



# Red Hat HyperConverged Infrastructure

RHUG Q3.2017

Marc Skinner

Principal Solutions Architect

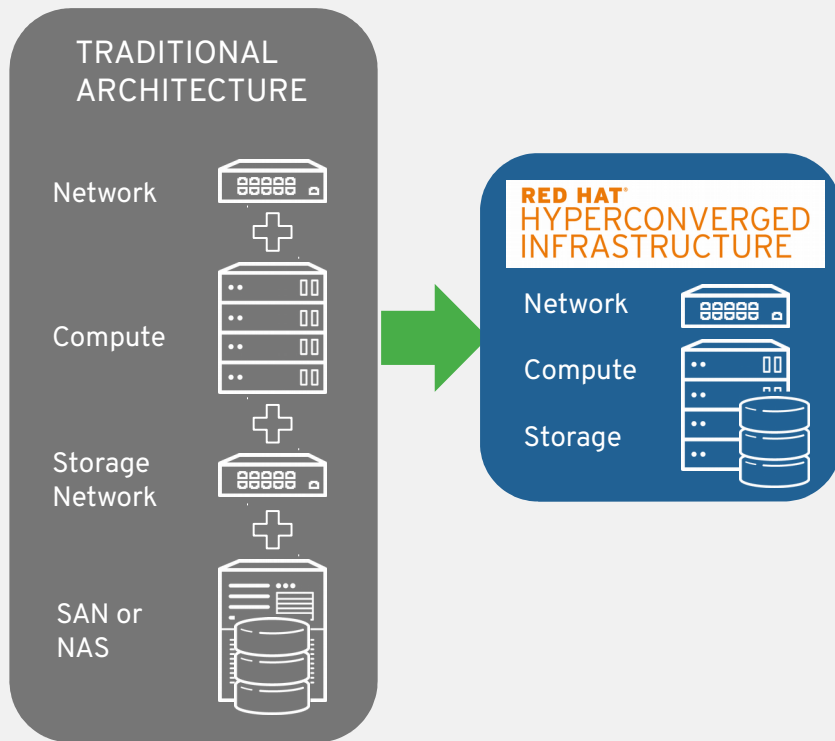
8/23/2017

# AGENDA

- What is RHHI?
- Use cases
- Technology
- Features
- Installation

What is RHHI?

# INFRASTRUCTURE CONSOLIDATION & OPERATIONAL EFFICIENCY

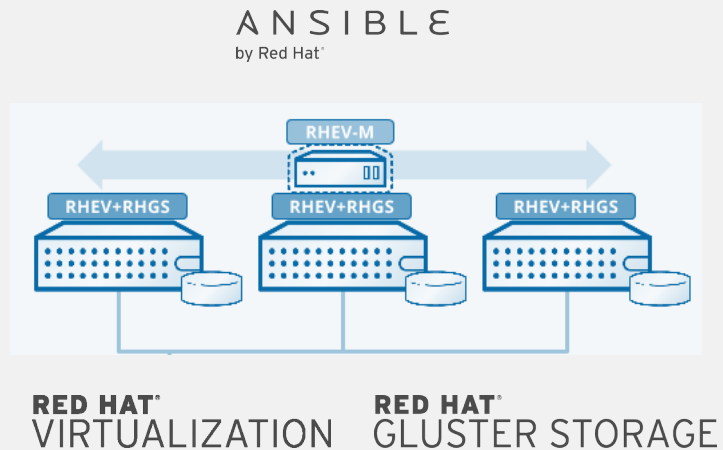


- **Eliminate** storage as a discrete tier
- **Easily virtualize** business applications, maximizing resource utilization
- **Single budget** for compute & storage
- **Single team** managing infrastructure
- **Simplified** planning & procurement
- **Streamlined** deployment & management
- **Single support** stack for compute & storage

# RHHI ARCHITECTURE

Self healing and highly available

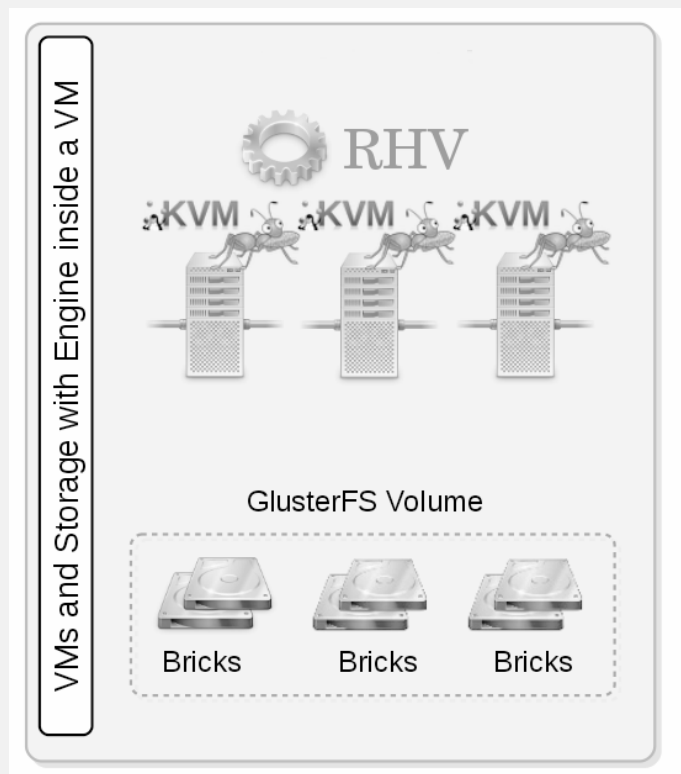
- 3-node base POD configuration
- RHGS with 3-way replica, Arbiter configs
- RHV-H with Self-Hosted Engine
- SSD cache fronting spinning media
- HW & SW monitoring, fault correlation
- Ansible based deployment tool & playbook



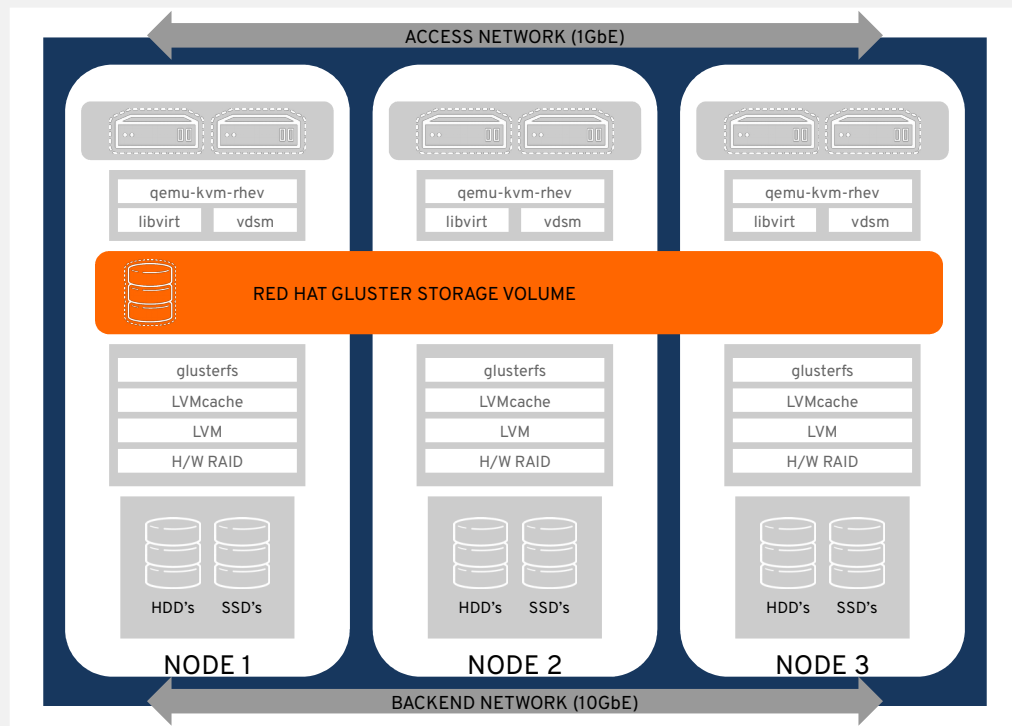
Software Solution with Reference Architecture

# RHHI ARCHITECTURE

- Hosted Engine + GlusterFS
- Same nodes used to:
  - Host the engine
  - Run VMs
  - Provide shared storage
- Storage is now highly available, redundant and **local**



# ARCHITECTURAL POD VIEW



3, 6 or 9 Node Pod Configuration

# SOFTWARE DEFINED INFRASTRUCTURE



## RED HAT HYPERCONVERGED INFRASTRUCTURE

Provides compute and storage in a single resource pool that is easily managed.



## ANSIBLE by Red Hat

Automates the deployment including software credentials & channels, package installation, security configuration, storage configuration, virtualization setup, high availability for RHV, replication for RHGS, and storage monitoring configuration.



## REST API

Provides means to integrate with other management tools including Red Hat CloudForms, Red Hat Satellite, or third-party tools. Additionally, software development kits are available (Java, Ruby, and Python) to extend functionality.




## Software Defined Networking (SDN) \*

Provides SDN capabilities by offering native tech preview\* support for Open Virtual Network (OVN) for Open vSwitch.

\* SDN is Tech Preview in RHV4



# CENTRALIZED MANAGEMENT OF YOUR RED HAT HYPERCONVERGED INFRASTRUCTURE INSTALLATIONS



RED HAT CLOUDFORMS (optional)  
Manage multiple sites centrally, from a  
single pane of glass.



Site 1



Site 2



Site 10



Site 20




Site 50

# Use cases

# REMOTE OFFICE/BRANCH OFFICE OR STORES

## PRIMARY USE CASE


- Seeking overall reduction in TCO
- Need infrastructure consolidation
- Need reduced footprint - power/cooling costs expanding with traditional models
- Dealing with too many vendors - ease of acquisition/support
- Need to keep key applications local to the remote site



# INTERNET of THINGS

## PRIMARY USE CASE

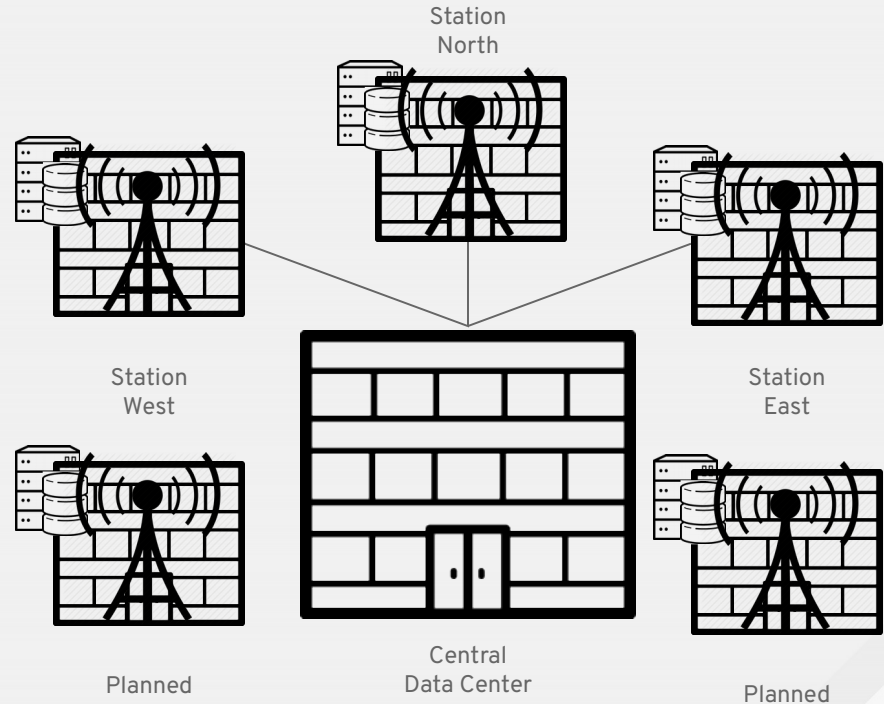
- Implement a robust intelligent gateway tier
- Deploy compute and storage resources closer to endpoints
- Red Hat Hyperconverged Infrastructure becomes a “micro-datacenter” for IoT



# EDGE COMPUTING


## PRIMARY USE CASE

- Deploy compute and storage resources closer to cellular customers
- Distributed infrastructure reduces cellular network congestion
- Enhance network performance and build additional resiliency



# Technology

# RED HAT HYPERCONVERGED INFRASTRUCTURE CORE COMPONENTS



## RED HAT VIRTUALIZATION


- Open source virtualization built on RHEL 7 and KVM
- High availability for VMs and RHV-M
- Security features like SELinux and sVirt inherited from RHEL



## RED HAT GLUSTER STORAGE

- Open, software-defined storage
- Modular architecture allows easy addition of features
- Data Replication with self healing features

# RHV SELF-HOSTED ENGINE



## Self-Hosted Engine

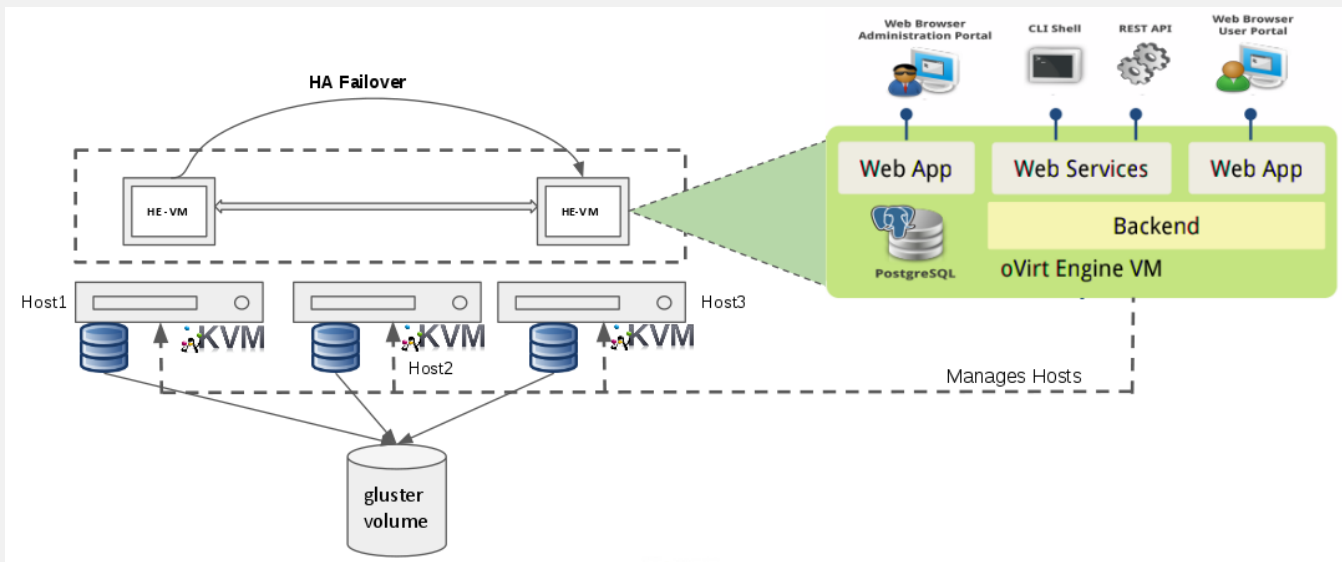
- High Availability for RHV-M
- Reduced hardware requirements for RHV
- RHV-M deployed as an appliance



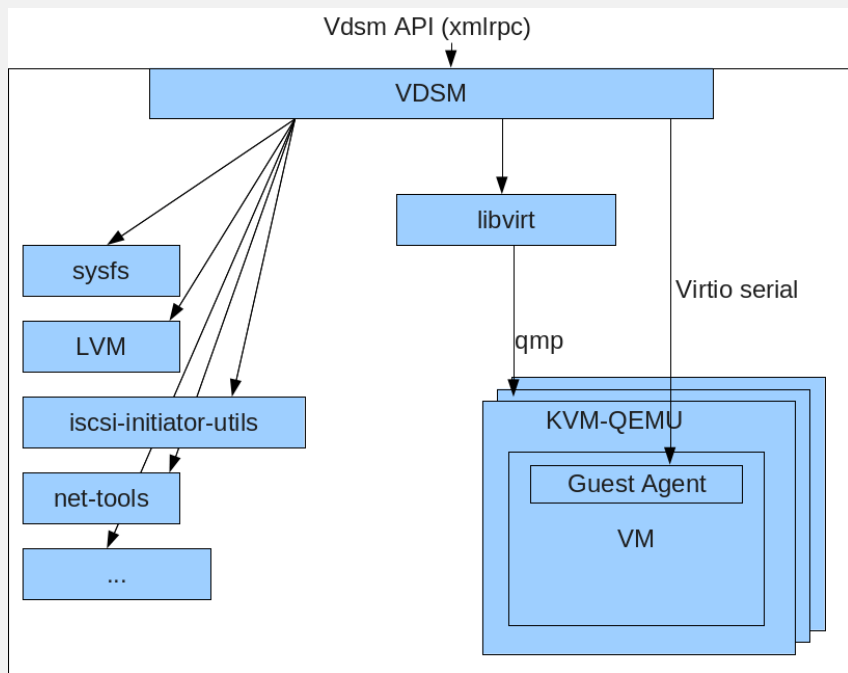
# RHV HOSTED ENGINE HA

## HA achieved using ovirt-ha-broker and ovirt-ha-agent

- Continuously monitors engine VM health
- Uses scoring mechanism to determine the best host to run
- Takes care of migrating and restarting the hosted engine in case of failures
- Engine VM image is stored on the Gluster volume

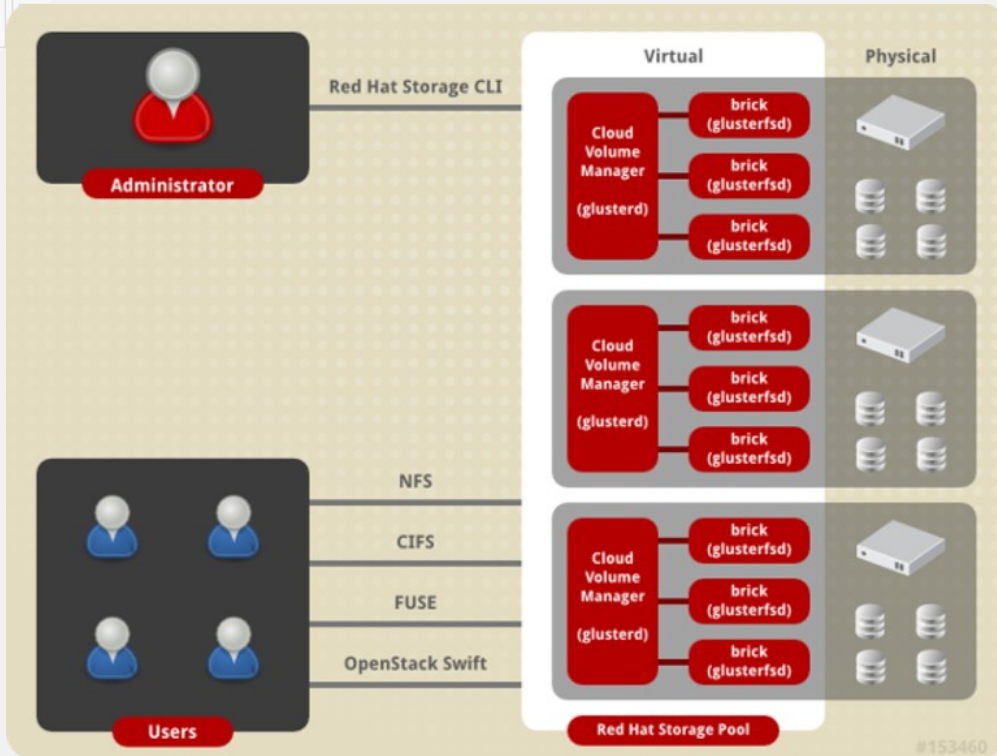


# VDSM :: daemon



- Agent that runs on all RHHI hosts
- Lifecycle managed by systemd
- Configures host, networking and shared storage
- Uses libvirt for VM lifecycle operations
- Multithreaded, can have multiple processes based on number of storage domains configured
- **Uses 2 CPU cores, pinned to cores 0 & 1 by default**


# GLUSTERED :: daemon




- Management daemon for Gluster, runs on all RHHI nodes
- Systemd based
- Spawns all gluster services
- Glusterfsd - for serving data to clients (1 per brick by default)
- Glustershd - Self-healing daemon
- **CGroups based resource allocation limits glusterd & related services to 4 cpu cores**

# REPLICA 3 VS ARBITER CONFIG

- GlusterFS replica 3 volumes keep 3 copies of data
- Only two copies of data are needed to make data redundant...
- ...while three participants are required to make a quorum




OR




- 2 bricks of data
- 1 brick quorum or arbiter
- Less space
- Better performance

# ROADMAP - PERFORMANT I/O :: libgfapi



**FUSE** access



**libgfapi** access

# Features

# RED HAT HYPERCONVERGED INFRASTRUCTURE 1.0

## FEATURES - CORE

Single Point of Management for Virtual Resources

HA for Virtual Machines & Management

Automated Resource Mgmt/Load Balancing

CPU Pinning

VM Templates

CPU QoS

RBAC & Tiered Access

Secure Browser Based Management

Hot Add Memory & CPU

Power Management

Streamlined Deployment & Operations

Advanced Live Migration Policies

Support for RHEL & Windows Workloads

Python, Ruby, & Java SDKs

Live Migration

Firewall/SELinux

REST API / Integrate w/ Red Hat Portfolio

# RED HAT HYPERCONVERGED INFRASTRUCTURE 1.0 FEATURES - NETWORK

VLAN Tagging

Open Virtual Network (Tech Preview)

Network QoS

IPv6 Support (guest)

NIC Bonding

Jumbo Frames

VM-FEX Support

Network Labels



# RED HAT HYPERCONVERGED INFRASTRUCTURE 1.0

## FEATURES - STORAGE

Geo-replication

REST API for backup/restore

Sharding support

3-way Data Replication

Live Snapshots/Merge

Thin & Thick Provisioning

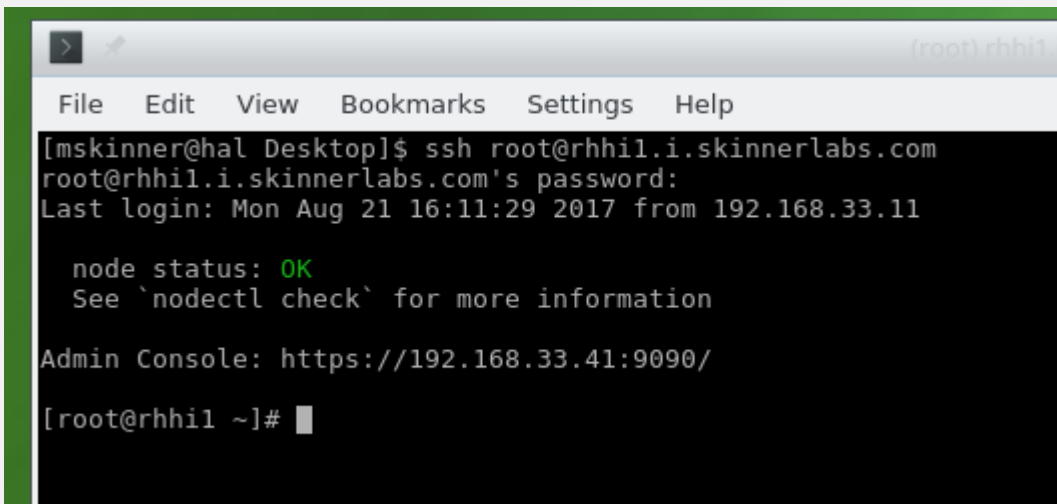
Block discard

Storage-based fencing

SSD Caching

# Installation

# RHV-H :: CONSOLE

A terminal window with a green title bar. The window title is "(root) rhhi1...". The menu bar contains "File", "Edit", "View", "Bookmarks", "Settings", and "Help". The terminal content shows a user logging in via SSH to root@rhhi1.i.skinnerlabs.com. The user is prompted for a password, and the login is successful. The terminal displays "node status: OK" and "See `nodectl check` for more information". It also shows the "Admin Console" URL: "https://192.168.33.41:9090/". The prompt is "[root@rhhi1 ~]#".


```
(root) rhhi1...  
File Edit View Bookmarks Settings Help  
[mskinner@hal Desktop]$ ssh root@rhhi1.i.skinnerlabs.com  
root@rhhi1.i.skinnerlabs.com's password:  
Last login: Mon Aug 21 16:11:29 2017 from 192.168.33.11  
  
node status: OK  
See `nodectl check` for more information  
  
Admin Console: https://192.168.33.41:9090/  
  
[root@rhhi1 ~]#
```

\* SSH keys need to be shared to all nodes, all interfaces, before installation

# ANSIBLE BASED DEPLOYMENT :: COCKPIT

https://rhh1.i.skinnerlabs.com:9090

Search

 redhat.

**RED HAT VIRTUALIZATION HOST 4.1 (EL7.4)**

User name

Password

Reuse my password for privileged tasks

▶ Other Options

Server: **rhh1.i.skinnerlabs.com**  
Log in with your server user account.

# ANSIBLE BASED DEPLOYMENT :: COCKPIT

The screenshot displays the Red Hat Virtualization Host 4.1 (EL7.4) Cockpit interface. The browser address bar shows `https://rhh1.i.skinnerlabs.com:9090/network#/ib1`. The main content area is titled "Virtualization Dashboard" and "Networking > ib1". A "Kbps Sending" graph is visible, showing data points at 16:13, 16:14, and 16:15. Below the graph, the network interface "ib1" is identified as "Mellanox Technologies MT25208 [InfiniHost III Ex] ib\_mthca". The status is "192.168.104.41/24, fe80:0:0:86f6:b44e:f9df:e843/64". The carrier is "Yes". Under the "General" section, there is a checkbox for "Connect automatically" which is currently unchecked. Below this, the IPv4 address is listed as "Address 192.168.104.41/24" and IPv6 as "Automatic".

An "IPv4 Settings" dialog box is open in the foreground, showing the following configuration:

- Addresses:** A dropdown menu is set to "Manual". There is a "+" button to the right. Below, there are three input fields: "192.168.104.41", "24", and "Gateway". A "-" button is to the right of the "Gateway" field.
- DNS:** Set to "Automatic" with an "ON" toggle and a "+" button.
- DNS Search Domains:** Set to "Automatic" with an "ON" toggle and a "+" button.
- Routes:** Set to "Automatic" with an "ON" toggle and a "+" button.

At the bottom of the dialog box, there are "Cancel" and "Apply" buttons.

# ANSIBLE BASED DEPLOYMENT :: COCKPIT

The screenshot shows a web browser window with the URL `https://rhh1.i.skinnerlabs.com:9090/ovirt-dashboard#/he`. The browser's address bar includes a search field and navigation icons. The page title is "RED HAT VIRTUALIZATION HOST 4.1 (EL7.4)" and the user is logged in as "root". The navigation menu on the left includes "Dashboard", "Hosted Engine", and "Virtual Machines". The main content area is titled "Hosted Engine Setup" and features a server rack icon. Below the icon, the text reads: "Configure and install a highly-available virtual machine which will run oVirt Engine to manage multiple compute nodes, or add this system to an existing hosted engine cluster." There are two radio button options: "Standard" (unselected) and "Hosted Engine with Gluster" (selected). A blue "Start" button is positioned below the options.

# ANSIBLE BASED DEPLOYMENT :: COCKPIT

Gluster Deployment ✕

Hosts Packages Volumes Bricks Review

1 2 3 4 5

Host1

Host2

Host3

**i** gdeploy will login to gluster hosts as root user using passwordless ssh connections. Make sure, passwordless ssh is configured for all gluster hosts from the first host.

Cancel < Back Next >

# ANSIBLE BASED DEPLOYMENT :: COCKPIT

### Gluster Deployment

Hosts Packages Volumes Bricks Review

1 2 3 4 5

Name	Volume Type	Arbiter	Brick Dirs	
engine	Replicate	<input checked="" type="checkbox"/>	/gluster_bricks/engine/engine	
data	Replicate	<input checked="" type="checkbox"/>	/gluster_bricks/data/data	
vmstore	Replicate	<input checked="" type="checkbox"/>	/gluster_bricks/vmstore/vmsto	

[Add Volume](#)

**3** First volume in the list will be used for hosted-engine deployment

Cancel < Back Next >



# ANSIBLE BASED DEPLOYMENT :: COCKPIT

### Gluster Deployment

Hosts Packages Volumes Bricks Review

1 2 3 4 5

#### Raid Information

Raid Type: JBOD

Stripe Size(KB): 256

Data Disk Count: 3

#### Brick Configuration

LV Name	Device Name	Size(GB)	Thinp	Mount Point
engine	md0	100	<input type="checkbox"/>	/gluster_bricks/engine
data	md0	1024	<input checked="" type="checkbox"/>	/gluster_bricks/data
vmstore	md0	1024	<input checked="" type="checkbox"/>	/gluster_bricks/vmstore

[Add Bricks](#)

**Arbiter bricks will be created on the third host in the host list.**

Cancel < Back Next >

# ANSIBLE BASED DEPLOYMENT :: COCKPIT

Gluster Deployment ✕

Hosts      Packages      Volumes      Bricks      Review

① ——— ② ——— ③ ——— ④ ——— ⑤

Generated Gdeploy configuration : /tmp/gdeployConfig.conf Edit Reload

```
#gdeploy configuration generated by cockpit-gluster plugin
[hosts]
192.168.103.41
192.168.103.42
192.168.103.43

[script1]
action=execute
ignore_script_errors=no
file=/usr/share/ansible/gdeploy/scripts/grafon-sanity-check.sh -d md0 -h 192.168.103.41,192.168.103.42,192.168.103.43

[disktype]
jbod
```


Cancel < Back Deploy

# ANSIBLE BASED DEPLOYMENT :: COCKPIT

Gluster Deployment ✕

Hosts Packages Volumes Bricks Review

① — ② — ③ — ④ — ⑤



Successfully deployed Gluster

[Continue to Hosted Engine Deployment](#)

Cancel < Back Close

# ANSIBLE BASED DEPLOYMENT :: COCKPIT

Continuing will configure this host for serving as hypervisor and create a VM where you have to install the engine afterwards.  
Are you sure you want to continue? (Yes, No)[Yes]:



THANK YOU



[plus.google.com/+RedHat](https://plus.google.com/+RedHat)



[facebook.com/redhatinc](https://facebook.com/redhatinc)



[linkedin.com/company/red-hat](https://linkedin.com/company/red-hat)



[twitter.com/RedHatNews](https://twitter.com/RedHatNews)



[youtube.com/user/RedHatVideos](https://youtube.com/user/RedHatVideos)