# Converged Infrastructure with Open Source

Theron Conrey

Open Source and Standards, Red Hat

@theronconrey

# Who is this guy?

- Hasn't taken a day off since joining Red Hat
- In a committed 10+ year relationship with virtualization software
- Career in helping customers deploy robust solutions
- LOVES helping people solve real world problems vs. deploying toolsets
- Virtualization community building for fun and ~~profit~~ parties.


theron@redhat.com

# compute virtualization today

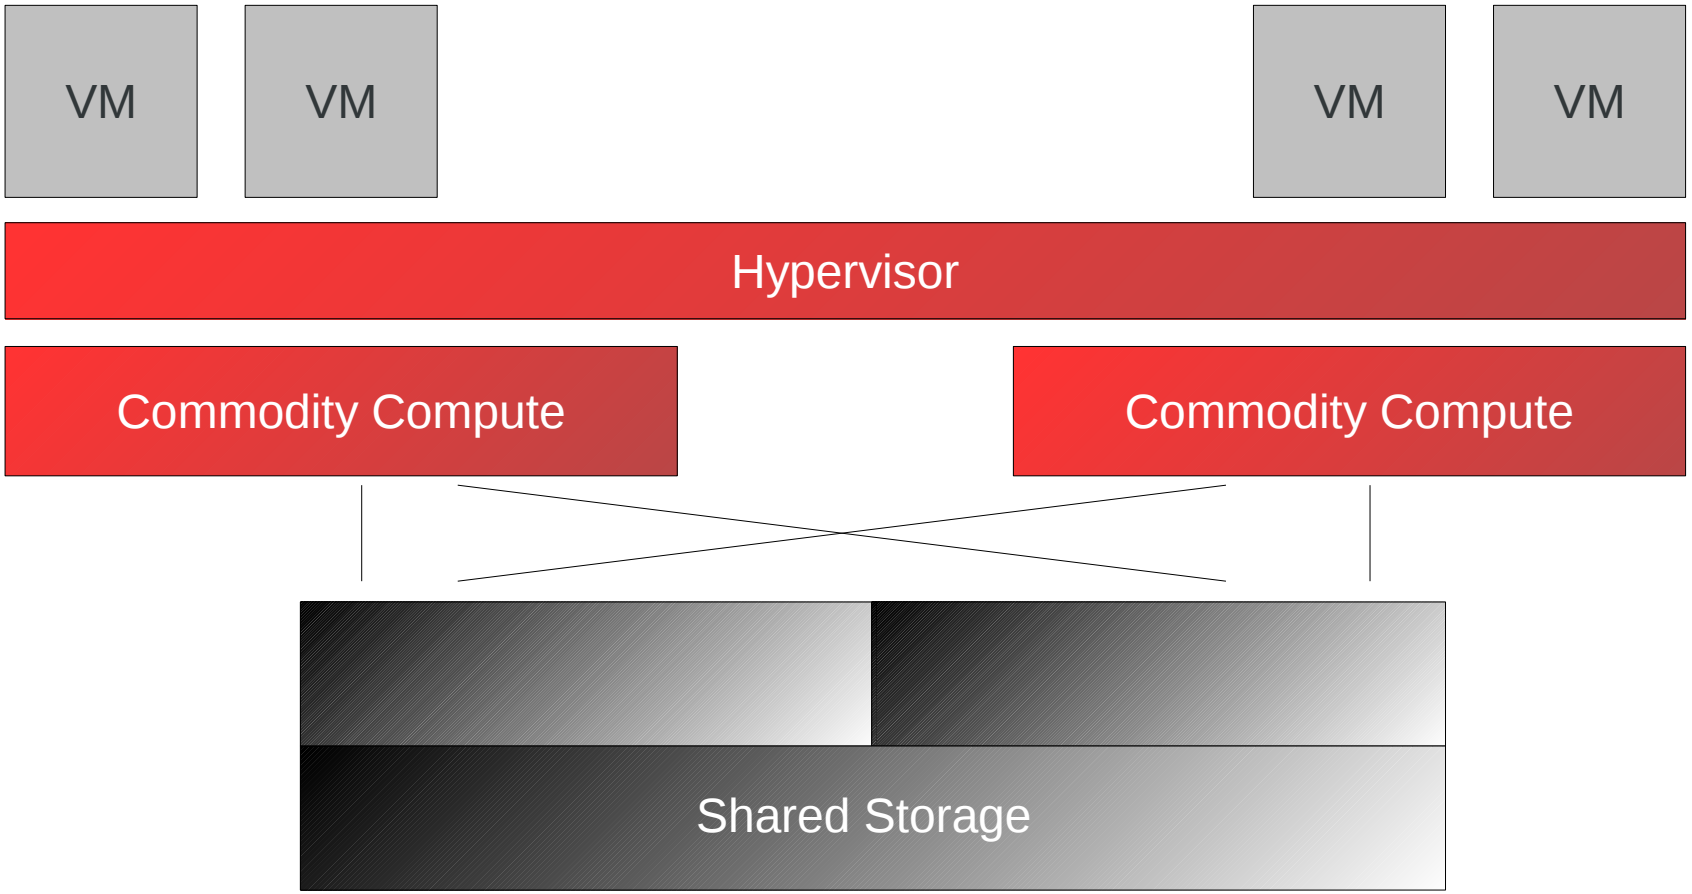




theron@redhat.com

# Advanced functionality sometimes isn't very useful


http://www.flickr.com/photos/mtaphotos/7337564952/



theron@redhat.com

# Problems with storage in virtualization platforms today

- Multiple choices for file system and virtualization management
- Lack of virtualization aware file systems
- No well defined interface points in the virtualization stack for storage integration
- No standard interface/APIs available for services like backup and restore
- Need for a single FS/storage solution that works for local, SAN and NAS storage
- Mixing different types of storage into a single filesystem namespace


theron@redhat.com

# What is converged infrastructure?

“Converged infrastructure packages multiple information technology (IT) components into a single, optimized computing solution. Components of a converged infrastructure solution include servers, data storage devices, networking equipment and software for IT infrastructure management, automation and orchestration”. - wikipedia


theron@redhat.com

# How is this different than just compute virtualization?

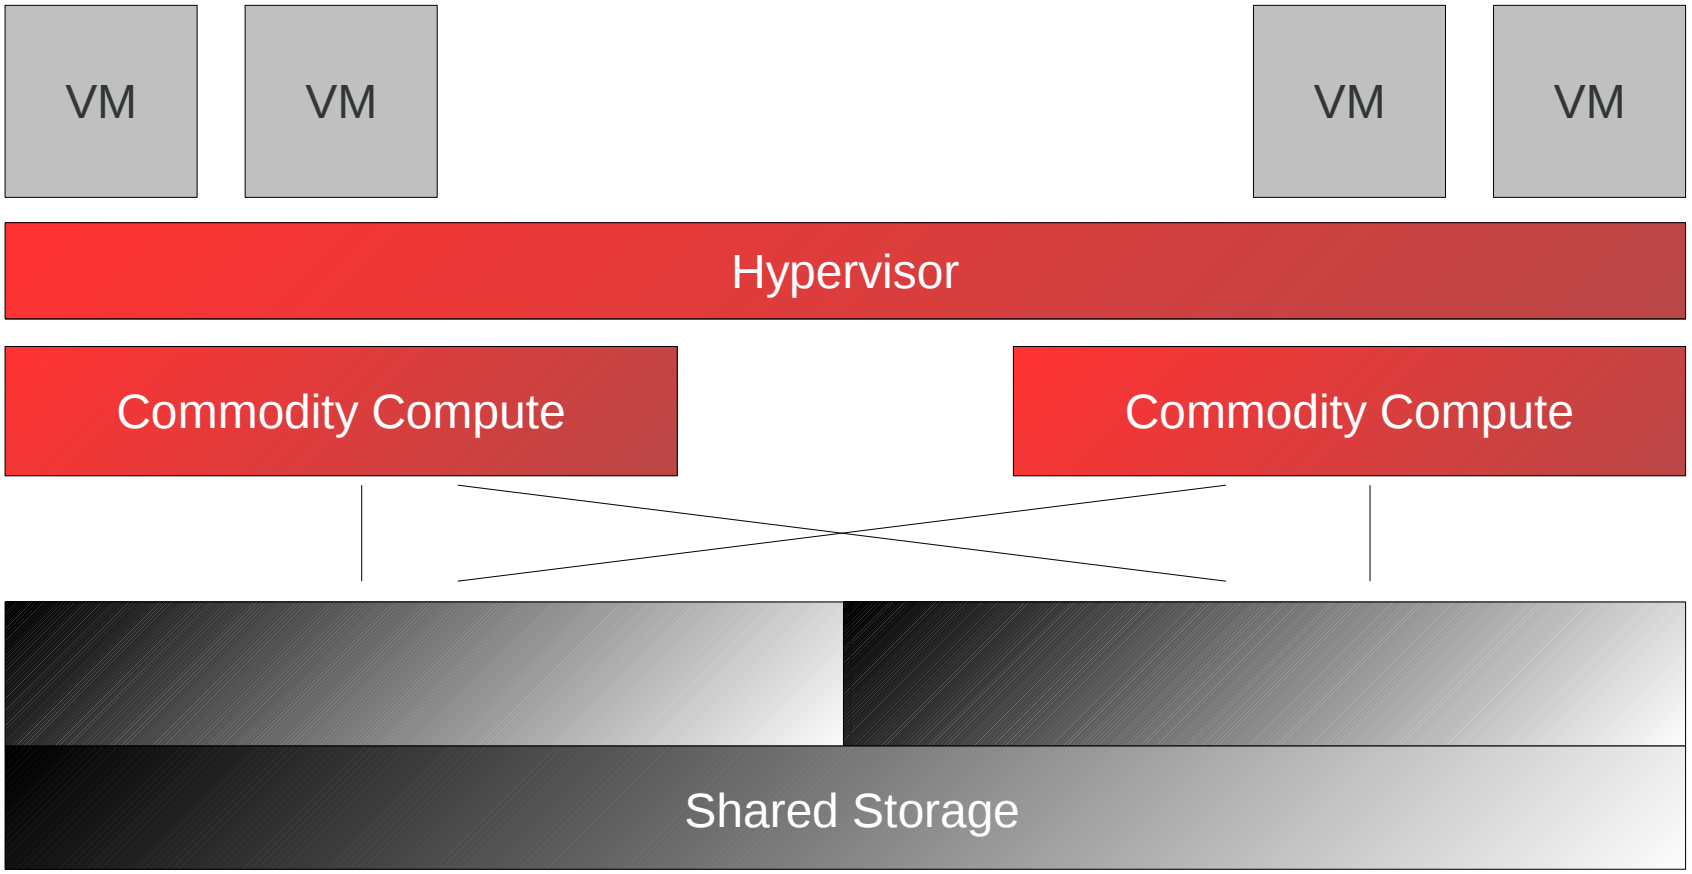




theron@redhat.com

# Who else in in this space?

- The established market
  - VCE
  - Dell
  - Nutanix
- All proprietary, all black box.


Image (c) nutanix

- Is this really a hard problem?


theron@redhat.com


http://www.flickr.com/photos/leehaywood/4481372869
YES,
It is that hard.

# How can open source software tackle such a diverse problem?

- oVirt


theron@redhat.com

# What is oVirt?

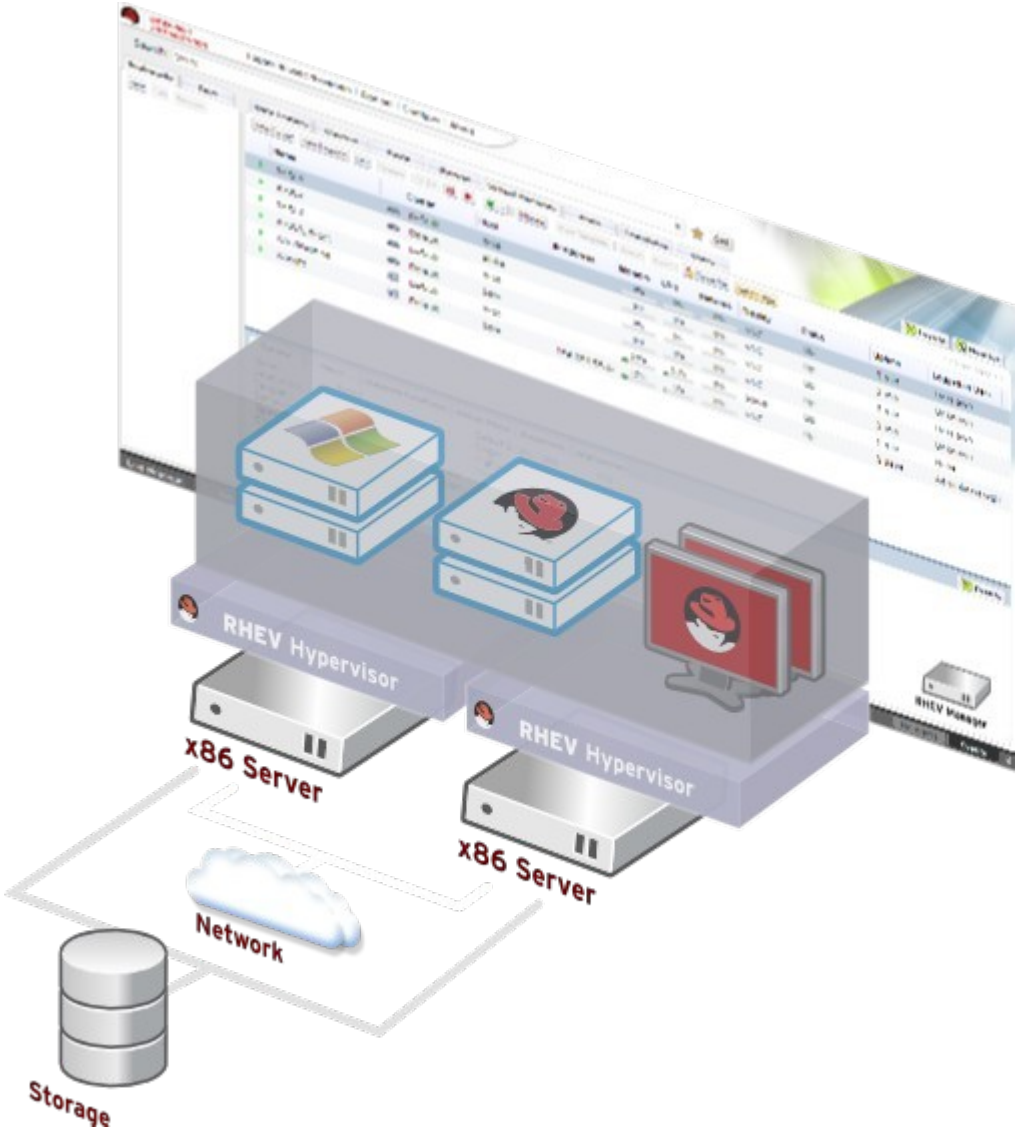


- Large scale, centralized management for server and desktop virtualization
- Open source alternative to vCenter / vSphere
- Focus on KVM
- www.ovirt.org
- #ovirt on OFTC


theron@redhat.com

# Traditional (base) oVirt architecture

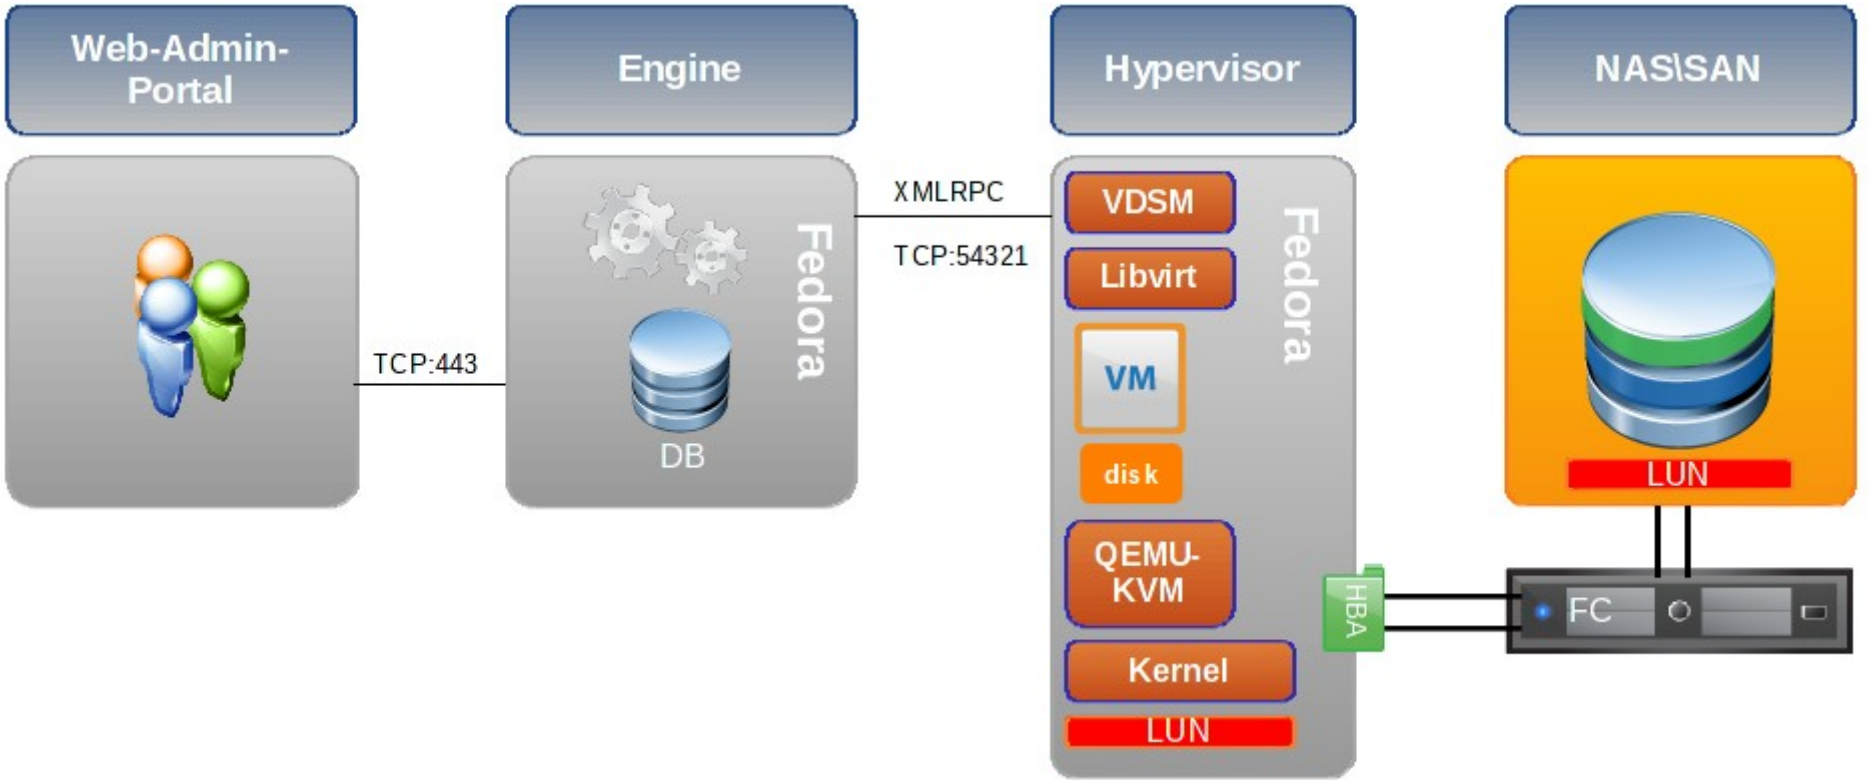




theron@redhat.com

# Overview of oVirt storage concepts

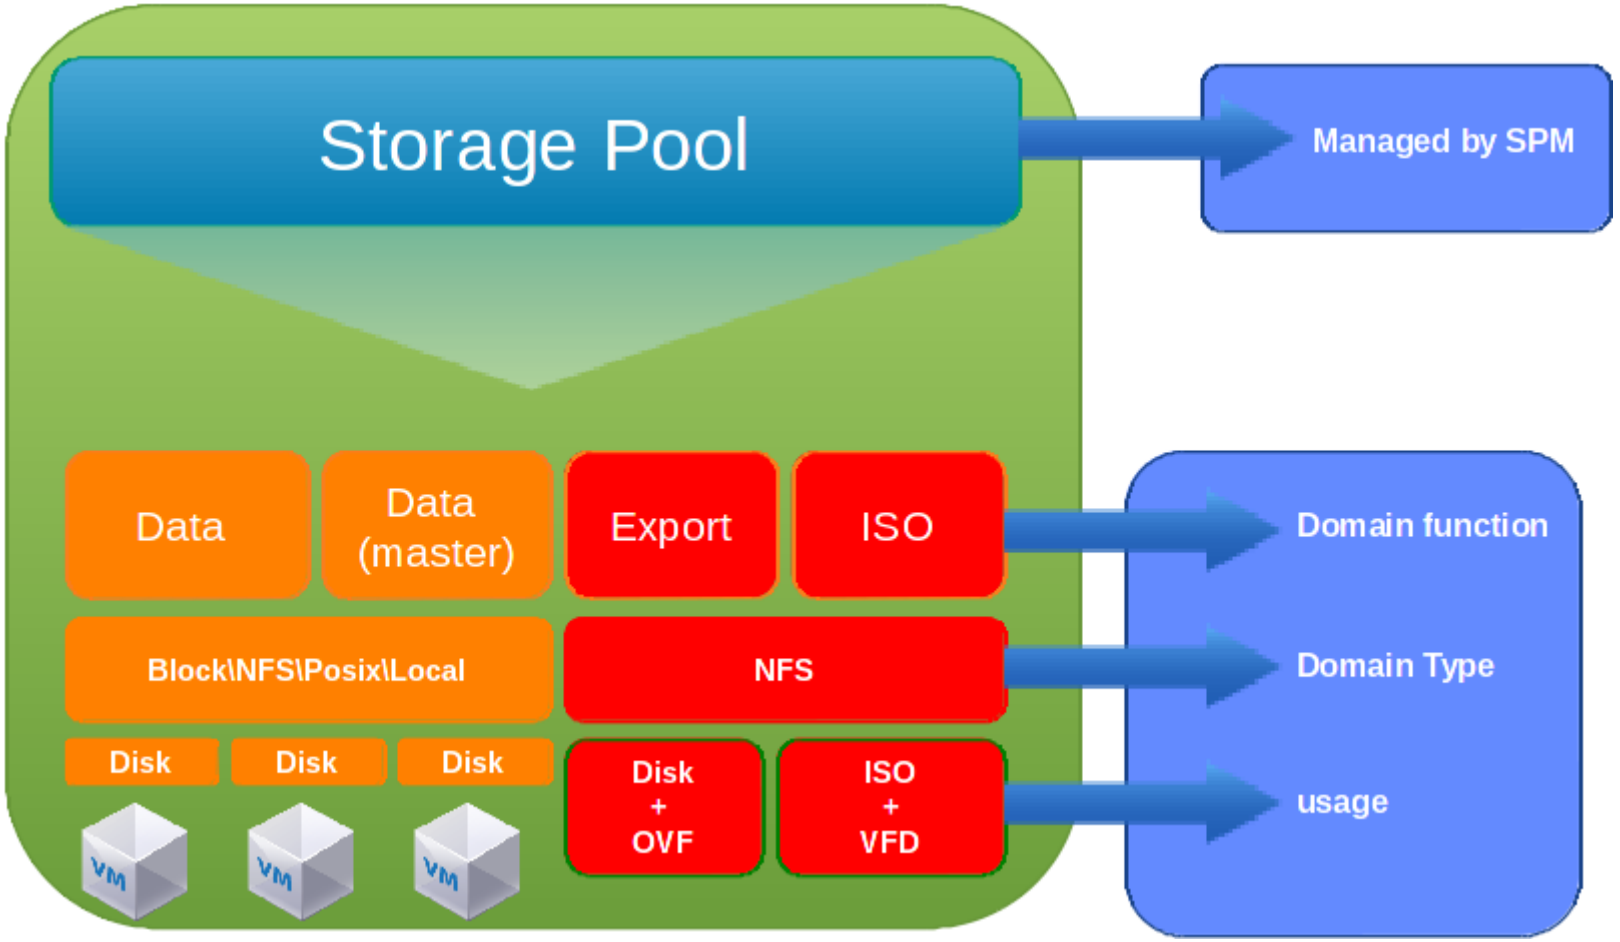




theron@redhat.com

# How can open source software tackle such a diverse problem?

- oVirt
- Gluster


theron@redhat.com

# Gluster

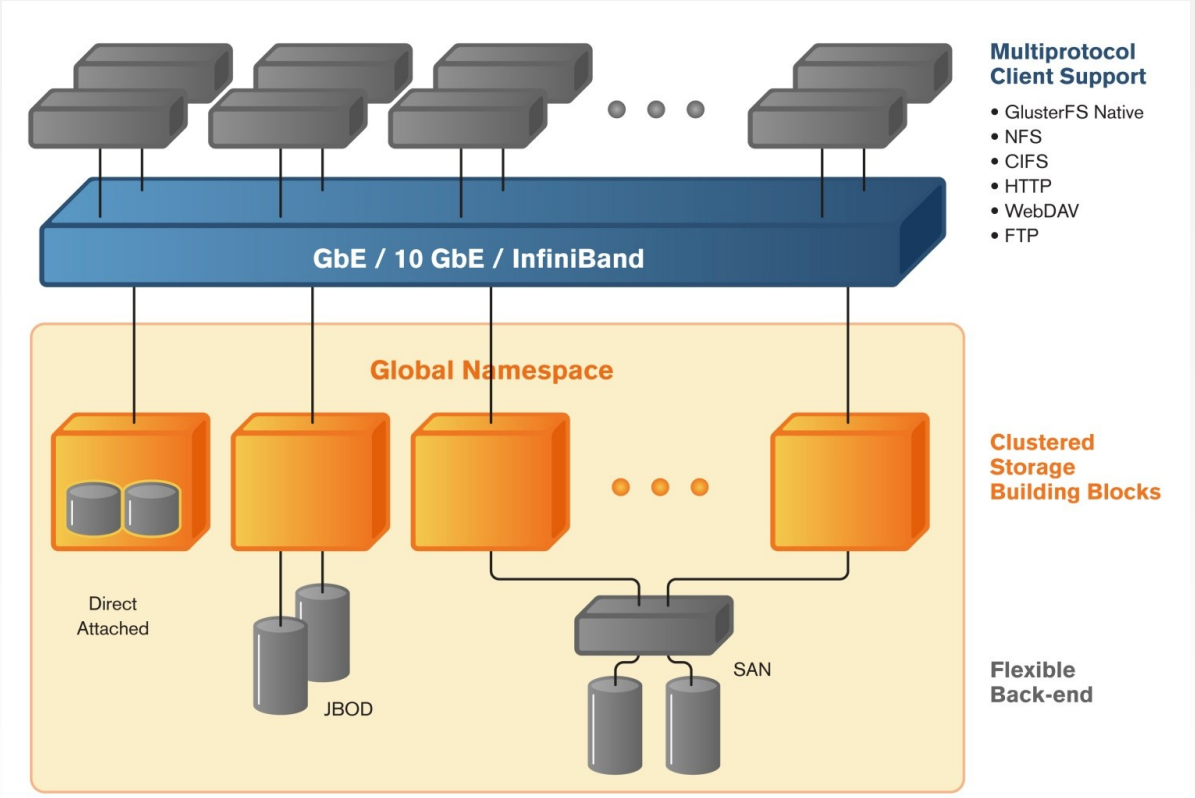


- User-space file system
- Global namespace
- Scale-out clustered storage building blocks
- Supports thousands of clients
- Access using various protocols
- Linear performance scaling
- www.gluster.org
- #gluster on freenode


theron@redhat.com

# GlusterFS Concepts

**VOLUME**

is a namespace presented as a POSIX mount point and is comprised of bricks.

**BRICK**

is the basic unit of storage, represented by an export directory on a server

**SERVER/NODES**

contain the bricks



theron@redhat.com

# GlusterFS Concepts

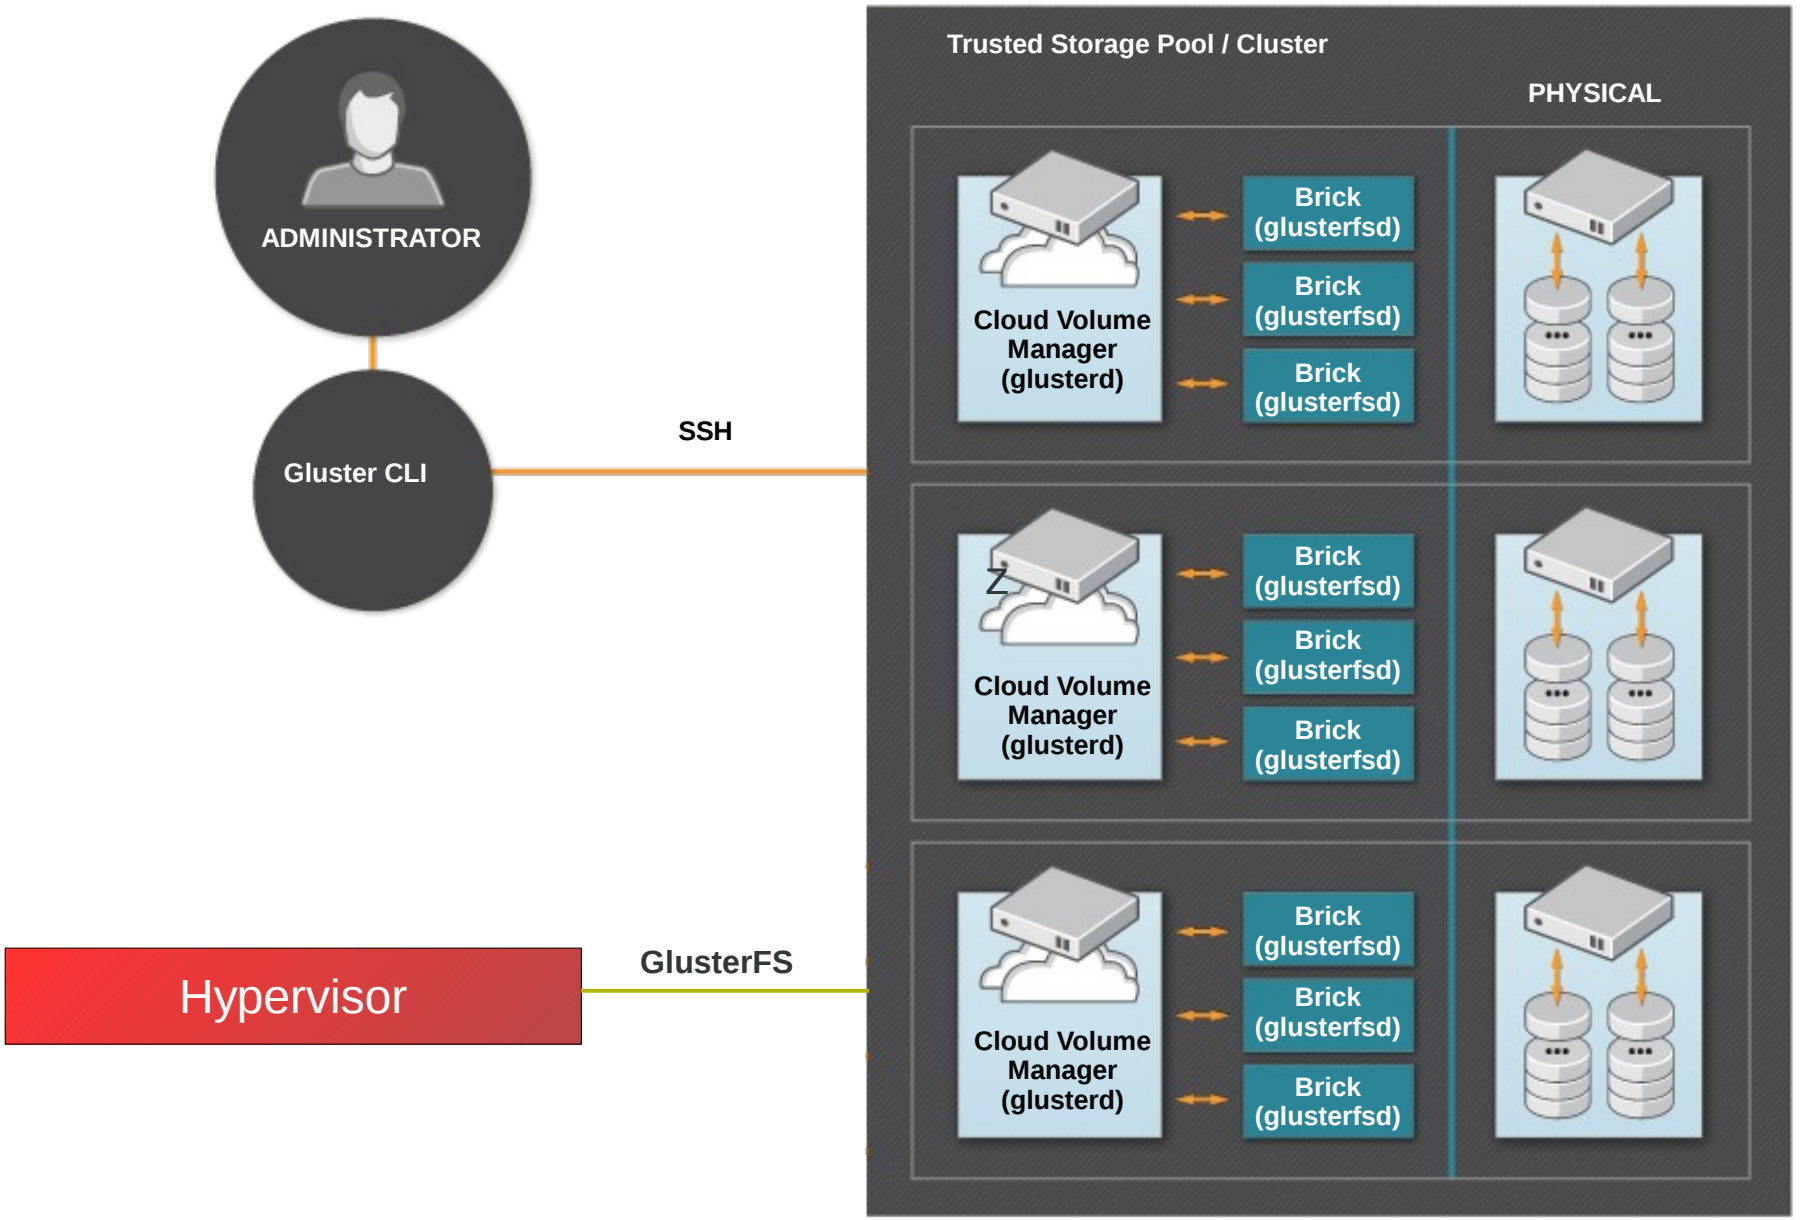




theron@redhat.com

# oVirt Gluster Integration

- Features added in oVirt 3.1[1]
- ApplicationMode configuration
  - 1 → Virtualization only (default)
  - 2 → Gluster only
  - 255 → Virtualization + Gluster
- Enable Gluster at cluster level
- New entities (Volumes, Bricks, Volume Options)
- VDSM verbs for gluster management
  - vdsm-gluster plug-in

[1] http://wiki.ovirt.org/wiki/Features/Gluster_Support


theron@redhat.com

# oVirt integrated Gluster Architecture

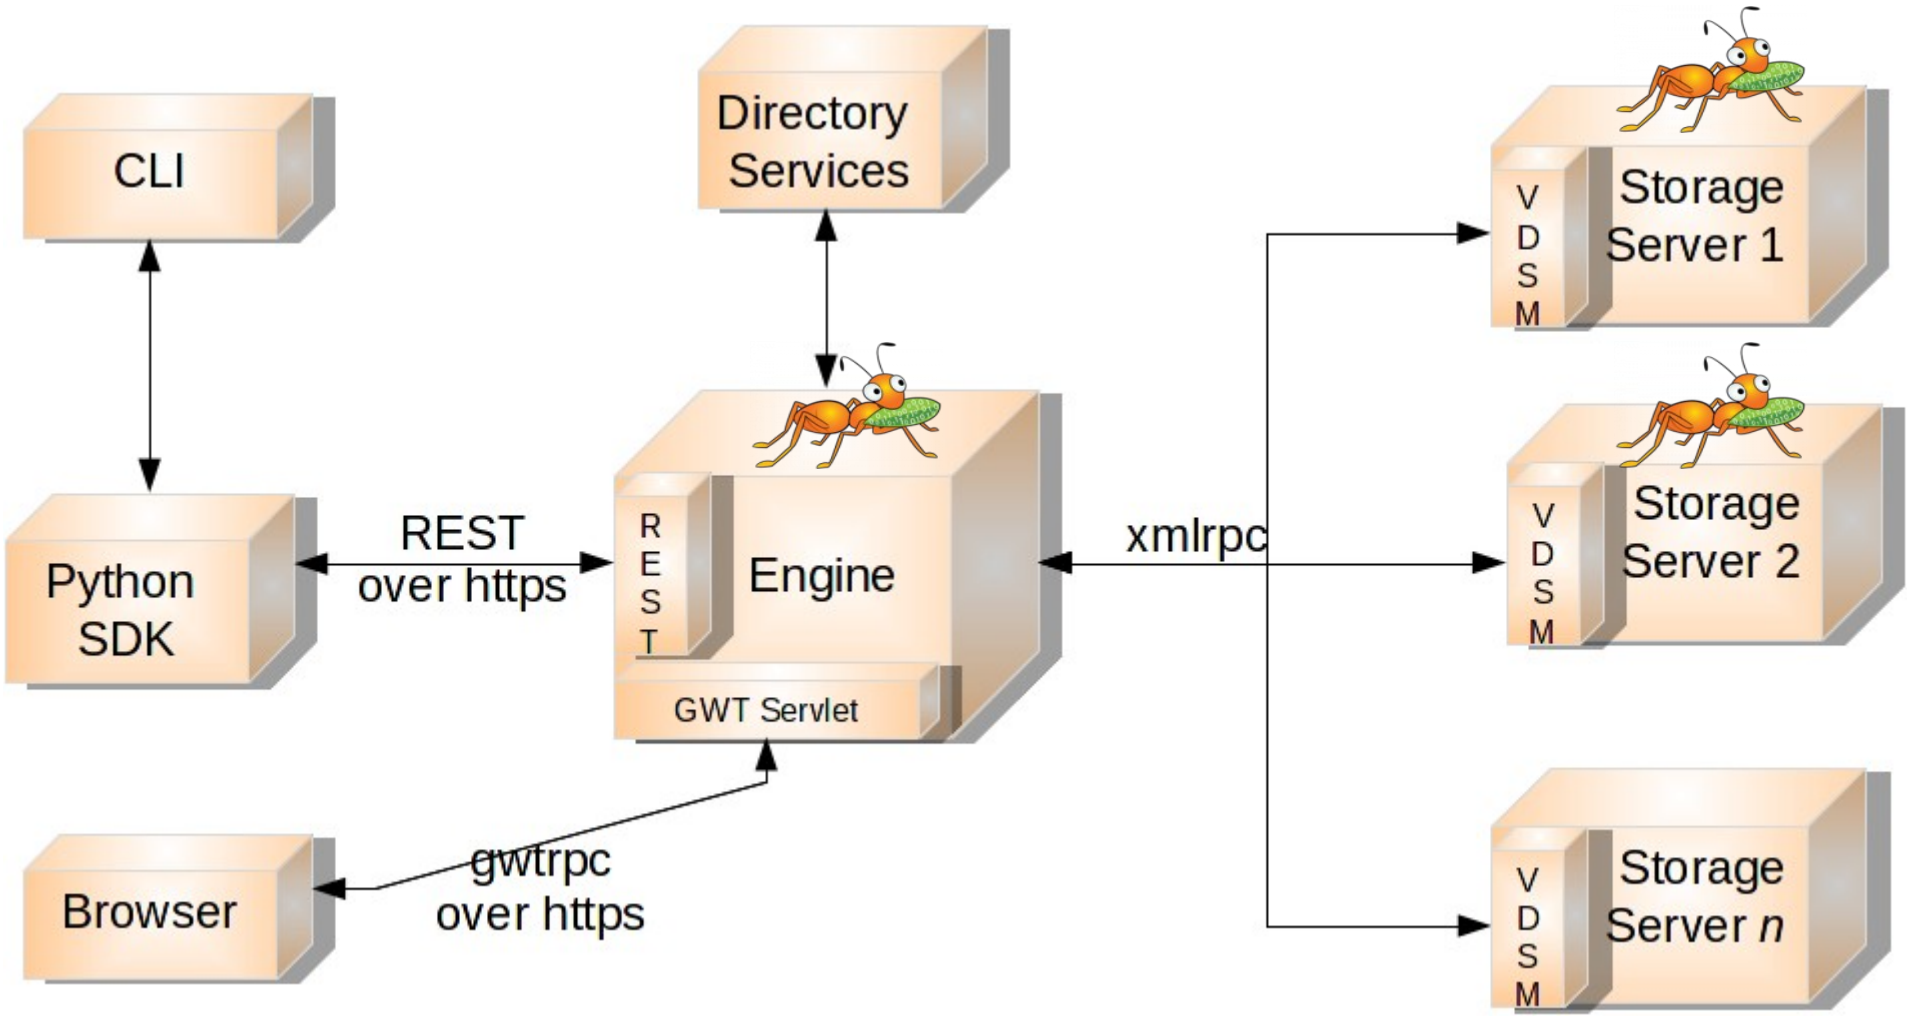




theron@redhat.com

# oVirt funtionality Gluster-ized!

- Cluster Management
  - Create Gluster Cluster
  - Add / Remove Storage Servers
  - Delete Cluster
- Volume Management
  - Create Volume
  - Add / Remove bricks
  - Start / Stop volume
  - Delete Volume

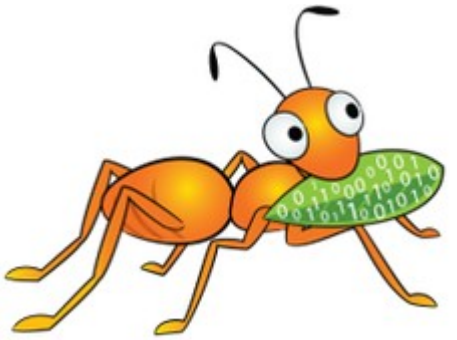


theron@redhat.com

# Gluster Storage Domains

General flow (fits all application modes)

1. Create Data Center (POSIX compliant)
2. Create Cluster
3. Add host (Hypervisor)
4. Create Volume
5. Add bricks
6. Start volume and optimize for virt
7. Create storage domain (POSIX compliant, VFS type glusterFS)


theron@redhat.com

# Gluster Storage Domains

- General flow
  - Create Data Center (POSIX compliant)
  - Create Cluster


  - Add host (Hypervisor / Storage server)
    - make sure your host has
    - Glusterfs-fuse-3.3.0.5+
    - Glusterfs-3.3.0.5+

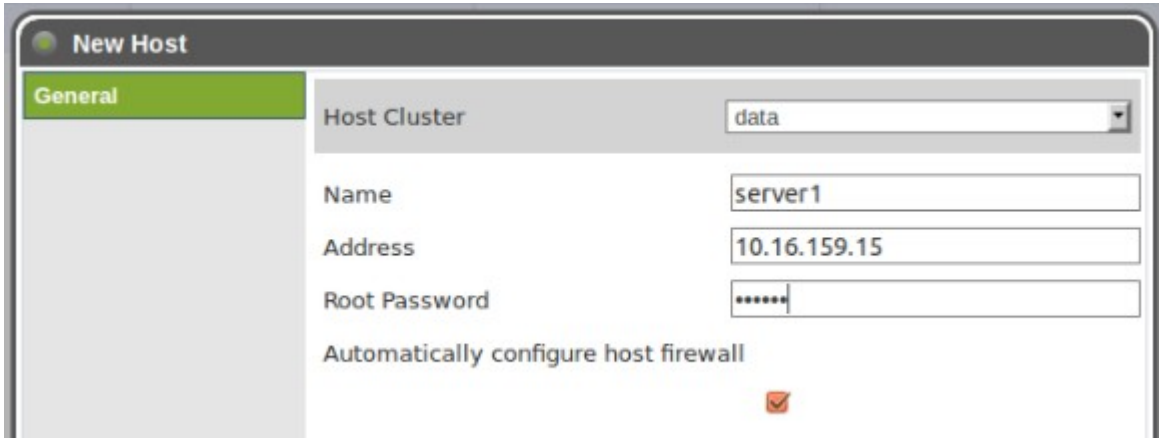



theron@redhat.com

# Gluster Storage Domains

- General flow
  - Create Volume
  - Add bricks

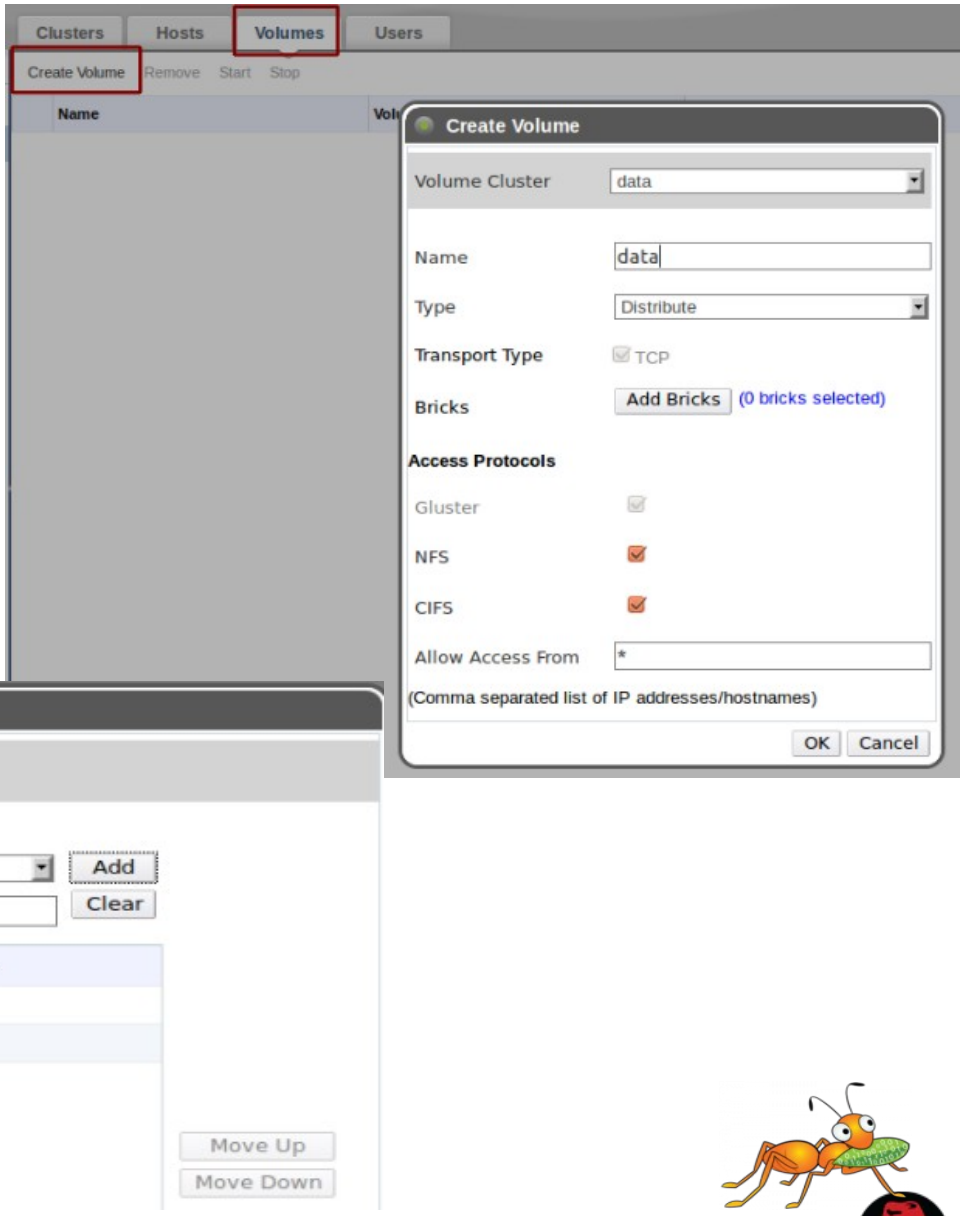




theron@redhat.com

# Gluster Storage Domains

- Create PosixFS storage domain (VFS type = glusterfs)

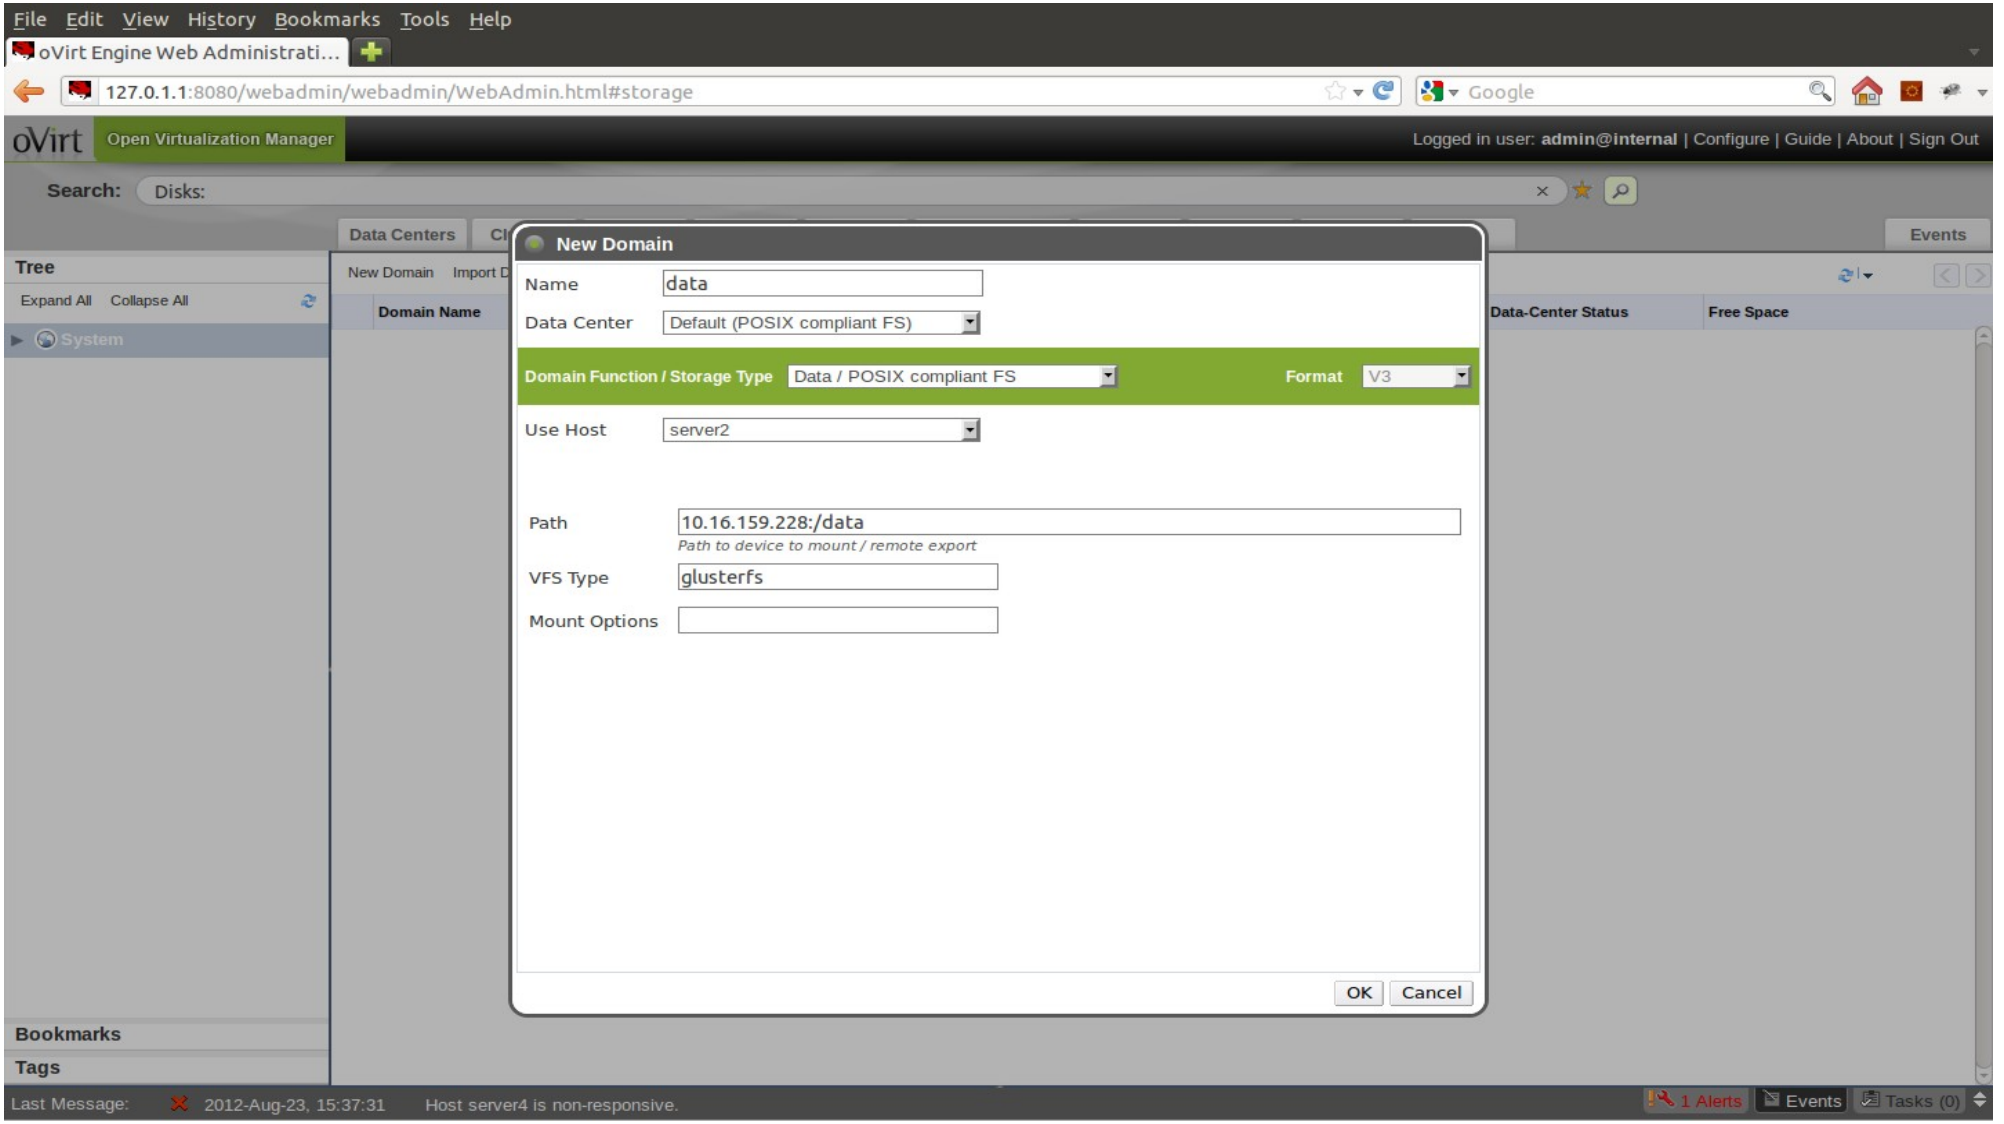




theron@redhat.com

# What's Coming in GlusterFS 3.4?

- QEMU-GlusterFS integration
- Native integration using libgfapi, No FUSE MOUNT
- QEMU leveraging GlusterFS natively on the back end

  **-drive file=gluster://server[:port]/volname/image[?transport=...]**
- Block device support in GlusterFS via Block Device translator
- Enhanced quorum
- Multi-threaded glusterd, NFSv3 ACLs
- In alpha (nudge, nudge), entering Beta soon


theron@redhat.com

# Demo time.....



theron@redhat.com

# Solid foundational elements.  What's left?

- oVirt
- Gluster
- Automation (command and control)
  - Ansible
  - Chef
  - Crowbar
  - Puppet


theron@redhat.com



theron@redhat.com

# A call to action!

- Git repository
  - git://gerrit.ovirt.org/ovirt-engine
  - git://gerrit.ovirt.org/vdsm
  - git://github.com/gluster/glusterfs.git
- IRC Channel
  - #ovirt on oftc
  - #gluster on freenode
- Websites
  - http://www.gluster.org
  - http://www.ovirt.org


theron@redhat.com


http://www.flickr.com/photos/mtaphotos/7337564952/



theron@redhat.com


http://www.flickr.com/photos/arnolddeleon/8009929099/



theron@redhat.com

http://www.flickr.com/photos/oneaustin/1261907803/



theron@redhat.com

# A call to action!

- Git repository
  - git://gerrit.ovirt.org/ovirt-engine
  - git://gerrit.ovirt.org/vdsm
  - git://github.com/gluster/glusterfs.git
- 
- IRC Channel
  - #ovirt on oftc
  - #gluster on freenode
- 
- Websites
  - http://www.gluster.org
  - http://www.ovirt.org


theron@redhat.com

# Thank you.

Theron Conrey

Open Source and Standards, Red Hat

@theronconrey