
cdn -device virtio-serial-pci,id=virtio-serial0,max_ports=16,bus=pci.0,addr=0x6  
-drive file=...
```

- Libvirt is even better:

To manage more than a handful of virtual machines you would need libvirt: virsh, virt-manager.

- Managing a virtualization node has much more than firing up a (K)VM.



Why Vdsm

Vdsm

Dan
Kenigsberg

Vdsm is Free

What is Vdsm
Responsibilities

Why Vdsm

Architecture

API

Storage

Architecture

Storage API

Thin

Provisioning

Bright Future

- qemu is good:

\$ qemu-kvm & yay! we have a virtual machine (but see the fine print).

```
/usr/libexec/qemu-kvm -S -M rhel6.0.0 -cpu Conroe -enable-kvm -m 2048
-smp 1,sockets=1,cores=1,threads=1 -name z-win7x86-1 -uuid e3e19b36-f6b7-4ab9-b604-1f8b5c471bda
-smbios type=1,manufacturer=Red Hat,product=RHEL,version=6Server-6.1.0.2.e16_1,serial=50C1C6F0-F
-nodefconfig -nodefaults -chardev socket,id=charmonitor,path=/var/lib/libvirt/qemu/z-win7x86-1.m
-mon chardev=charmonitor,id=monitor,mode=control -rtc base=2011-08-04T06:17:36 -boot
cdn -device virtio-serial-pci,id=virtio-serial0,max_ports=16,bus=pci.0,addr=0x6
-drive file=...
```

- Libvirt is even better:

To manage more than a handful of virtual machines you would need libvirt: `virsh`, `virt-manager`.

- Managing a virtualization node has much more than firing up a (K)VM.



Why Vdsm

Vdsm

Dan
Kenigsberg

Vdsm is Free

What is Vdsm
Responsibilities

Why Vdsm

Architecture

API

Storage

Architecture

Storage API

Thin

Provisioning

Bright Future

- qemu is good:

\$ qemu-kvm & yay! we have a virtual machine (but see the fine print).

```
/usr/libexec/qemu-kvm -S -M rhel6.0.0 -cpu Conroe -enable-kvm -m 2048  
-smp 1,sockets=1,cores=1,threads=1 -name z-win7x86-1 -uuid e3e19b36-f6b7-4ab9-b604-1f8b5c471bda  
-smbios type=1,manufacturer=Red Hat,product=RHEL,version=6Server-6.1.0.2.e16_1,serial=50C1C6F0-B  
-nodefconfig -nodefaults -chardev socket,id=charmonitor,path=/var/lib/libvirt/qemu/z-win7x86-1.m  
-mon chardev=charmonitor,id=monitor,mode=control -rtc base=2011-08-04T06:17:36 -boot  
cdn -device virtio-serial-pci,id=virtio-serial0,max_ports=16,bus=pci.0,addr=0x6  
-drive file=....
```

- Libvirt is even better:

To manage more than a handful of virtual machines you would need libvirt: virsh, virt-manager.

- Managing a virtualization node has much more than firing up a (K)VM.



Why Vdsm

Vdsm

Dan
Kenigsberg

Vdsm is Free

What is Vdsm
Responsibilities

Why Vdsm

Architecture

API

Storage

Architecture

Storage API

Thin

Provisioning

Bright Future

- qemu is good:

\$ qemu-kvm & yay! we have a virtual machine (but see the fine print).

```
/usr/libexec/qemu-kvm -S -M rhel6.0.0 -cpu Conroe -enable-kvm -m 2048
-smp 1,sockets=1,cores=1,threads=1 -name z-win7x86-1 -uuid e3e19b36-f6b7-4ab9-b604-1f8b5c471bda
-smbios type=1,manufacturer=Red Hat,product=RHEL,version=6Server-6.1.0.2.e16_1,serial=50C1C6F0-B
-nodefconfig -nodefaults -chardev socket,id=charmonitor,path=/var/lib/libvirt/qemu/z-win7x86-1.m
-mon chardev=charmonitor,id=monitor,mode=control -rtc base=2011-08-04T06:17:36 -boot
cdn -device virtio-serial-pci,id=virtio-serial0,max_ports=16,bus=pci.0,addr=0x6
-drive file=...
```

- Libvirt is even better:

To manage more than a handful of virtual machines you would need libvirt: virsh, virt-manager.

- Managing a virtualization node has much more than firing up a (K)VM.



Why Vdsm(-18)

Vdsm

Dan
Kenigsberg

Vdsm is Free

What is Vdsm
Responsibilities
Why Vdsm
Architecture
API
Storage
Architecture
Storage API
Thin
Provisioning
Bright Future

With thousands of machines, a cluster of hundreds of hosts, multiple storage exports — what to do?

- One solution is Red Hat Enterprise Virtualization (RHEV).
- Vdsm is its node agent, tailored for its needs.
- Agile, KVM-specific.
- RHEV is being open sourced.
- Vdsm is moving to a more general use case, applicable to other management platforms.



Why Vdsm(-18)

Vdsm

Dan
Kenigsberg

Vdsm is Free

What is Vdsm
Responsibilities

Why Vdsm
Architecture
API

Storage
Architecture
Storage API

Thin
Provisioning
Bright Future

With thousands of machines, a cluster of hundreds of hosts, multiple storage exports — what to do?

- One solution is Red Hat Enterprise Virtualization (RHEV).
- Vdsm is its node agent, tailored for its needs.
- Agile, KVM-specific.
- RHEV is being open sourced.
- Vdsm is moving to a more general use case, applicable to other management platforms.



Why Vdsm(-18)

Vdsm

Dan
Kenigsberg

Vdsm is Free

What is Vdsm
Responsibilities

Why Vdsm
Architecture
API

Storage
Architecture
Storage API

Thin
Provisioning
Bright Future

With thousands of machines, a cluster of hundreds of hosts, multiple storage exports — what to do?

- One solution is Red Hat Enterprise Virtualization (RHEV).
- Vdsm is its node agent, tailored for its needs.
- Agile, KVM-specific.
- RHEV is being open sourced.
- Vdsm is moving to a more general use case, applicable to other management platforms.



Why Vdsm(-18)

Vdsm

Dan
Kenigsberg

Vdsm is Free

What is Vdsm
Responsibilities

Why Vdsm

Architecture

API

Storage

Architecture

Storage API

Thin

Provisioning

Bright Future

With thousands of machines, a cluster of hundreds of hosts, multiple storage exports — what to do?

- One solution is Red Hat Enterprise Virtualization (RHEV).
- Vdsm is its node agent, tailored for its needs.
- Agile, KVM-specific.
- RHEV is being open sourced.
- Vdsm is moving to a more general use case, applicable to other management platforms.



Why Vdsm(-18)

Vdsm

Dan
Kenigsberg

Vdsm is Free

What is Vdsm
Responsibilities

Why Vdsm
Architecture
API

Storage
Architecture
Storage API

Thin
Provisioning
Bright Future

With thousands of machines, a cluster of hundreds of hosts, multiple storage exports — what to do?

- One solution is Red Hat Enterprise Virtualization (RHEV).
- Vdsm is its node agent, tailored for its needs.
- Agile, KVM-specific.
- RHEV is being open sourced.
- Vdsm is moving to a more general use case, applicable to other management platforms.



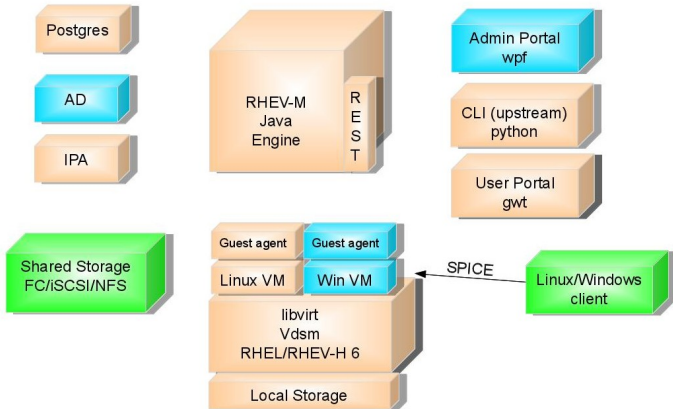
Place within oVirt

Vdsm

Dan Kenigsberg

Vdsm is Free

- What is Vdsm
- Responsibilities
- Why Vdsm
- Architecture
- API
- Storage
- Architecture
- Storage API
- Thin
- Provisioning
- Bright Future



Freeing Vdsm is another step in freeing RHEV-M within the oVirt framework.



Architecture and Implementation

Vdsm

Dan
Kenigsberg

Vdsm is Free

- What is Vdsm
- Responsibilities
- Why Vdsm
- Architecture**
- API
- Storage
- Architecture
- Storage API
- Thin
- Provisioning
- Bright Future



By Daniel Schwen, from Wikipedia



Architecture and Implementation

Vdsm

Dan
Kenigsberg

Vdsm is Free

What is Vdsm
Responsibilities
Why Vdsm

Architecture

API
Storage
Architecture
Storage API
Thin
Provisioning
Bright Future

- Written in Python.
- Multithreaded, Multi-processes.
- Used to speak directly to qemu, now uses libvirt for that.
- Speaks with its guest agent via virtio-serial.
- Implements a distributed image repository over the supported storage types: Local directory, FCP, FCoE, iSCSI, NFS...

Manages all the virtual disks, fully up down
and ready to go
Scales linearly in data writes
Integrates with the different hypervisors, qemu, xen



Architecture and Implementation

Vdsm

Dan
Kenigsberg

Vdsm is Free

What is Vdsm
Responsibilities
Why Vdsm

Architecture

API
Storage
Architecture
Storage API
Thin
Provisioning
Bright Future

- Written in Python.
- Multithreaded, Multi-processes.
- Used to speak directly to qemu, now uses libvirt for that.
- Speaks with its guest agent via virtio-serial.
- Implements a distributed image repository over the supported storage types: Local directory, FCP, FCoE, iSCSI, NFS...

Managing state over a distributed system, fully distributed
APIs for storage
Scalable, distributed
Scalable, distributed



Architecture and Implementation

Vdsm

Dan
Kenigsberg

Vdsm is Free

What is Vdsm
Responsibilities
Why Vdsm
Architecture
API
Storage
Architecture
Storage API
Thin
Provisioning
Bright Future

- Written in Python.
- Multithreaded, Multi-processes.
- Used to speak directly to qemu, now uses libvirt for that.
- Speaks with its guest agent via virtio-serial.
- Implements a distributed image repository over the supported storage types: Local directory, FCP, FCoE, iSCSI, NFS...



Architecture and Implementation

Vdsm

Dan
Kenigsberg

Vdsm is Free

What is Vdsm
Responsibilities
Why Vdsm
Architecture
API
Storage
Architecture
Storage API
Thin
Provisioning
Bright Future

- Written in Python.
- Multithreaded, Multi-processes.
- Used to speak directly to qemu, now uses libvirt for that.
- Speaks with its guest agent via virtio-serial.
- Implements a distributed image repository over the supported storage types: Local directory, FCP, FCoE, iSCSI, NFS...



Architecture and Implementation

Vdsm

Dan
Kenigsberg

Vdsm is Free

What is Vdsm
Responsibilities
Why Vdsm
Architecture
API
Storage
Architecture
Storage API
Thin
Provisioning
Bright Future

- Written in Python.
- Multithreaded, Multi-processes.
- Used to speak directly to qemu, now uses libvirt for that.
- Speaks with its guest agent via virtio-serial.
- Implements a distributed image repository over the supported storage types: Local directory, FCP, FCoE, iSCSI, NFS...
 - Maintains that on a multihost system, only one concurrent metadata writer exists.
 - Scales linearly in data writers.
 - Image-aware API. (import, templates, qcow, raw)



Architecture and Implementation

Vdsm

Dan
Kenigsberg

Vdsm is Free

What is Vdsm
Responsibilities
Why Vdsm
Architecture
API
Storage
Architecture
Storage API
Thin
Provisioning
Bright Future

- Written in Python.
- Multithreaded, Multi-processes.
- Used to speak directly to qemu, now uses libvirt for that.
- Speaks with its guest agent via virtio-serial.
- Implements a distributed image repository over the supported storage types: Local directory, FCP, FCoE, iSCSI, NFS...
 - Maintains that on a multihost system, only one concurrent metadata writer exists.
 - Scales linearly in data writers.
 - Image-aware API. (import, templates, qcow, raw)



Architecture and Implementation

Vdsm

Dan
Kenigsberg

Vdsm is Free

What is Vdsm
Responsibilities
Why Vdsm
Architecture
API
Storage
Architecture
Storage API
Thin
Provisioning
Bright Future

- Written in Python.
- Multithreaded, Multi-processes.
- Used to speak directly to qemu, now uses libvirt for that.
- Speaks with its guest agent via virtio-serial.
- Implements a distributed image repository over the supported storage types: Local directory, FCP, FCoE, iSCSI, NFS...
 - Maintains that on a multihost system, only one concurrent metadata writer exists.
 - Scales linearly in data writers.
 - Image-aware API. (import, templates, qcow, raw)



Architecture and Implementation

Vdsm

Dan
Kenigsberg

Vdsm is Free

What is Vdsm
Responsibilities
Why Vdsm
Architecture
API
Storage
Architecture
Storage API
Thin
Provisioning
Bright Future

- Written in Python.
- Multithreaded, Multi-processes.
- Used to speak directly to qemu, now uses libvirt for that.
- Speaks with its guest agent via virtio-serial.
- Implements a distributed image repository over the supported storage types: Local directory, FCP, FCoE, iSCSI, NFS...
 - Maintains that on a multihost system, only one concurrent metadata writer exists.
 - Scales linearly in data writers.
 - Image-aware API. (import, templates, qcow, raw)



Architecture and Implementation (-18)

Vdsm

Dan
Kenigsberg

Vdsm is Free

What is Vdsm
Responsibilities
Why Vdsm

Architecture

API

Storage
Architecture
Storage API

Thin
Provisioning
Bright Future

Robustness as a design goal:

- Evaporated NFS exports.
- Faulty multipaths.
- Node crash.
- Live-locked qemu.
- Internal Python exceptions.
- Self-fencing of metadata writer.



Architecture and Implementation (-18)

Vdsm

Dan
Kenigsberg

Vdsm is Free

What is Vdsm
Responsibilities
Why Vdsm
Architecture
API
Storage
Architecture
Storage API
Thin
Provisioning
Bright Future

Robustness as a design goal:

- Evaporated NFS exports.
- Faulty multipaths.
- Node crash.
- Live-locked qemu.
- Internal Python exceptions.
- Self-fencing of metadata writer.



Architecture and Implementation (-18)

Vdsm

Dan
Kenigsberg

Vdsm is Free

What is Vdsm
Responsibilities
Why Vdsm
Architecture
API
Storage
Architecture
Storage API
Thin
Provisioning
Bright Future

Robustness as a design goal:

- Evaporated NFS exports.
- Faulty multipaths.
- Node crash.
- Live-locked qemu.
- Internal Python exceptions.
- Self-fencing of metadata writer.



Architecture and Implementation (-18)

Vdsm

Dan
Kenigsberg

Vdsm is Free

What is Vdsm
Responsibilities
Why Vdsm
Architecture
API
Storage
Architecture
Storage API
Thin
Provisioning
Bright Future

Robustness as a design goal:

- Evaporated NFS exports.
- Faulty multipaths.
- Node crash.
- Live-locked qemu.
- Internal Python exceptions.
- Self-fencing of metadata writer.



Architecture and Implementation (-18)

Vdsm

Dan
Kenigsberg

Vdsm is Free

What is Vdsm
Responsibilities
Why Vdsm
Architecture
API
Storage
Architecture
Storage API
Thin
Provisioning
Bright Future

Robustness as a design goal:

- Evaporated NFS exports.
- Faulty multipaths.
- Node crash.
- Live-locked qemu.
- Internal Python exceptions.
- Self-fencing of metadata writer.



Architecture and Implementation (-18)

Vdsm

Dan
Kenigsberg

Vdsm is Free

What is Vdsm
Responsibilities
Why Vdsm
Architecture
API
Storage
Architecture
Storage API
Thin
Provisioning
Bright Future

Robustness as a design goal:

- Evaporated NFS exports.
- Faulty multipaths.
- Node crash.
- Live-locked qemu.
- Internal Python exceptions.
- Self-fencing of metadata writer.



Host Monitoring API

Vdsm

Dan
Kenigsberg

Vdsm is Free

What is Vdsm
Responsibilities
Why Vdsm
Architecture
API
Storage
Architecture
Storage API
Thin
Provisioning
Bright Future

- `getVdsCapabilities`
- `getVdsStats`
- `ping`

- `fenceNode`




VM Lifecycle API

Vdsm

Dan
Kenigsberg

Vdsm is Free

- What is Vdsm
- Responsibilities
- Why Vdsm
- Architecture
- API
- Storage
- Architecture
- Storage API
- Thin
- Provisioning
- Bright Future



Il était une fois... la Vie, Albert Barillé



VM Lifecycle API

Vdsm


Dan
Kenigsberg

Vdsm is Free

What is Vdsm
Responsibilities
Why Vdsm
Architecture
API
Storage
Architecture
Storage API
Thin
Provisioning
Bright Future

- create
- destroy
- pause
- continue
- setVmTicket
- changeCD
- changeFloppy
- migrate (downtime, timeout)
- hibernate

▪ hooks





VM Lifecycle API (agent-dependent)

Vdsm

Dan
Kenigsberg

Vdsm is Free

What is Vdsm
Responsibilities
Why Vdsm
Architecture
API
Storage
Architecture
Storage API
Thin
Provisioning
Bright Future

- shutdown
- desktopLogin
- desktopLogoff
- desktopLock



VM Monitoring API

Vdsm

Dan
Kenigsberg

Vdsm is Free

What is Vdsm
Responsibilities
Why Vdsm
Architecture
API
Storage
Architecture
Storage API
Thin
Provisioning
Bright Future

- list
- getAllVmStats
- getVmStats
 - Interesting applications installed
 - Logged in users
 - CPU consumption
 - Memory usage
 - ...



Network Configuration API (agent-dependent)

Vdsm

Dan
Kenigsberg

Vdsm is Free

What is Vdsm
Responsibilities

Why Vdsm
Architecture

API

Storage
Architecture

Storage API

Thin
Provisioning

Bright Future

- addNetwork
- delNetwork
- editNetwork
- setSafeNetworkConfig
- setupNetworks
 - connectivityCheck



Async Tasks API

Vdsm

Dan
Kenigsberg

Vdsm is Free

What is Vdsm
Responsibilities
Why Vdsm
Architecture
API
Storage
Architecture
Storage API
Thin
Provisioning
Bright Future

- getAllTasksStatuses
- getTaskStatus
- clearTask
- revertTask
- stopTask

(currently used only by the storage subsystem)



Storage Architecture

Vdsm

Dan
Kenigsberg

Vdsm is Free

What is Vdsm
Responsibilities
Why Vdsm
Architecture
API
Storage
Architecture
Storage API
Thin
Provisioning
Bright Future

- A *Storage Domain* is one of: Local directory, NFS export, VG over iSCSI LUNs, VG over FCP LUNs. Vdsm (tries to) abstract the differences away.
- A *Storage Pool* aggregates several storage domains. We try to remove this concept from the Vdsm level.
- Single meta data writer:
 - SPM lease mechanism (Chockler and Malkhi 2004, Light-Weight Leases for Storage-Centric Coordination).
 - Storage-centric mailbox.



Storage API

Vdsm

Dan
Kenigsberg

Vdsm is Free

What is Vdsm
Responsibilities
Why Vdsm
Architecture
API
Storage
Architecture
Storage API
Thin
Provisioning
Bright Future

- connectStorageServer
 - getDeviceList
 - createStorageDomain
 - attachStorageDomain
 - createImage
 - prepareVolume
 -
 -
 - (and many, many more)
- repoStats
 - getSpmStatus
 - spmStart
 - extendVolume



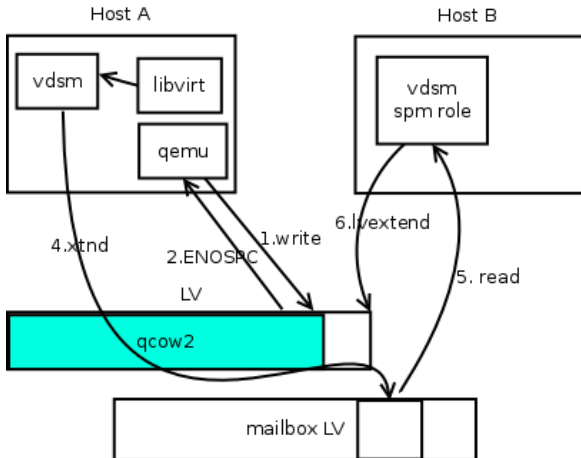
Thin Provisioning and LV extend

Vdsm

Dan
Kenigsberg

Vdsm is Free

What is Vdsm
Responsibilities
Why Vdsm
Architecture
API
Storage
Architecture
Storage API
Thin
Provisioning
Bright Future



- Disk image is created as a big empty qcow2 file within a small LV.
- When the guest writes, qcow2 grows.
- If qcow2 file fills the LV, guest stops.
- Local Vdsm senses that, and issues an extend request
- SPM extends LV, and returns success message
- Local Vdsm continue guest execution.



Future

Vdsm

Dan
Kenigsberg

Vdsm is Free

What is Vdsm
Responsibilities
Why Vdsm
Architecture
API
Storage
Architecture
Storage API
Thin
Provisioning
Bright Future

- New, stable API. RESTful(?), similar to that of rhevm-api
 - Current API is not very clean (createVG, createStorageDomain)
- Split Vdsm up to reusable autonomous parts.
 - Spin off as a generic image repository.
- SANLock, Vepa, VN-Link, storage network, storage migration, ??
- Other use cases:
 - Single-host
 - Linux containers?



Contribution

Vdsm

Dan
Kenigsberg

Vdsm is Free

What is Vdsm
Responsibilities
Why Vdsm
Architecture
API
Storage
Architecture
Storage API
Thin
Provisioning
Bright Future

- vdsm-devel@fedorahosted.org
- <http://git.fedorahosted.org/git/?p=vdsm.git>
- Ping us on #vdsml at irc.freenode.org:
Saggi Mizrahi, Federico Simoncelli, Igor Lvovsky, Eduardo Warszawasky, Yotam Oron, Ayal Baron, yours truly.
- Come over to see our demo at Andy's!



Vdsm

Dan
Kenigsberg

Vdsm is Free

What is Vdsm

Responsibilities

Why Vdsm

Architecture

API

Storage

Architecture

Storage API

Thin

Provisioning

Bright Future

The end.

Thanks for listening.