

RED HAT :: CHICAGO :: 2009

SUMMIT

FOLLOW US:

[TWITTER.COM/REDHATSUMMIT](https://twitter.com/redhatsummit)

TWEET ABOUT US:

ADD #SUMMIT AND/OR #JBOSSWORLD TO THE END
OF YOUR EVENT-RELATED TWEET

presented by



RED HAT :: CHICAGO :: 2009

SUMMIT

Red Hat Enterprise Linux 5 Virtualization in Real Life

Vinny Valdez, RHCA
Sr. Product Architect
Red Hat Consulting
September 3, 2009

presented by



Agenda

- Virtualization Benefits Brief
- Red Hat Enterprise Linux 5.{0-3} Virtualization
- Red Hat Enterprise Linux 5.4+ Virtualization
- The Linux Kernel – Scalability, Security, Performance
- Real Life Use Cases
- Best Practices
- Managing your Virtualization Infrastructure
- Red Hat Enterprise Virtualization

RED HAT :: CHICAGO :: 2009

SUMMIT

Virtualization Benefits Brief

presented by



Benefits of Virtualization

- Consolidation
- Increased utilization
- Power/Cooling Savings (Green IT)
- Leverage new technologies (FCoE, 10G)
- Rapid provisioning
- Manage risks and minimize costs
- Extend legacy software lifecycle
- Dynamic software fault tolerance
- Hardware fault tolerance through live migration

RED HAT :: CHICAGO :: 2009

SUMMIT

Red Hat Enterprise Linux 5.{0-3} Virtualization

presented by



Xen Virtualization

Red Hat Enterprise Linux 5 includes Integrated virtualization

- Xen Hypervisor
 - Included in 5.0
 - Fully Supported during product lifetime
 - Until at least 2014 (later with Mission Critical program)
 - Using Xen 3.1.2 (with selected backports)
 - Available for x86, x86_64 and IA64
- Red Hat Enterprise Linux Advanced Platform
 - Includes GFS Cluster file system and Red Hat Cluster Suite
 - Unlimited virtualized guests

RED HAT :: CHICAGO :: 2009

SUMMIT

Red Hat Enterprise Linux 5.4+ Virtualization

presented by



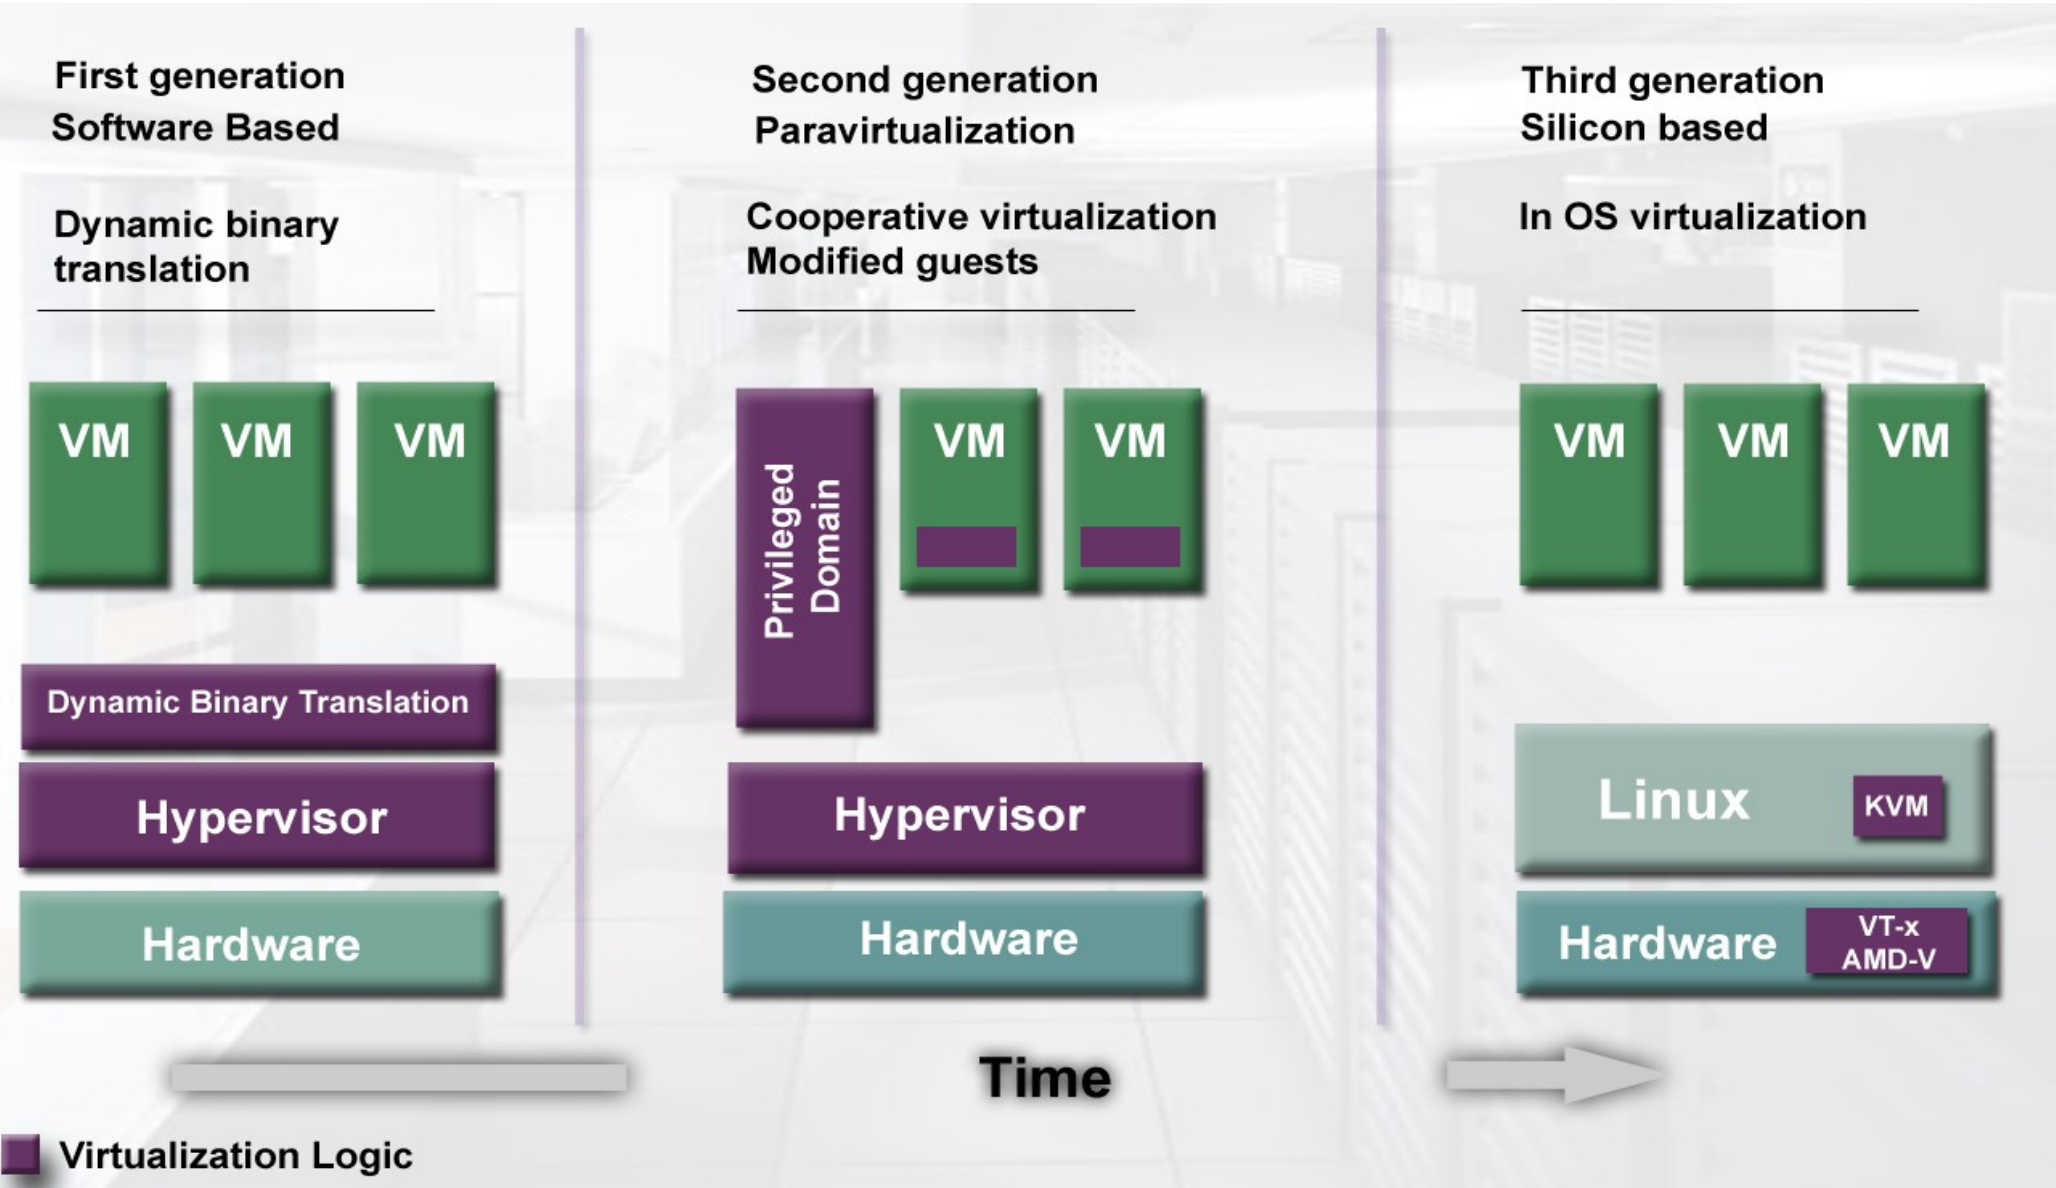
Kernel Virtual Machine

- KVM Hypervisor
 - Added in RHEL 5.4
 - Included in bare metal kernel
 - Hypervisor available for **x86_64 only** (guests can be mixed)
 - Requires Intel VT-x (VMX) or AMD AMD-V (SVM)
 - Support for RHEL {3-5} guests, other Linux distros
 - Support for Microsoft Windows Servers guests (Windows Server 2003, 2008), and Windows XP, Windows Vista
 - Paravirtualized drivers for network and disk (WHQL certified)
 - Microsoft SVVP Certification
- Xen will continue to be shipped and supported in RHEL 5.4+ (2014)

RHEL 5.4 Virtualization

- Support for advanced hardware features for both Xen and KVM
 - VT-d/IOMMU for secure PCI Pass-through
 - SR-IOV for PCI device sharing
 - Hardware assisted virtualization
- Performance and scalability out-of-the-box
- Management compatibility via libvirt
- Virt-tools allow transparent transition
- Tools to transition from Xen to KVM (para or HVM)

Evolution of x86 Virtualization



RED HAT :: CHICAGO :: 2009

SUMMIT

The Linux Kernel – Scalability, Security, Performance

presented by



Virtualization with the Linux Kernel

- All certifications immediately apply
- Performance tuning done as normal kernel process
- Offload work to hardware
- NUMA, huge pages, Kernel Shared Memory, CPU scheduling, power management, SELinux, etc.
- 256 cores, 1TB RAM
- Xen makes sense for some workloads (hardware without virtualization extensions)
- Xen always needs specialized knowledge to code to
- KVM is just another process

RED HAT :: CHICAGO :: 2009

SUMMIT

Real Life Use Cases

Typical Use Cases

- Training
- Support
- Development/testing/demos
- Leverage new technologies (FCoE, 10G)
- Kiosks
- VDI
- Compatibility
- High Availability for non-cluster-aware Applications
- Cloud Computing

Customer Deployments

- Amazon EC2
- Electronic Arts
- Dreamworks
- Chicago Tribune (RHCS/GFS)
- US Courts (dom0 mount offline LVM for backup)
- Booze-Allen-Hamilton (Innovation Winner 2008)
- What are YOU doing?

RED HAT :: CHICAGO :: 2009

SUMMIT

Best Practices

Best Practices

- For performance, use pre-allocated disks
- LVM snapshots to quickly clone
- Shared storage (NFS/iSCSI/Fibre) with GFS for ease of management and live migration
- Use symlinks to common config files
- Red Hat Cluster Suite for resilience and fencing
- Use Satellite 5.3 with cobbler/koan for provisioning
- Script live migrations based on utilization
- Use para-virt drivers for best performance on HVM guests

RED HAT :: CHICAGO :: 2009

SUMMIT

Managing Your Virtualization Infrastructure

libvirt: Stable and Open API

- Management of Hypervisor(s) using standard APIs
- Hypervisor / Platform agnostic
- Stable API for ISV and application developers
 - Isolation from Xen HV instability
 - Isolation from XenD protocol changes
 - Formalized error reporting/handling
- XML definitions for VMs
- Contribution of partners (Fujitsu, SUN, IBM, Scalent, Platform, Enomaly, etc..)
- Core API in C, Python/Perl bindings for scripting
- CLI access via virsh
- Remote management and storage APIs in development
- Addition of CIM capabilities (IBM contribution)



RHEL5 Virt Tools

- Addressing short/mid-term holes in tool suite
 - Focus on “low hanging fruit” tools/capabilities to address common issue
 - Can be CLI or GUI
 - Longer term focus on GUI interface and Virtual Center-like capabilities
 - Tools developed so far
 - “RHEL5 for Dummies” Guide
 - Red Hat Magazine articles
 - Virtualization tools for RHEL5
 - virt-top , top like performance monitoring utility
 - virt-p2v , convert physical instance to virtual machine
 - nagios-virt , integration of virtual machine monitoring into NAGIOS
 - virt-image in progress (to convert VMware images to RHEL5)

RHEL5/Virt Tools (<http://et.redhat.com/~jmh/virt-tools/>)

- All tools available in Fedora and via EPEL for RHEL
- virt-top
 - Performance monitoring using a 'top' like tool
- virt-p2v
 - Convert physical to virtual instances
- libguestfs / guestfish
 - Library and tools to modify guest images offline
- virt-inspector
 - Generate XML configuration file from existing guest image
- virt-what
 - Determine guest type and hosted Hypervisor
- virt-df
 - 'df' for virtual machines, examine file system usage from Hypervisor
- virt-mem / virt-ps / virt-dmesg / virt-uname etc
 - Tools to inspect running guest from Hypervisor
- Cobbler/Koan (integrated into Satellite 5.3)

virsh

- Exposes libvirt functionality
- Command line administration tool
- Interactive administrator shell (**# virsh**)
- Scriptable batch operations
- Stable and dependable interface to HV administration
- Distributed with libvirt in Fedora and RHEL5
- Should be used for all interactive/remote/shell/script management tasks
 - Can be used for all management tasks supported by xm(8)
 - Support for remote management via ssh or TLS

virt-manager

- Desktop application for VM management
- Summary view for
 - Running domains
 - Performance & Resource statistics
- Support for domain management
 - Guest Creation, deletion, suspend/resume/start/stop/save/restore
 - Configuration of domain resources
 - vCPUs, Memory, Storage, Network, Bridging
- Live resource utilization graphs
 - CPU, Memory and Networking
- Available in Fedora and RHEL5
- Graphical guest OS console
- Secure console access
- Support for remote management via SSH or TLS



virt-manager

- Initial Overview window
 - Non-root connections open a read-only connection
 - Connection to localhost and/or multiple remote hosts in parallel
- Host Summary Window
 - List of all running Virtual Machines and resource usage
 - Dynamic hiding/showing of fields
 - Detailed domain information via 'Details' tab
 - Domain console via 'Open' or click onto domain entry
 - Create new virtual machine using the 'New' tab
- Virtual Machine Graphical Console
 - Graphical/serial interface for guest console via VNC for PV and FV
 - Supports authentication

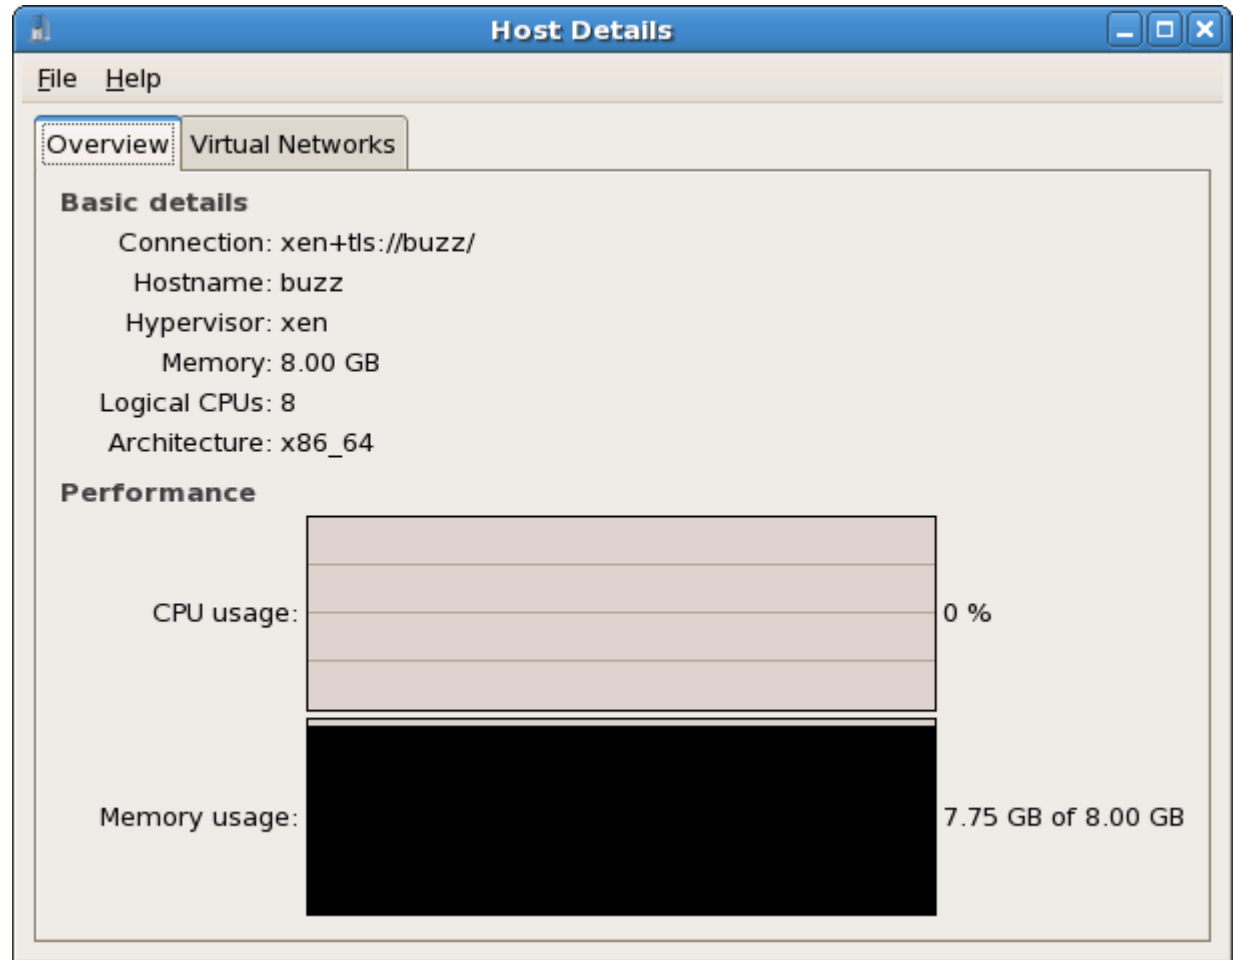


virt-manager

- Virtual Machine Performance Details
 - Live graphs of various machine and HyperVisor resources
 - Network, CPU, Memory, Disk usage
 - Current view is from the HyperVisor, will include guest OS view
 - Allows CPU time to be broken down to user,system,idle,iowait and nice categories
 - UUID is a globally unique identifier
 - Persistent across migrations
- Virtual Machine Configuration
 - Allows changes to various resources
 - Memory, Virtual CPU count
 - Dynamic tunables for immediate changes
 - Static tunables will be saved and applied during next reboot of virtual machine
 - Network and Storage management (adding/removing components)



virt-manager / Connection & Host Overview



virt-manager / VM Overview

The image displays two screenshots of the Virtual Machine Manager (virt-manager) interface. The left screenshot shows a list of virtual machines, and the right screenshot shows a detailed view of the 'buzz' host.

Left Screenshot: Virtual Machine Manager (Xen: woodie)

ID	Name	Status	CPU usage	VCPUs	Memory usage
-	CaneRHEL	Shutoff	0.00 %	16	2.00 GB
0	Domain-0	Running	6.09 %	4	2.00 GB
-	FreeBSD	Shutoff	0.00 %	1	500.00 MB
-	HastaVista	Shutoff	0.00 %	1	1024.00 MB
-	HastaVista64	Shutoff	0.00 %	1	1024.00 MB
-	NetBoot	Shutoff	0.00 %	1	1024.00 MB
-	OEL5x86_64	Shutoff	0.00 %	1	2.00 GB
-	OpenSOLfv	Shutoff	0.00 %	1	1024.00 MB
-	OracleVM-Manager	Shutoff	0.00 %	1	1024.00 MB
46	OracleVM-Server	Running	0.09 %	1	2.00 GB
-	TulsaRHEL	Shutoff	0.00 %	4	2.00 GB
-	VMwareVC	Shutoff	0.00 %	1	2.00 GB
2	WinXP	Running	0.02 %	1	512.00 MB
-	WinXPpv	Shutoff	0.00 %	1	512.00 MB

Right Screenshot: Virtual Machine Manager

View: All virtual machines

Name	ID	Status	CPU usage	VCPUs	Memory usage
buzz	xen	Active	0.33 %	8	7.75 GB
Domain-0	0	Running	4.50 %	8	14.75 GB
rhel4u4fvpv	-	Shutoff	0.00 %	4	4.00 GB
rhel4u5x86_64PV	-	Shutoff	0.00 %	4	4.00 GB
rhel52pv	-	Shutoff	0.00 %	2	512.00 MB
vBlade5	-	Shutoff	0.00 %	1	512.00 MB
localhost	xen	Disconnected	0.00 %	0	0.00 MB
woodie	xen	Active	4.54 %	8	15.76 GB
CaneRHEL	-	Shutoff	0.00 %	16	2.00 GB
Domain-0	0	Running	3.46 %	8	14.75 GB
HastaVista	-	Shutoff	0.00 %	1	1024.00 MB
HastaVista64	-	Shutoff	0.00 %	1	1024.00 MB
NetBoot	-	Shutoff	0.00 %	1	1024.00 MB
OEL5x86_64	-	Shutoff	0.00 %	1	2.00 GB
OracleVM-Manager	-	Shutoff	0.00 %	1	1024.00 MB
OracleVM-Server	-	Shutoff	0.00 %	1	2.00 GB
TulsaRHEL	-	Shutoff	0.00 %	4	2.00 GB
VMwareVC	-	Shutoff	0.00 %	1	2.00 GB
WinXP	-	Shutoff	0.00 %	1	512.00 MB
WinXP64bit	-	Shutoff	0.00 %	1	1024.00 MB
WinXPpv	-	Shutoff	0.00 %	1	512.00 MB

virt-manager / Host Details

Host Details

File Help

Overview Virtual Networks

Basic details

Hostname: woodie
Hypervisor: Xen
Memory: 16.00 GB
Logical CPUs: 8
Architecture: x86_64

Performance

CPU usage: 19 %

Memory usage: 7.02 GB of 16.00 GB

Host Details

File Help

Overview Virtual Networks

Basic details

Name: dummy12
UUID: 0483010e-67b7-0ecb-6d42-1e286e2780aa
Device: vnet16
State: Active
Autostart: On boot

IPv4 configuration

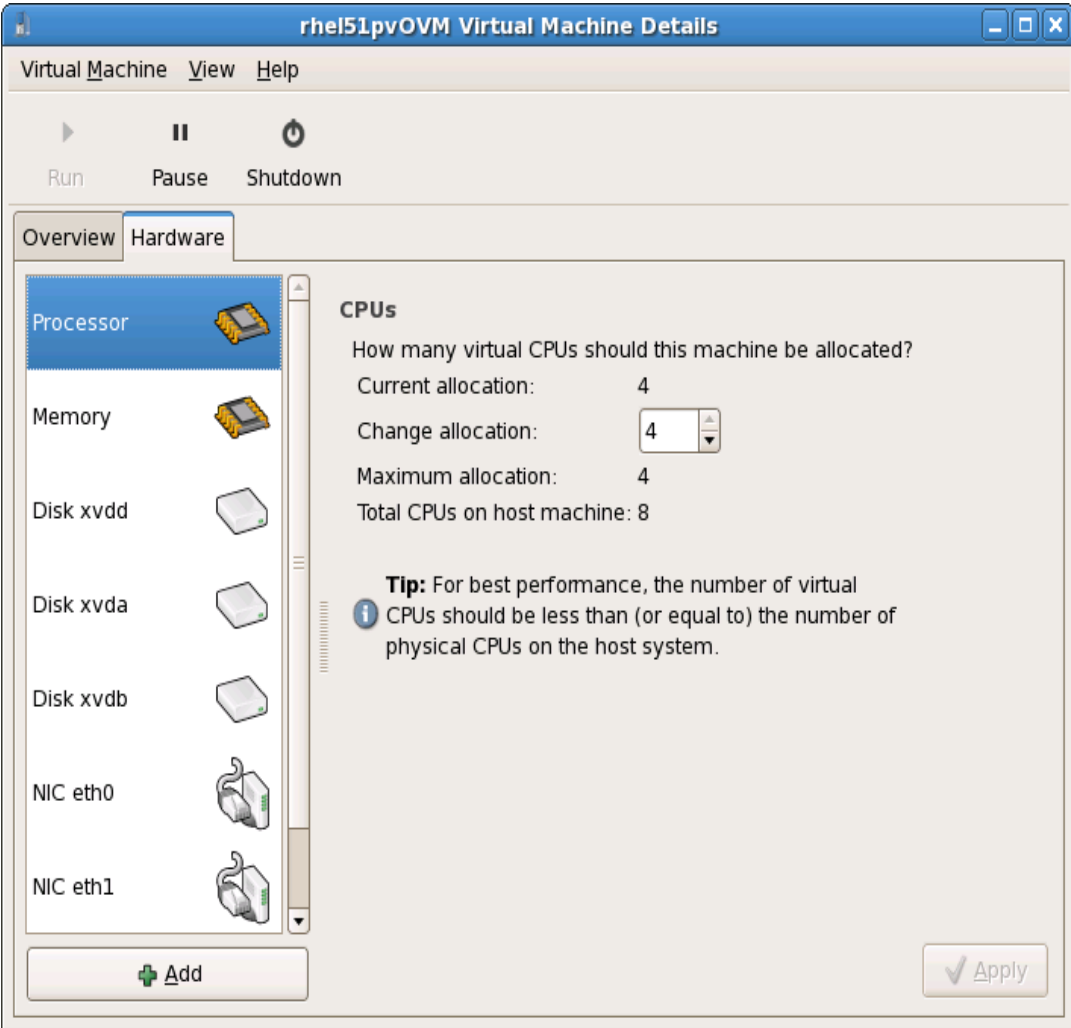
Network: 10.1.12.0/24
DHCP start: 10.1.12.128
DHCP end: 10.1.12.254
Forwarding: NAT to physical device peth0

dummy12
dummy11
dummy5
dummy6
dummy15
dummy2

+ Add Delete Start Stop

virt-manager / Resource Management

PV Guest

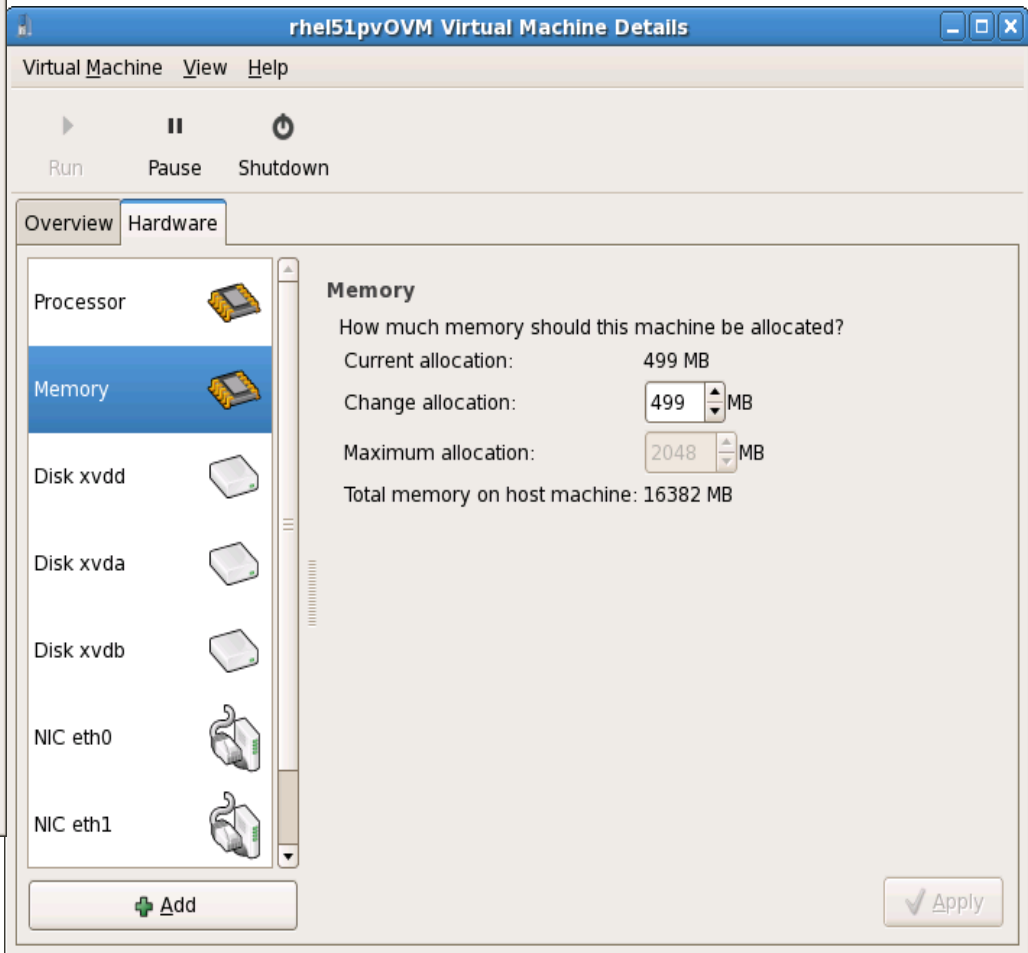


The screenshot shows the 'rhel51pvOVM Virtual Machine Details' window with the 'Hardware' tab selected. The 'CPU' section is active, displaying the following information:

- How many virtual CPUs should this machine be allocated?
- Current allocation: 4
- Change allocation: 4 (with a spin button)
- Maximum allocation: 4
- Total CPUs on host machine: 8

A tip is provided: **Tip:** For best performance, the number of virtual CPUs should be less than (or equal to) the number of physical CPUs on the host system.

The hardware list on the left includes: Processor, Memory, Disk xvdd, Disk xvda, Disk xvdb, NIC eth0, and NIC eth1. There are 'Add' and 'Apply' buttons at the bottom.



The screenshot shows the 'rhel51pvOVM Virtual Machine Details' window with the 'Hardware' tab selected. The 'Memory' section is active, displaying the following information:

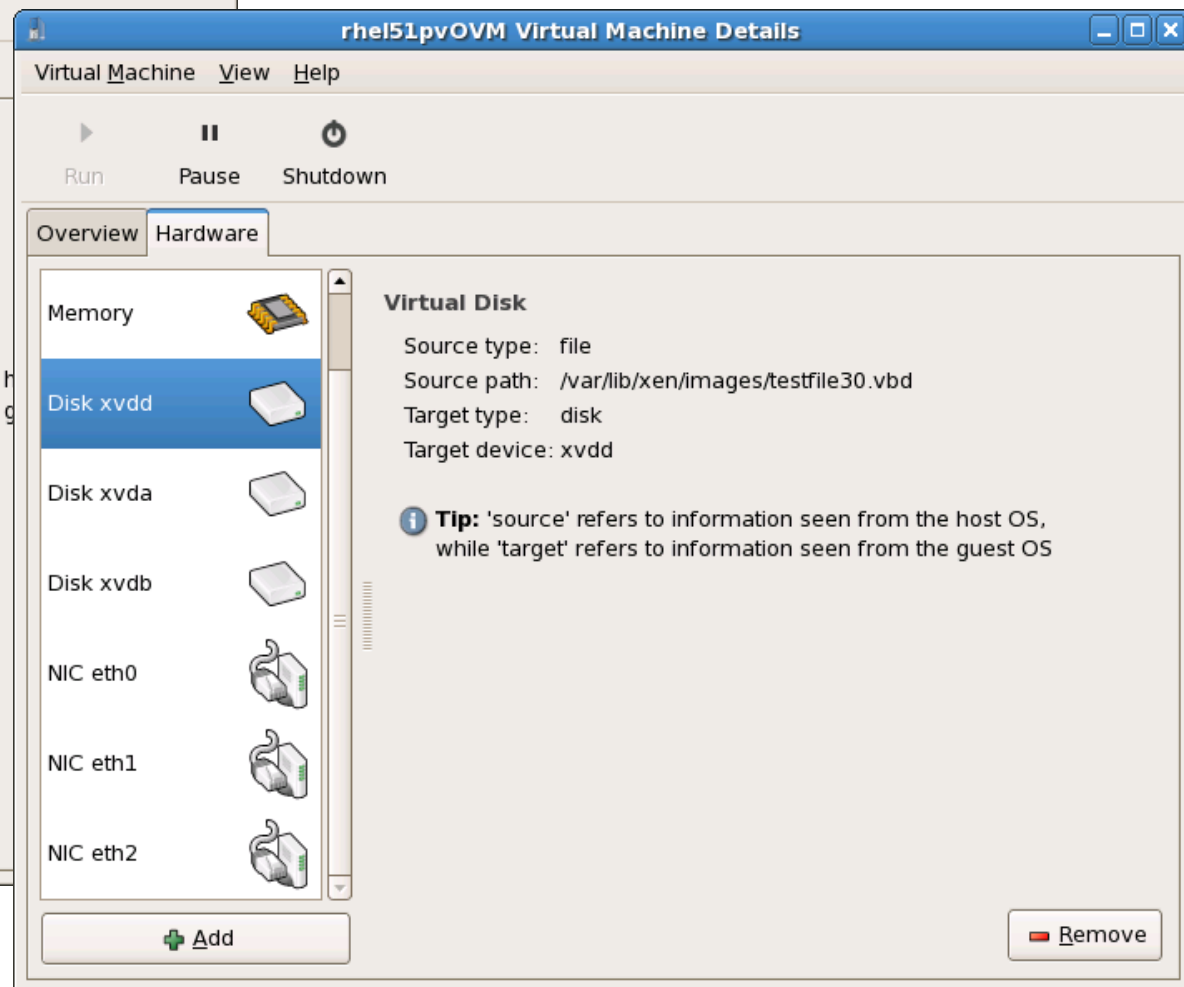
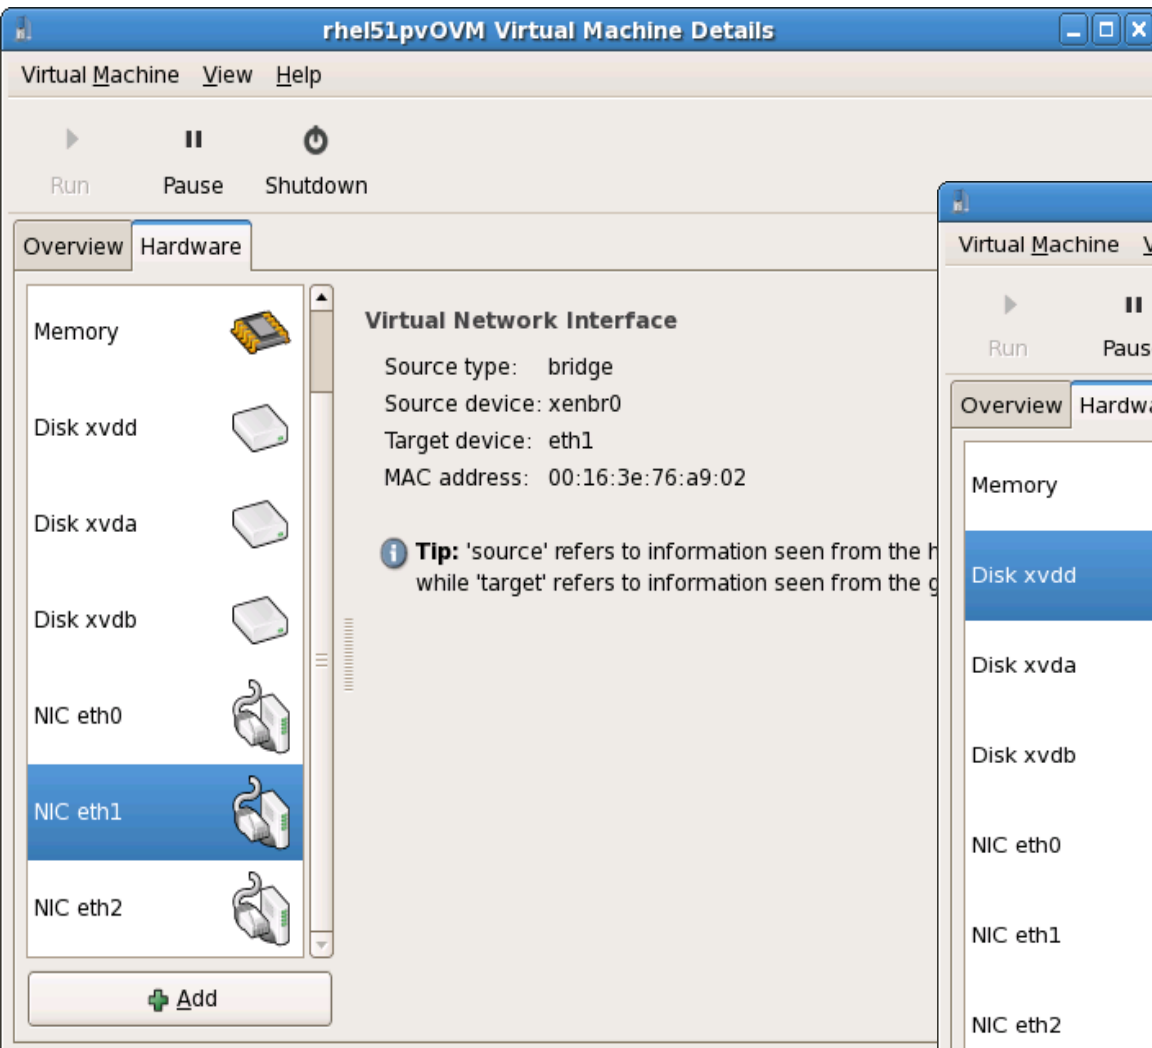
- How much memory should this machine be allocated?
- Current allocation: 499 MB
- Change allocation: 499 MB (with a spin button)
- Maximum allocation: 2048 MB (with a spin button)
- Total memory on host machine: 16382 MB

The hardware list on the left includes: Processor, Memory, Disk xvdd, Disk xvda, Disk xvdb, NIC eth0, and NIC eth1. There are 'Add' and 'Apply' buttons at the bottom.



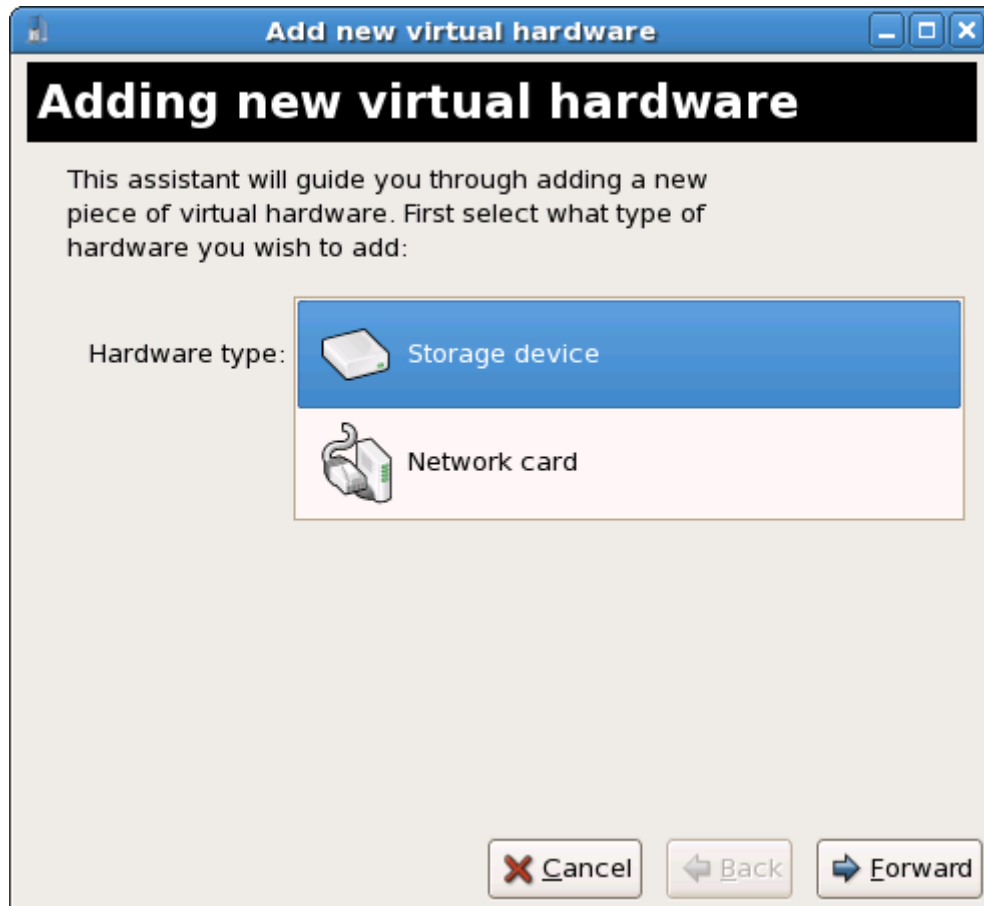
virt-manager / Hardware Management

Guest hardware resources



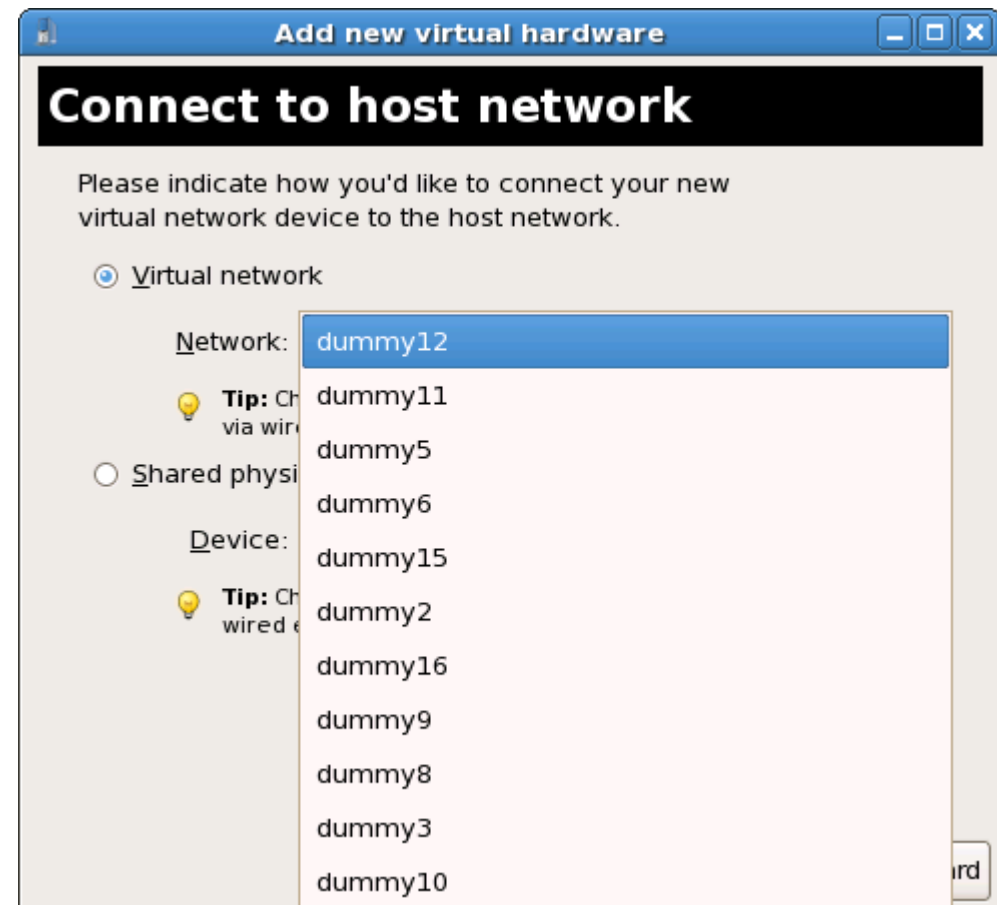
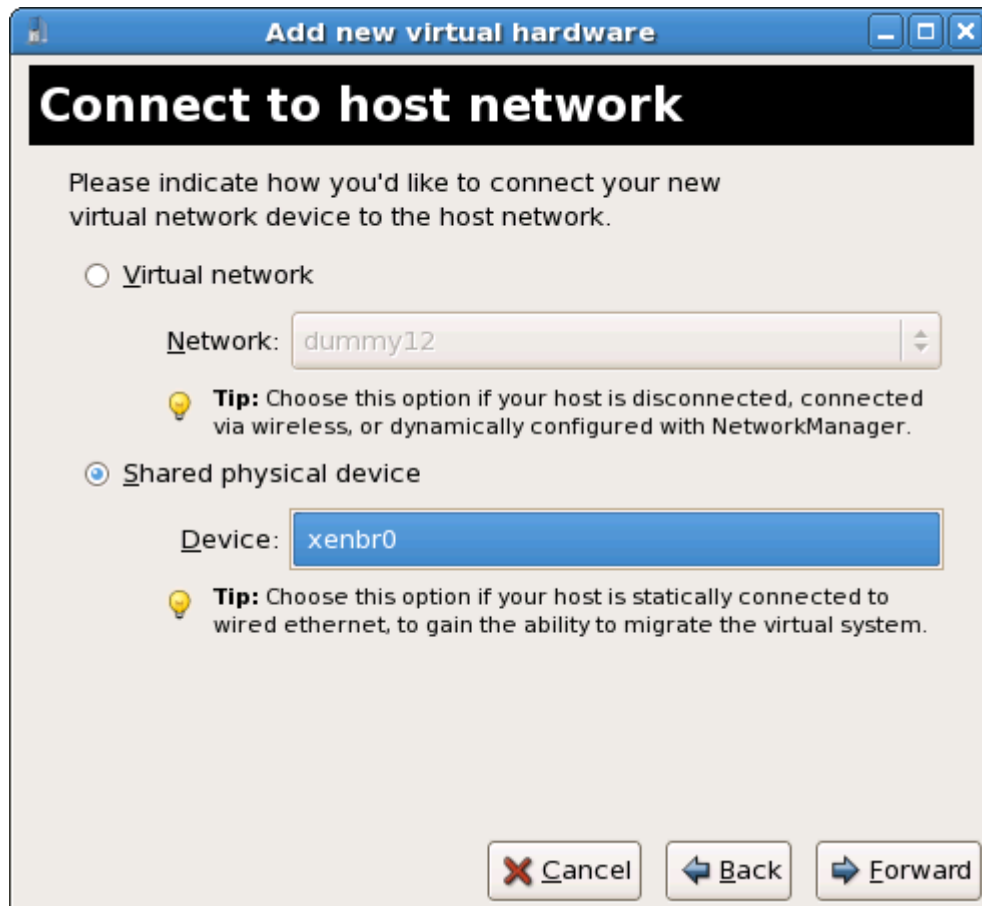
virt-manager / Hardware Management

Guest hardware resources



virt-manager / Hardware Management

Guest hardware resources



virt-manager / Resource Management

HVM Guest

rhel3u9x86ptest Virtual Machine Details

Virtual Machine View Help

Run Pause Shutdown

Overview Hardware

Basic details

Name: rhel3u9x86ptest

UUID: 39082041-950e-9f88-8cb5-5ae1c0782482

Status: Running

Performance

CPU usage:	0 %
Memory usage:	1031 MB of 16382 MB

rhel3u9x86ptest Virtual Machine Details

Virtual Machine View Help

Run Pause Shutdown

Overview Hardware

Processor

Memory

Disk hda

Disk xvdb

NIC eth0

CPUs

How many virtual CPUs should this machine be allocated?

Current allocation: 2

Change allocation:

Maximum allocation: 2

Total CPUs on host machine: 8

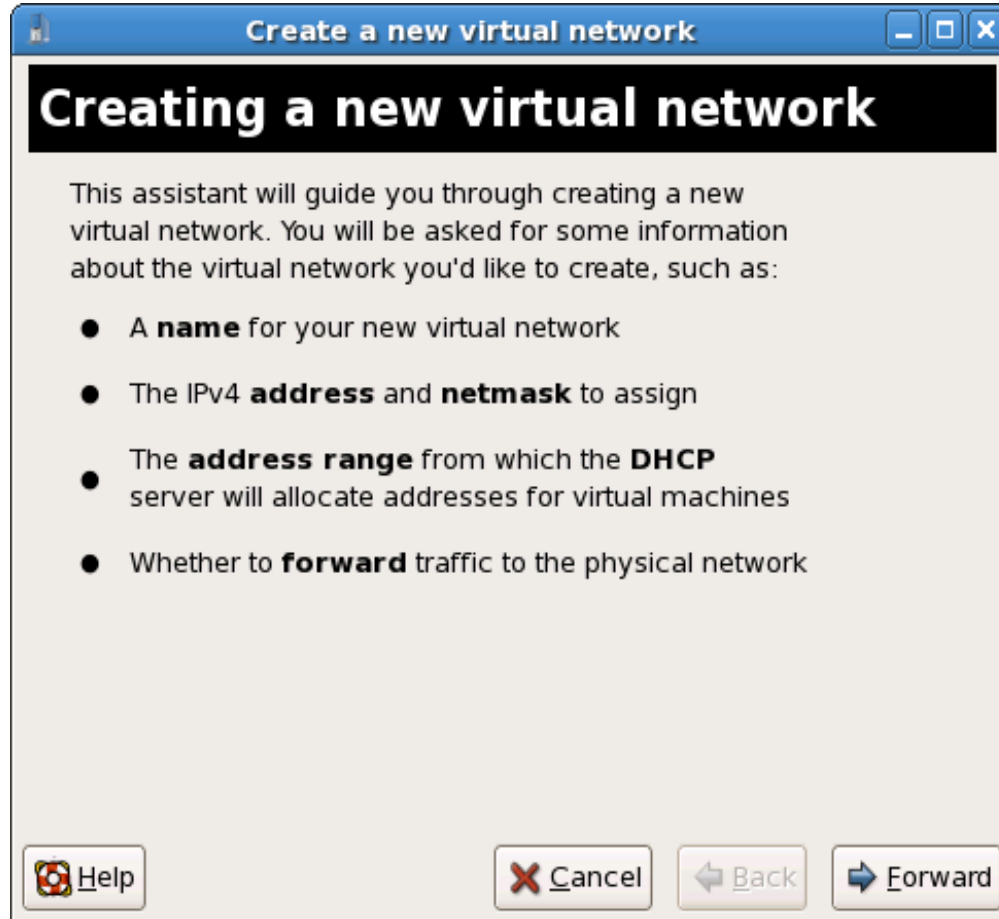
Tip: For best performance, the number of virtual CPUs should be less than (or equal to) the number of physical CPUs on the host system.

Add

Apply

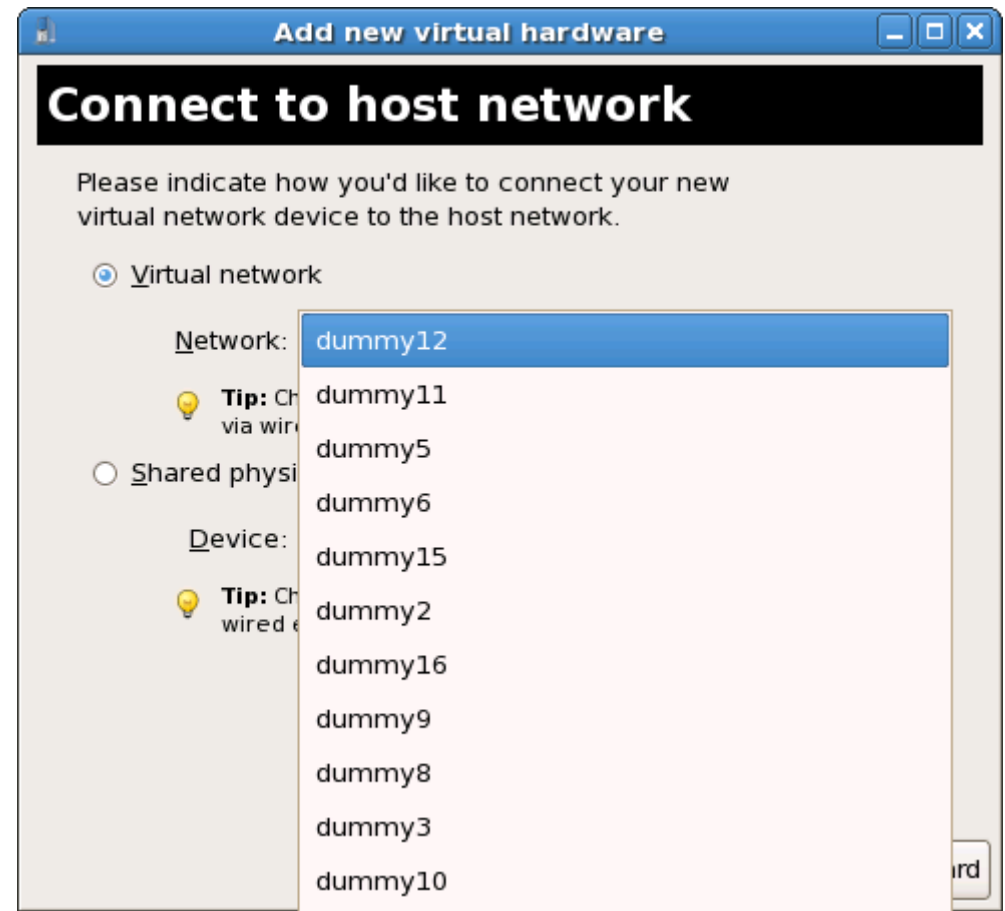
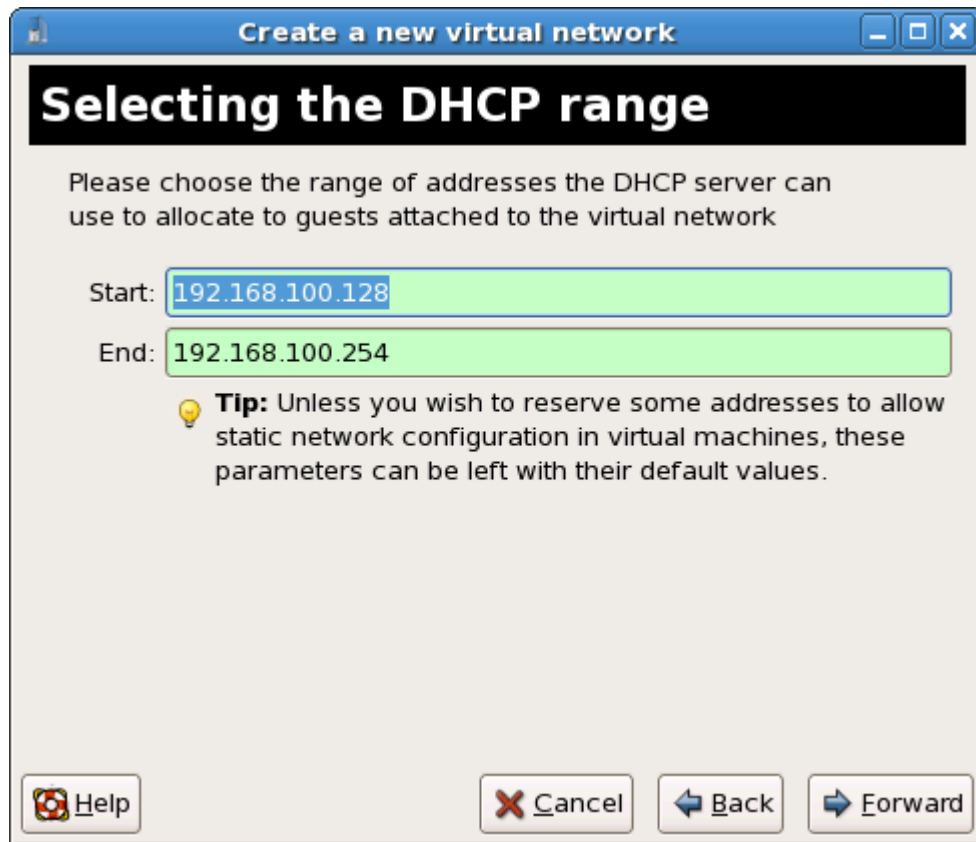


virt-manager / Virtual Network Management HyperVisor



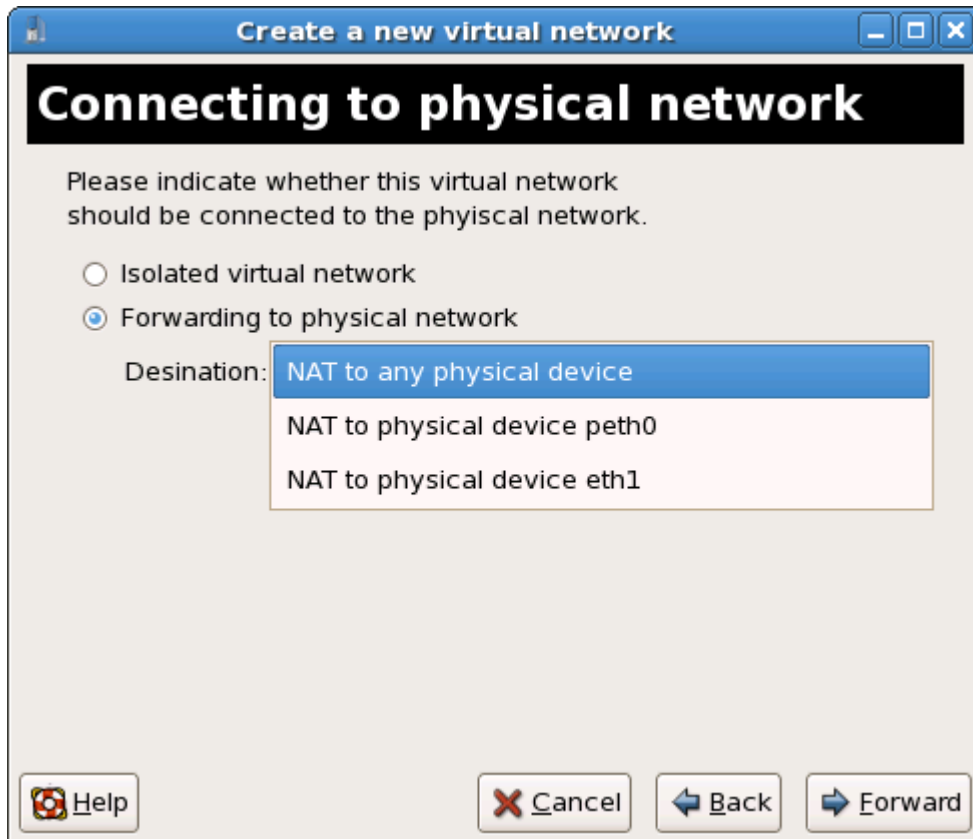
virt-manager / Virtual Network Management

HyperVisor



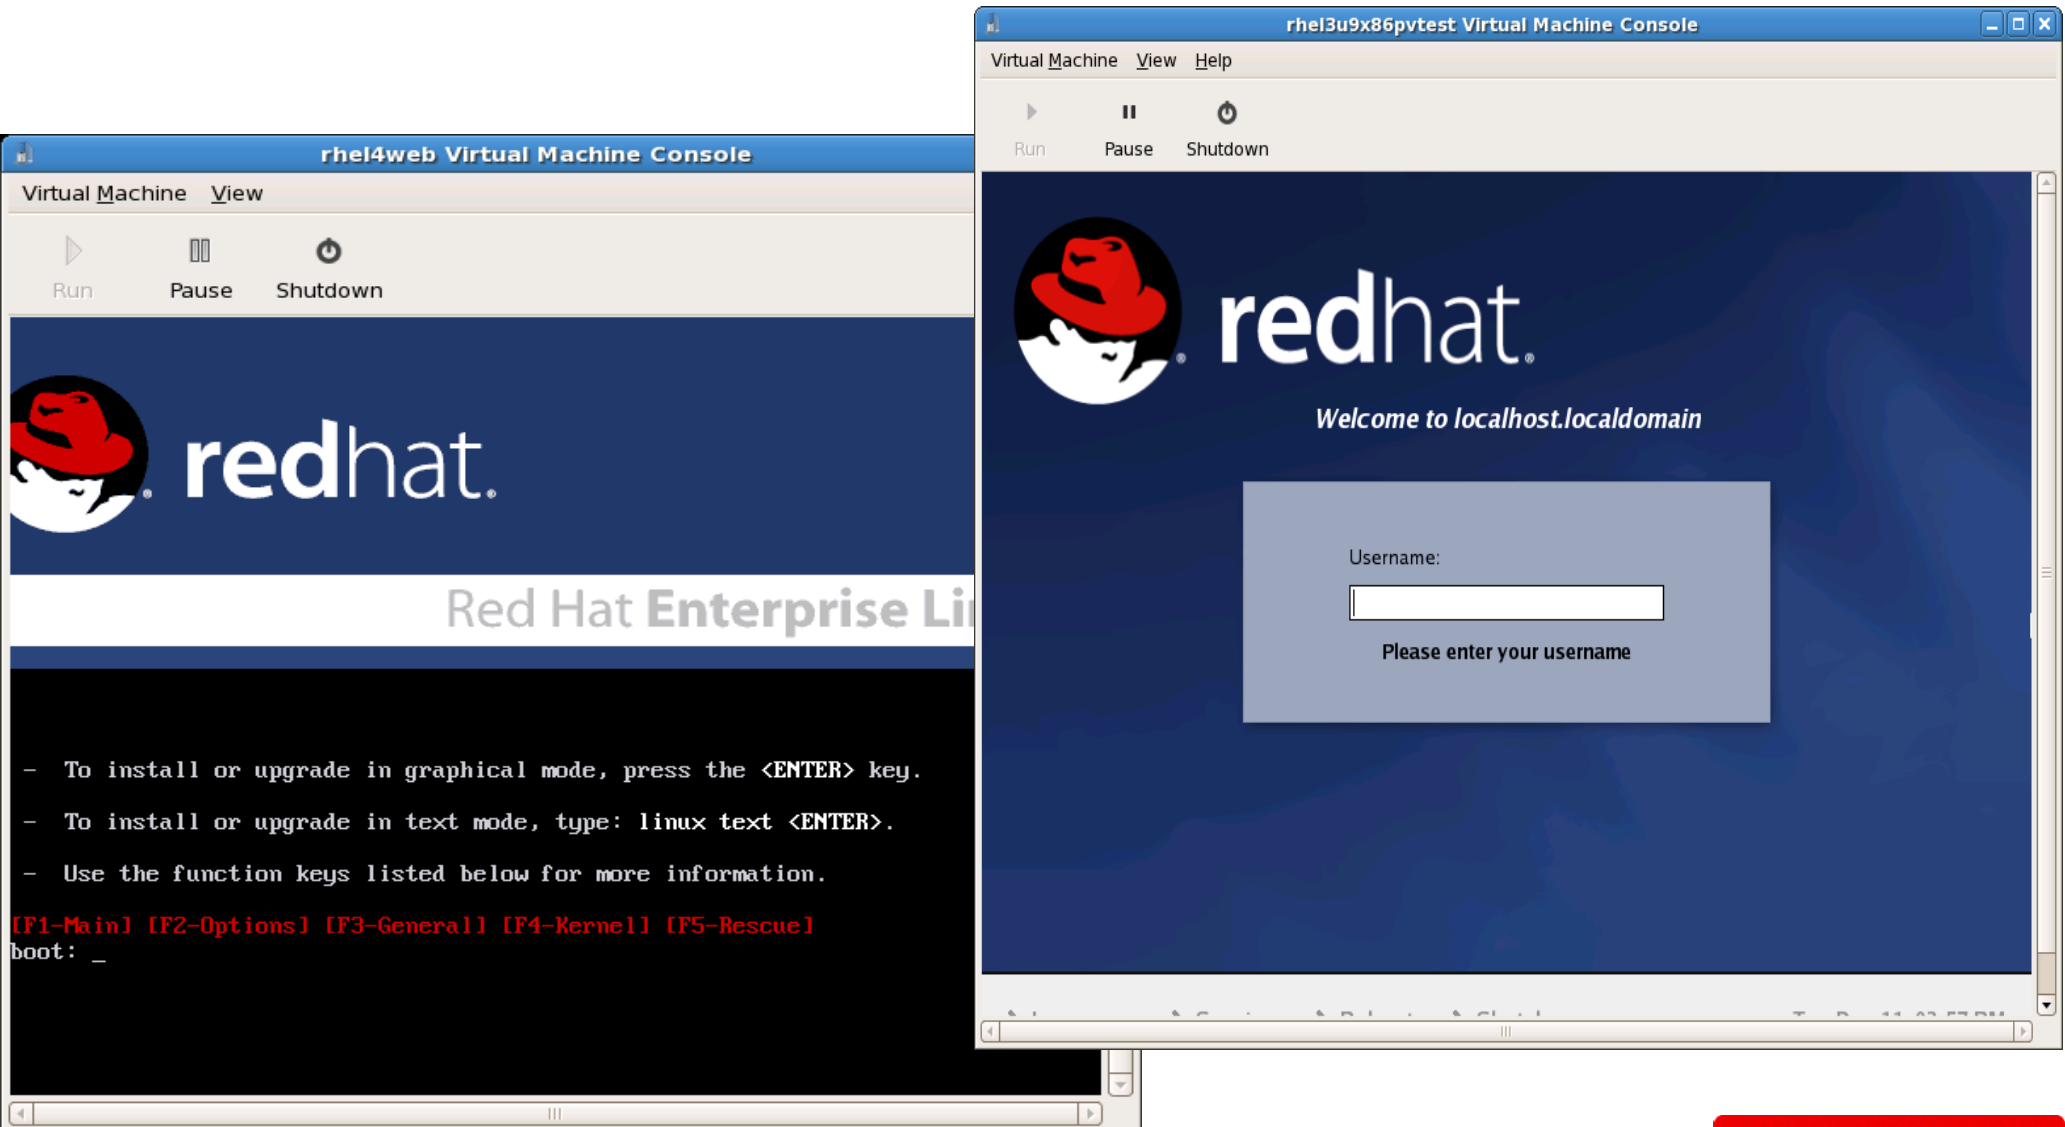
virt-manager / Virtual Network Management

HyperVisor

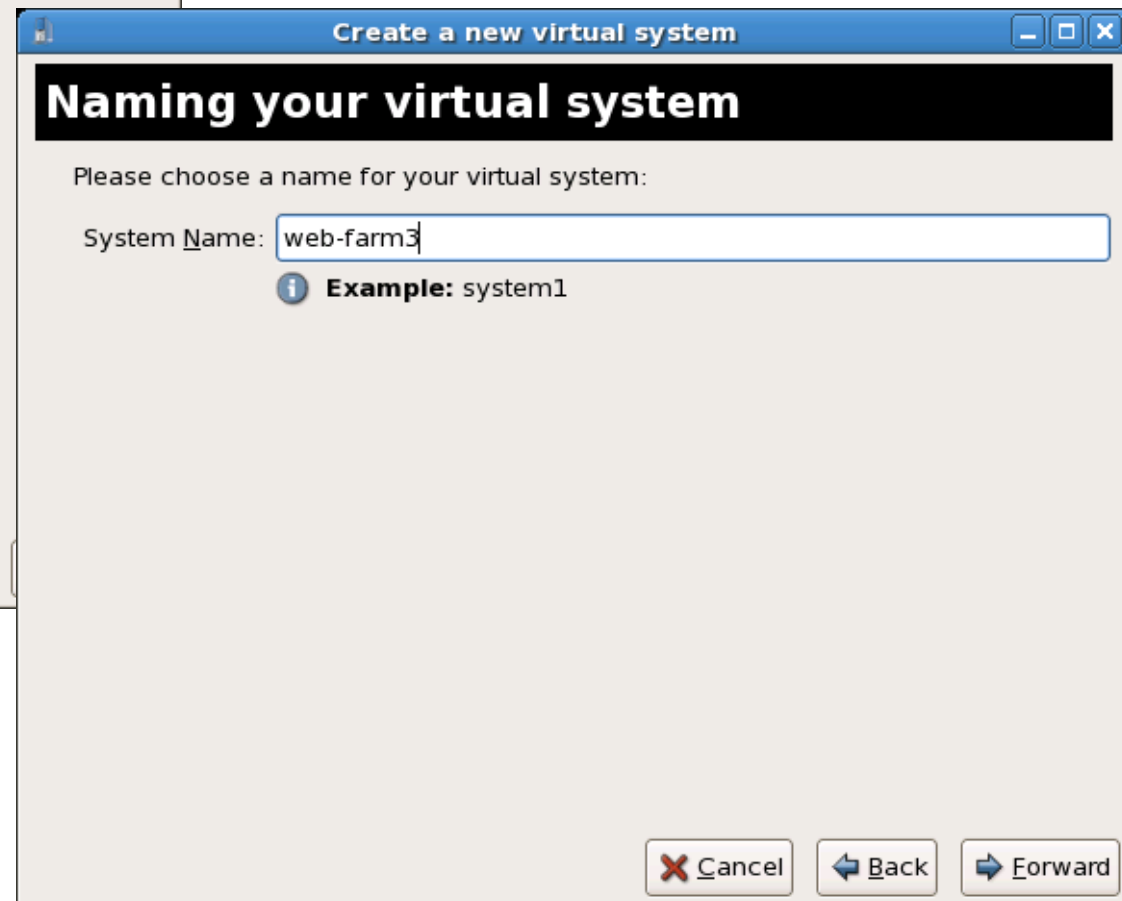
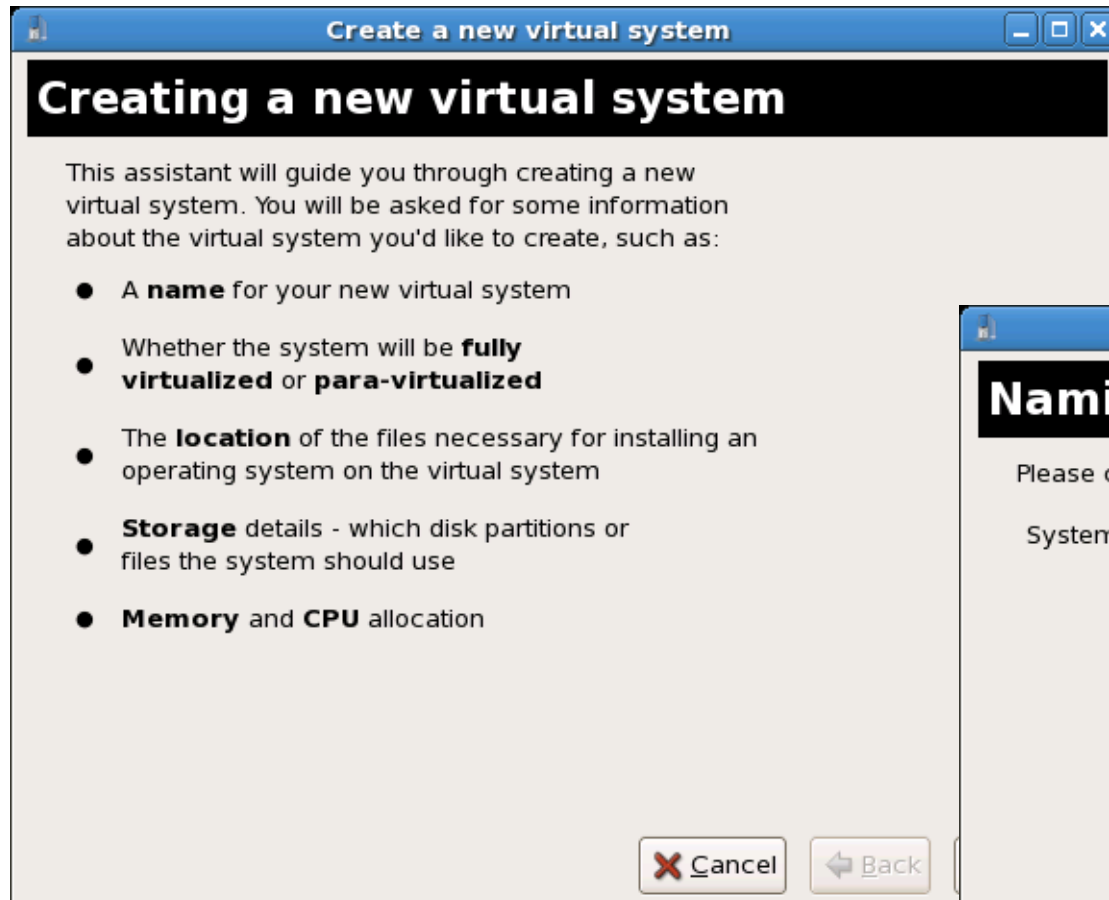


virt-manager / Console Window

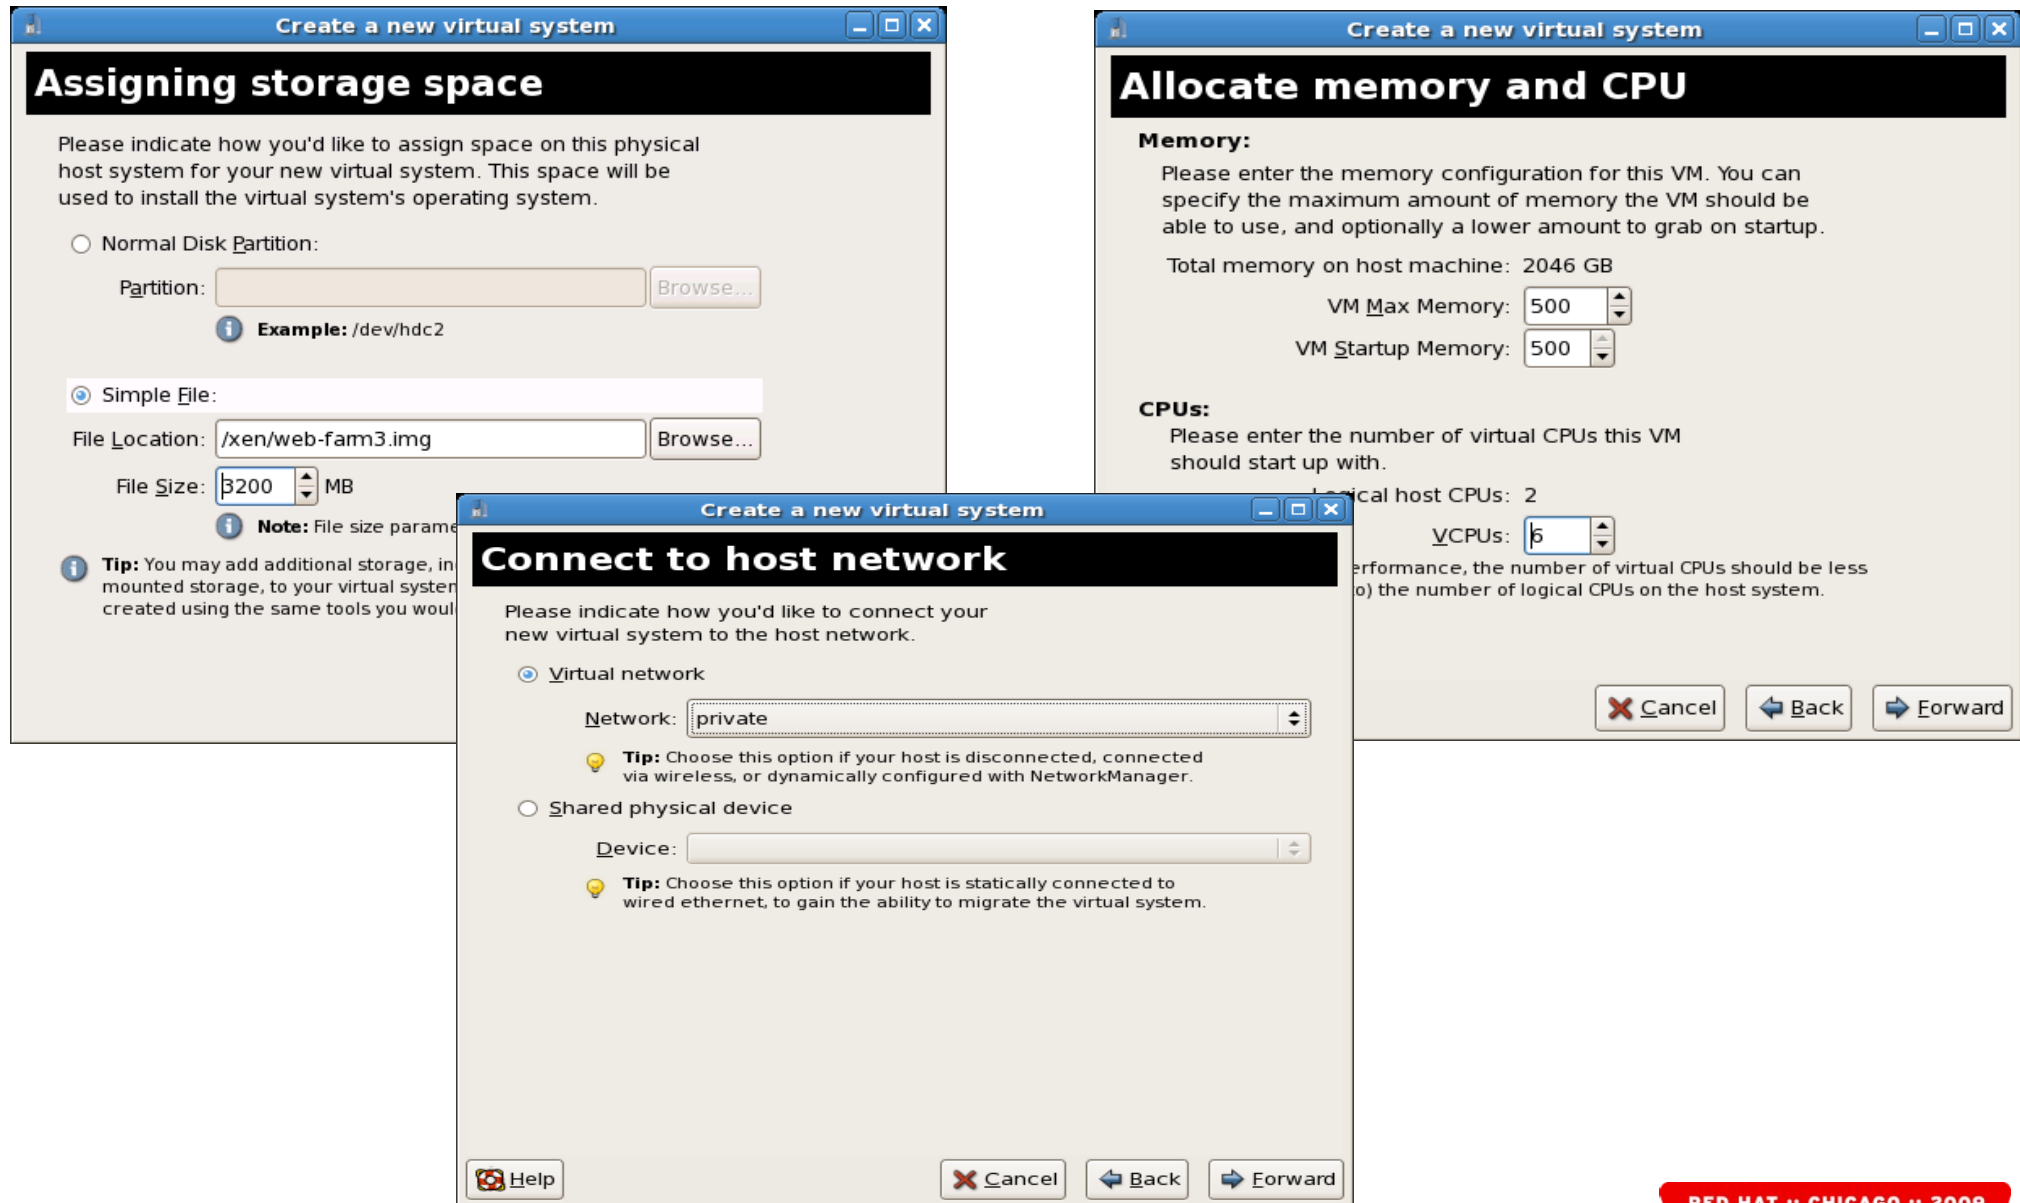
Guest



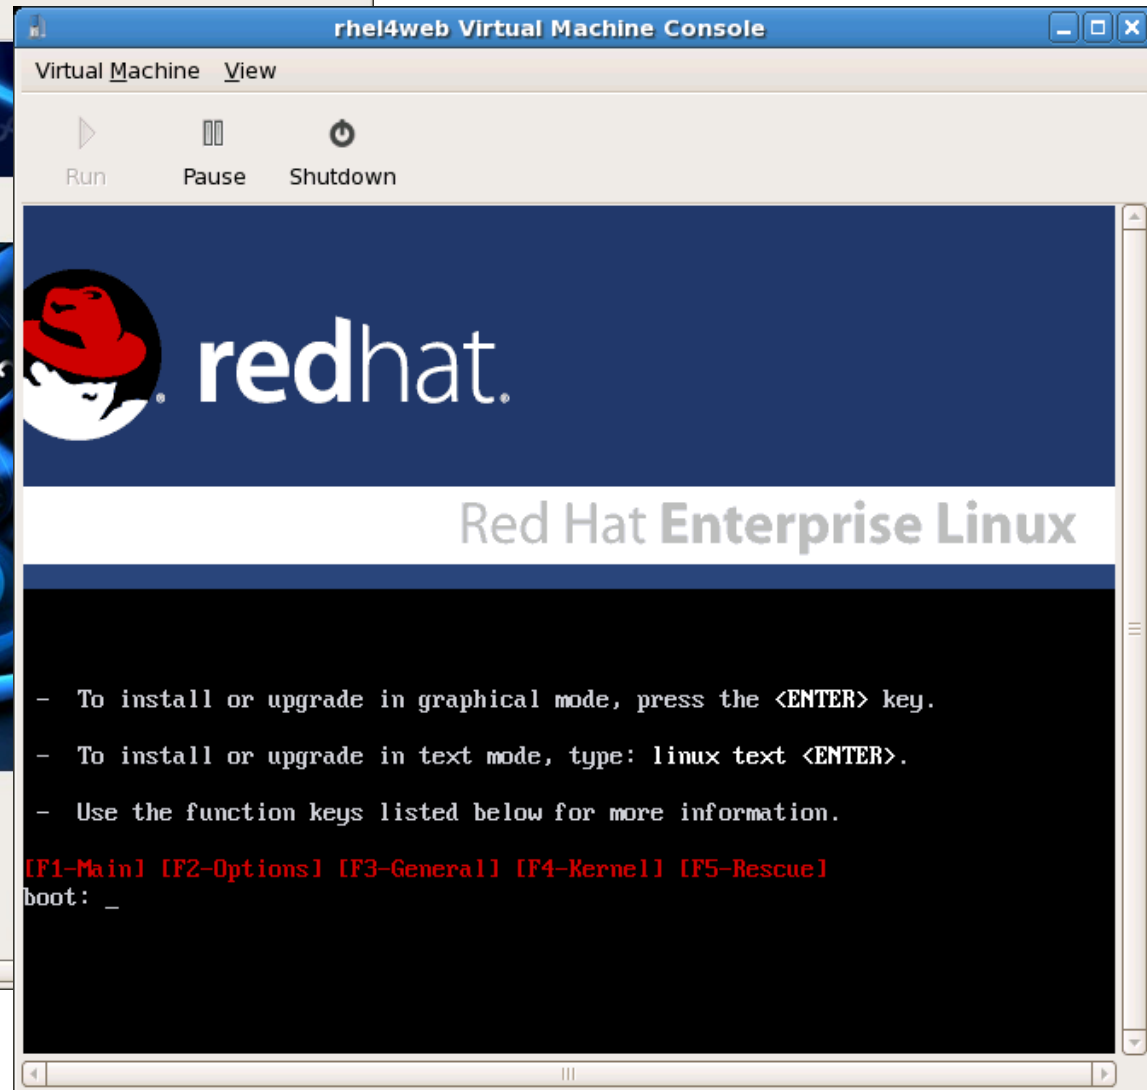
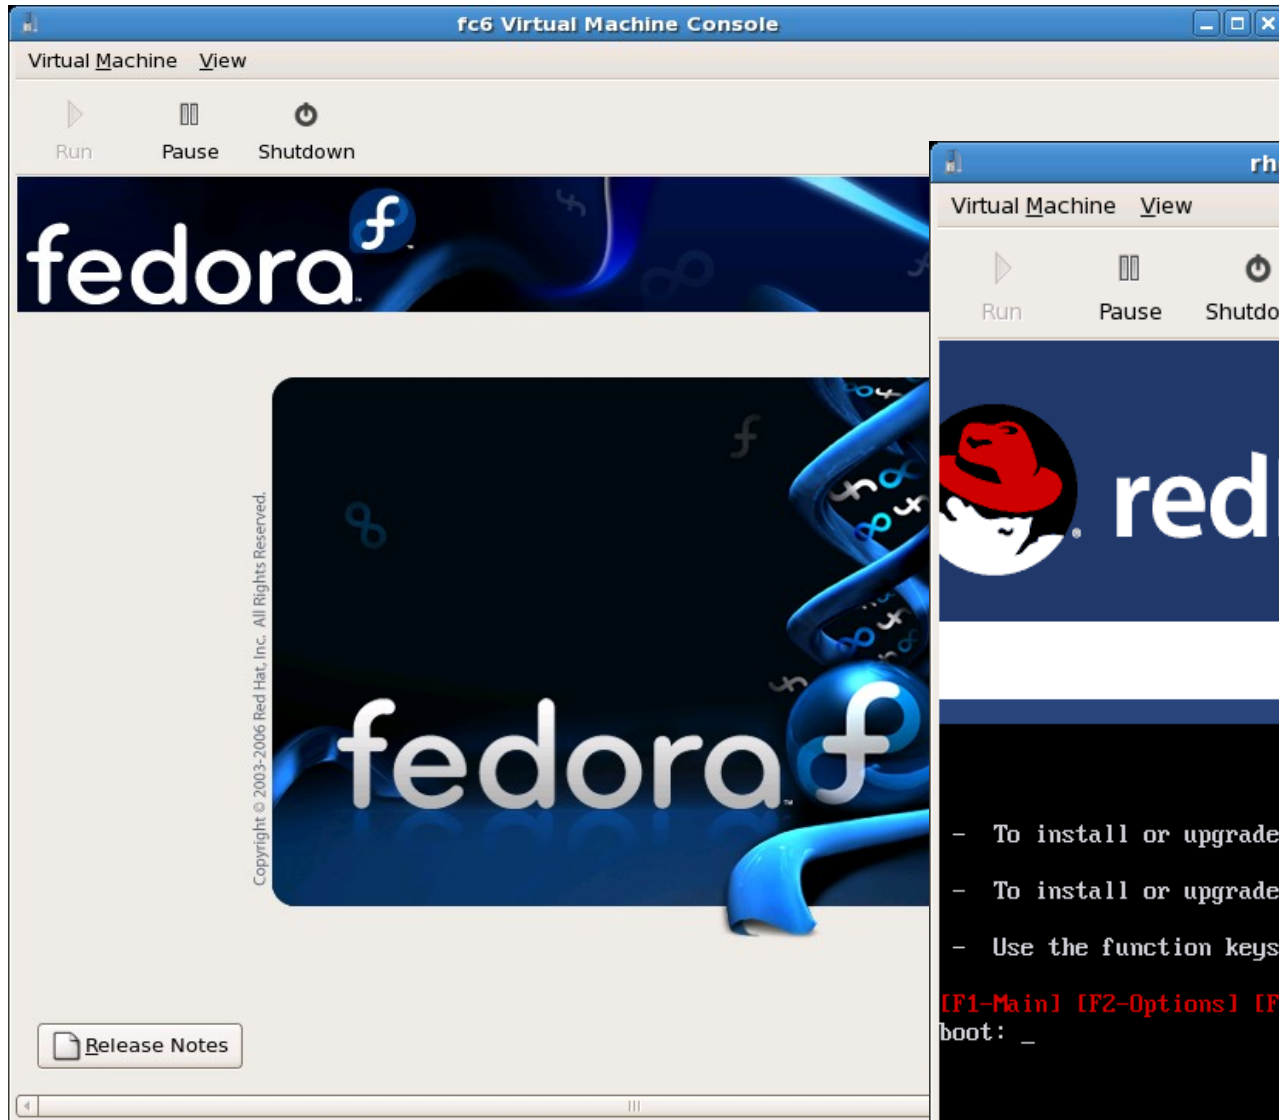
virt-manager / VM Creation Wizzard



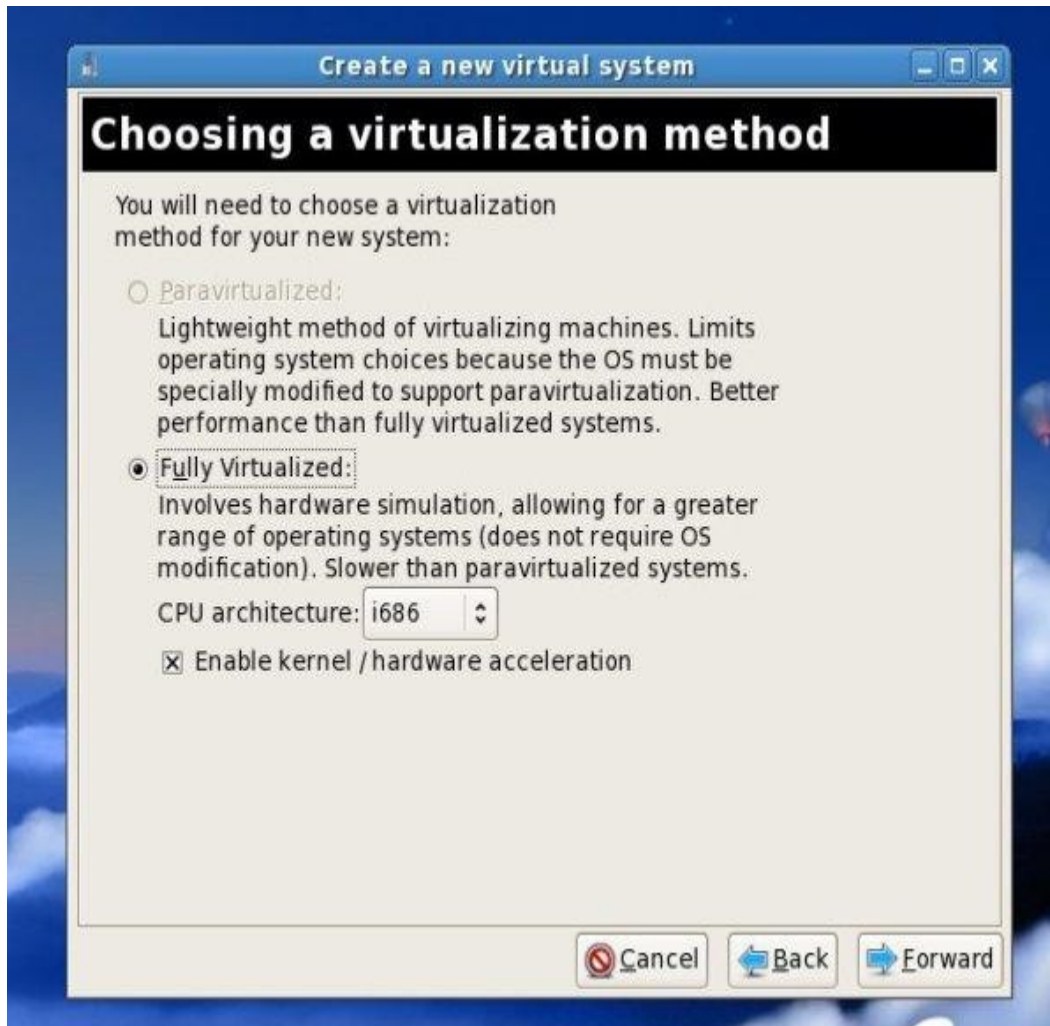
virt-manager Install Wizard (Resources)



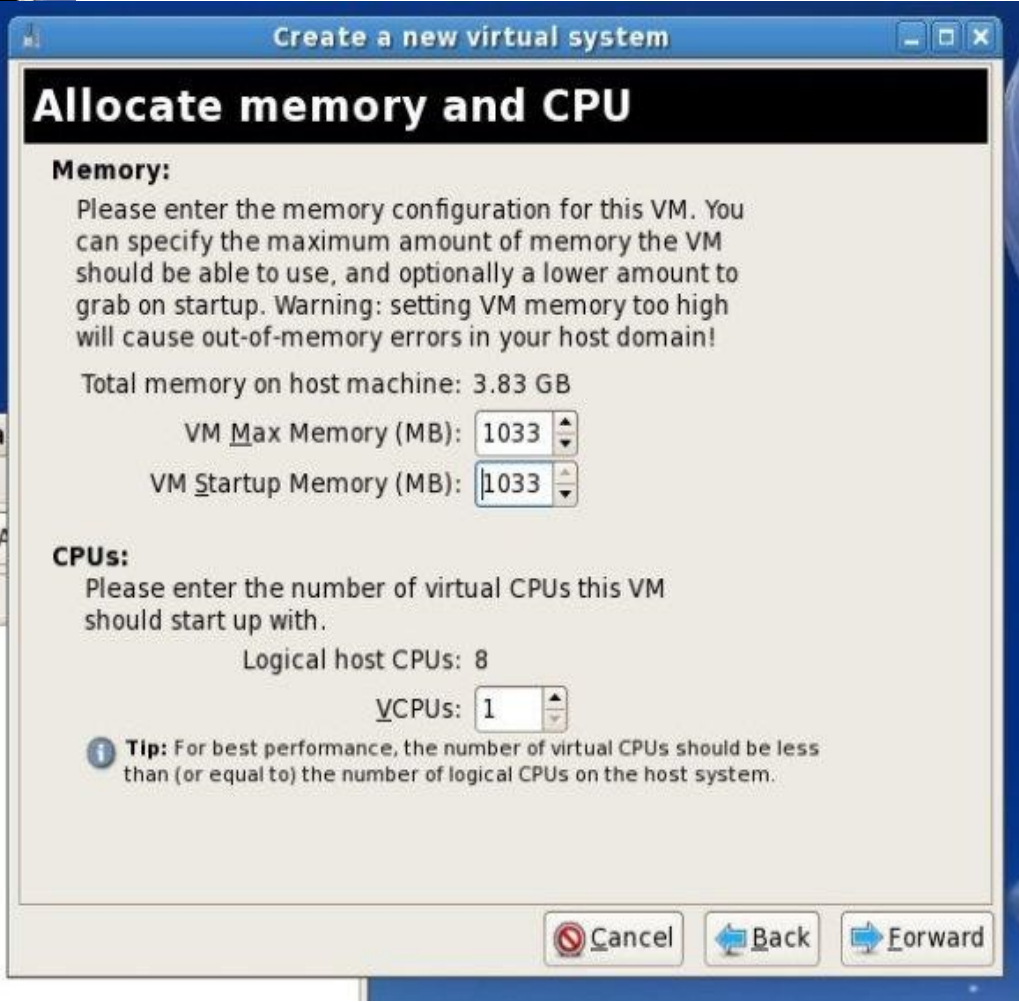
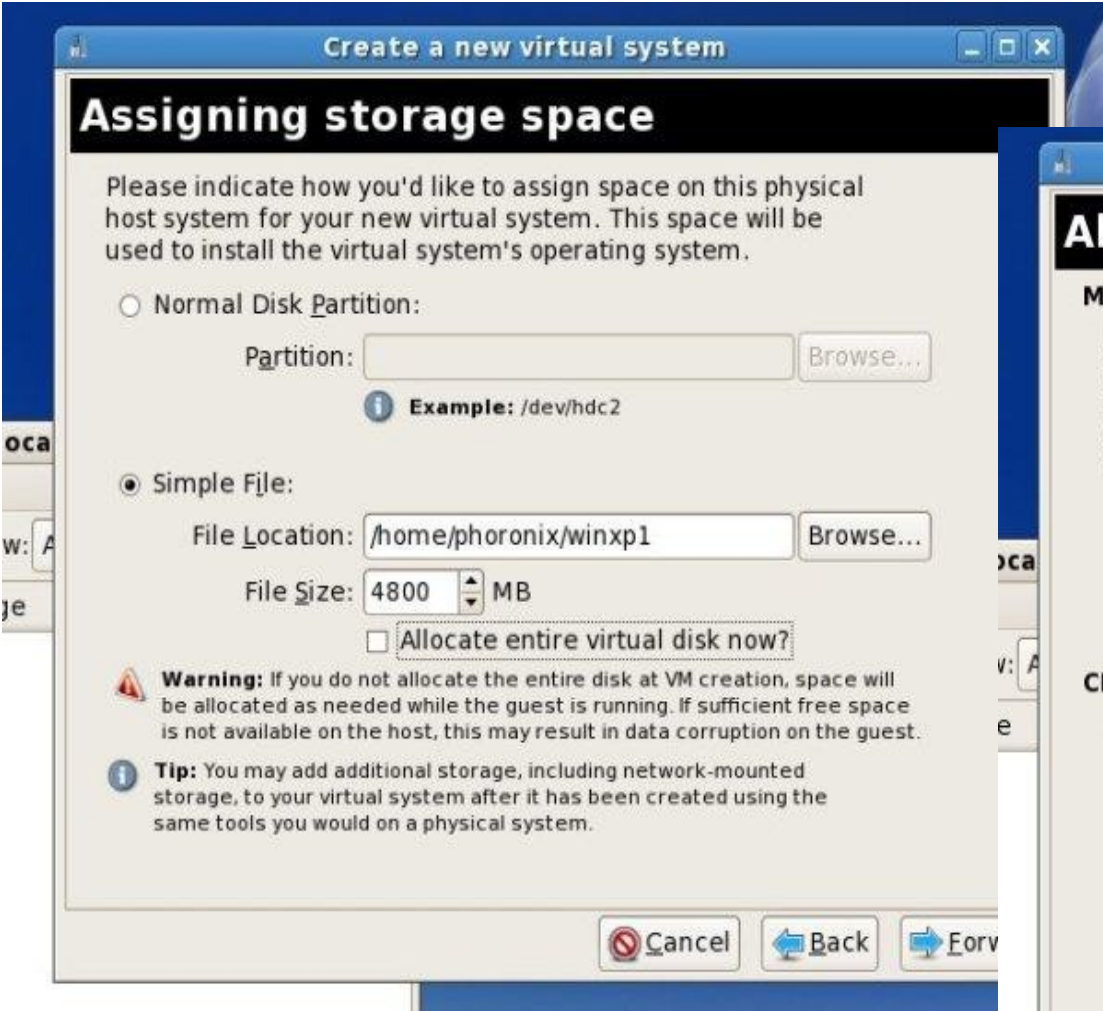
virt-manager Install (console)



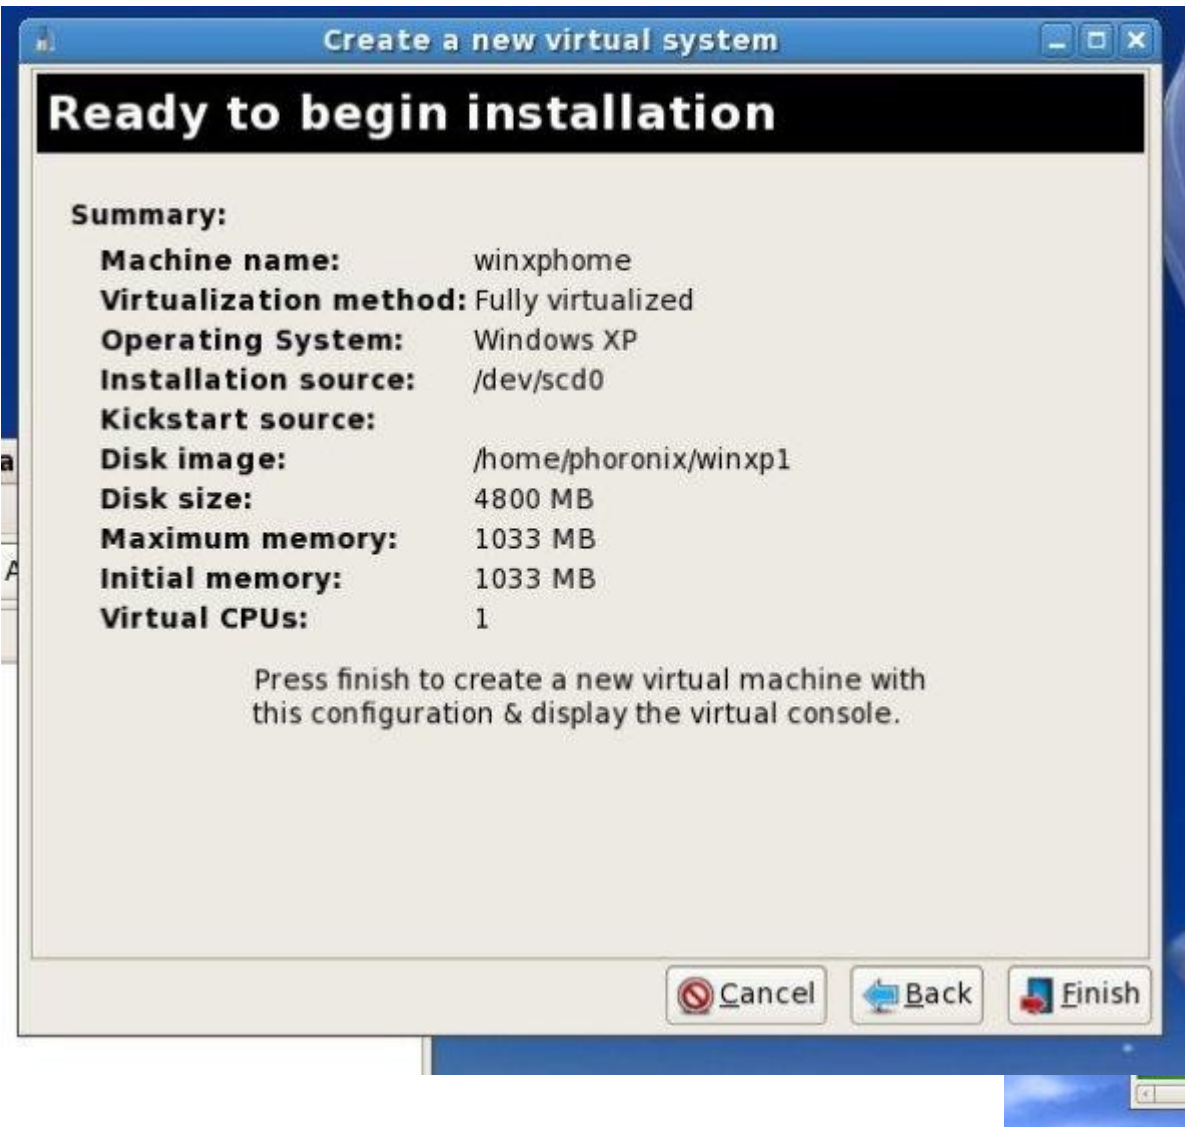
KVM virt-manager (Hardware Selection)



KVM virt-manager (Storage/Resource Selection)




KVM virt-manager (Summary)



Management Best Practices

- As of RHEL5.2 remote management via virt-manager/virsh
 - Can implement a “virtual appliance” style management instance
 - Can trigger live migration from remote , as well as all other virsh commands
- Can consolidate all management interfaces for virtualization and clustering onto a single management “appliance”/host
- If clustering is deployed use a bonded network interface for the cluster traffic and “virtualization” (live migration) traffic
 - Currently the clustering software assumes the network associated to the hostname is used for cluster and “virtualization” traffic
- Even if no clustering is used it is recommended to use a dedicated network for migration traffic (and use at least a GbE connection)

You need a centralized UI?



-
-
-
-



RED HAT :: CHICAGO :: 2009

SUMMIT

Red Hat Enterprise Virtualization

presented by



Red Hat Enterprise Virtualization Manager

ENTERPRISE VIRTUALIZATION

Logged in user: rhevadmin | Sign out | Configure | About

Search: Vms: [x] [★] [GO]

Bookmarks Tags

Data Centers Clusters Hosts Storage Virtual Machines Pools Templates Users Events Monitor

New Edit Remove

New Server New Desktop Edit Remove Run Once [▶] [⏸] [⏹] [🖥] Migrate Make Template Custom Actions Assign tags

Name	Cluster	Host	IP Address	Memory	CPU	Network	Display	Status
RHEL-53	Default			0%	0%	0%		Down
Window-2003	Default	station	.1	11%	0%	1%	VNC	Up
Windows-XP	Default			0%	0%	0%		Down

- New
- Edit
- Remove
- Run
- Suspend
- Shut down
- Stop
- Migrate
- Make Template
- Run Once
- Console: VNC
- Change CD

General Network Interfaces Virtual Disks Snapshots Ap History

Create Preview Commit Undo

June, 2009

Su	Mo	Tu	We	Th	Fr	Sa	Time	Description	Disks	Installed Applications
31	1	2	3	4	5	6	11:41 AM	Snapshot-01	1	
7	8	9	10	11	12	13				
14	15	16	17	18	19	20				
21	22	23	24	25	26	27				
28	29	30	1	2	3	4				
5	6	7	8	9	10	11				

Red Hat Enterprise Virtualization Hypervisor

Scalability

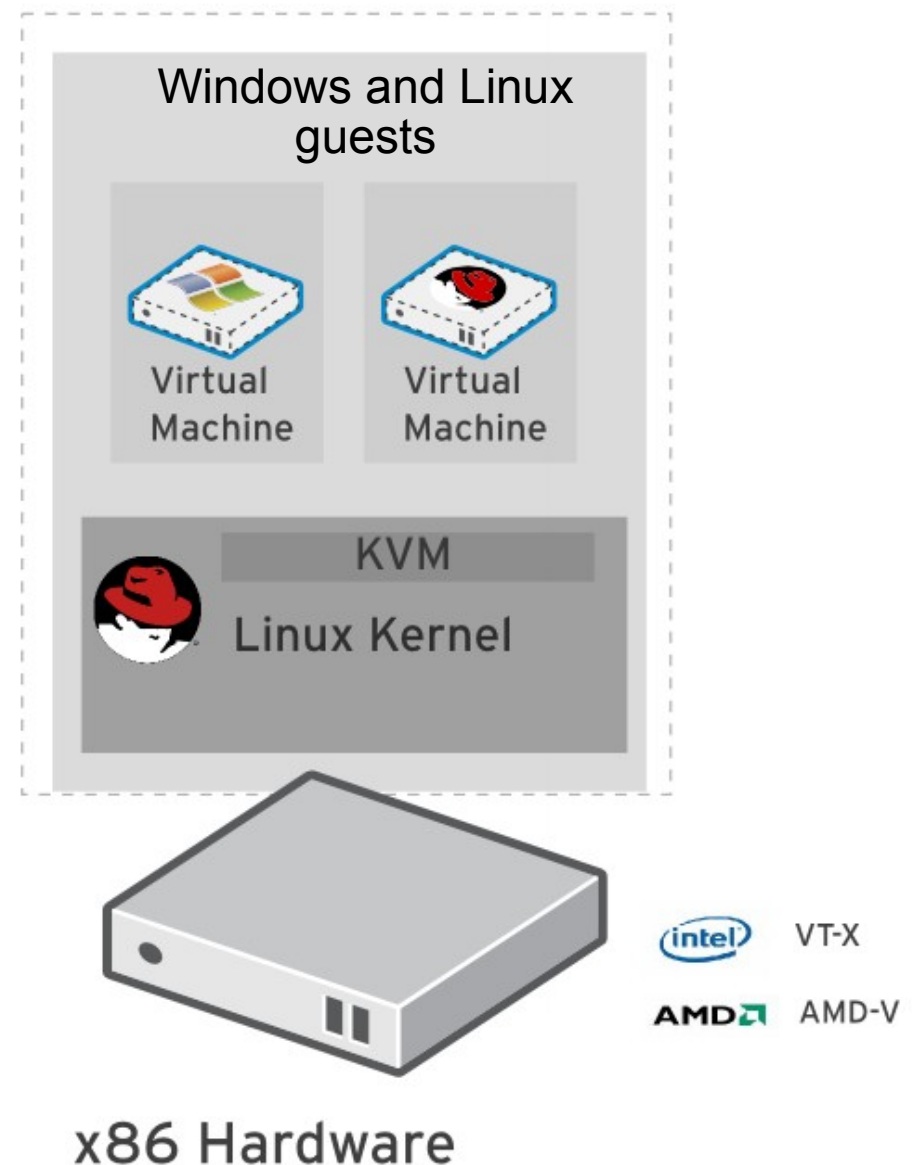
- Host: 96 cores, 1 TB RAM
- Guest: 16vCPU, 64 GB RAM

Industry Standards

- Trusted RHEL kernel + KVM
- High performance VirtIO drivers
- Libvirt management interface

Advanced Features

- Memory Page sharing
- SELinux for high security and isolation
- Live migration
- Snapshots
- Thin provisioning

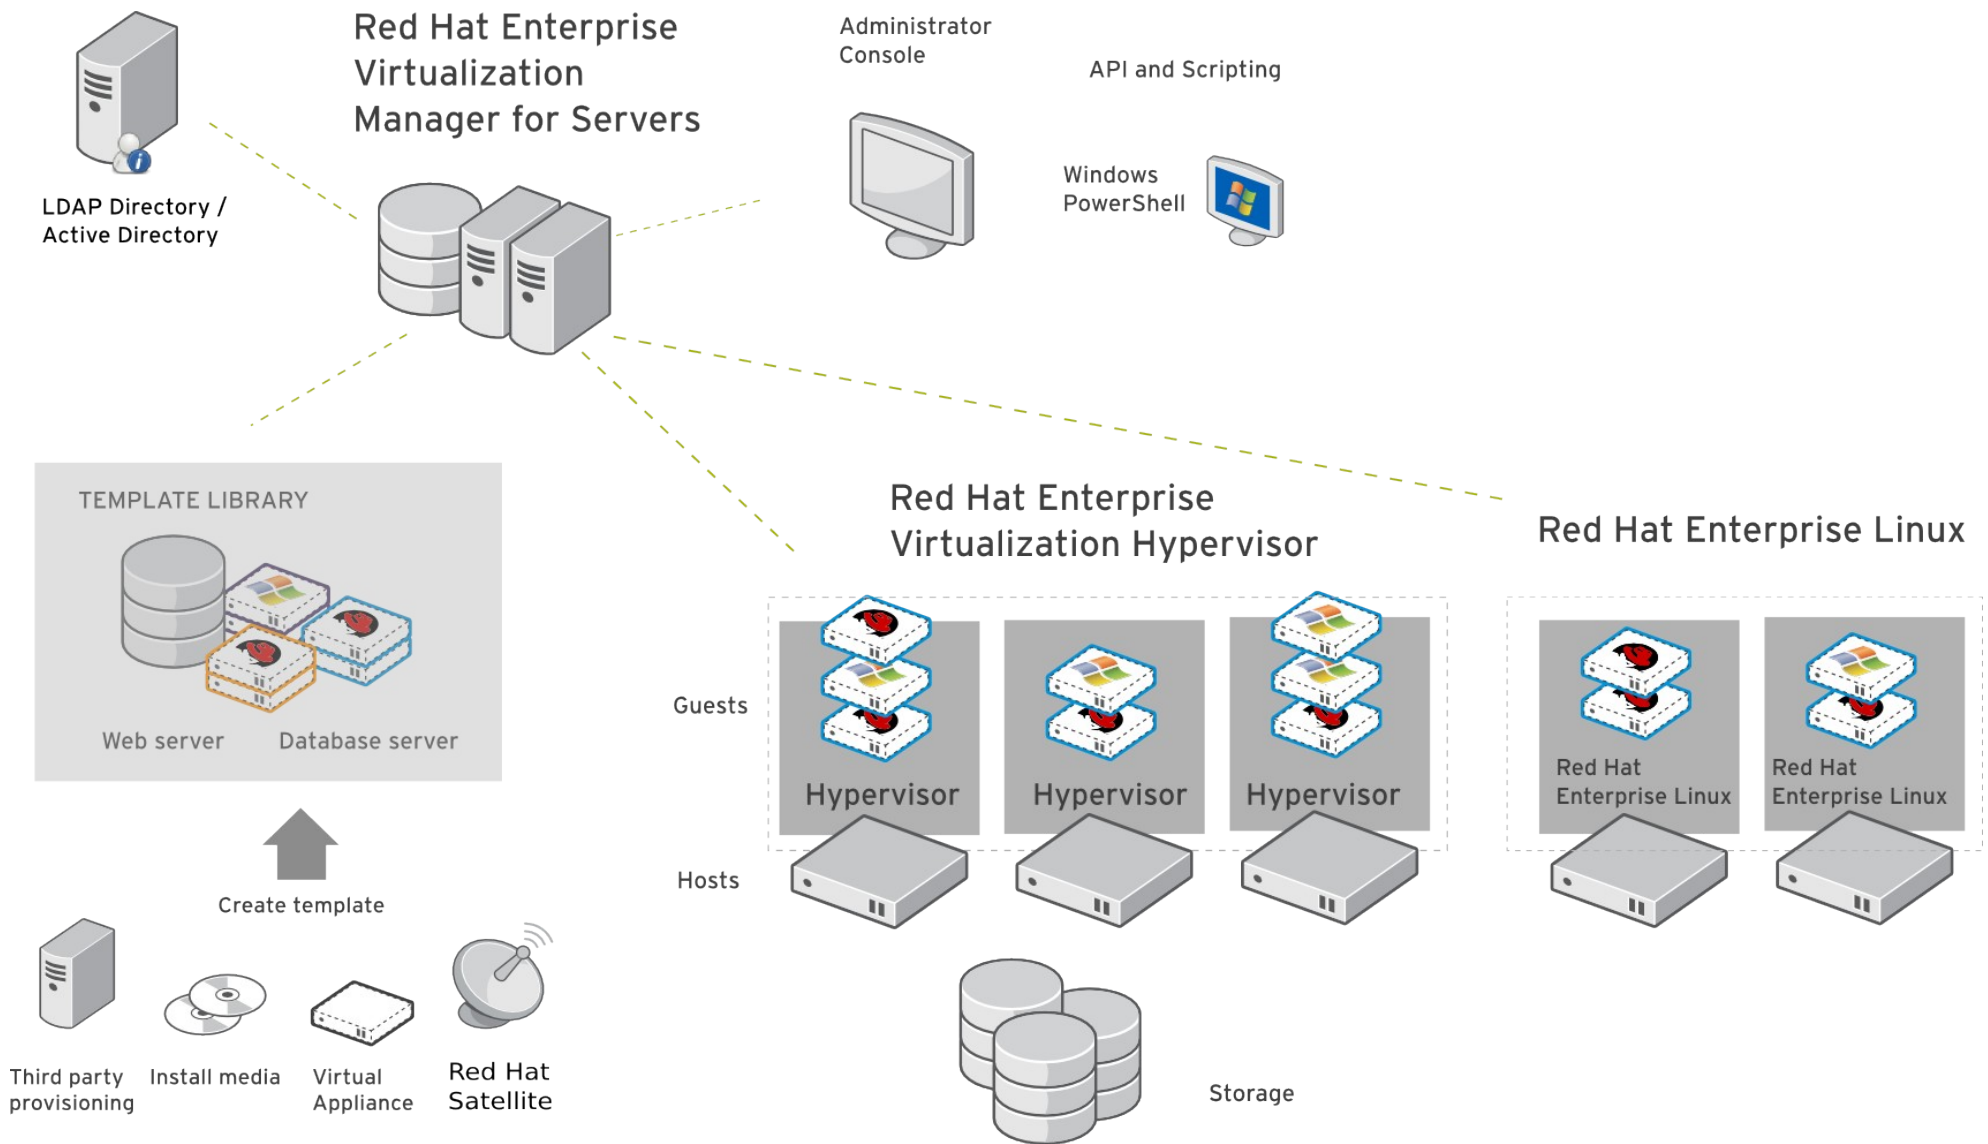


RED HAT :: CHICAGO :: 2009

SUMMIT

RHEV Server Virtualization

Red Hat Enterprise Virtualization - Servers

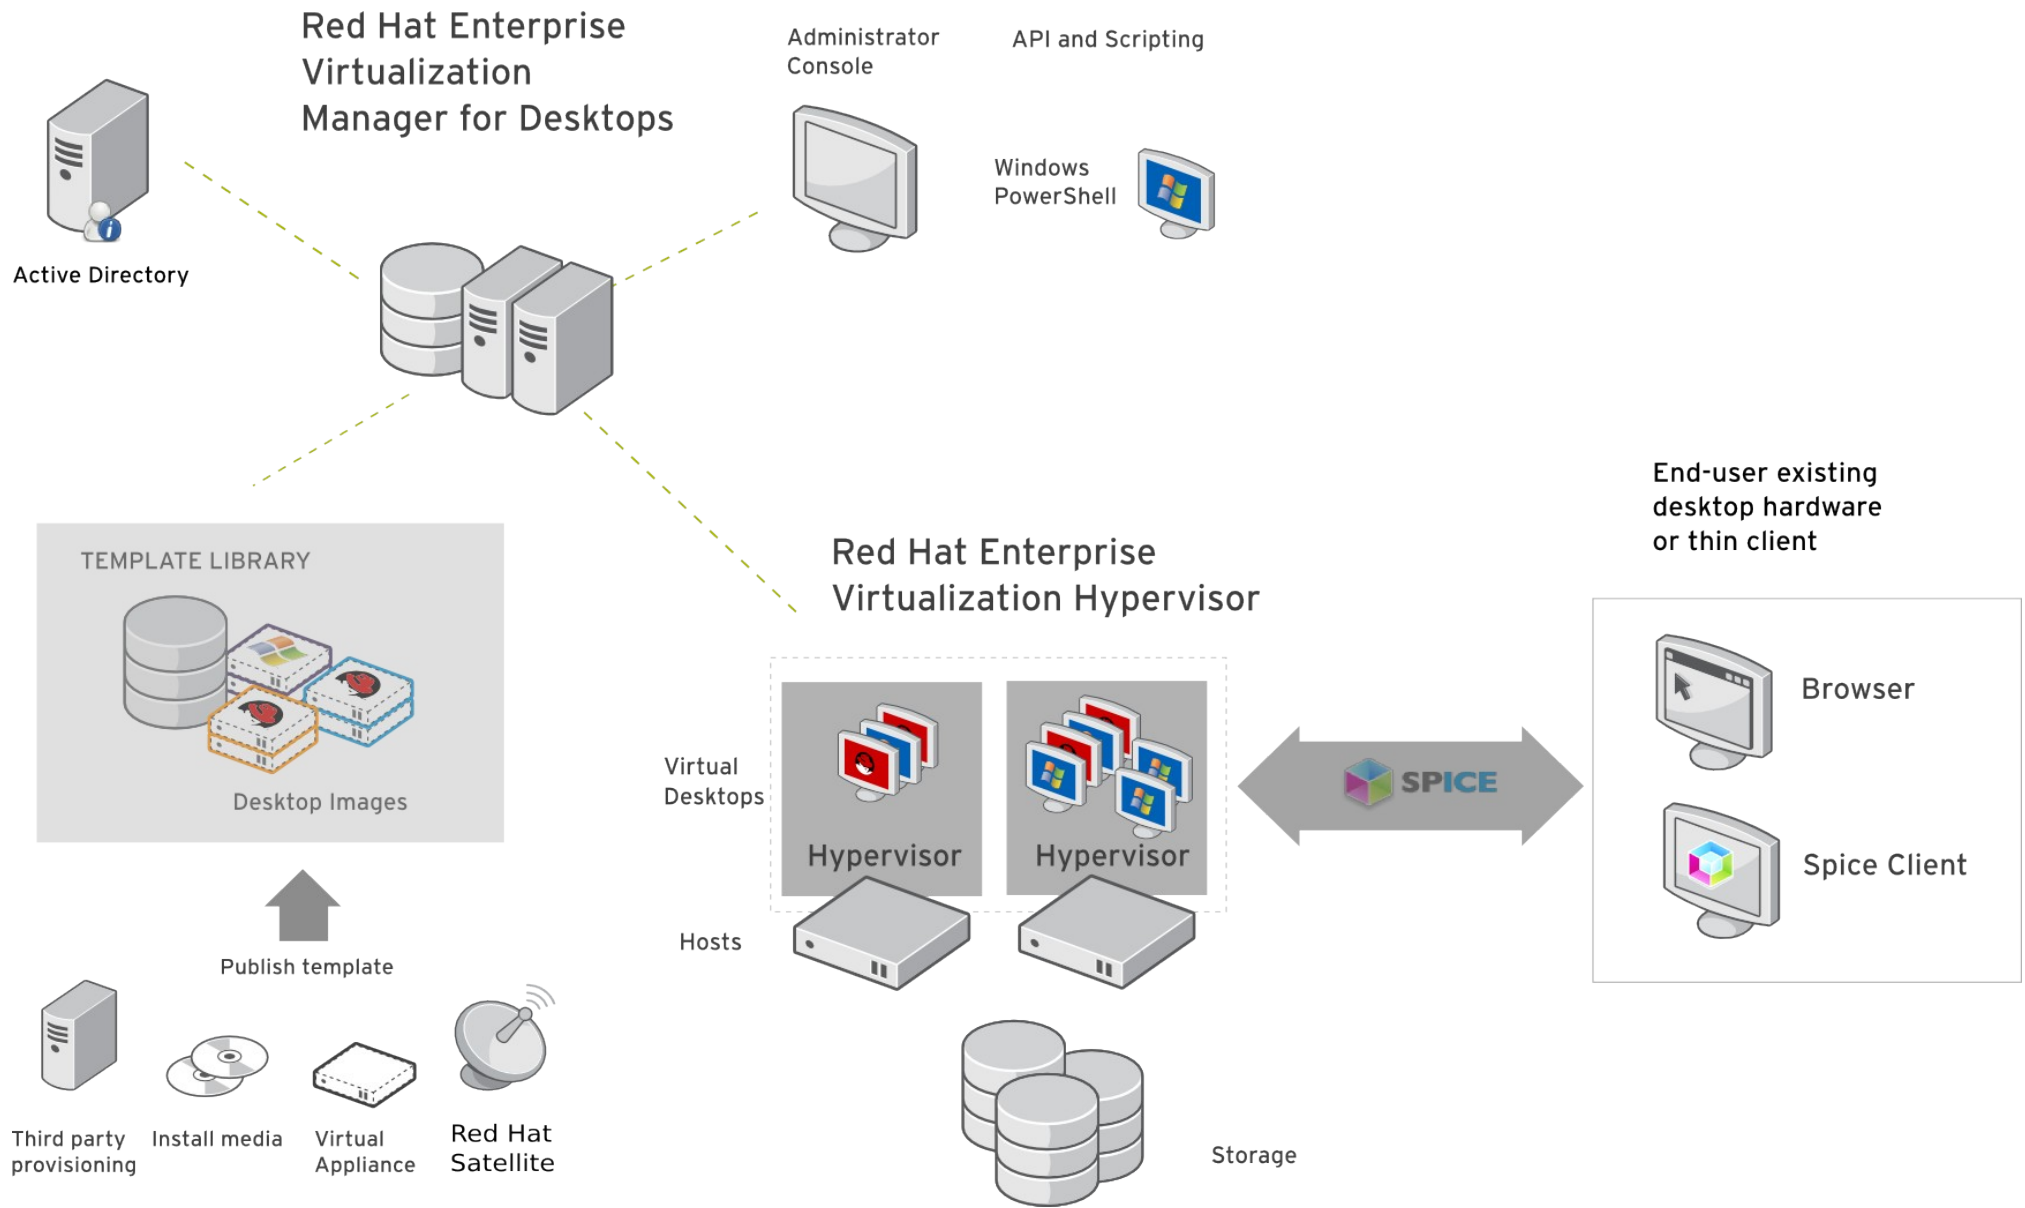


RED HAT :: CHICAGO :: 2009

SUMMIT

RHEV Desktop Virtualization

Red Hat Enterprise Virtualization - Desktops



QUESTIONS?

**TELL US WHAT YOU THINK:
[REDHAT.COM/SUMMIT-SURVEY](https://redhat.com/summit-survey)**