



Mike Christie  
Senior Software Engineer, Red Hat

# Red Hat Enterprise Linux iSCSI Update

- Status of iSCSI in Red Hat Enterprise Linux.
- New Feature Examples.
- RHEL 5.x and RHEL 6.
- Questions.



# Status of iSCSI in Red Hat Enterprise Linux

- Red Hat Enterprise Linux 3.
  - Initiator based on Cisco/sourceforge linux-iscsi-3 code base.
  - Frozen.
- Red Hat Enterprise Linux 4.
  - Initiator based on Cisco/sourceforge linux-iscsi-4 code base.
  - Semi-Feature Frozen.
    - IPv6 and configuration enhancements are the only planned features.

# Status of iSCSI in Red Hat Enterprise Linux Continued

- Red Hat Enterprise Linux 5.
  - Initiator.
    - Based on upstream open-iscsi project.
    - New Features in 5.1 and 5.2.
      - iSCSI root/boot.
      - Host based multipath using dm-multipath.
  - iSCSI Target.
    - Tech Preview in 5.1 and 5.2.
  - iSNS client and server.
    - Tech Preview in 5.2.

## iSCSI Root/Boot in RHEL 5.2

- Support for iSCSI root using netboot added in 5.1.
  - Must have local /boot or use network boot loader like pxelinux.
  - Several bugs including CDROM based install network setup and shutdown hangs fixed in 5.2.
- Support for iBFT and the Open Firmware iSCSI Boot in 5.2.
  - Mini iSCSI initiator lives on NIC.
  - Anaconda (installer) supports installation to a iSCSI disk.
  - Limitations:
    - CHAP only supported on boot.
    - Does not pick up iBFT network settings.
- All the kinks out in 5.3

# iSCSI Boot Installation is Simple

Click the “Advanced storage configuration” button.



**RED HAT  
ENTERPRISE LINUX 5**

Installation requires partitioning of your hard drive. By default, a partitioning layout is chosen which is reasonable for most users. You can either choose to use this or create your own.

Remove linux partitions on selected drives and create default layout. ▾

Select the drive(s) to use for this installation.

<input checked="" type="checkbox"/>	sda	2047 MB	VMware, VMware Virtual S
-------------------------------------	-----	---------	--------------------------

+ Advanced storage configuration

Review and modify partitioning layout

[Release Notes](#) [Back](#) [Next](#)

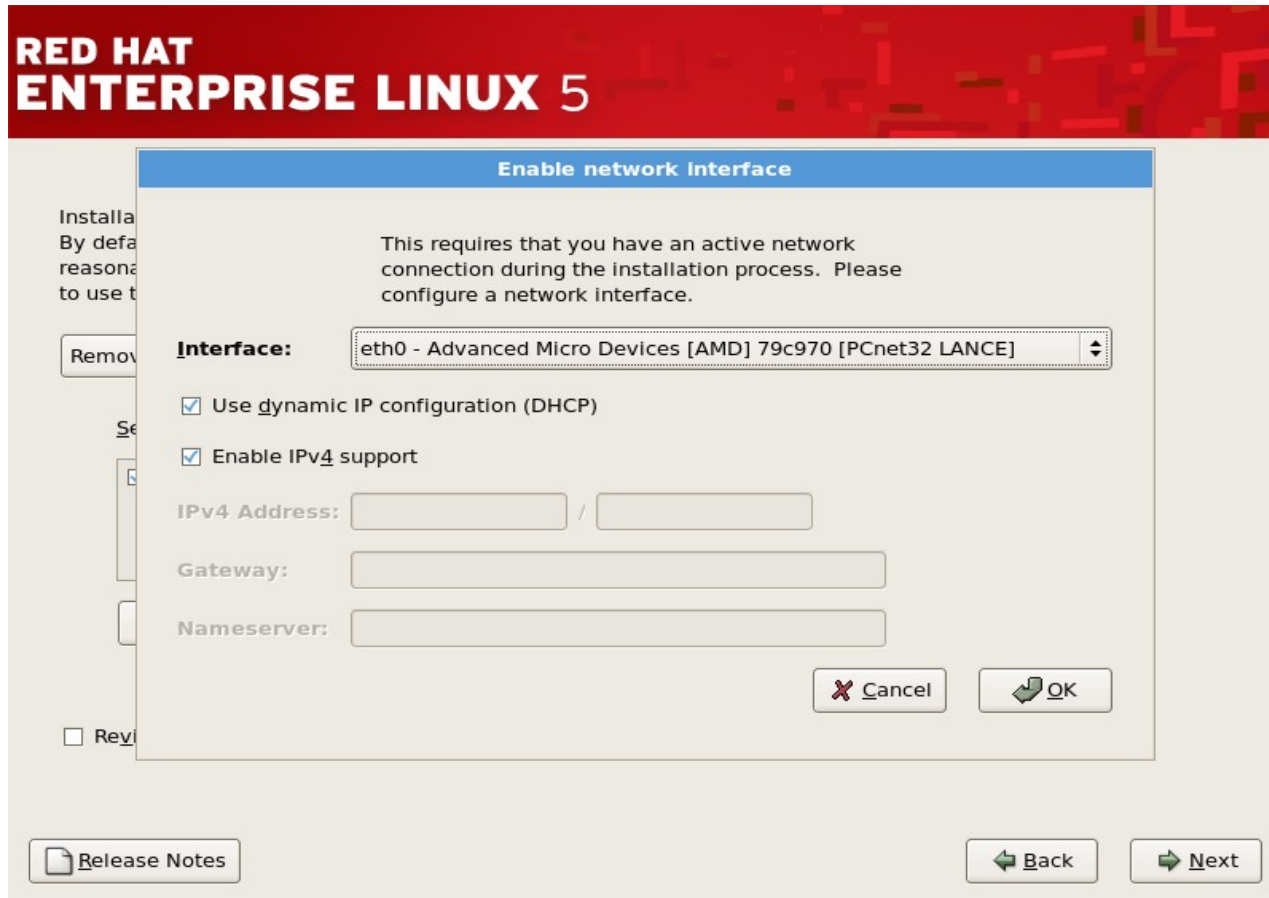
# iSCSI Boot Installation Continued

Select “Add iSCSI target”, and click “Add drive”.



# iSCSI Boot Installation Continued

Enter network info.



**RED HAT  
ENTERPRISE LINUX 5**

**Enable network interface**

This requires that you have an active network connection during the installation process. Please configure a network interface.

**Interface:** eth0 - Advanced Micro Devices [AMD] 79c970 [PCnet32 LANCE]

Use dynamic IP configuration (DHCP)

Enable IPv4 support

IPv4 Address:  /

Gateway:

Nameserver:

Review

# iSCSI Boot Installation Continued

Enter iSCSI information.

**RED HAT  
ENTERPRISE LINUX 5**

Installation  
By default,  
reasonable  
to use this

Remove I

Select

### Configure iSCSI Parameters

To use iSCSI disks, you must provide the address of your iSCSI target and the iSCSI initiator name you've configured for your host.

**Target IP Address:**

**iSCSI Initiator Name:**

Review and modify partitioning layout

# Installation was Simple, but Optimizing the Root Session will Make You Cry

- Timers not optimized for root setups.
- If using iSCSI with dm-multipath you want short iSCSI timers and lots of dm-multipath retries.
  - *node.session.timeo.replacement\_timeout = 5*
  - *node.conn[0].timeo.noop\_out\_timeout = 5*
  - *node.conn[0].timeo.noop\_out\_interval = 5*
  - *no\_path\_retry = queue*
- If using iSCSI without multipath you want long iSCSI timers.
  - *node.session.timeo.replacement\_timeout = 172800*
  - *node.conn[0].timeo.noop\_out\_timeout = 0 # turns off*
  - *node.conn[0].timeo.noop\_out\_interval = 0 # turns off*

# How to Get the New Settings Picked Up

- The iSCSI tools store config settings in a data base (DB).
- When the discovery command is run, a portal record is created using the */etc/iscsi/iscsid.conf* values.

```
# iscsiadm -m discovery -t st -p 10.15.85.19  
10.15.84.19:3260,2 iqn.1992-08.com.redhat  
10.15.85.19:3260,3 iqn.1992-08.com.redhat
```

```
# iscsiadm -m node  
10.15.85.19:3260,3 iqn.1992-08.com.redhat  
10.15.84.19:3260,2 iqn.1992-08.com.redhat
```

# Managing a Node/Portal Record

Use *iscsiadm* in node mode.

- “-o show” displays settings.
- “-o update” changes a setting.
- “-o delete” removes a record.

```
# iscsiadm -m node -o show -T iqn.1992-08.com.redhat -p  
10.15.84.19:3260 | grep  
node.session.timeo.replacement_timeout  
node.session.timeo.replacement_timeout = 120
```

```
# iscsiadm -m node -o update -T iqn.1992-08.com.redhat -p  
10.15.84.19:3260 -n node.session.timeo.replacement_timeout  
-v 5
```

# iSCSI Setting Update Pain is Almost Over

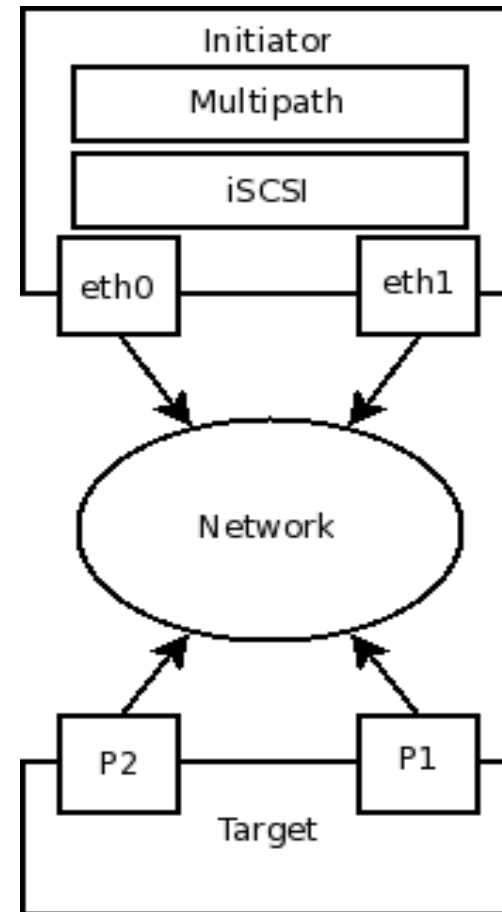
- Updating record does not affect a running session.
- When *iscsiadm* logs into a portal it will read the record info.
  - For the root session you must reboot.
  - For all other sessions you need to restart the session by logging the session out then logging back in.

# iSCSI dm-multipath Enhancements

- Root on a dm-multipath device consisting of iSCSI paths.
  - Qla4xxx supported since 5.1.
  - Software iSCSI support added in 5.2.
    - One path is used during the early boot stage.
      - Will be fixed in RHEL 5.3.
- Improved performance and availability using host based multipath.
  - Added ability to bind iSCSI sessions to specific network devices.

# iSCSI iface Binding

- Used to have to use trunking, subnets, etc.
- Iface binding allows host based multipathing at the block layer.
  - Create 4 paths/sessions.
    - session1 through eth0 to P1 and session2 through eth0 to P2
    - Session3 through eth1 to P1 , and session4 through eth1 to session P2



# Setup iSCSI Iface

- Create iSCSI iface binding for eth0.
- Can bind by network device name (eth0, alias, vlan name) or MAC address.

```
# ifconfig | grep eth
```

```
eth0      Link encap:Ethernet  HWaddr 00:80:3F:2B:79:AE
```

```
eth1      Link encap:Ethernet  HWaddr 00:80:3F:2B:79:AF
```

```
# iscsiadm -m iface -I iscsi-eth0 -o new
```

```
New interface iscsi-eth0 added
```

```
# iscsiadm -m iface -I iscsi-eth0 -o update -n  
iface.net_ifacename -v eth0
```

# Bind Iface to Existing Portal/Target

- Can be done using the “*-o new*” command in node mode if you know the target and portal information.

```
# iscsiadm -m node -T iqn.1992-08.com.redhat -I iscsi-eth0 -o  
new
```

**New iSCSI node [tcp:**

```
[hw=default,ip=,net_if=eth0,iscsi_if=iscsi-eth0]  
10.15.85.19,3260,3 iqn.1992-08.com.redhat added
```

**New iSCSI node [tcp:**

```
[hw=default,ip=,net_if=eth0,iscsi_if=iscsi-eth0]  
10.15.84.19,3260,2 iqn.1992-08.com.redhat] added
```

# Bind Iface During Discovery

- Pass interfaces to bind to portals that are found during discovery.

```
# iscsiadm -m discovery -t st -p 10.15.85.19 -I iscsi-eth0 -I  
iscsi-eth1 -P 1
```

```
Target: iqn.1992-08.com.redhat
```

```
Portal: 10.15.85.19:3260,3
```

```
Iface Name: iscsi-eth1
```

```
Iface Name: iscsi-eth0
```

```
Portal: 10.15.84.19:3260,2
```

```
Iface Name: iscsi-eth0
```

```
Iface Name: iscsi-eth1
```

# Optimizing the Session

- Block Layer Settings.
  - Experiment with different IO Schedulers.
    - *echo noop > /sys/block/sdX/queue/scheduler*
  - Experiment with different IO sizes.
    - *echo N > /sys/block/sdX/queue/max\_sectors\_kb*
    - Values from 64 to 512 work best, but this depend on the target and workload.
- iSCSI Settings.
  - *node.conn[0].iscsi.MaxRecvDataSegmentLength = N*
    - Values from 65536 (64 kb) to 524288 (512 kb) work best, but it also depends on the work load (target will negotiate for its optimal value).
    - Try aligning with the block layer's *max\_sectors\_kb* limit.

# Logging into Bound Records

- *service iscsi start*
- *iscsiadm -m node -T target -p ip:port -l iface -l*
- *iscsiadm -m node -l iface -l*

```
# iscsiadm -m node -I iscsi-eth1 -l
```

```
Logging in to [iface: iscsi-eth1, target:  
iqn.1992-08.com.redhat, portal: 10.15.85.19,3260]
```

```
Logging in to [iface: iscsi-eth1, target:  
iqn.1992-08.com.redhat, portal: 10.15.84.19,3260]
```

```
Login to [iface: iscsi-eth1, target: iqn.1992-08.com.redhat,  
portal: 10.15.85.19,3260]: successful
```

```
Login to [iface: iscsi-eth1, target: iqn.1992-08.com.redhat,  
portal: 10.15.84.19,3260]: successful
```

# Display Session Information

- *iscsiadm -m session -P N*
  - *N can be from 0 to 2 and will show more information at higher values.*

```
# iscsiadm -m session -P 3
```

```
iSCSI Transport Class version 2.0-870
```

```
iscsiadm version 2.0-869
```

```
Target: iqn.1992-08.com.redhat
```

```
Current Portal: 10.15.85.19:3260,3
```

```
....
```

```
iSCSI Connection State: LOGGED IN
```

```
....
```

```
scsi6 Channel 00 Id 0 Lun: 0
```

```
Attached scsi disk sdd
```

```
State: running
```

# SCSI Target Support

- Added to RHEL 5.1 as a tech preview.
  - Package named scsi-target-utils.
- Based on upstream TGT project.
  - iSCSI engine from iSCSI Enterprise Target (IET).
  - Can support multiple types of targets.
    - FCP – qla2xxx, lpfc.
    - iSCSI and iSER – Software iSCSI and iSER, qla4xxx.
    - VSCSI – IBM target based on SRP used in PPC boxes.
  - Controversial design.
    - SCSI processing engine is in userspace, but data paths are in the kernel.

# SCSI Target Support Continued

- For the initial RHEL release we are concentrating on software iSCSI and iSER.
  - Will be solid, but few features.
    - By reusing IET code, its feature set is short, but it's iSCSI processing engine has received a lot of testing.
    - Can use any block device or normal File for Logical Unit backing, and supports block layer multipath.
    - ERL0, MaxOutstandingR2T=1, No MC/s.

# SCSI Target Setup

- tgtadm - SCSI target administration utility.
  - Can manage any target type.
  - Setup is not persistent.
    - What do people prefer?
      - IET config file or Opensolaris DB.

```
# tgtadm --lld iscsi --mode target --op new --tid 1  
--targetname iqn.2001-04.redhat
```

```
# tgtadm --lld iscsi --mode logicalunit --op new --tid 1 --lun  
1 --backing-store /var/lib/tgt/lun1
```

# SCSI Target Information

```
# tgtadm --lld iscsi --op show --mode target
```

```
Target 1: iqn.2001-04.redhat
```

```
System information:
```

```
Driver: iscsi
```

```
State: ready
```

```
LUN information:
```

```
.....
```

```
LUN: 1
```

```
Type: disk
```

```
SCSI ID: deadbeaf1:1
```

```
SCSI SN: beaf11
```

```
Size: 8389 MB
```

```
Online: Yes
```

```
Backing store: /var/lib/lun1
```

# SCSI Session Information

```
# tgtadm --lld iscsi --op show --mode target
```

```
Target 1: iqn.2001-04.redhat
```

```
System information:
```

```
Driver: iscsi
```

```
State: ready
```

```
I_T nexus information:
```

```
I_T nexus: 1
```

```
Initiator: iqn.2005-06.com.redhat:max
```

```
Connection: 0
```

```
.....
```

# iSNS Support

- Added to RHEL 5.2 as a tech preview.
  - Package named isns-utils.
  - Includes iSNS server and command line client tool.
- Server and client tool based on upstream open-isns project.
- iscsi-initiator-utils in 5.2 includes iSNS client support.
  - Can only use iSNS for discovery.
    - Will support more iSNS features when it uses the isns-utils library.

# iSNS Server Setup

- Simple setups using the DefaultDomain should be easy to configure.
  - Just install isns-utils rpm and start the service by running:
    - *service isnsd start*
  - Point targets and initiators to the server.
- For security and domain management see the setup guide under `/usr/share/docs/isns-utils-*/README.redhat.setup`

# iSNS Initiator Setup

- 1) In */etc/iscsi/iscsid.conf*, set the *isns.address* to the iSNS server's IP address.
- 2) Start *iscsi* service with “*service iscsi start*”.
- 3) Run *iscsiadm* discovery command with type as “*isns*”

```
# iscsiadm -m discovery -t isns
```

```
10.15.84.19:3260,2 iqn.1992-08.com.redhat
```

# Future Features for Red Hat Enterprise Linux 5/6

- iSCSI Library.
  - Based on SNIA iSCSI Management Interface (IMA).
- Better hardware support.
  - iscsiadm and anaconda support for Qlogic qla4xxx driver.
    - No more islci and Qlogic ioctl module hassles!
  - Broadcom iSCSI offload card support.
  - Intel IOAT integration.
- Extend iSCSI ifaces to support virtualization needs.
  - Create ifaces based on different initiator names.
- iSCSI over 10 gig ethernet performance improvements?

# iSCSI Management Improvements GUI – conga integration.

The screenshot displays the Red Hat Cluster and Storage Systems GUI. The top navigation bar includes the Red Hat logo, the text "CLUSTER AND STORAGE SYSTEMS", and tabs for "homebase", "cluster", and "storage". The "storage" tab is active. The main content area shows the URL "tng3-4.lab.msp.redhat.com" and a "Volume Group new\_vg" configuration page. A sidebar on the left lists navigation options: "System List", "Hard Drives", "Partition Tables", "Software RAID", "Volume Group", "New Volume Group", and "new\_vg". The "Graphical View" is checked, showing two horizontal bars representing physical volumes: a blue bar for "Lvg1/Extents" and a red bar for "Phys1/Extents". Below the graphical view is a table of properties for the "Volume Group 'new\_vg'":

Property	Value
Volume Group Name	new_vg
Extent Size	4.0 MB
Total Extents	13164
Free Extents	13164
Size	51.42 GB
Used Extents	0
Maximum Physical Volumes	256
Maximum Logical Volumes	256
Attributes	vg=lv
Clustered	false
UUID	jkQJGzZKXO-QpMIO-0118-htwQ-wWgd#C5C02

At the bottom of the configuration panel are buttons for "Remove", "Add Physical Volumes", "New Logical Volume", "Reset", and "Apply". A footer at the bottom of the page reads: "The Conga Cluster and Storage Management System is Copyright © 2000-2006 by Red Hat, Inc. Distributed under the GNU GPL license."

# Where to Find More Info

- iscsi-initiator-utils
  - man iscsiadm
  - /usr/share/docs/iscsi-initiator-utils-\*
  - <http://www.open-iscsi.org>
  - <http://people.redhat.com/mchristi/iscsi/RHEL5/readme>
- scsi-target-utils
  - man tgtadm
  - /usr/share/docs/scsi-target-utils-\*
  - <http://stgt.berlios.de/>
- isns-utils
  - man isnsadm
  - /user/share/docs/isns-utils-\*
  - <http://oss.oracle.com/~okir/open-isns/>

# Where to Find More Info Continued

- Red Hat Kbase
  - <http://kbase.redhat.com/faq>

# Questions, Feature Requests, Suggestions?

- Make a bugzilla request at <https://bugzilla.redhat.com>.
- Send email to Mike Christie <[mchristi@redhat.com](mailto:mchristi@redhat.com)>

It became very clear to me sitting out there today, that every decision I've ever made, in my entire life, has been wrong.

- George Costanza

- Iscsiadm was made by a Bad Storage Admin (Me).
- Common user mistakes due to iscsiadm being written by me.
  - If you change a value in /etc/iscsi/iscsi.conf, it does not get picked up until you rerun discovery or update a record manually.
  - If you run discovery the node records are overwritten with values from iscsi.conf so you lose all your previous config data.
- What can we do to make the user experience better?
- What are your priorities?