oVirt Node

May 9, 2013

Mike Burns
Software Engineer
Tech Lead for oVirt Node
Red Hat, Inc

Agenda

- Introduction
- Architecture and Packaging
- Configuration Persistence
- Installation and Configuration
- Plugins
- Stateless
- Future Features
- Discussion

Introduction

oVirt Architecture

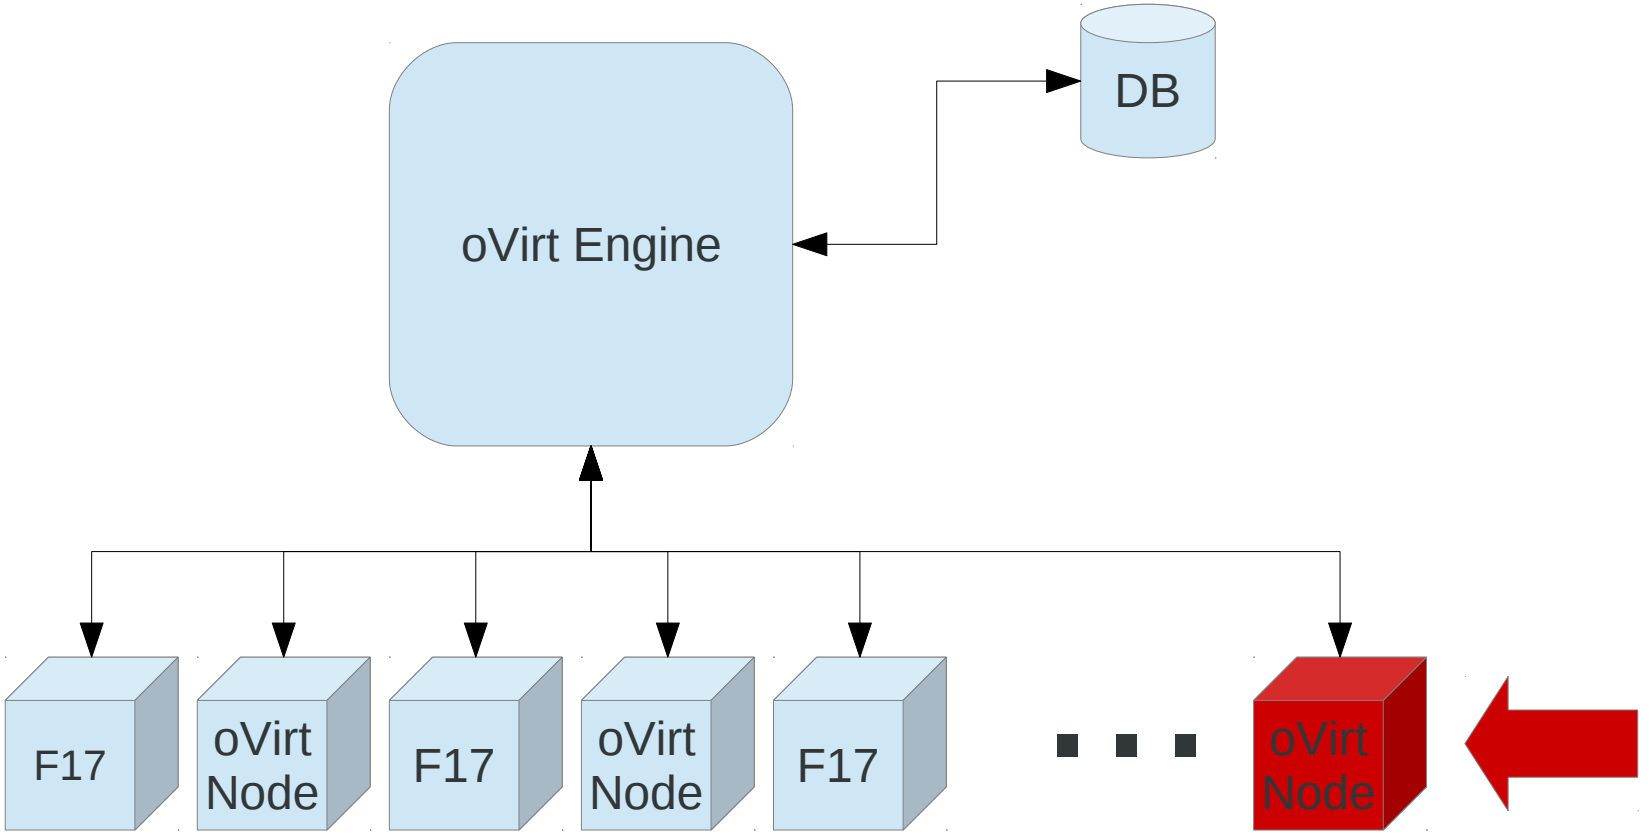

Virtual Machine Hosts
oVirt Node
Fedora

Is your datacenter this diverse?

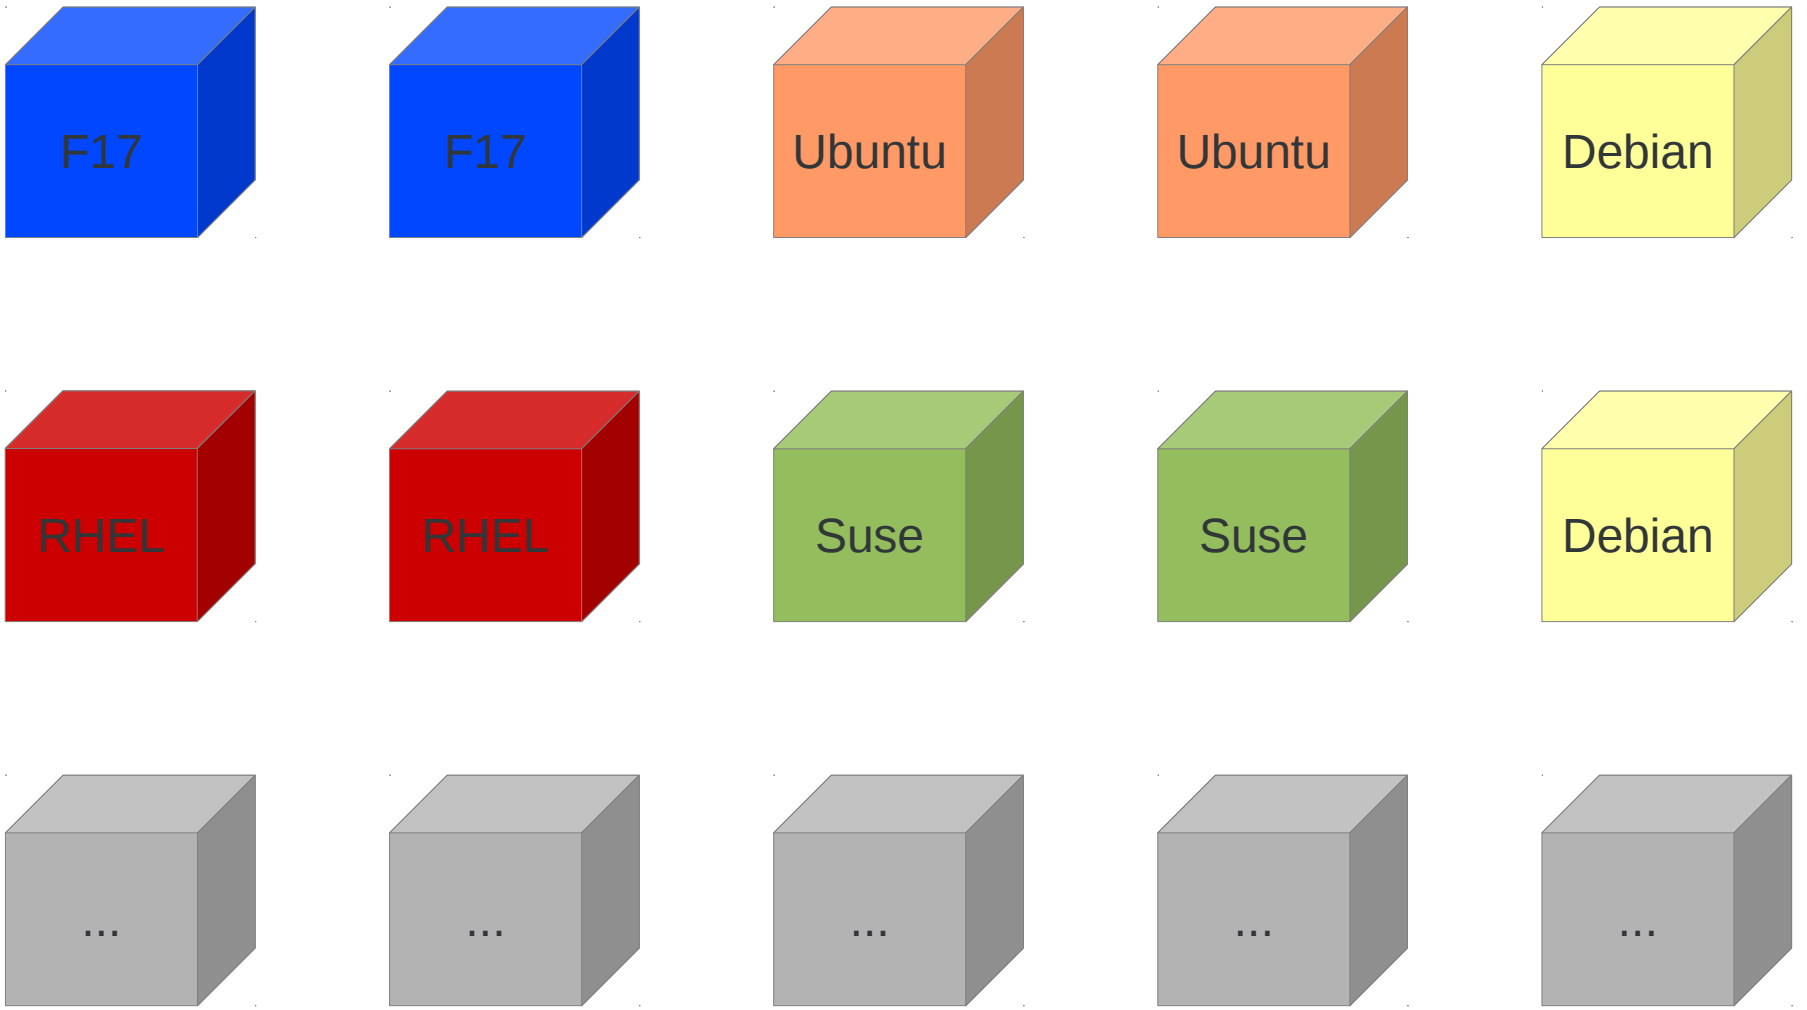

What if it could look like this?

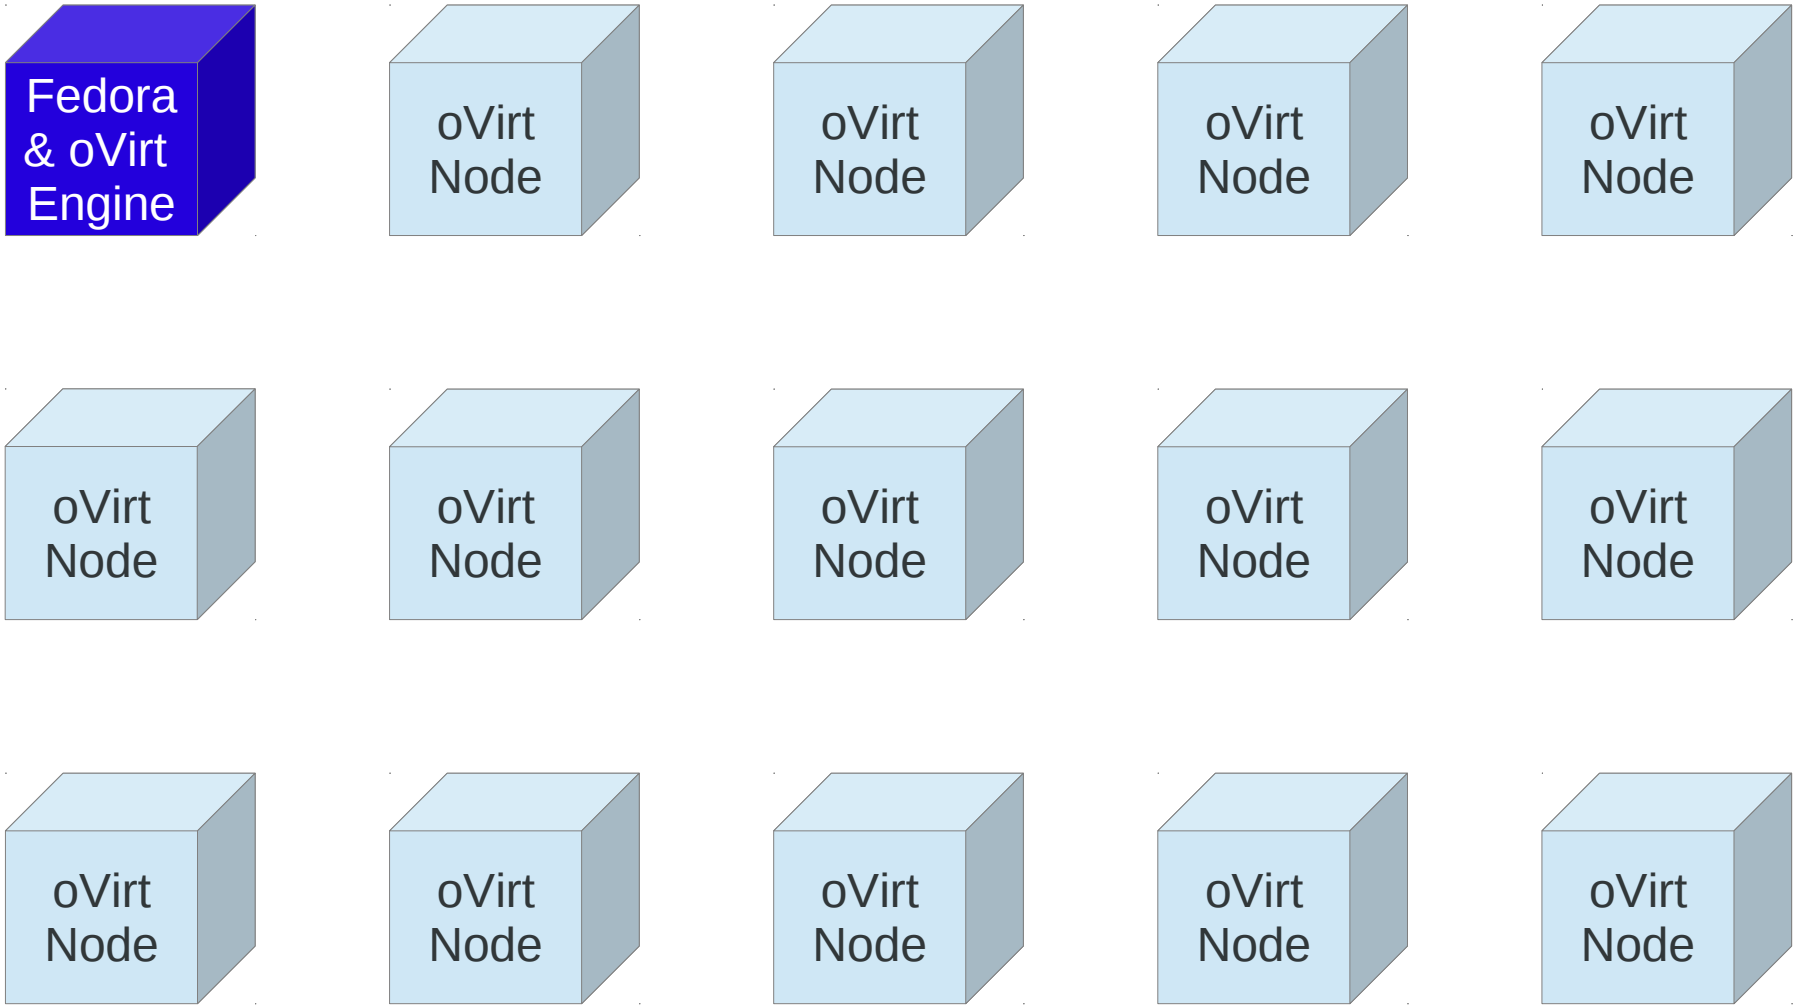

What is oVirt Node?

- Dedicated hypervisor
- JEOS
- livecd
- Built on Fedora
- Firmware
  - Install and forget about it
  - Similar to ESXi or OpenWRT
- Small Footprint (< 200MB)

Advantages and Disadvantages

- Advantages
  - Single image
  - Easy Upgrades
  - No managing individual package updates
  - Upgrade directly from management
- Disadvantages
  - Lack of Customization
  - No easy shell access
  - More difficult to debug problems

Architecture and Packaging

Architecture

- ISO Image
- Created using standard packages in Fedora
- Generated based on layered kickstarts
- %post scripts handle default configuration and setup
- TUI provided for installation and configuration post install
- Source Repositories
  - ovirt-node
  - ovirt-node-iso

Packaging

- Packaged into 4 distinct RPMs
  - ovirt-node – configuring the ISO
  - ovirt-node-tools – working with ISO
  - ovirt-node-recipe – building the ISO
  - ovirt-node-iso – wraps the ISO
- Additional RPMs to be added in the future
  - ovirt-node-plugin-* - customizations for the ISO

Key Technologies

- qemu-kvm - http://www.linux-kvm.org/
- libvirt - http://libvirt.org/
- spice - http://spice-space.org/
- device-mapper-multipath
  - http://christophe.varoqui.free.fr/
- newt/snack
  - http://en.wikipedia.org/wiki/Newt_(programming_library)
- Livecd-tools
  - http://fedoraproject.org/wiki/FedoraLiveCD

Configuration Persistence

Configuration Persistence

- Root FS is non-persistent by default
  - On reboot, the original filesystem is loaded
- Root FS is mounted readonly
- Some things do need to be persisted across reboots
  - persist and unpersist commands added
- Persistent changes are stored in /config
  - Limited space available by default (8 MB)
  - bind-mounted automatically at boot time
- Some packages will handle this automatically
  - vdsm and ovirt-node

Installation and Configuration

Deployment Modes

- oVirt Node can be installed using a variety of methods
  - CD/DVD-Rom (including virtual CD)
  - Flash Memory (USB or SD Card)
  - Network (PXE)
- Limited stateless support
- Install to disk
  - Can be either HDD or Flash disk (USB or SD Card)
- Installation Methods
  - Automatic
  - Manual

Automatic Installation

- Triggered using kernel command line parameters
  - Requires storage_init and BOOTIF
  - Should include additional options or adminpw if you want to configure things later
- All configurations done on the TUI can be done through kernel command line options
- After installation completes, machine will reboot automatically

Manual Installation

- Done using a TUI interface
- Keyboard driven

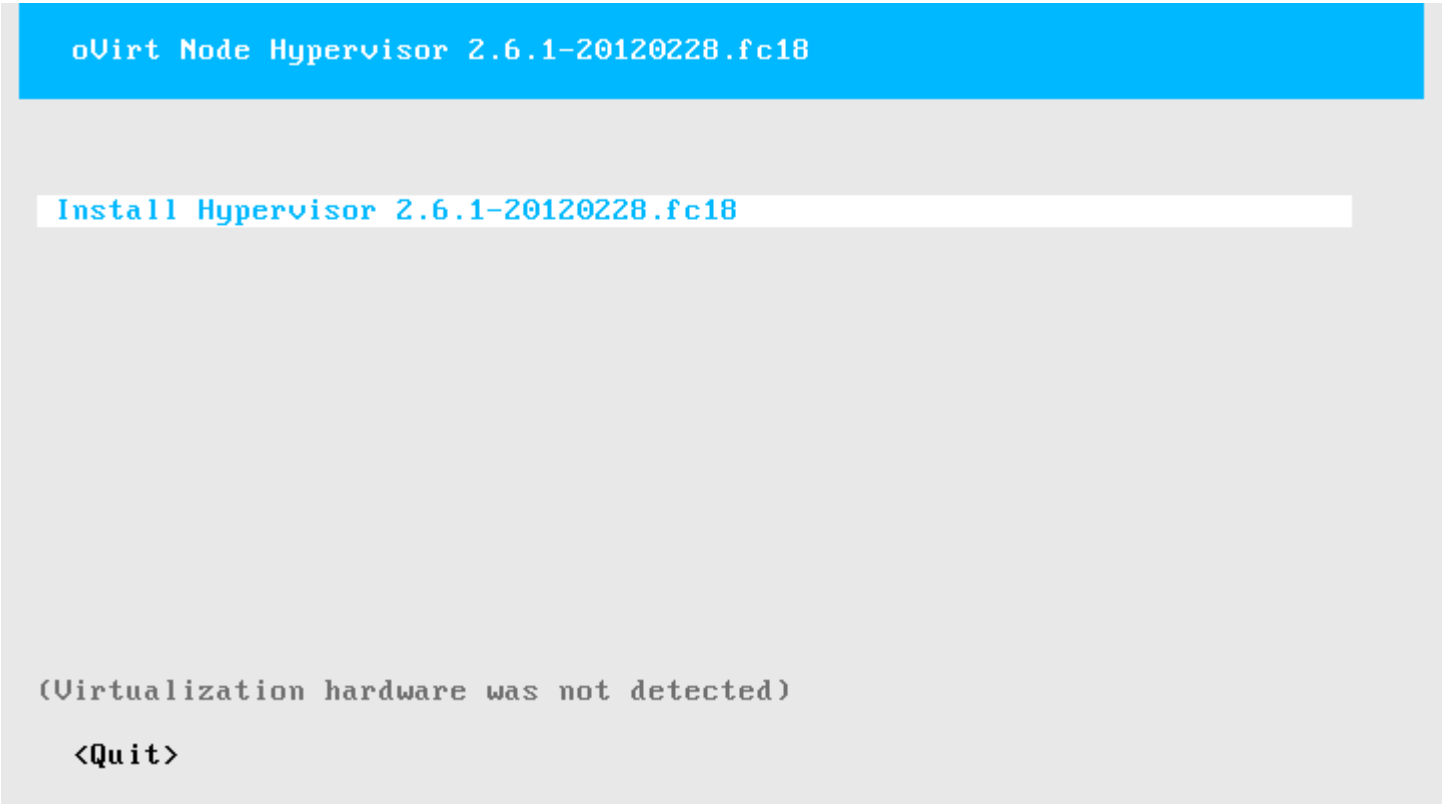

Configuration

- After installation or stateless boot
  - Login as admin to access the TUI to make changes

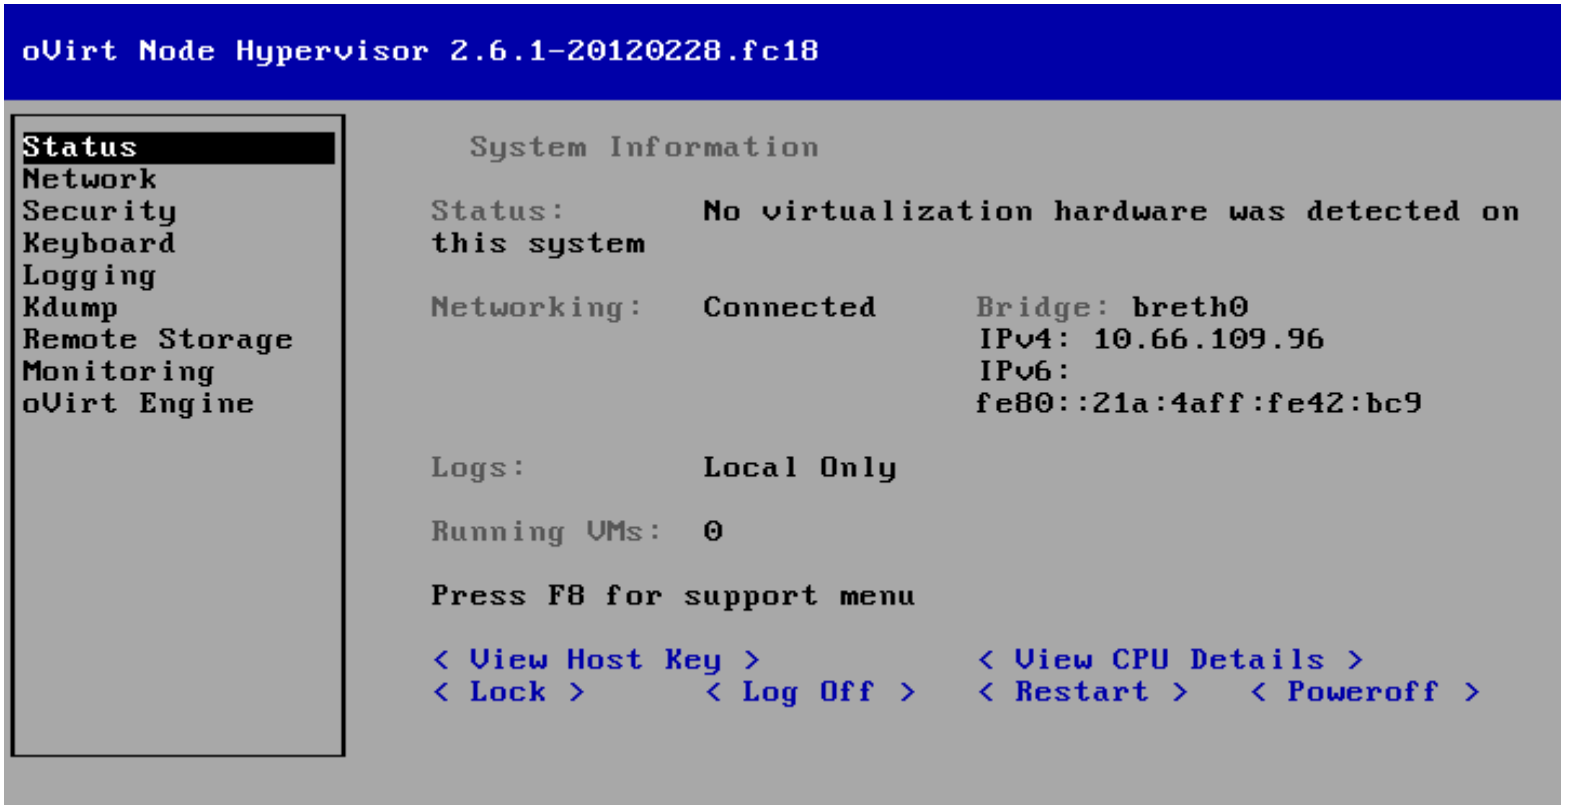

Installation Disk Layout

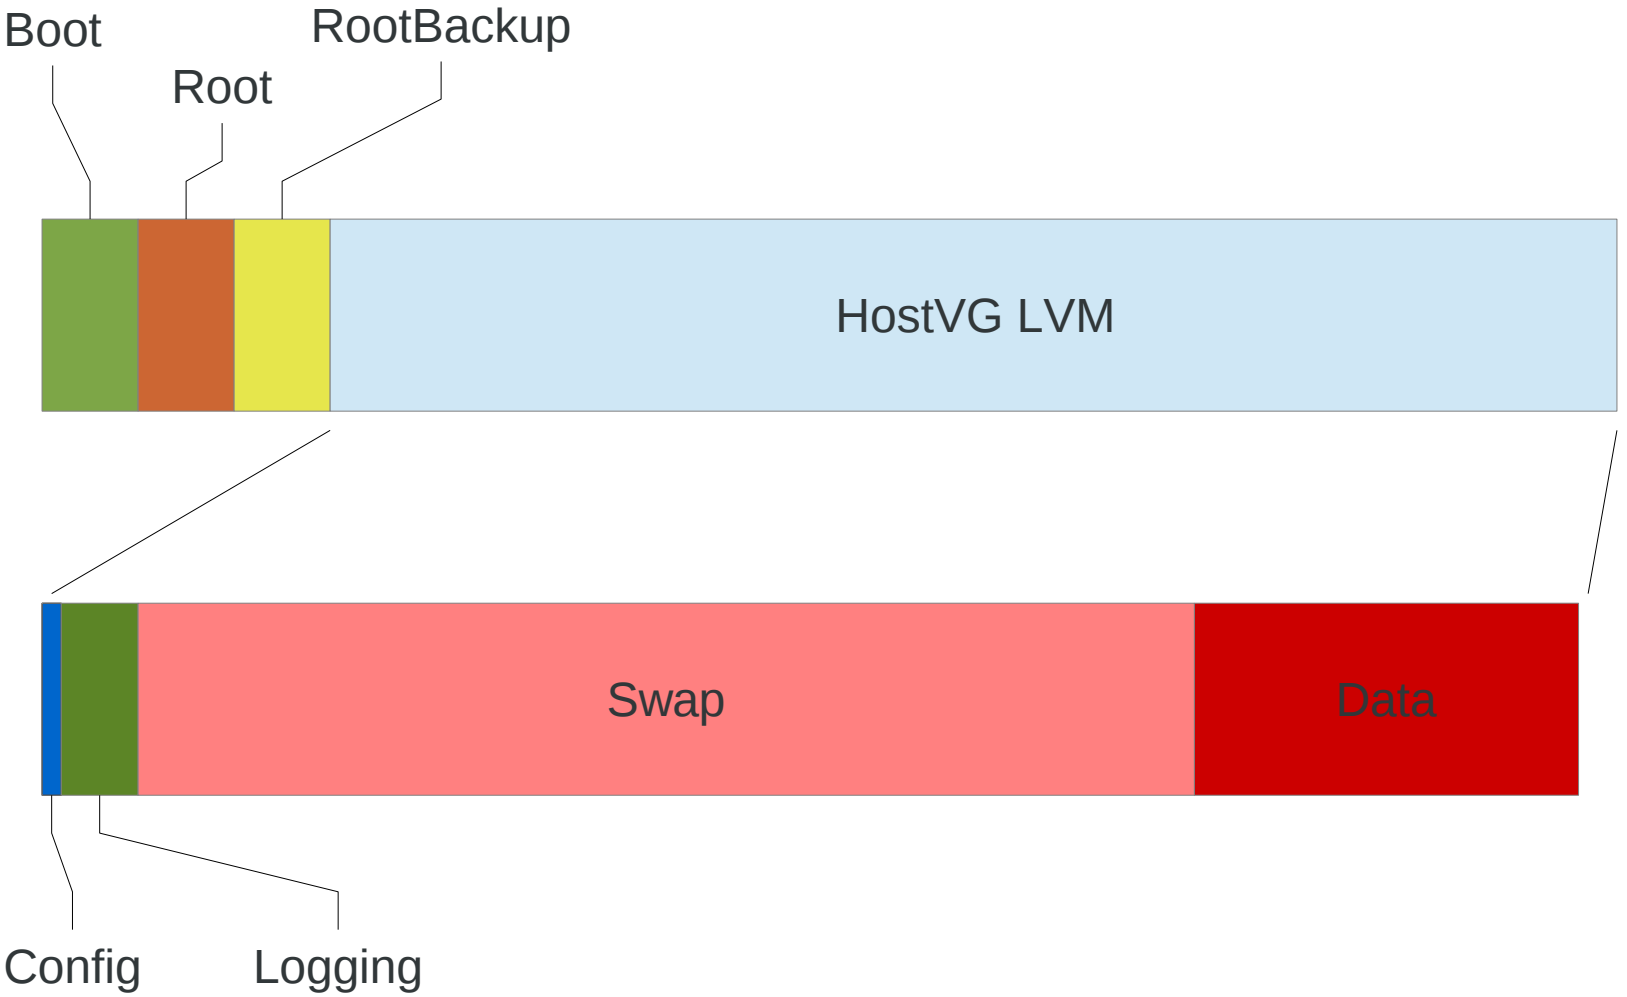

Upgrades

- Usually as simple as booting the new image
  - Update the PXE image
  - Boot new CD/USB/SD
  - In Place Upgrade
    - Upload new image to running system
    - Trigger Upgrade logic
    - Used by oVirt Engine
- Can be done automatically using the command line
- Can be done through TUI
- Clean installs can be triggered with a command line option as well

RootBackup

- Provide roll back capability in the case of a bad upgrade
- Using Grub savedefault
- Upgrade ISO gets installed into RootBackup partition
- Partitions are renamed
  - RootBackup -> RootNew
  - Root -> RootBackup
  - RootNew to Root
- If machine fails to boot, it rolls back to RootBackup

Plugins

What are Plugins?

- Preview in oVirt Node 2.5.1
- Add functionality not included in the base image
- Packaged as RPMs
- Installed offline using edit-node tool (ovirt-node-tools)
  - Start with oVirt Node ISO image
  - Run edit-node tool
  - Get a different oVirt Node ISO image
- Can install arbitrary number of plugins

Plugin Examples

- Update default passwords
- Install or update packages
- Install new kernel modules
- Add vdsm hooks

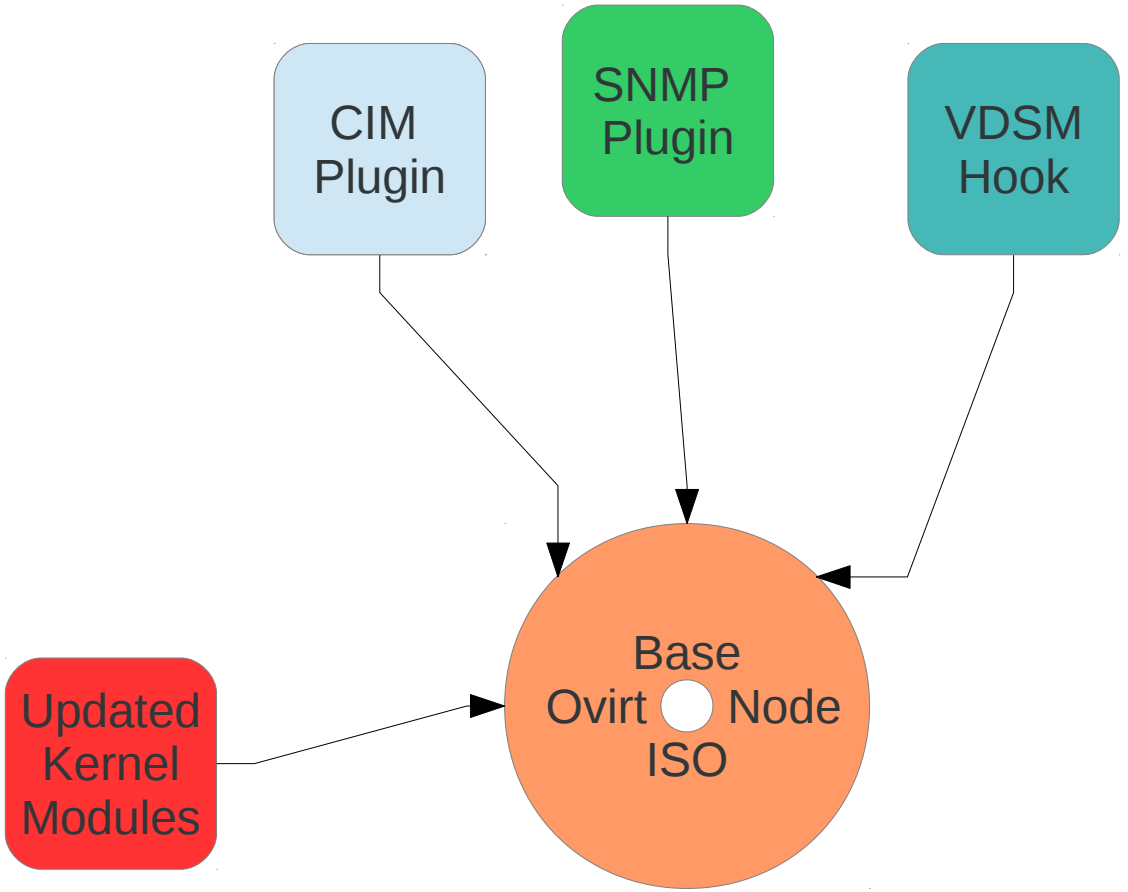

Stateless

Current Support

- Enabled by passing “stateless” on the command line
- Ignores all local storage
- persist and unpersist commands do nothing
- Configuration TUI works the same as in an installed system
- Honors all regular kernel command line options except those having to do with local storage

Current Limitations

- No support for swap
  - This means that you need to disable overcommit in oVirt Engine
- No local storage partition
  - Local Storage Domains in oVirt Engine are not supported
- oVirt Engine does not understand the concept of stateless nodes
  - Have to re-register and approve every reboot

Roadmap

Additional Use Cases

- Non-oVirt use cases
- Can be done by utilizing Plugins
- Investigating OpenStack and Gluster as possible consumers
- Steps needed
  - Remove vdsm from ovirt-node default build and make it a plugin
  - Develop plugins that would pull in appropriate packages for OpenStack and/or Gluster
- Base image is generic and not used for anything
- Admin would use edit-node to install their plugin(s)

Plugin Enhancements

- Add some net-new plugins
  - OpenStack, Gluster, etc...
- Remove some functionality from base image to minimize size
  - SNMP, CIM, etc...
- Additional supported features
  - Service enablement
  - Firewall configuration
  - Auto-installation process
  - Security Profiles

Other Future Features

- Software iSCSI Root Support
- Network Manager
  - Move away from using ifcfg scripts directly to using NM libraries
  - Depends on NetworkManager support for bridges, bonds, and vlans
- Stateless
  - Fix the limitations on swap and local storage domain
- UI Enhancements
  - Make code more re-usable to ease TUI Plugin design

    Allow different size screens (currently only 80x20)

Discussion and Questions

More information

- http://www.ovirt.org/get-ovirt/ (Installation guide available)
- Mailing Lists:
  - node-devel@ovirt.org
  - users@ovirt.org
- IRC: #ovirt on OFTC
- Web Site: http://www.ovirt.org
- Git Repository: git://gerrit.ovirt.org/ovirt-node.git
- Documents: http://www.ovirt.org/wiki/Special:AllPages
- Bugzilla: https://bugzilla.redhat.com (Community->oVirt)

THANK YOU !

http://www.ovirt.org