

SUMMIT

**JBoss
WORLD**

PRESENTED BY RED HAT

**LEARN. NETWORK.
EXPERIENCE OPEN SOURCE.**

www.theredhatsummit.com

STORAGE RECONFIGURATION WITH RED HAT ENTERPRISE LINUX AND RED HAT ENTERPRISE VIRTUALIZATION

Ian Pilcher, RHCA – Sr. Solution Architect
Tom Coughlan – Sr. Engineering Manager
Red Hat, Inc.

June 25, 2010

SUMMIT

**JBoss
WORLD**

PRESENTED BY RED HAT

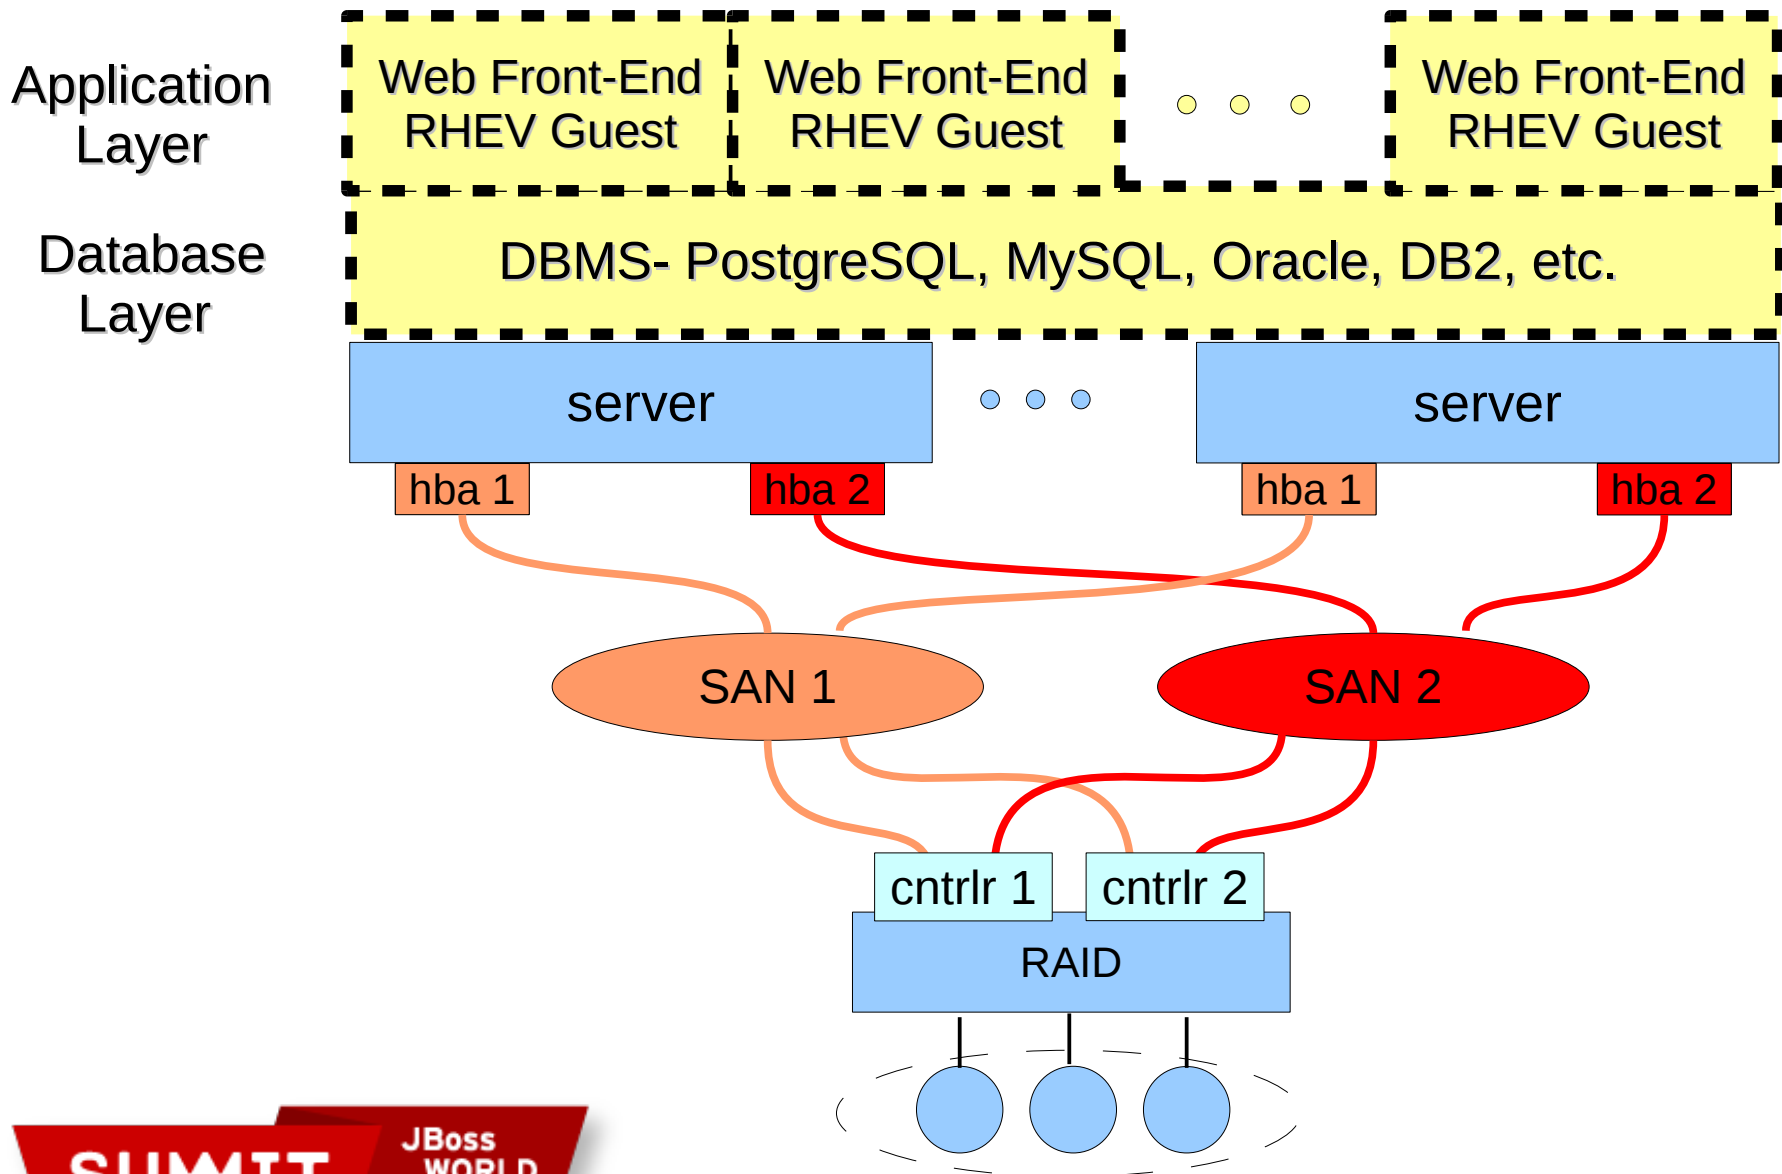


Objectives

- Add additional LUNs to a Red Hat Enterprise Linux (RHEL) server
- Expand LUNs used by a RHEL server
- Make use of additional/expanded LUNs
- Add additional LUNs to a Red Hat Enterprise Virtualization (RHEV) storage pool
- Overcommit RHEV guest storage with “sparse provisioning”
- Use RHEV templates to ease deployment of guests
- Manage guest updates and testing with RHEV snapshots

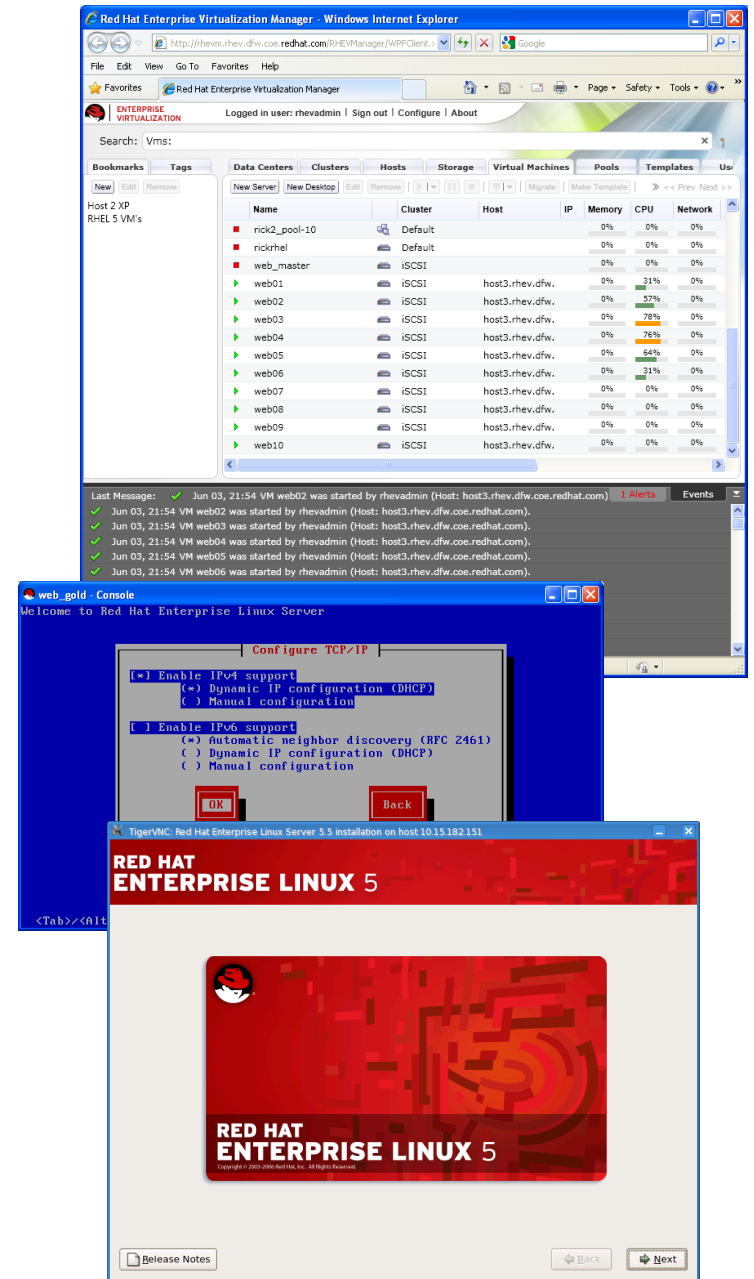


Environment



Application Tier Setup

- Create “master” guest
- “Unconfigure” master guest
- Shut down master guest
- Create template
- Create guests from template
- Start guests



SUMMIT

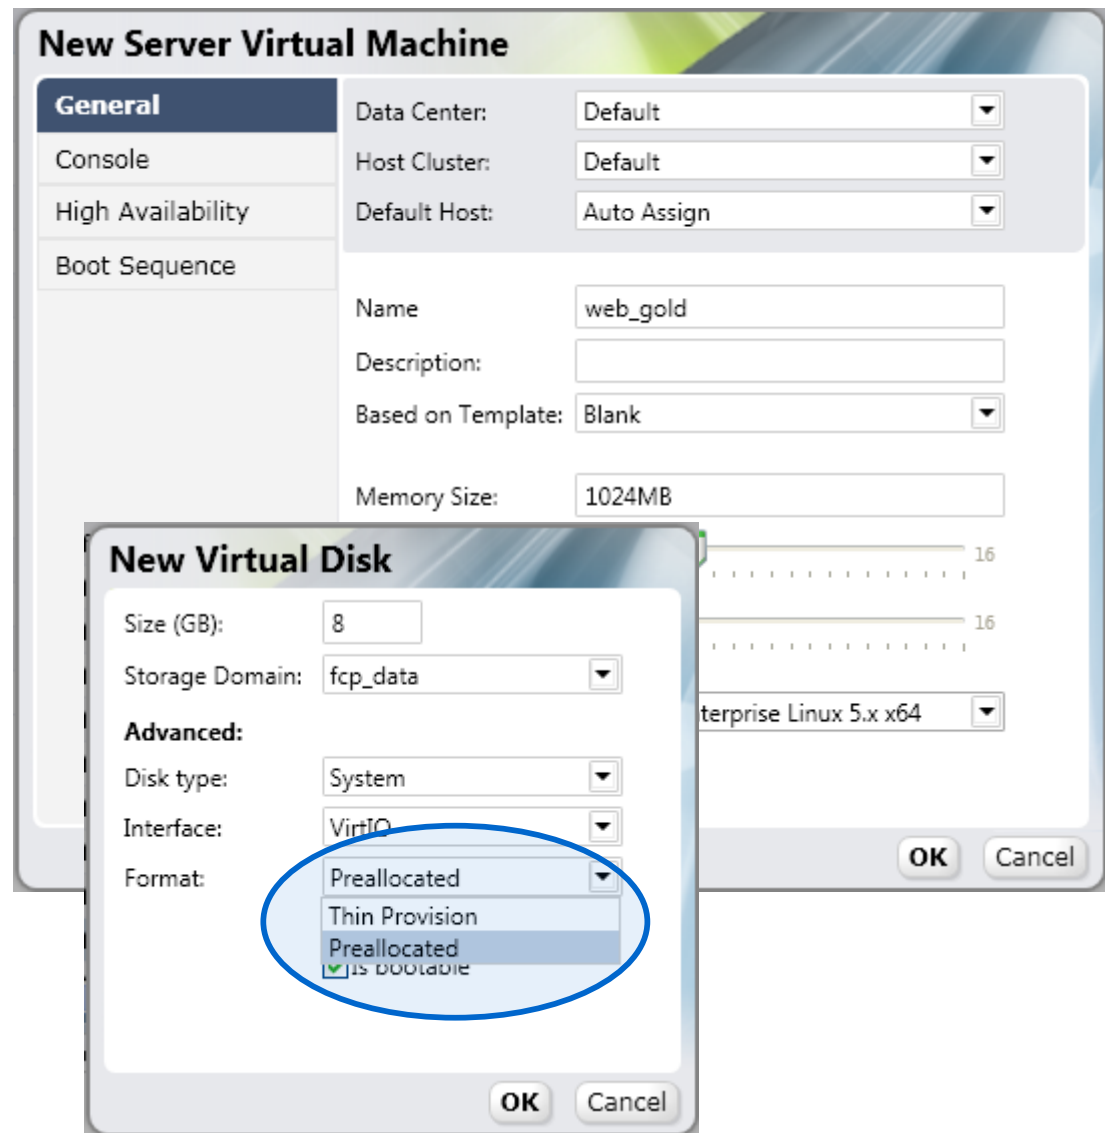
JBoss
WORLD

PRESENTED BY RED HAT



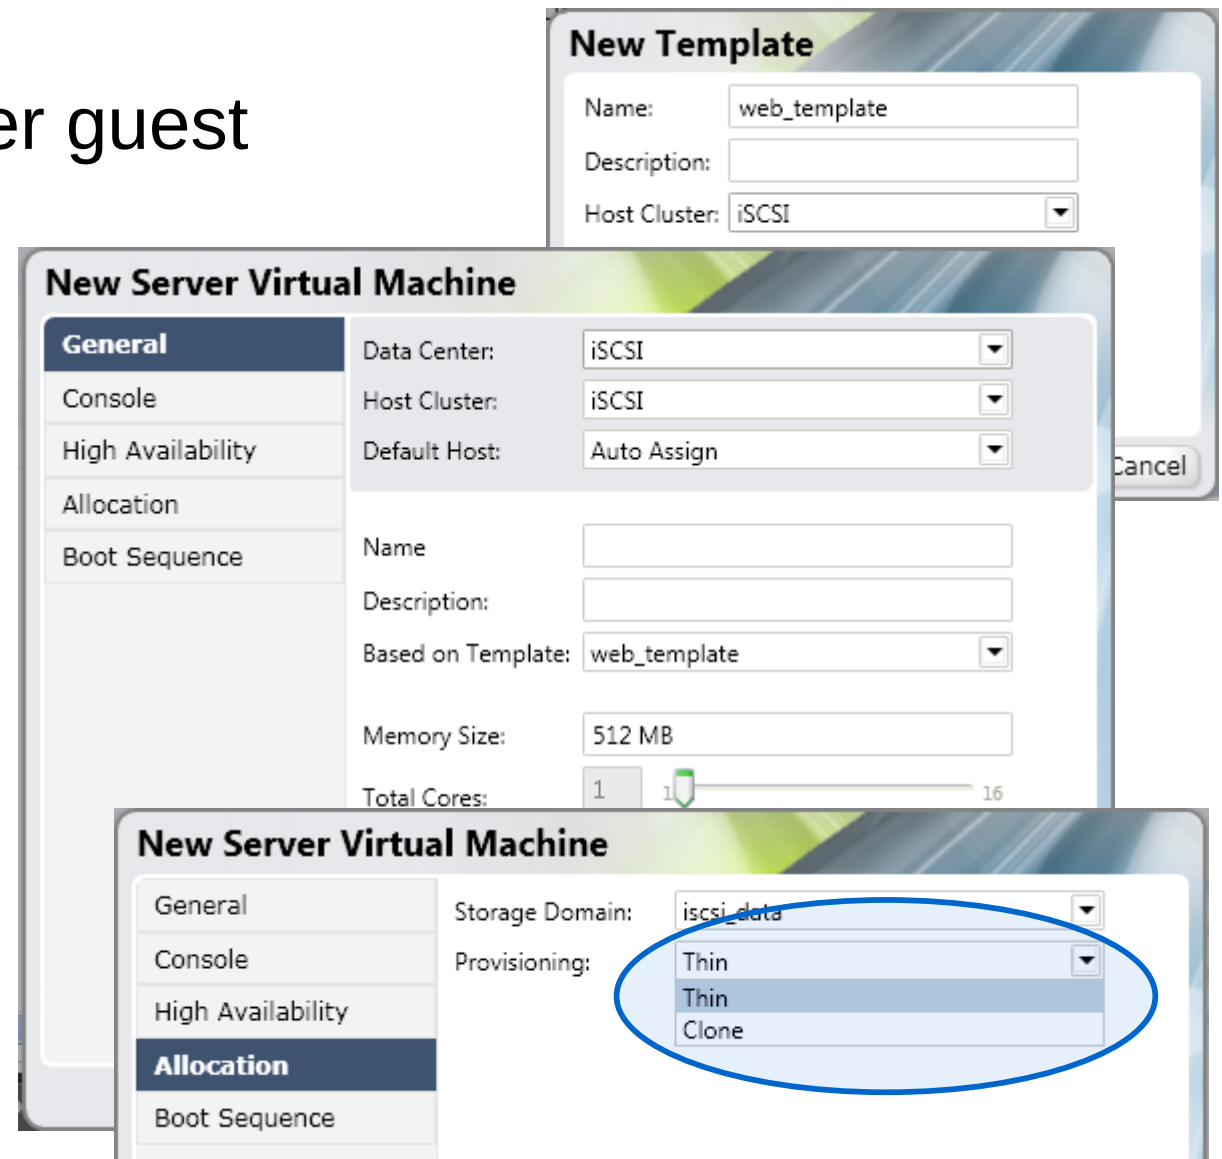
Creating the “Master” Guest

- Define the guest
 - CPUs
 - Memory
 - Disks
 - NICs
- Set up operating system
- Set up applications



Creating the Template and Guests

- “Unconfigure” master guest
 - SSH host keys
 - MAC address(es)
 - Use DHCP
 - “Firstboot” scripts?
 - Shut down
- Create template
- Create guests




Guest Virtual Disks

- Storage types
 - NFS – files (possibly sparse)
 - iSCSI and Fibre-Channel – logical volumes
- “Preallocated” – raw image
- “Thinly provisioned” – qcow2
 - Holds changes relative to “backing file”
 - Changes to backing file will corrupt qcow2 data



Templates

- Complete copy of guest disk
- Independent of guest disk once created
- Derived guests
 - “Cloned” – full copy
 - “Thin” (shown)



Thin Provisioning in Action

Name	Cluster	Host	IP Address	Memory	CPU	Network	Display	Status	Uptime
web_master	iSCSI			0%	0%	0%		Down	
web01	iSCSI	host3.rhev.dfw.		0%	77%	0%	VNC	Powering Up	< 1 min
web02	iSCSI	host3.rhev.dfw.		0%	99%	0%	VNC	Powering Up	< 1 min
web03	iSCSI	host3.rhev.dfw.		0%	95%	0%	VNC	Powering Up	< 1 min
web04	iSCSI	host3.rhev.dfw.		0%	82%	0%	VNC	Powering Up	< 1 min
web05	iSCSI	host3.rhev.dfw.		0%	48%	0%	VNC	Powering Up	< 1 min
web06	iSCSI	host3.rhev.dfw.		0%	30%	0%	VNC	Powering Up	< 1 min
web07	iSCSI	host3.rhev.dfw.		0%	0%	0%	VNC	Powering Up	< 1 min
web08	iSCSI	host3.rhev.dfw.		0%	0%	0%	VNC	Powering Up	< 1 min
web09	iSCSI	host3.rhev.dfw.		0%	0%	0%	VNC	Powering Up	< 1 min
web10	iSCSI	host3.rhev.dfw.		0%	79%	0%	VNC	Up	< 1 min

Name	Type	Cross Data-Center Status	Avail.
fcg_data	Data (Master)	Active	2,734 GB
iscsi_data	Data (Master)	Active	10 GB
nfs_iso	ISO	Active	198 GB

General	Data Center	Virtual Machines	Templates
Size:		39 GB	
Available:		10 GB	
Used:		29 GB	
Over allocation Ratio:		700%	


SUMMIT

**JBoss
WORLD**

PRESENTED BY RED HAT




Environment



Objective: Enlarge the existing LUN, and add a new LUN



The Names of Things



“Attached scsi disk sdc at scsi3, channel 0, id 2, lun 5”...

“Attached scsi disk sdd at scsi4, channel 0, id 8, lun 5”...



The mpath virtual device:

Multipath will notice the matching WWID for sdc and sdd:

```
multipath -ll
```

```
mpath3 (3600508b400105e210000900000490000)  
[size=100 GB][features="0"][hwhandler="0"]  
\_ round-robin 0 [prio=1][active]  
\_ 3:0:2:5 sdc 8:32 [active][ready]  
\_ round-robin 0 [prio=1][enabled]  
\_ 4:0:8:5 sdd 8:48 [active][ready]
```

Be aware: The WWID is the only persistent name for the device !



Add New Storage

You: “Hello, I would like to request a 100GB increase in the size of LUN 5 (WWID=3600508b400105e210000900000490000), and I'd like a second LUN at 300 GB.”

Storage Admin: “I am *Mordac* the preventer of information services...”

<much later>

Storage Admin: “Okay, it's done.”

You: “Ummm, would you please tell me the WWNN of the storage server, so I can identify the paths to it, and the new LUN number?”

Storage Admin: <more abuse...you try to straighten your tie...>

Storage Admin: “5000-1FE1-5009-7080”, “LUN 16”

You: “Thank-you.”



Make the system aware of the new storage

We will add a LUN while the system is running.

– so use the least disruptive scan possible, to probe just the device we want:

```
# find the paths to the storage server whose WWNN=50001FE150097080
```

```
grep -i 50001FE150097080 /sys/class/fc_transport/*/node_name
```

```
/sys/class/fc_transport/target3:0:2/node_name:0x50001fe150097080
```

```
/sys/class/fc_transport/target4:0:8/node_name:0x50001fe150097080
```

```
# probe each path for LUN=16
```

```
echo "0 2 16" > /sys/class/scsi_host/host3/scan
```

```
echo "0 8 16" > /sys/class/scsi_host/host4/scan
```

```
# Check /var/log/messages...
```



A new multipath LUN is added

kernel: SCSI device sdk: 629145600 512-byte hdwr sectors (322123 MB)

kernel: sd 3:0:2:16: Attached scsi disk sdk

kernel: sd 3:0:2:16: Attached scsi generic sg12 type 0

...

kernel: SCSI device sdl: 629145600 512-byte hdwr sectors (322123 MB)

kernel: sd 4:0:8:16: Attached scsi disk sdl

kernel: sd 4:0:8:16: Attached scsi generic sg13 type 0

multipath -ll

mpath5 (3600508b400105e210000c000003c0000) dm-9

[size=300G][features=1 queue_if_no_path][hwhandler=0][rw]

_ round-robin 0 [prio=100][enabled]

_ 3:0:2:16 sdk 8:128 [active][ready]

_ 4:0:8:16 sdl 8:144 [active][ready]

SUMMIT

JBoss
WORLD

PRESENTED BY RED HAT



Increase the size of an existing LUN

```
echo 1 > /sys/block/sdc/device/rescan  
echo 1 > /sys/block/sdd/device/rescan
```

```
kernel: sdc: detected capacity change from 107374182400 to 214748364800  
kernel: SCSI device sdc: 419430400 512-byte hdwr sectors (214748 MB)
```

...

```
kernel: sdd: detected capacity change from 107374182400 to 214748364800  
kernel: SCSI device sdd: 419430400 512-byte hdwr sectors (214748 MB)
```

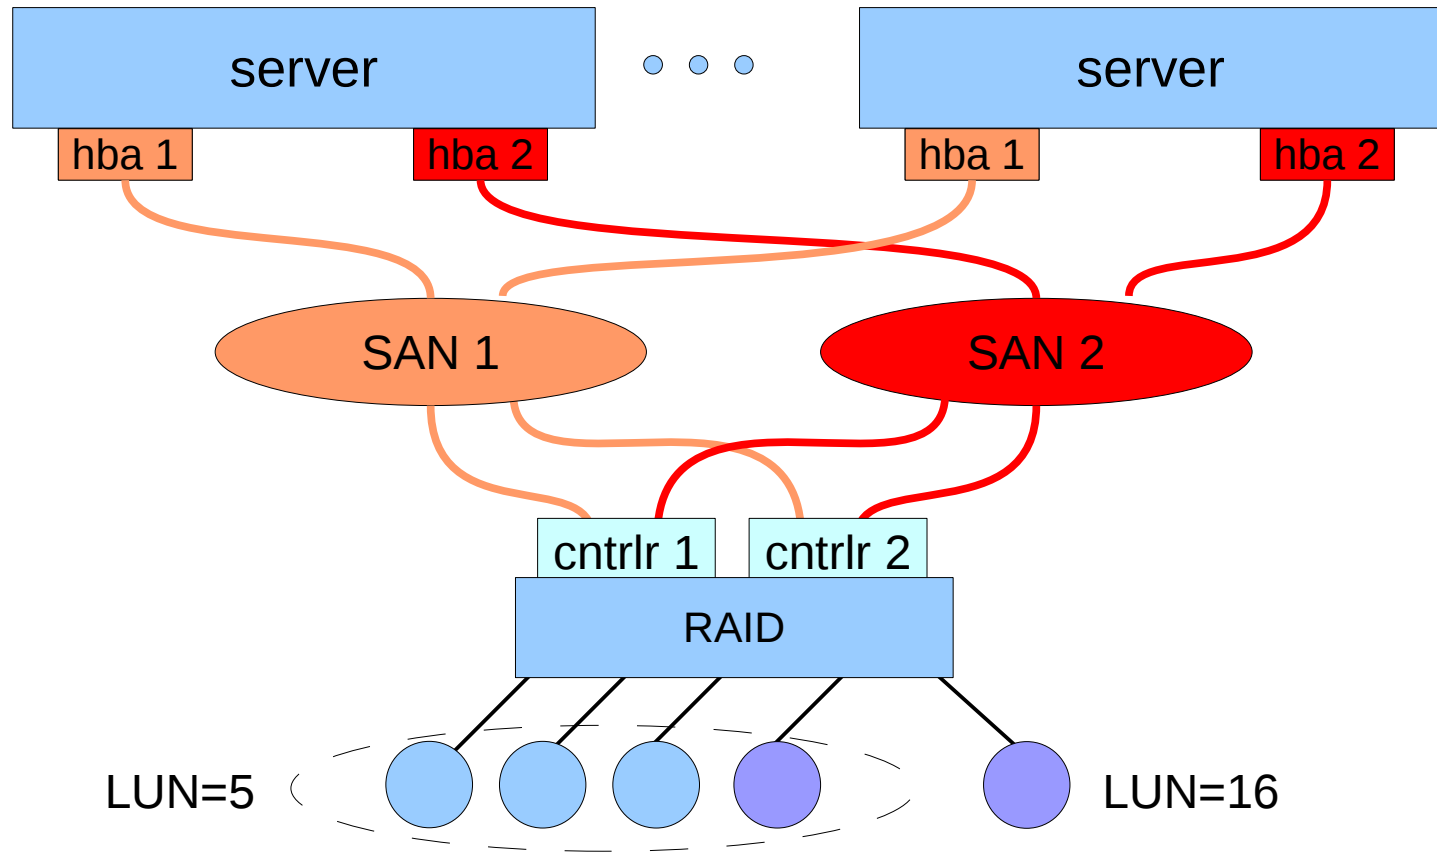
```
multipathd -k"resize map mpath4"
```

```
multipath -ll
```

```
mpath3 (3600508b400105e210000900000490000)  
[size=200 GB][features="0"][hw_handler="0"]  
\_ round-robin 0 [prio=1][active]  
  \_ 3:0:2:5 sdc 8:32 [active][ready]  
\_ round-robin 0 [prio=1][enabled]  
  \_ 4:0:8:5 sdd 8:48 [active][ready]
```



We've doubled the size of LUN 5 and added LUN 16




Application Tier Upgrade

- Create master guest snapshot
- Add additional virtual disk to master guest
 - Expand storage pool
- Boot master guest
- Set up new virtual disk
 - fdisk, pvcreate, vgextend, lvextend, resize2fs
- Upgrade master guest
 - Operating system, applications, etc.
- Test
 - Revert to snapshot if necessary
- Unconfigure
 - SSH host keys
- Shutdown
- Create **new** template
- Create **new** guests from new template
- Transition from old guests to new guests

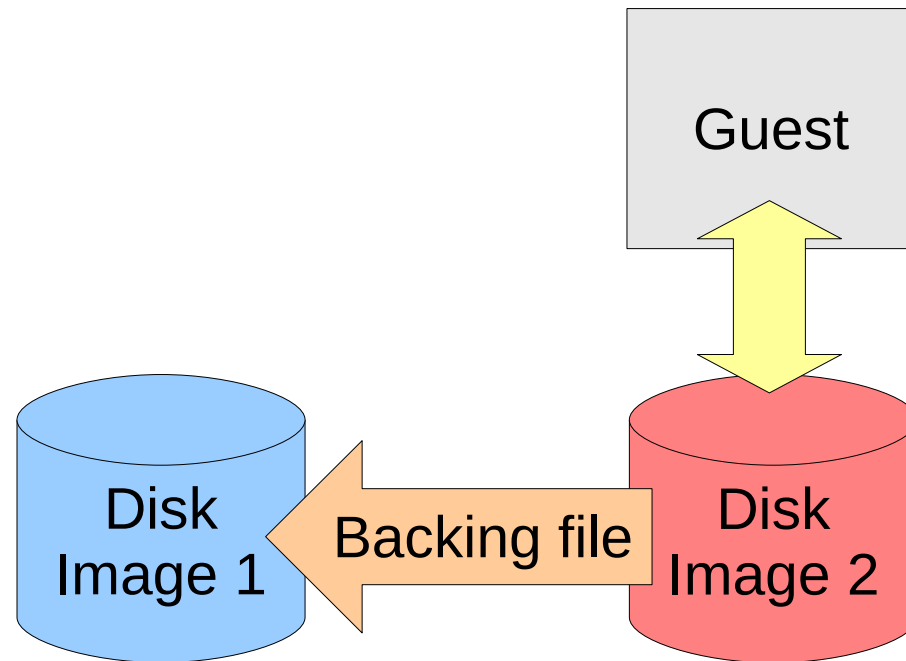


Creating a Snapshot

Before



After



Original "base image" is now snapshot


SUMMIT

**JBoss
WORLD**

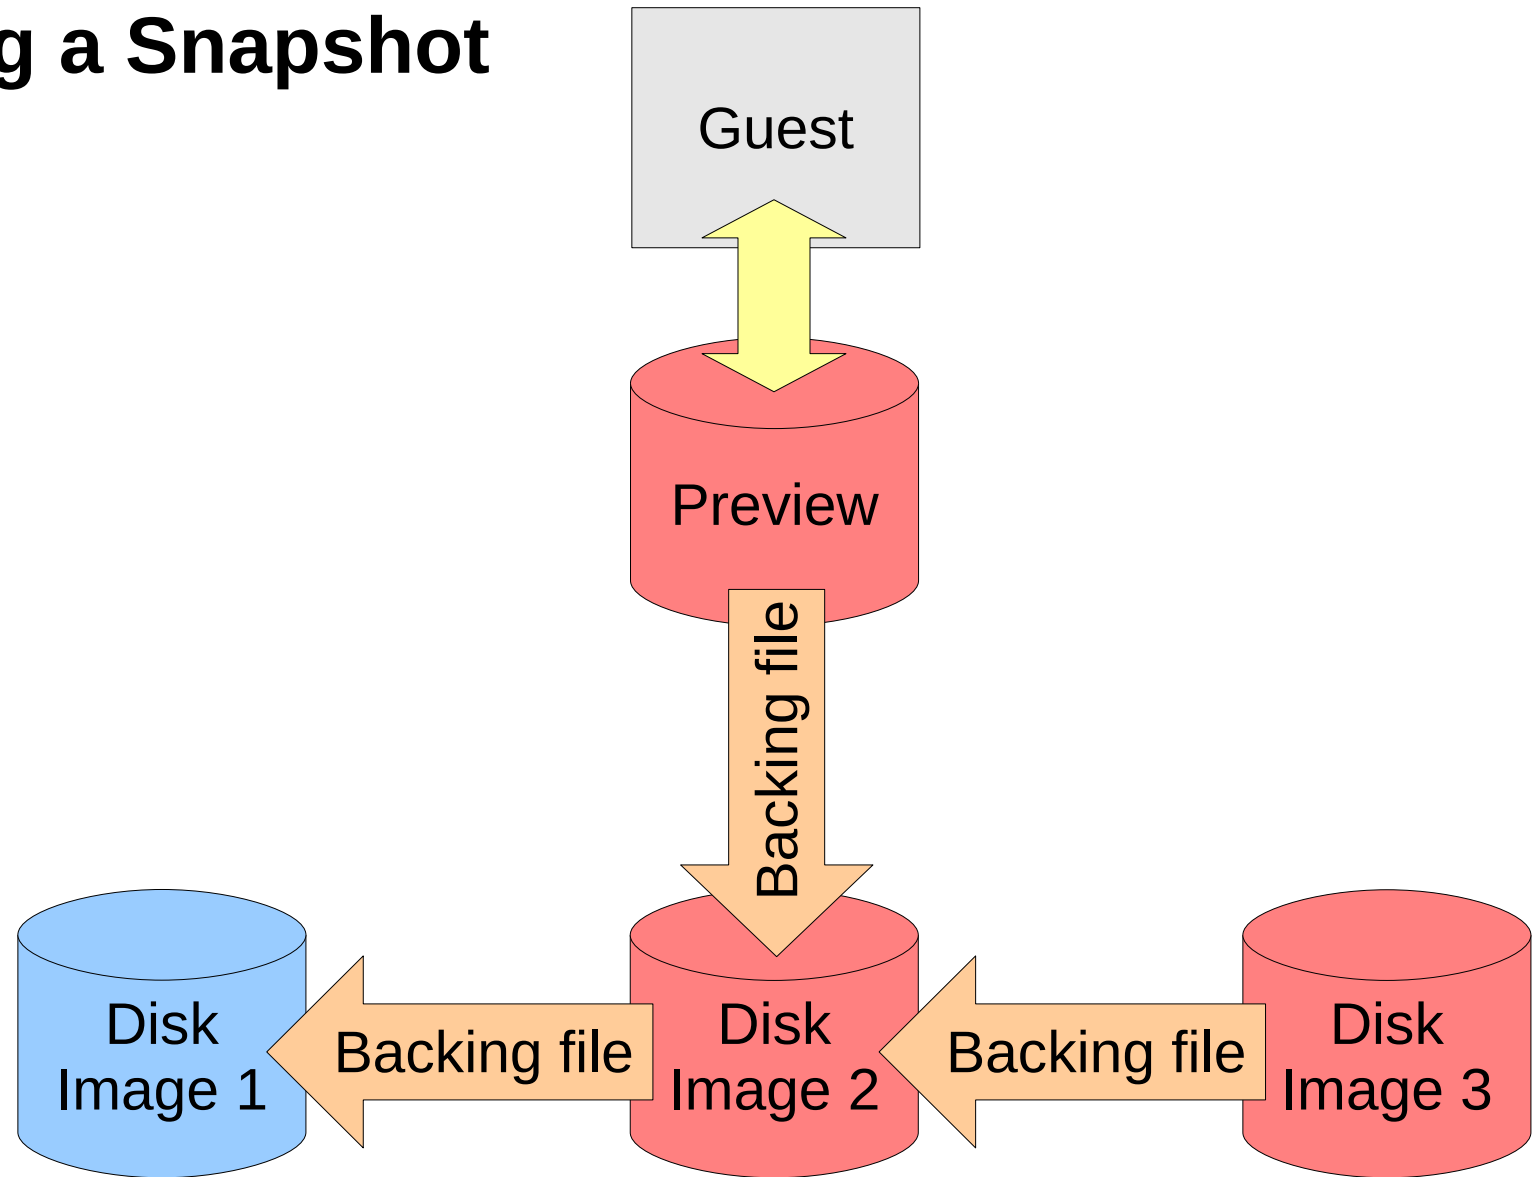
PRESENTED BY RED HAT



Creating Another Snapshot



Previewing a Snapshot



SUMMIT

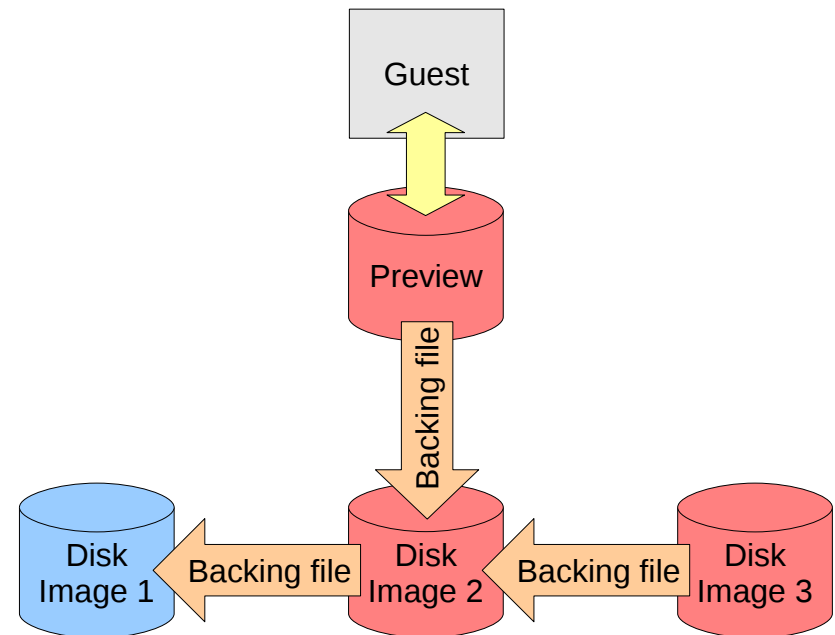
JBoss
WORLD

PRESENTED BY RED HAT

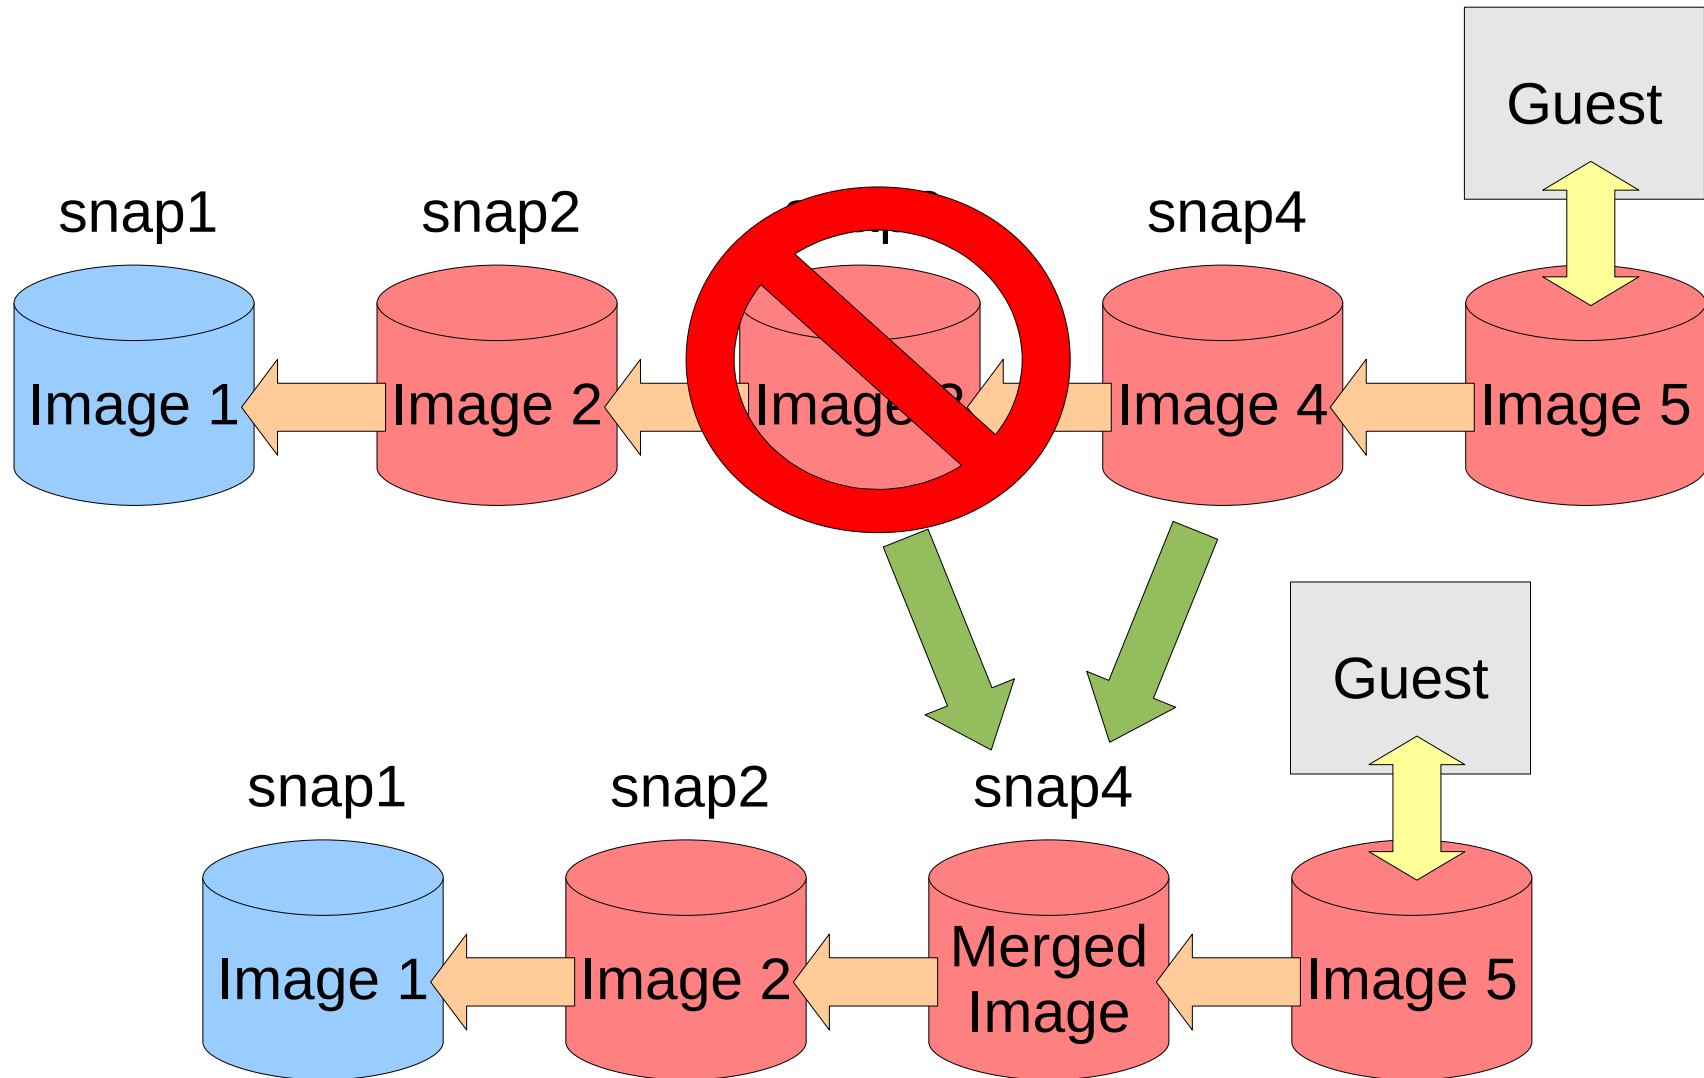


Snapshot Undo/Commit

- Available when snapshot is previewed
- “Undo”
 - Preview is discarded
- “Commit”
 - Subsequent snapshots discarded
 - Guest “attached” to preview



Deleting a Snapshot



Expanding the Storage Pool

New Virtual Disk

Size (GB):

Advanced:

Disk type:

Interface:

Format:

Thin Provision

Preallocated

iS bootable

Edit Domain

Name:

Domain function: Data ISO Export

Storage type: iSCSI

UUID: 8464c3fa-287d-4df7-9396-493ed08c758c

Use host:

Discovered LUNs

	LUN ID	UUID	Multipathing	Dev. Size
<input type="checkbox"/>	+ 11ET_00010003		1	13GB

Selected LUNs

	LUN ID	UUID	Multipathing	Dev. Size
<input checked="" type="checkbox"/>	+ 11ET_00010002		1	40GB

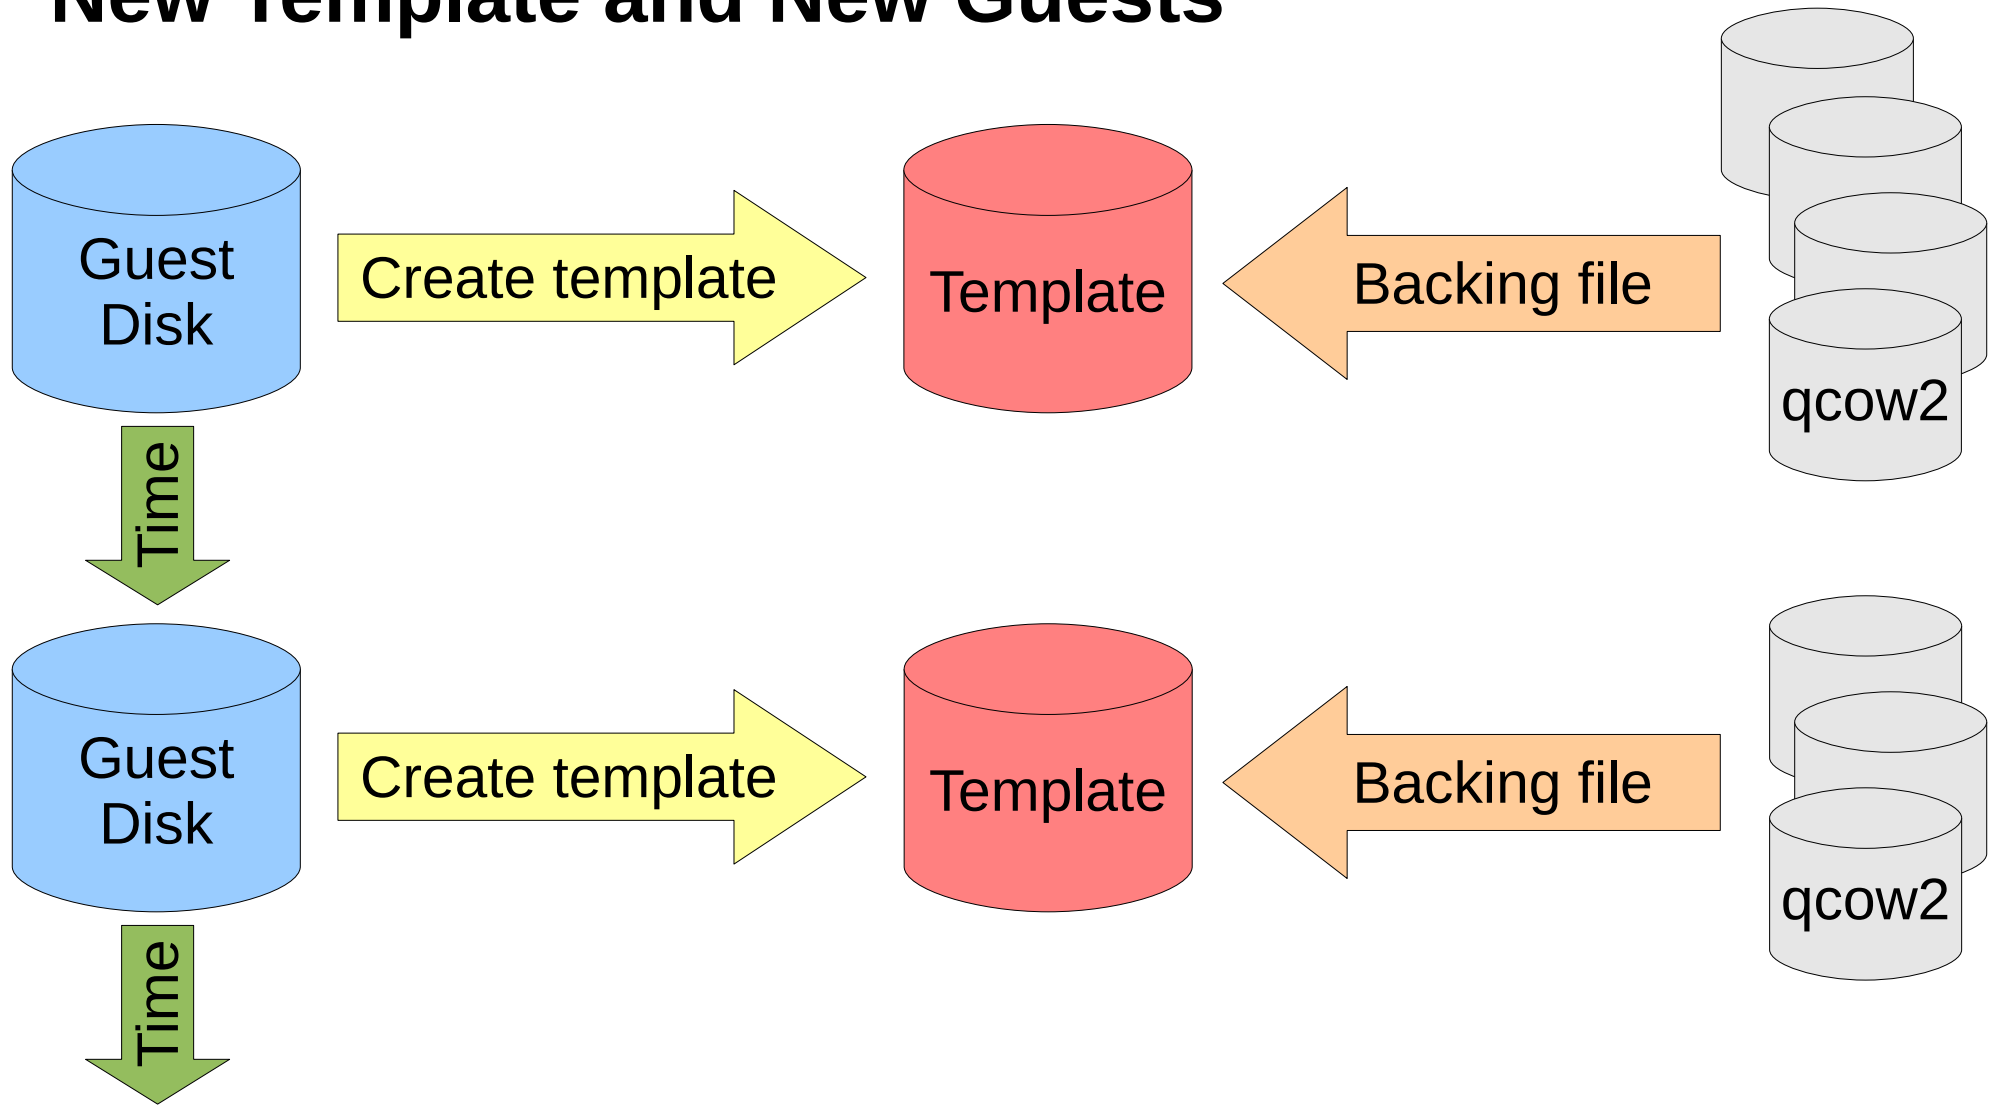
Operation Canceled

web_master:

- Cannot add Virtual Machine Disk. Low disk space on relevant Storage Domain.



New Template and New Guests



Summary

- **Red Hat Enterprise Linux**
 - LUN(s)
 - Expand
 - Add
 - Rescanning
 - Multipath
- **Red Hat Enterprise Virtualization**
 - Guest disks
 - Pre-allocated
 - Thin
 - Templates & guests
 - Cloned
 - Thin
 - Snapshots
 - Storage pools



For More Information – <http://www.redhat.com/docs>

- Red Hat Enterprise Linux
 - Red Hat Enterprise Linux 5
 - “Online Storage Reconfiguration Guide”
 - Red Hat Enterprise Linux 6 Beta
 - “Storage Administration Guide”
 - Chapter 22 - “Online Storage Management”
- Red Hat Enterprise Virtualization
 - “Administration Guide”
 - Chapter 4 - “Managing Storage”
 - Chapter 6 - “Managing Virtual Resources”
 - 6.1 - VM Storage
 - 6.2 - Creating VMs
 - 6.5 - Snapshots
 - 6.6 - Templates



Please Turn In Your Evaluations

Thank



You!

SUMMIT

**JBoss
WORLD**

PRESENTED BY RED HAT



FOLLOW US ON TWITTER

www.twitter.com/redhatsummit

TWEET ABOUT IT

[#summitjbw](https://twitter.com/summitjbw)

READ THE BLOG

<http://summitblog.redhat.com/>

SUMMIT

**JBoss
WORLD**

PRESENTED BY RED HAT

