oVirt Node

Mar 21, 2012

Mike Burns

Agenda

- Introduction
- Architecture Overview
- Deployment Modes
- Installation and Configuration
- Upgrading
- Configuration Persistence
- Future Features
- Discussion

Introduction to oVirt Node

- Dedicated Hypervisor
  - Minimum OS needed to run and manage VMs
  - Well defined management interfaces and APIs
- Small Footprint
  - Less than 150MB image size
  - ~850MB disk space required for installation
  - 512MB runtime RAM footprint
- Built from Fedora components
  - Supports same hardware as Fedora
- Utilizes KVM
- Includes VDSM for VM Management

Architecture – RPMs

- ovirt-node
  - A TUI interface for installation and configuration
  - Automatic installation and configuration
  - Modifications to sysinit scripts
- ovirt-node-tools
  - A set of kickstart files
  - Minimal package list
  - Blacklisting for image minimization
  - Used to create an image

Architecture - Image

- Built using kickstarts from ovirt-node-tools package
- Provided in ISO format
- Installs to local hard drive, flash drive, SD card
- Minimal state information persisted to a config partition
- Multiple upgrade paths
  - Booting a new image
  - In place upgrade

livecd-tools

- Utility for create an ISO image containing a LiveOS
  - http://fedoraproject.org/wiki/FedoraLiveCD
- Uses the following as input
  - Kickstart file to automate OS installation
    - package list (RPMs)
    - custom %post script for configuration
    - scripts for minimizing image size using file blacklisting and forced package removal
  - Yum Repositories for OS packages
  - Output is an ISO image

Key Packages

- qemu-kvm – provides KVM virtualization platform
- qemu-kvm-tools – kvmtrace and kvm_stat for debugging utilities
- vdsm – daemon for managing the node from oVirt Engine
- vdsm-cli – command line interface to VDSM daemon
- libvirt – virtualization API and VM control daemon
- spice-server – Provides guest remote connections

Manifest Files

- The produced image file contains manifest files in the top-level isolinux folder
- Manifests document the content of the final image after minimization (blacklisting)
  - manifest-deps.txt.bz2 – dependency (stderr of “rpm -e --test” result)
  - manifest-dir.txt.bz2 – directories in the image
  - manifest-file.txt.bz2 – files in the image
  - manifest-license.txt – licenses for all installed RPMs
  - manifest-owns.txt.bz2 – file ownership by RPM
  - manifest-rpm.txt – installed RPMs
  - manifest-srpm.txt – source RPMs for installed RPMs

Deployment Modes

- CD-ROM
  - Burn ISO image to writable CD with standard tools
  - Boot server from CD and install to local hard disk
- Flash Memory (USB stick or SD card)
  - Use livecd-iso-to-disk to copy image to USB or SD card
  - Boot from USB/SD and install to local hard disk
- Network (PXE) Boot
  - Use livecd-iso-to-pxeboot on the iso
  - Deploy generated vmliuz0/initrd0 files to PXE/tftp server
  - Boot server from PXE

Installation

- The oVirt Node image currently needs to be installed to a hard disk or flash drive to run
- After installation, boot method should be changed to hard disk
- There are two modes of installation
  - Booting to installation TUI
  - Autoinstallation via kernel command line arguments
- Booting from CD or Flash memory will bring you to a Boot Menu

Boot Menu

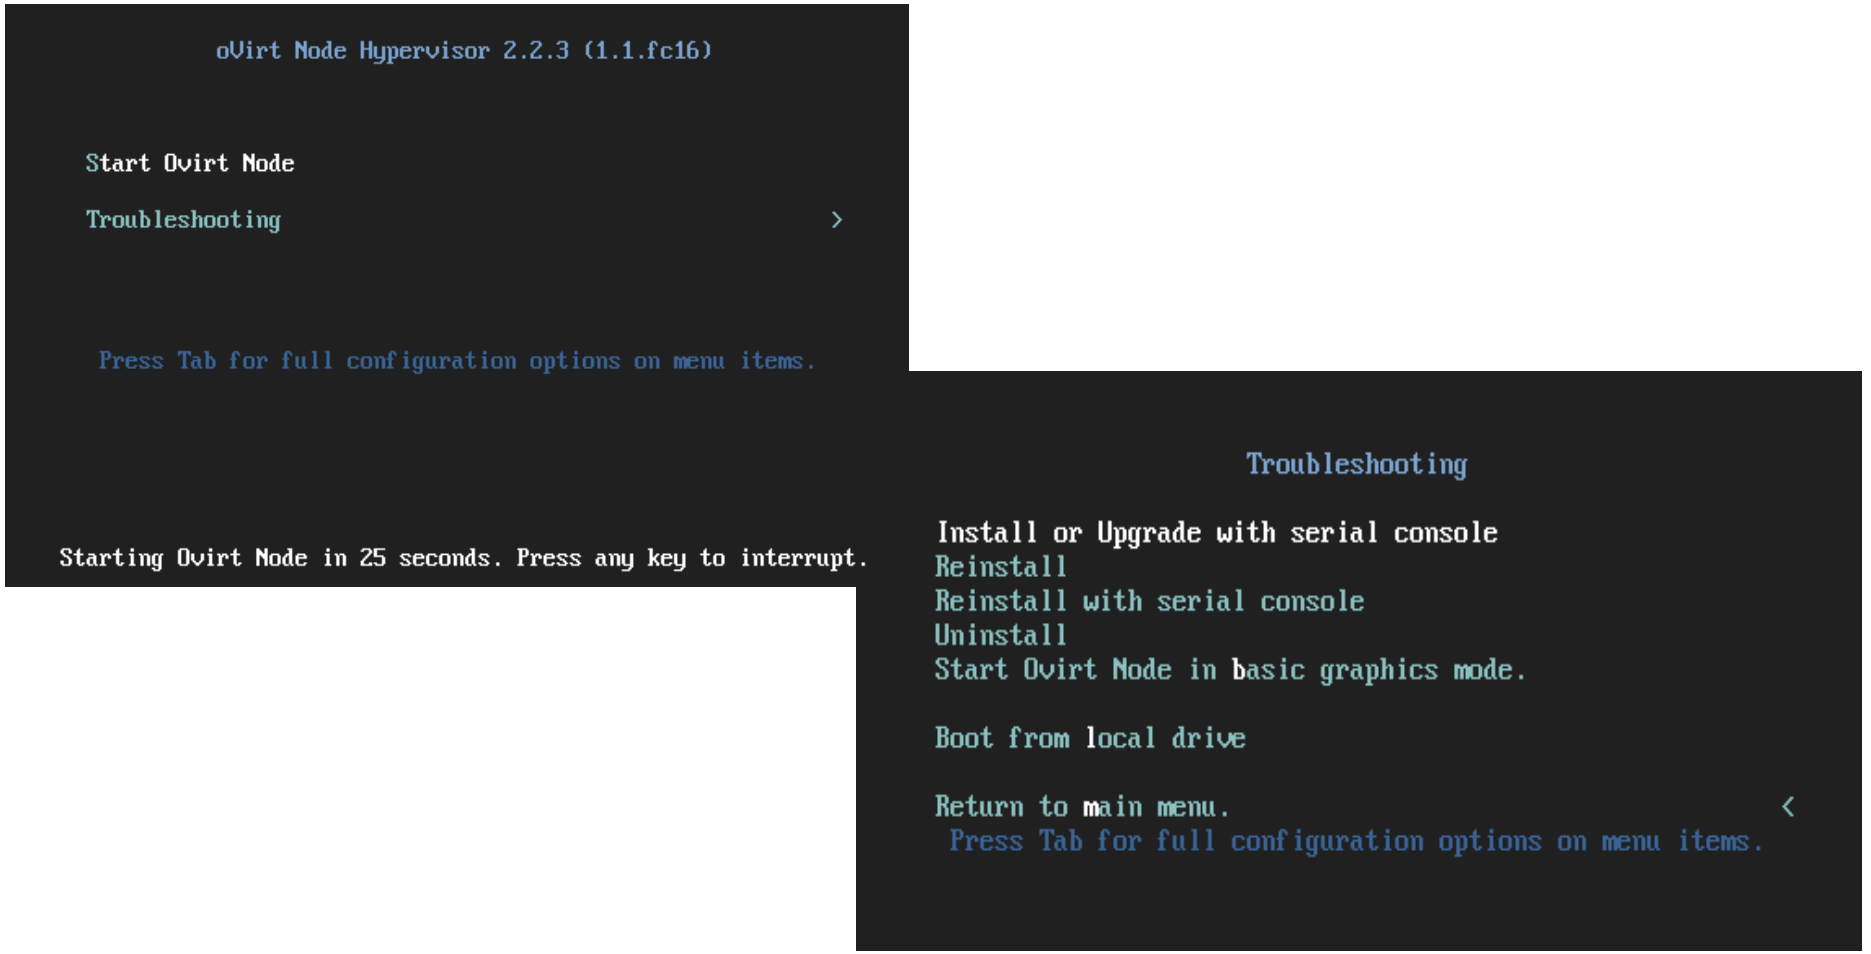

Note: linux rescue can be passed to the boot: parameter to boot in rescue mode

TUI Installation

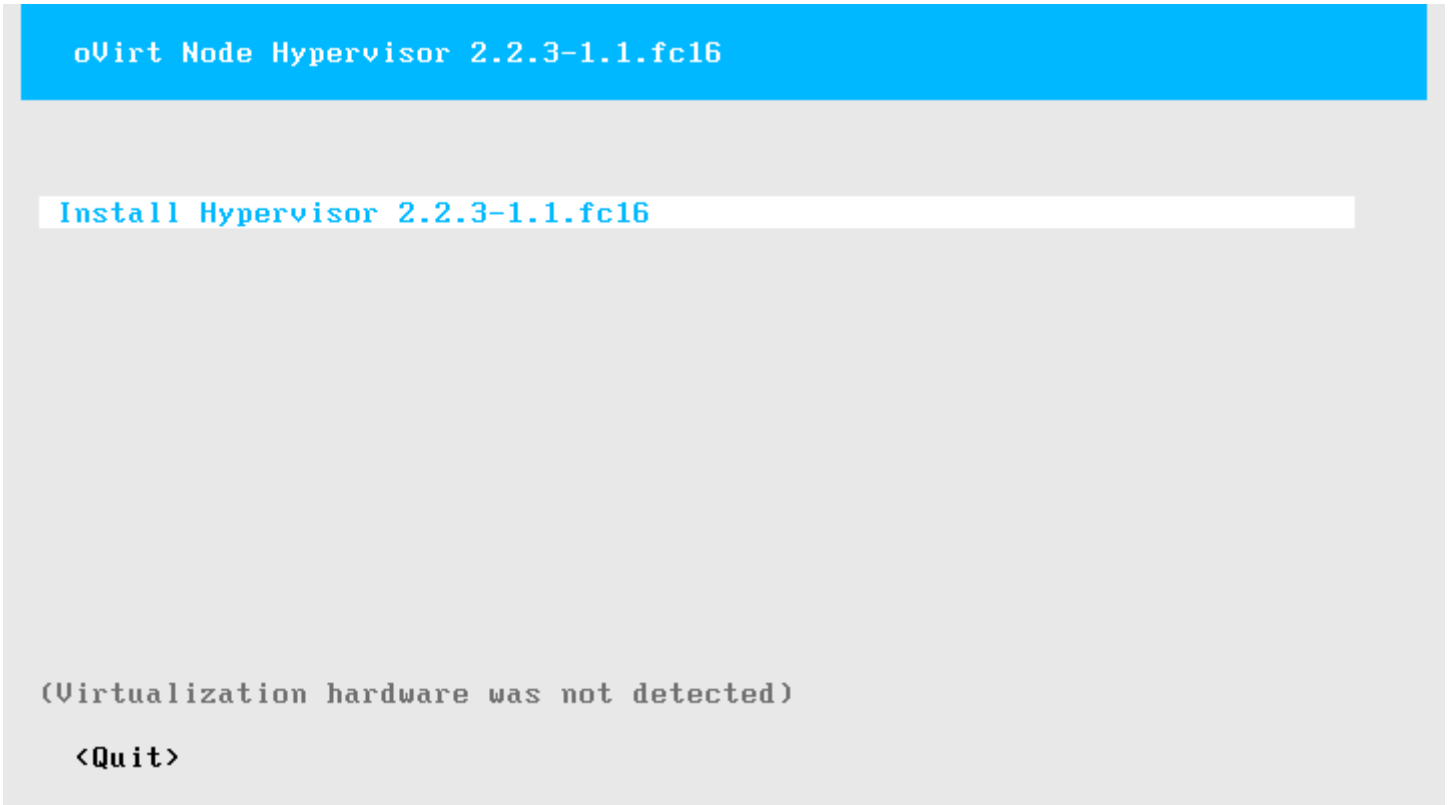

This screen will warn if Hardware Virtualization is not enabled on the host (see bottom on screen shot above)

TUI Installation – Disk Selection

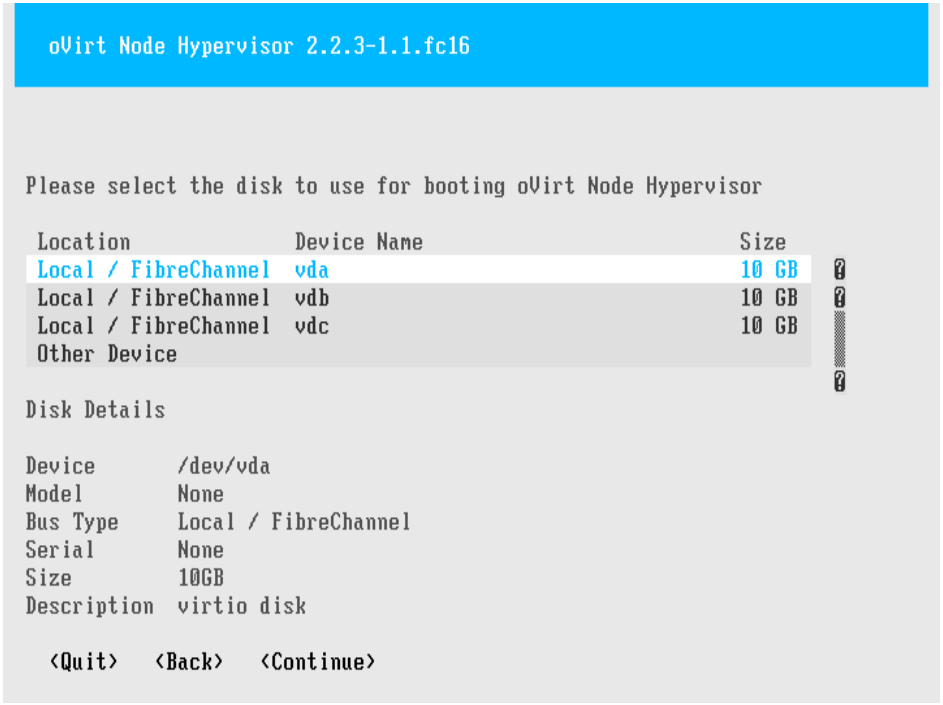

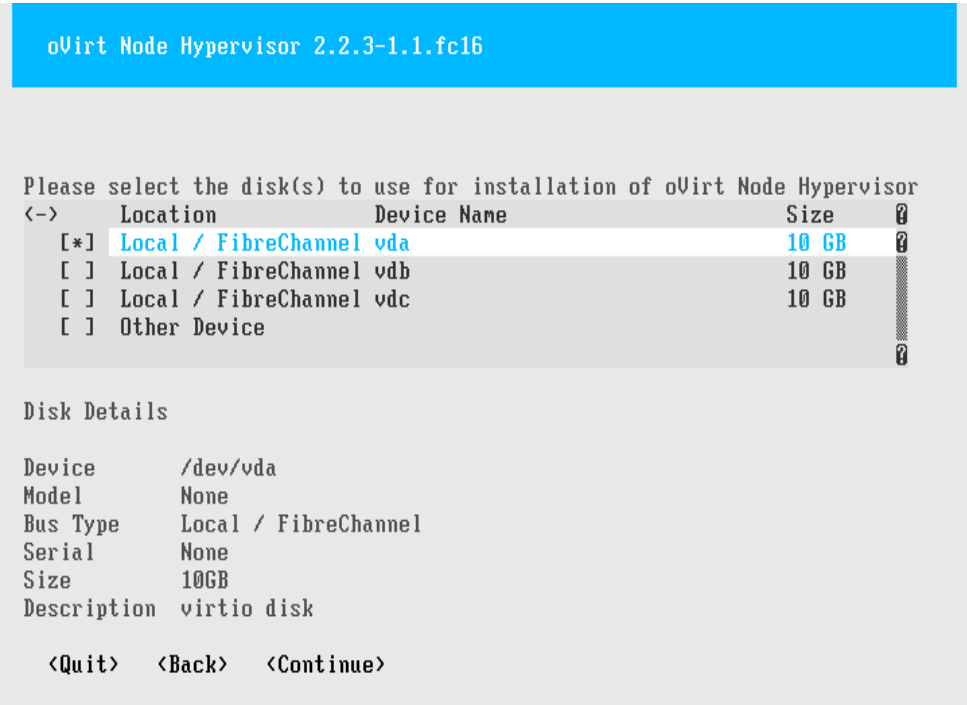

- Select a device to boot from (left screen)
- Select 1 or more disks for storing configuration data and swap (right screen)

TUI Installation – Admin Password

- Set a password for the admin user
- Proceeding from this screen starts the installation

TUI Installation -- Finishing

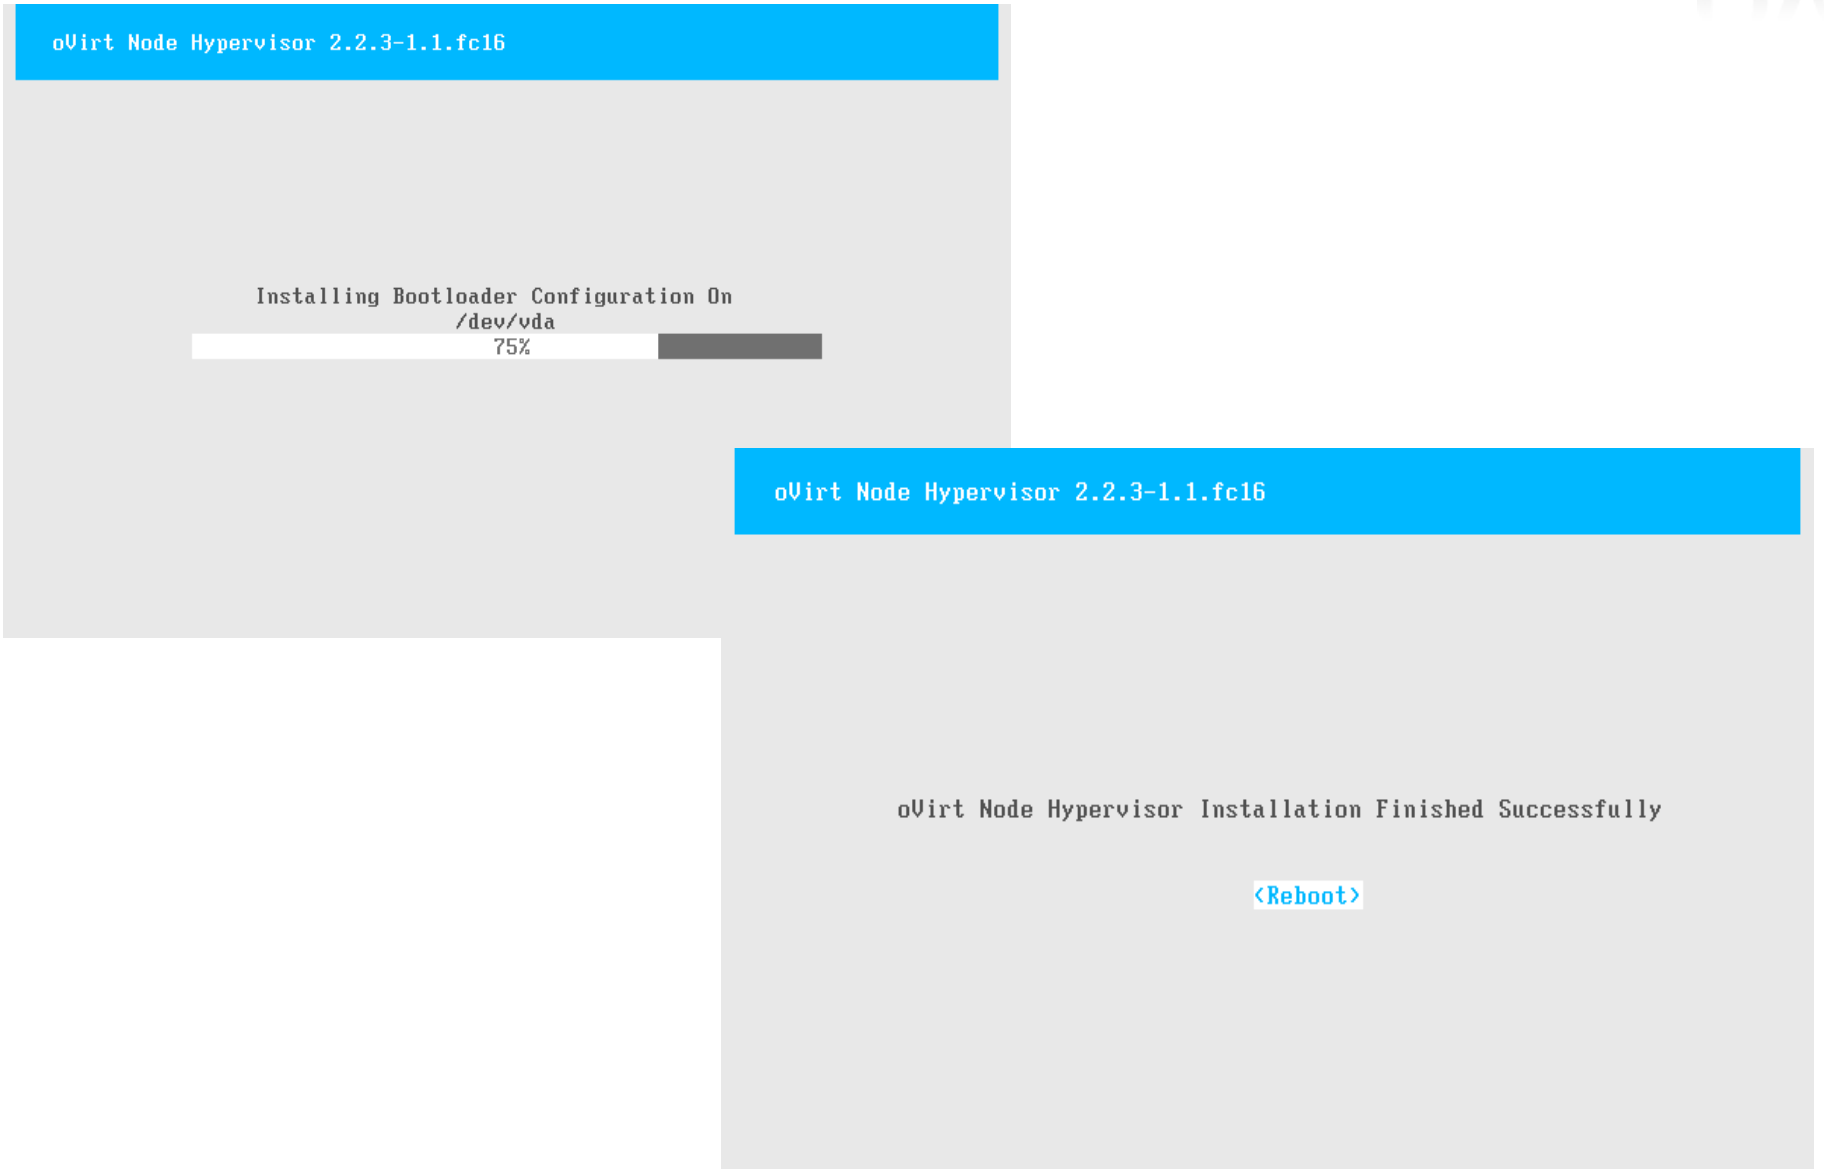

Configuration - Status

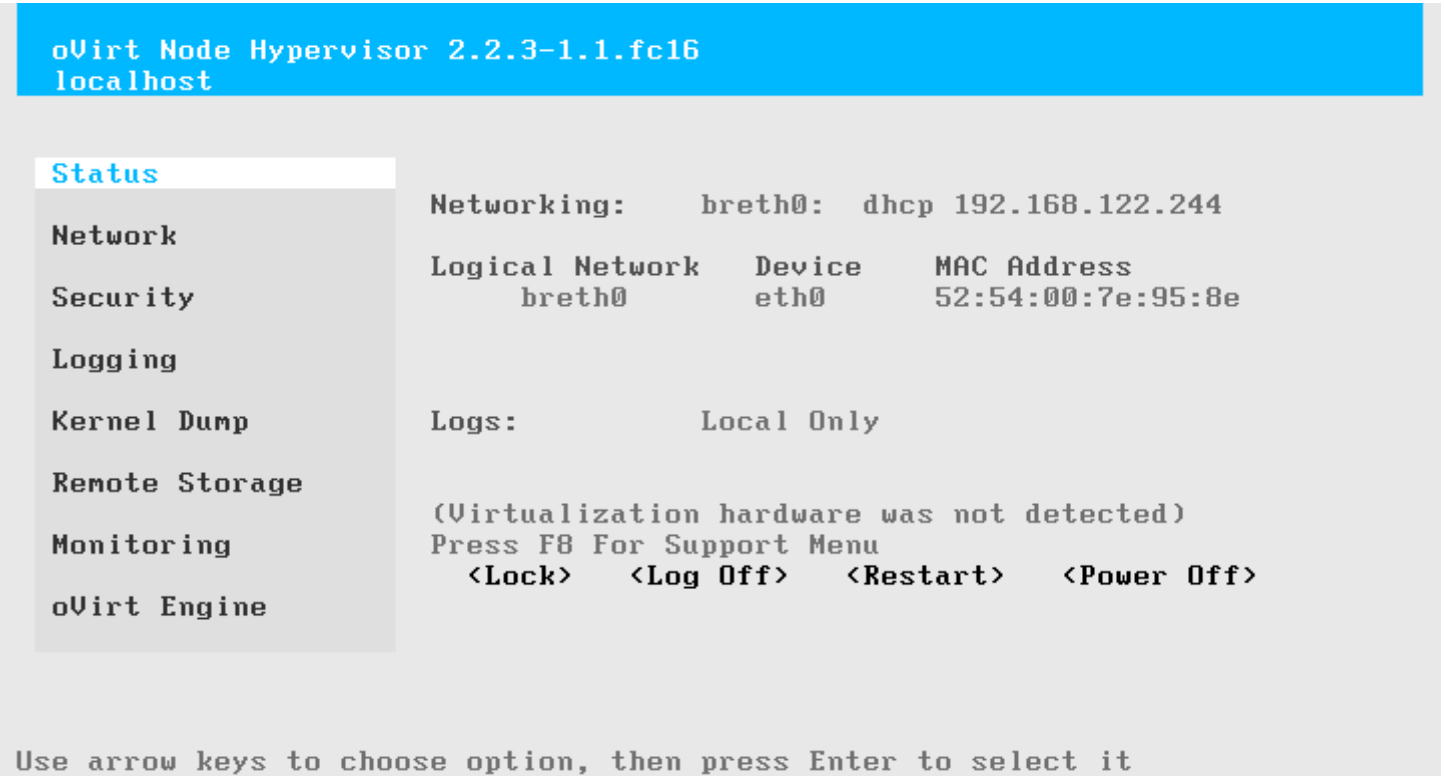

- Get overall status of the system
- See number of VMs runing (not available in the screen shot, because we're installing the hypervisor in a kvm guest)
- Support menu

Configuration – Networking

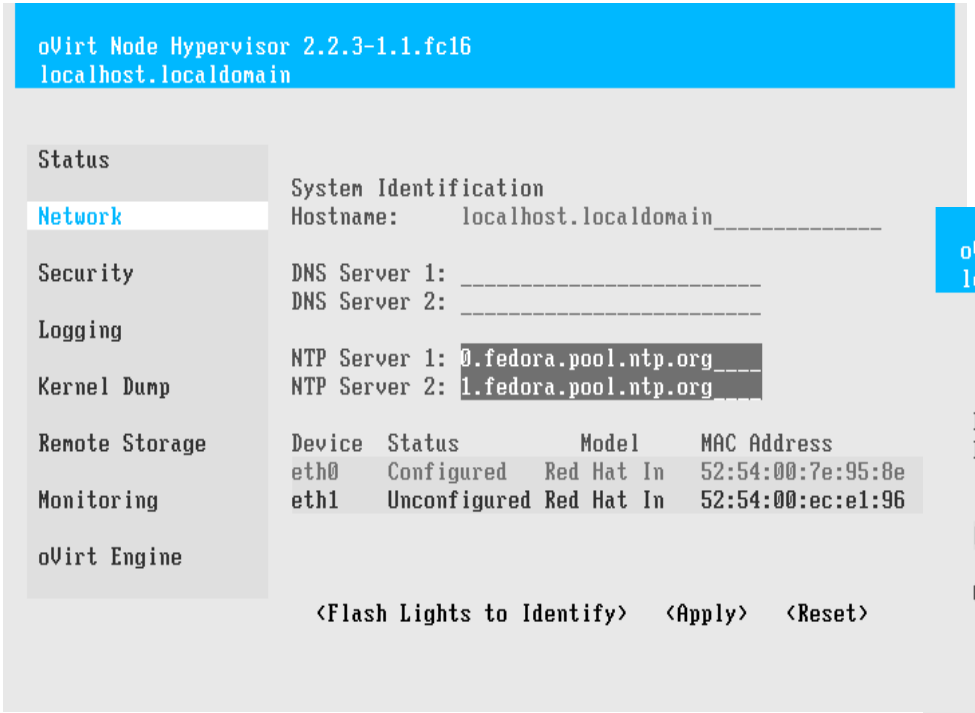

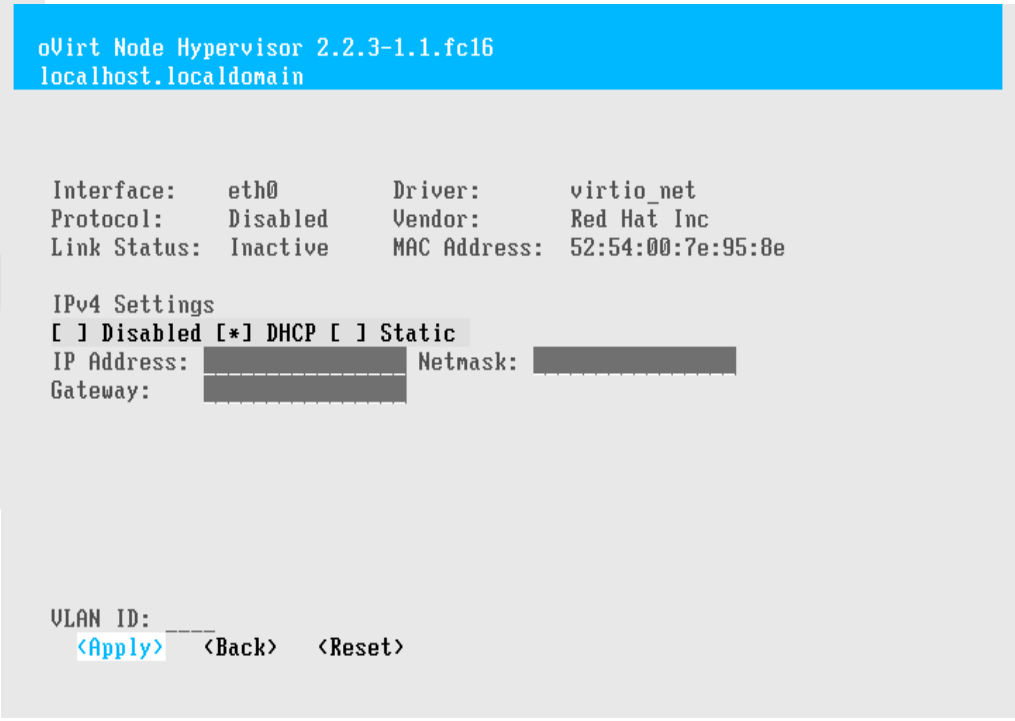

- Choose device to setup for management bridge
- Supports vlans

Configuration - Security

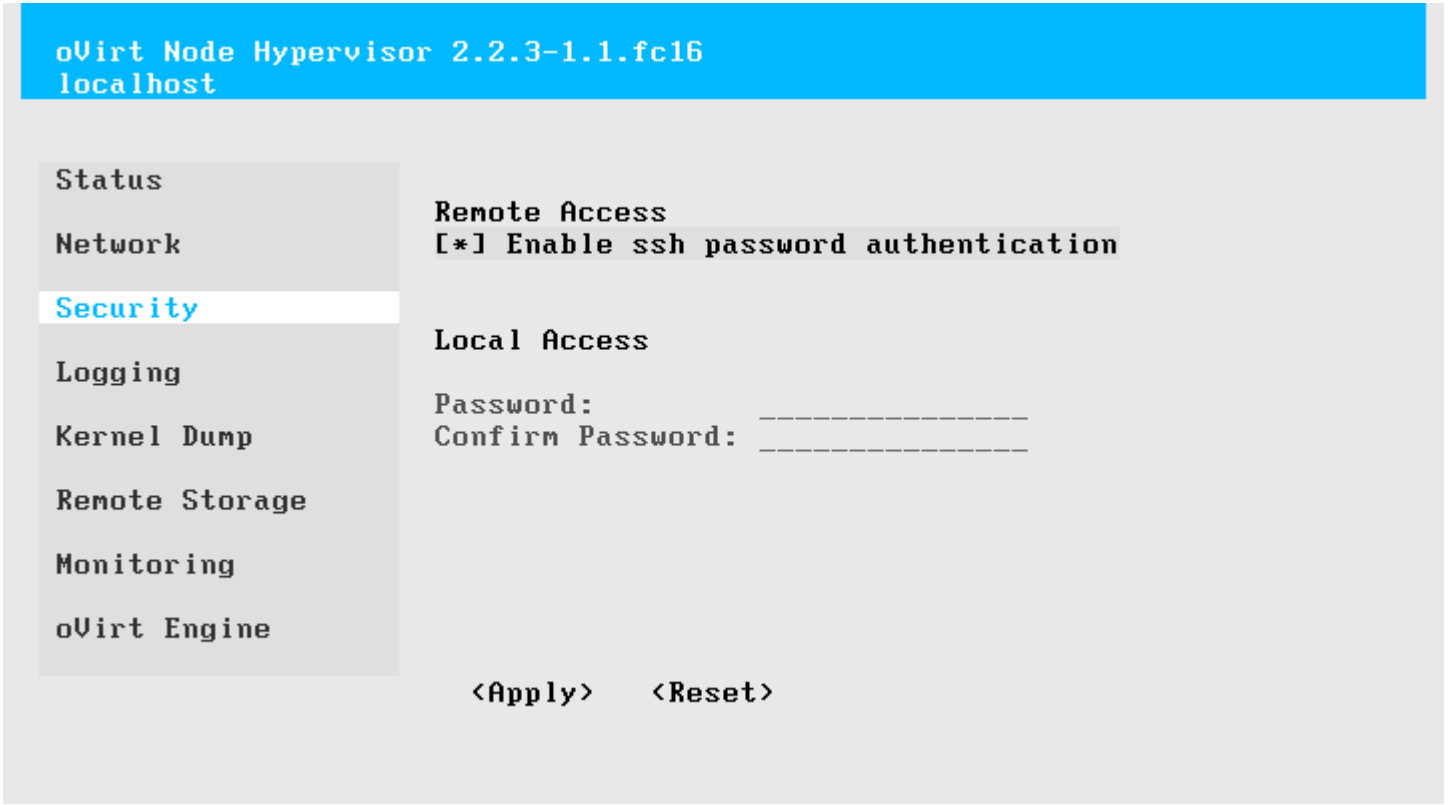

- Enable password based ssh authentication
- Reset admin password

Configuration - Logging

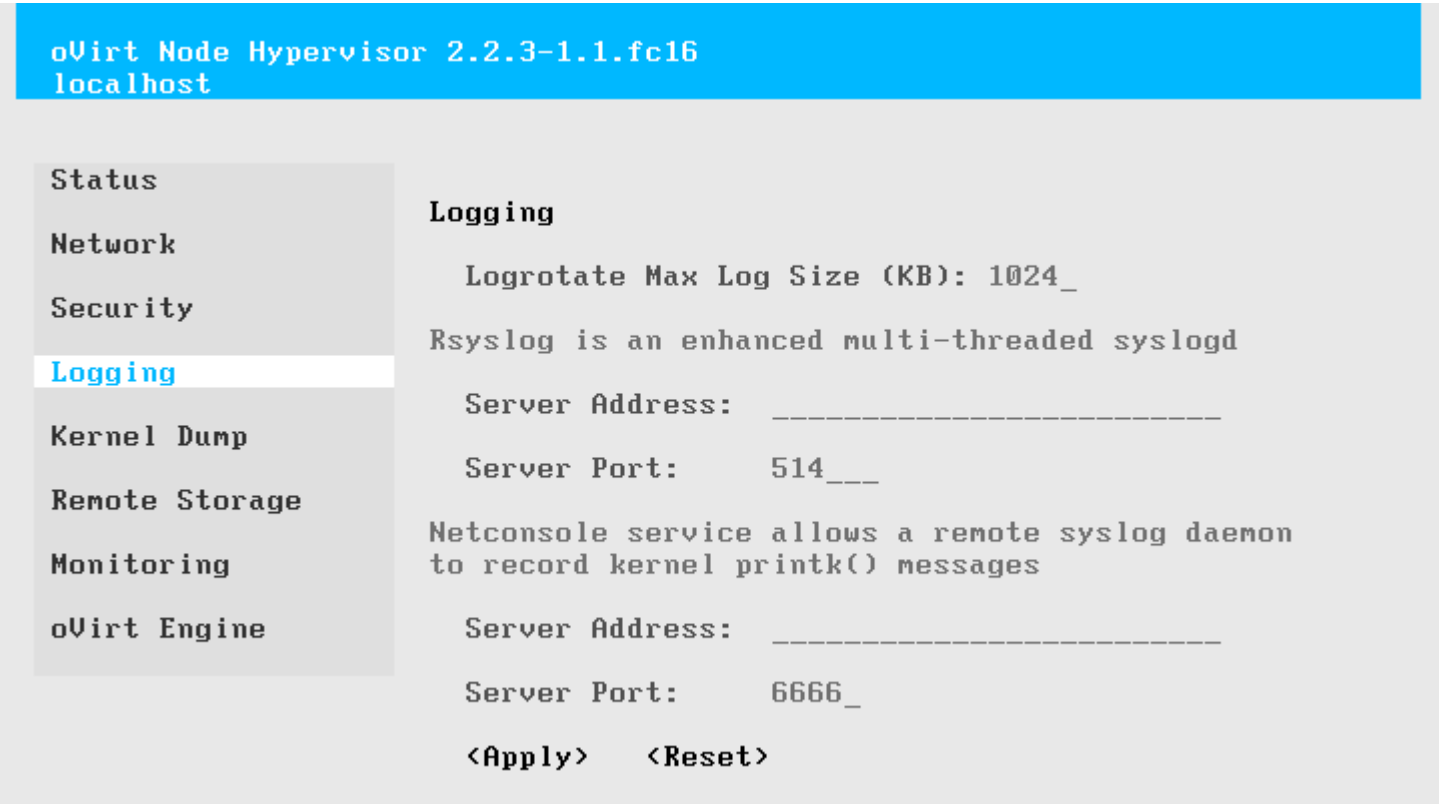

- Logrotate
- Remote logging server
- Netconsole

Configuration - kdump

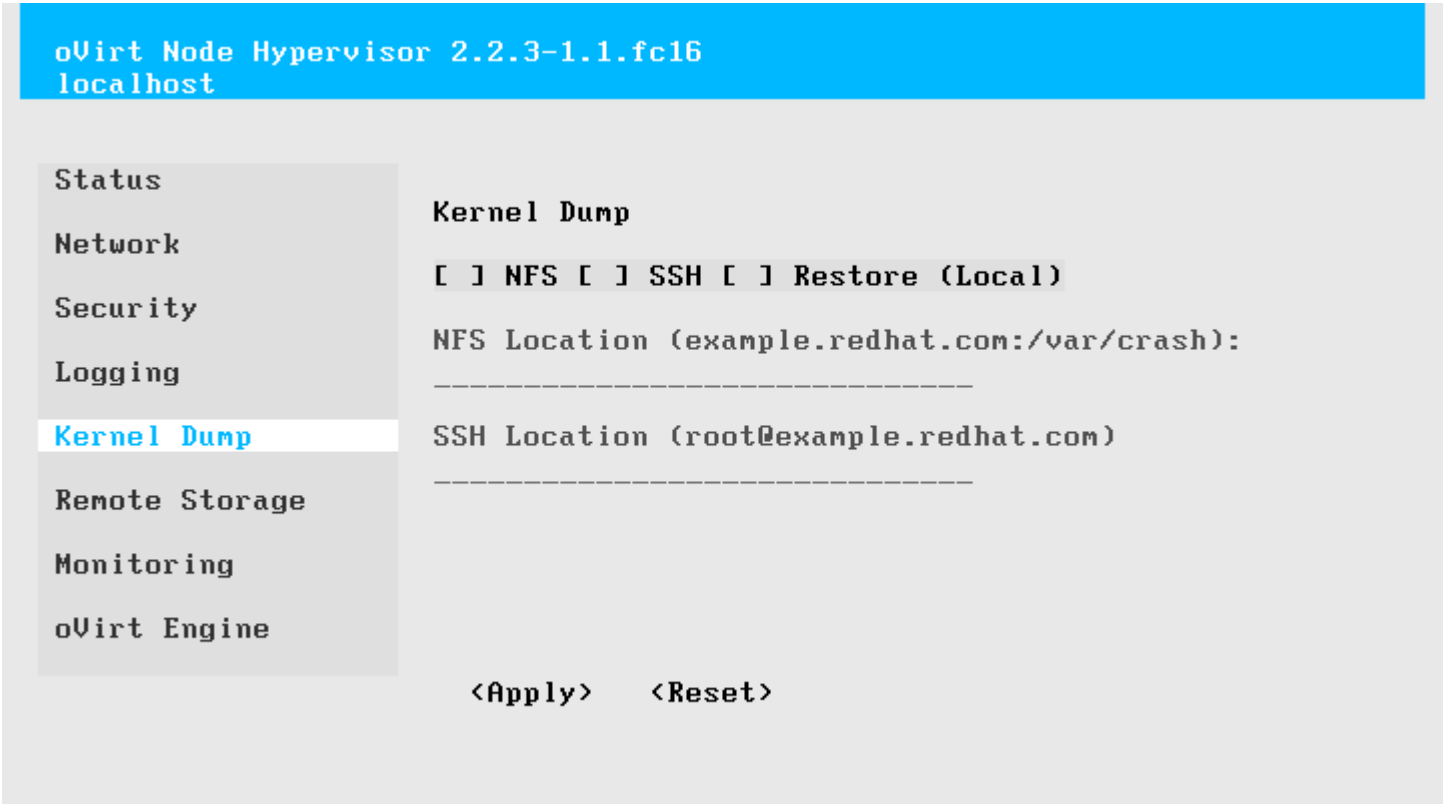

- Configure kdump server

Configuration – Remote Storage

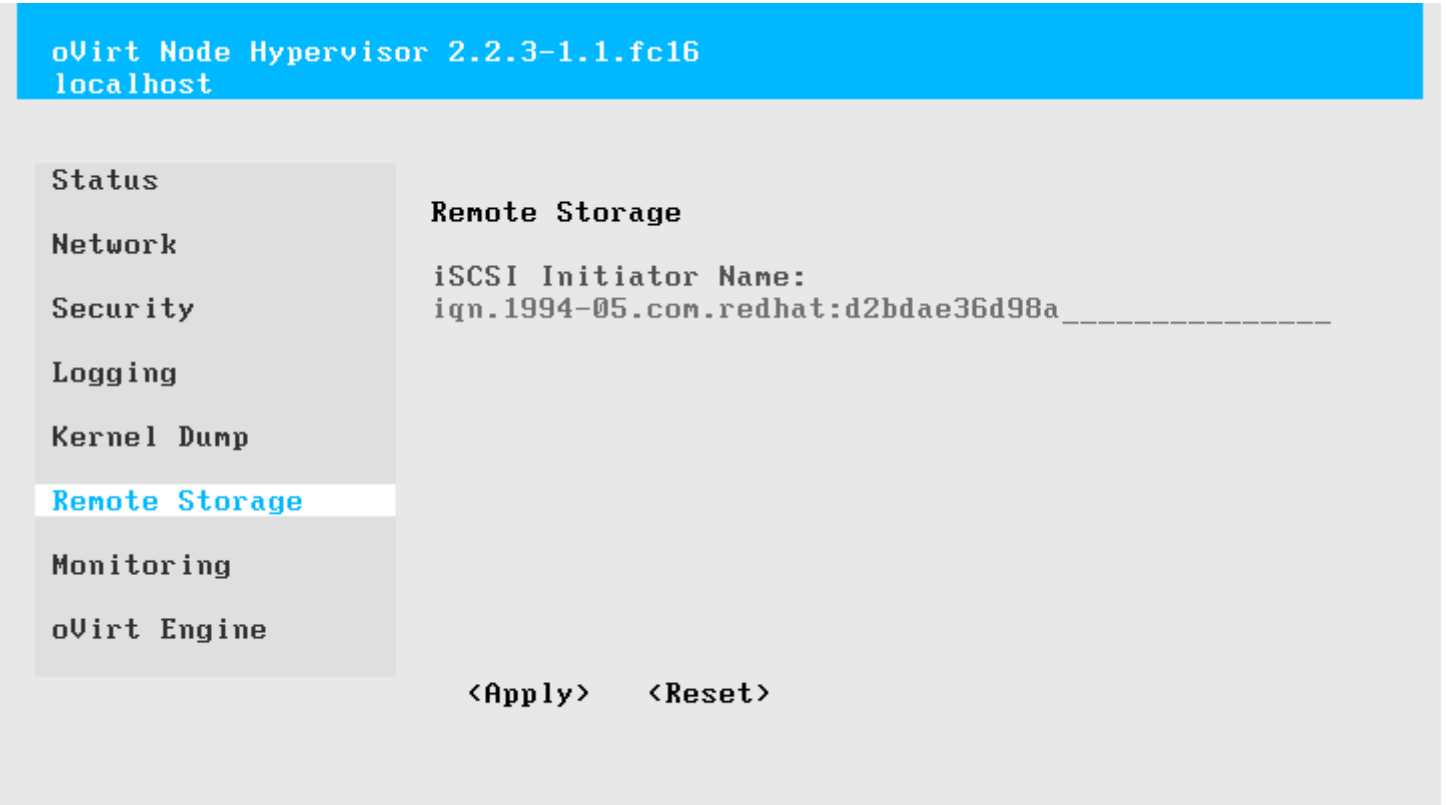

- Setup an iSCSI Initiator Name
- One is randomly generated during installation

Configuration – Monitoring

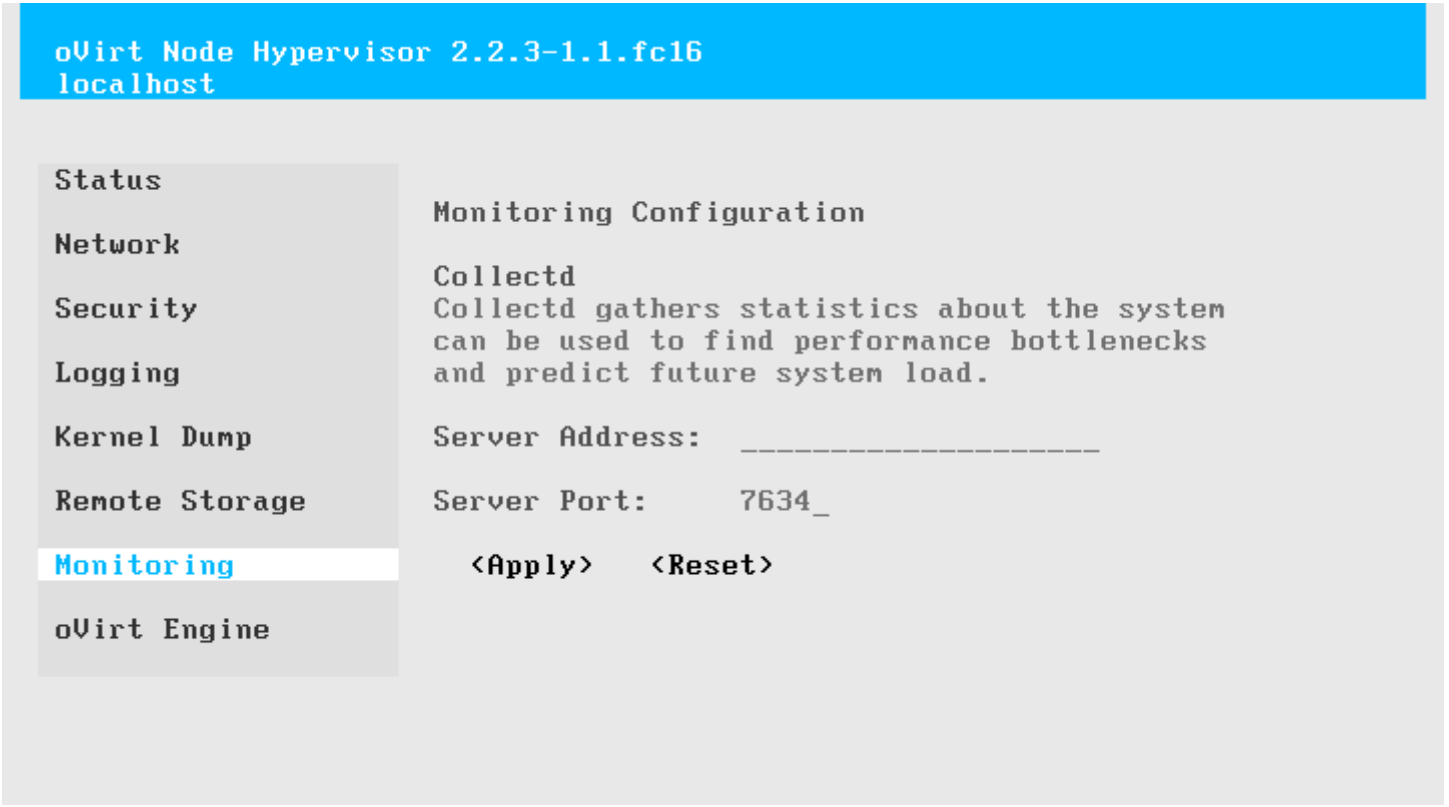

- Connect with a remote collectd server

Configuration – oVirt Engine

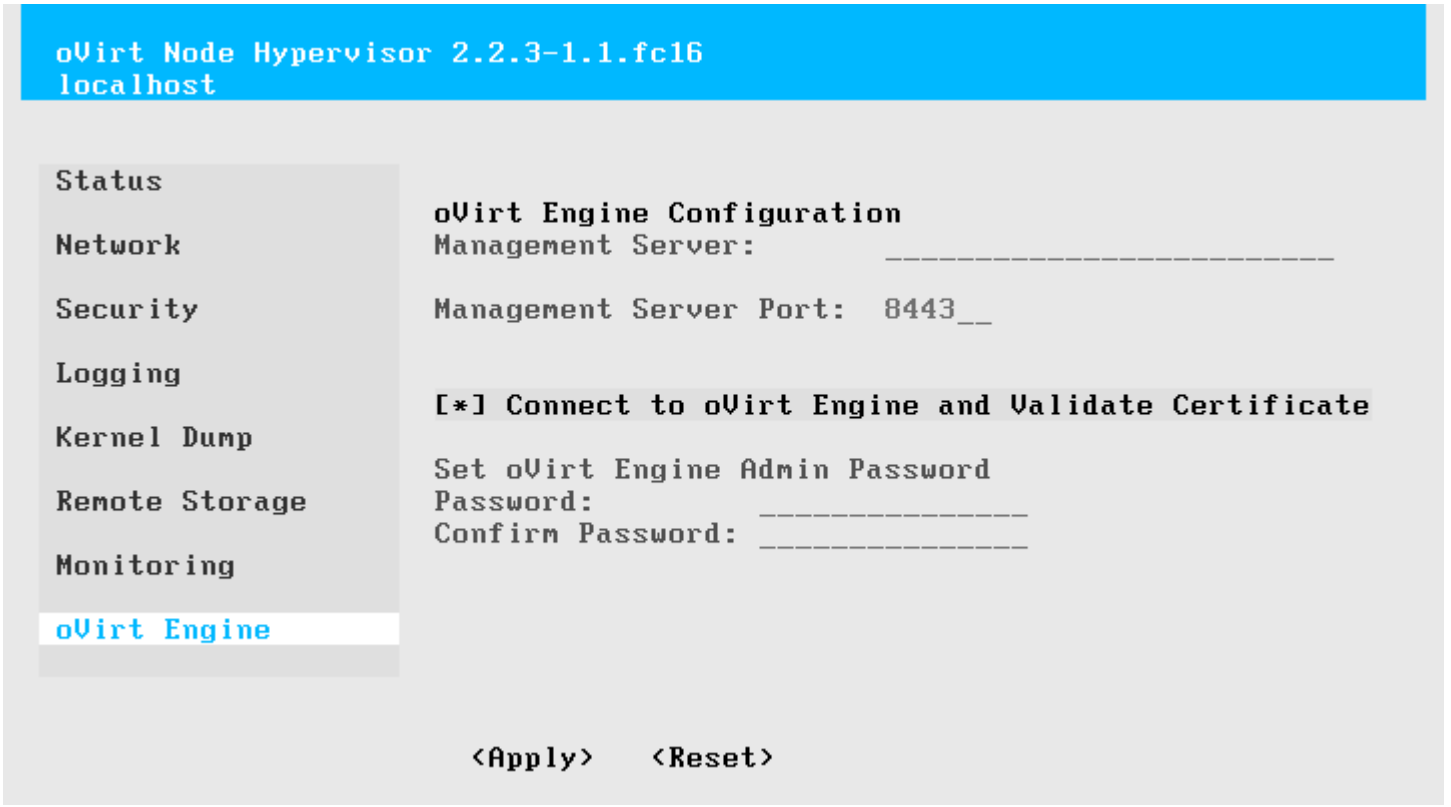

- Register to the oVirt Engine management server

Automatic Installation and Configuration

- All configuration in the TUI can be automated with kernel command line parameters
- Ideal for PXE boot environments
- Requires storage_init and BOOTIF parameters
- adminpw parameter recommended for management after installation

Upgrading

- oVirt Node image is a dedicated appliance
  - no yum/rpm upgrading in the live image
  - Rootfs is non-persistent so upgrades are lost
    - Warning: Runtime rootfs (/) is in-memory overlay. Writing excessive amounts of data to it can cause out of space issues
    - It is mounted read-only by default to avoid this issue
- Three upgrade paths supported
  - Update the PXE server and set host to network boot
  - Boot from new media (CD, USB, SD)
  - In-place upgrade
- ISO/USB/PXE upgrades must specify **upgrade** on the kernel command line to trigger upgrade logic

Configuration Persistence

- Root FS is mounted read-only
  - even if remounted RW, changes are not persisted
- Current Persitence uses rc.sysinit stateless support
- Important files are persisted automatically by oVirt and VDSM as needed
- To manually persist a file, use the persist command:
  - # persist /etc/hosts
- The /config partition is only a few MB by default, so use sparingly.

Roadmap – Stateless

- Feature
  - http://ovirt.org/wiki/Node_Stateless
  - Be able to boot from media/PXE and run completely stateless
  - Kernel Commandline Parameters used for configuration
- Status
  - Booting stateless and all oVirt Node functionality should work currently
  - oVirt Engine does not support Stateless nodes currently
    - Nodes can register fine, but a reboot (reinstall) will require re-registration

Roadmap – Plugins

- Feature
  - http://ovirt.org/wiki/Node_plugins
  - Ability to add custom packages and functionality
  - Added offline to the ISO image
- Status
  - Work in progress
  - First patches are posted, but work is still ongoing

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