Red Hat Reference Architecture Series

Red Hat Enterprise Virtualization (RHEV) 3.1 Backup and Recovery

Using Symantec Netbackup 7.5 (7.5.0.5)

Balaji Jayavelu

Principal Software Engineer

Version 2.0

March 2013

redhat.

1801 Varsity Drive™
Raleigh NC 27606-2072 USA
Phone: +1 919 754 3700
Phone: 888 733 4281
Fax: +1 919 754 3701
PO Box 13588
Research Triangle Park NC 27709 USA

The GPG fingerprint of the security@redhat.com key is:
CA 20 86 86 2B D6 9D FC 65 F6 EC C4 21 91 80 CD DB 42 A6 0E

Send feedback to refarch-feedback@redhat.com

Table of Contents

1 Executive Summary

Virtualization is being rapidly adopted in all business segments big, medium or small. Initially, poor stability and reliability of Virtualization limited its usage to low impact applications. However, over the years the technology has evolved and matured, rendering Virtualization suitable for mission critical applications.

Despite low cost, efficient usage of resources and faster provisioning, Virtualization, like any other technology brings challenges that can erode its benefits and leave the infrastructure vulnerable. As organizations embrace virtualization, it is increasingly important to protect data and ensure business continuity is not compromised. By incorporating and validating backup and restore capabilities, IT departments can work to mitigate this risk.
Backup and recovery consists of multiple layers of data protection to include:

- Application level- Typically consists of using a client or agent which has the ability to quiesce a particular application.
- Operating system level- Includes the ability to backup specific operating system details.
- File level-refers to backing up and restoring a file or set of files as opposed to an entire operating system, application or disk drive.
- Snapshot- a point in time, disk based copy of a client volume(s).
- Bare Metal- provides the ability to take a backup of a machine and restore to a previous state without the need for a previously installed software.

Deciding upon and implementing the proper policies is just as important as determining the data content that needs backing up. This may include:

- Type of backup- differential, incremental or full
- SLA- recovery point objective (RPO) or recovery time objective (RTO)
- Government/Business regulations- HIPAA, SOX or PCI

The goal of this paper is not to cover all the backup and recovery scenarios. However it is meant to provide guidance for proper data backup and recovery for customers utilizing Red Hat Enterprise Virtualization 3.1 to include the following:

- Virtual Machine
    - File level backup and restore
    - Virtual Machine recovery
- Red Hat Enterprise Virtualization Manager (*henceforth referred to as* RHEV-M *in this paper)*
    - Application backup and recovery
    - Database backup and recovery
    - Full system recovery (Using Bare Metal Restore)

2 Red Hat Enterprise Virtualization (RHEV)

Red Hat Enterprise Virtualization provides IT departments with the tools to meet the challenges of managing complex environments. Red Hat's state-of-the-art virtualization platform enables administrators to reduce the cost and complexity of large deployments. Red Hat Enterprise Virtualization platform provides:

- High availability to quickly configure virtual machines for fault tolerance.
- Live migration to move virtual machines between physical hosts without interruption.
- System scheduler to create policies to dynamically balance compute resources.
- Power saver to create policies to conserve power and cooling costs.
- Image manager to create, manage and provision virtual machines.
- Storage virtualization to consistently access common storage from any server.
- Multi-level administration to enable administration of physical infrastructure as well as administration of virtual objects.
- Ability to convert existing virtual machines on foreign hypervisors to Red Hat Enterprise Virtualization platform.
- A range of reports either from the reports module based on JasperReports, or from the data warehouse. The reports enable administrators to monitor and analyze information on virtual machines, hosts and storage usage and performance.

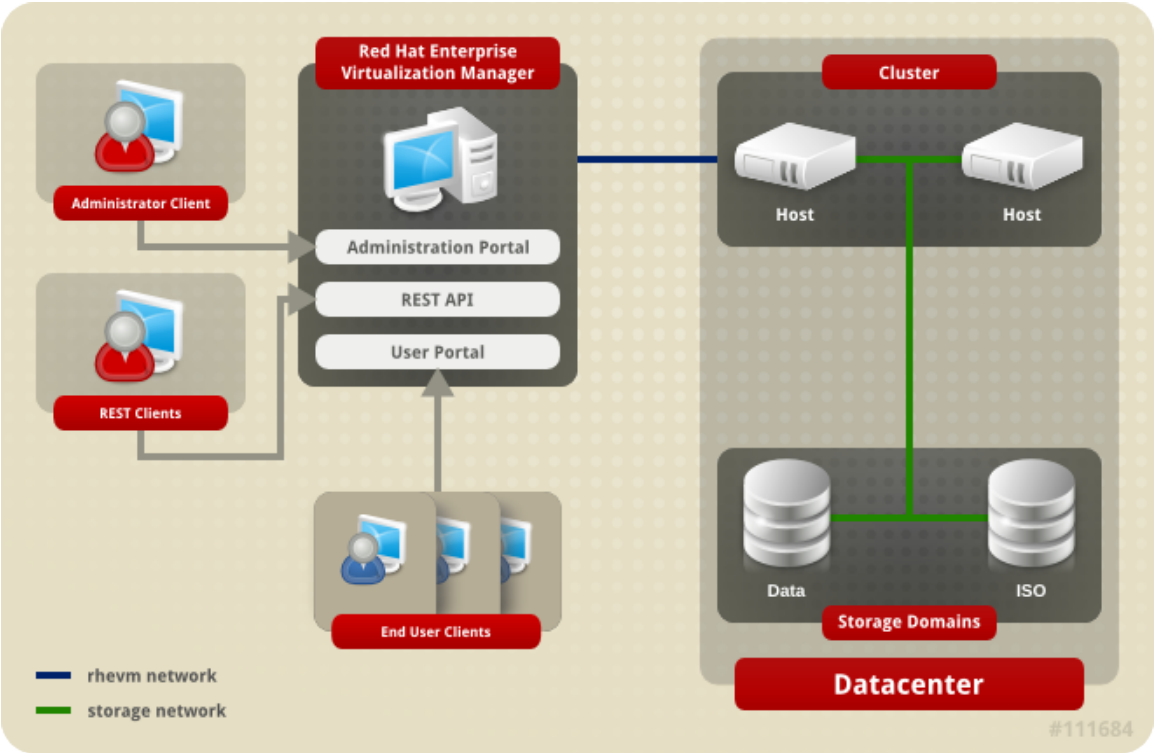

Figure 2-1: Red Hat Enterprise Virtualization Component Overview

Red Hat Enterprise Virtualization Architecture

A Red Hat Enterprise Virtualization environment consists of:

- Virtual machine hosts using the Kernel-based Virtual Machine (KVM).
- Agents and tools running on hosts including VDSM, QEMU, and libvirt. These tools provide local management for virtual machines, networks and storage.
- The Red Hat Enterprise Virtualization Manager (RHEV-M); a centralized management platform for the Red Hat Enterprise Virtualization environment. It provides a graphical interface where one can view, provision and manage resources.
- Storage domains to hold virtual resources like virtual machines, templates, ISOs.
- A database to track the state of and changes to the environment.
- Access to an external Directory Server to provide users and authentication.
- Networking to link the environment together- physical and logical networks.

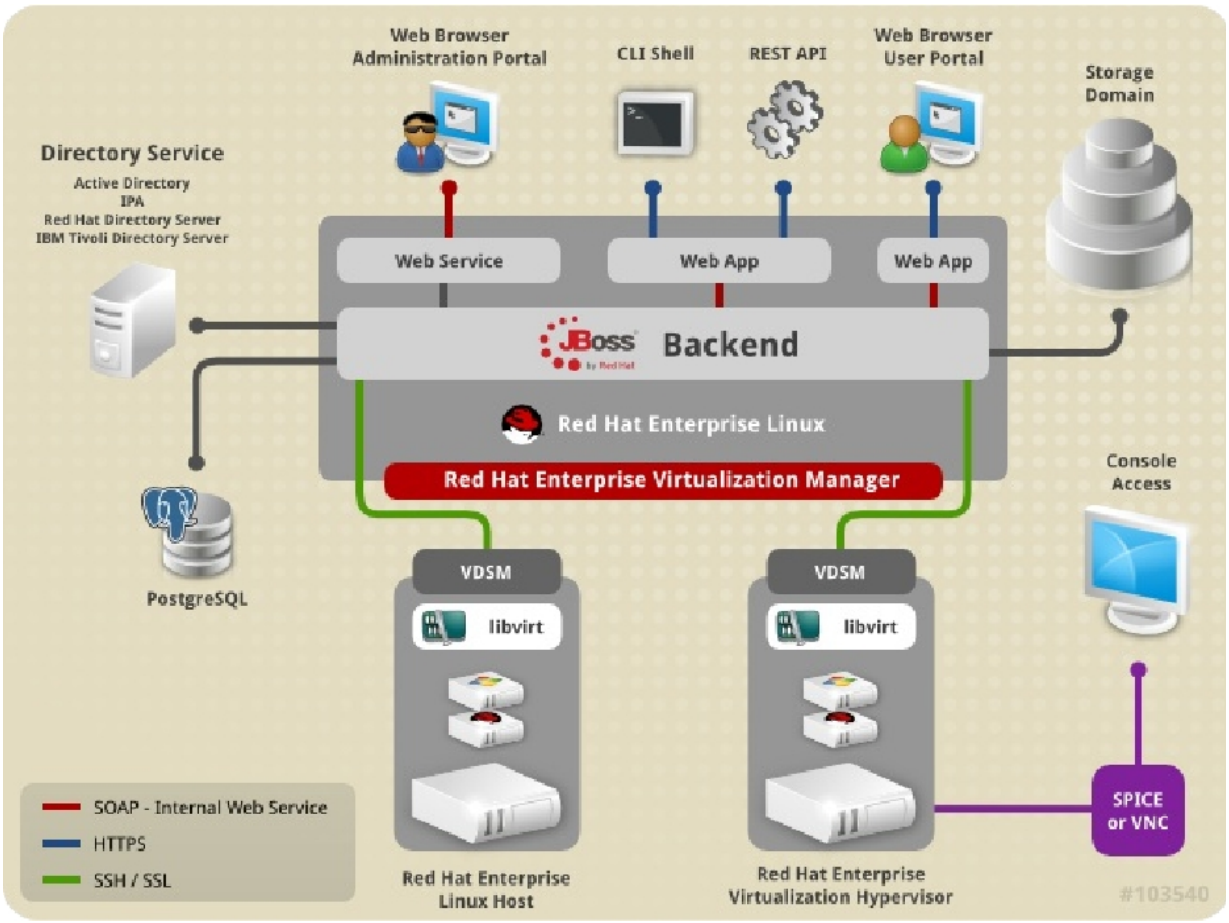

Figure 2-2: Red Hat Enterprise Virtualization Platform Overview

2.1 Red Hat Enterprise Virtualization Manager (RHEV-M)

Red Hat Enterprise Virtualization Manager provides a centralized management system to provision, administer and control the virtualized infrastructure (hosts and virtual machines) managing storage, network, high availability (live migrations) and reporting. The Red Hat Enterprise Virtualization Manager itself executes on Red Hat Enterprise Linux 6.

RHEV-M provides a user friendly interface that allows an administrator to manage their virtual infrastructure from a web browser. Utilizing SPICE or VNC, RHEV-M provides graphical console to the virtual machines.
The Red Hat Enterprise Virtualization Manager exposes an Administration Portal, a User Portal, and an Application Programming Interface (API).

- The ***Administration Portal*** is used to perform setup, configuration, and management of the Red Hat Enterprise Virtualization environment.
- The ***User Portal*** is used to start, stop, reboot, and connect to virtual machines.
- The ***REST API*** provides and interface for automation of tasks normally accomplished manually by users. Scripts that make use of the REST API can be written in any language which supports accessing HTTP and HTTPS resources.

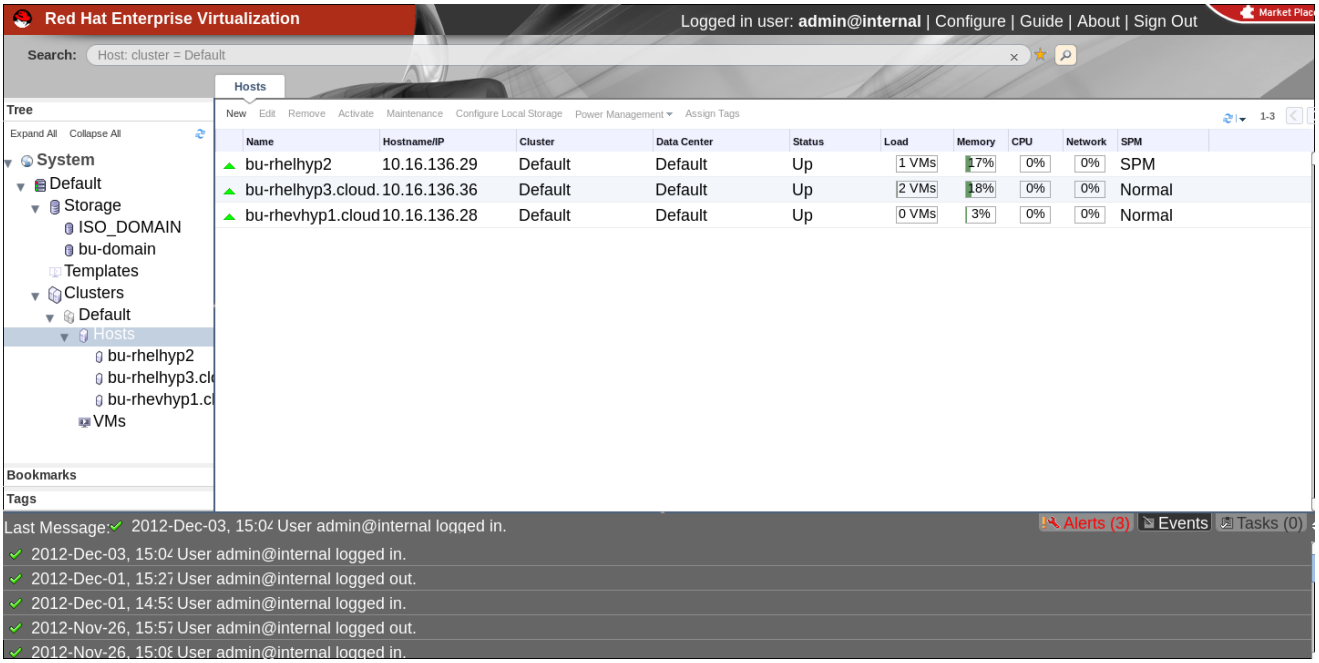

Figure 2.1-1: Red Hat Enterprise Virtualization Platform Overview

2.2 Red Hat Enterprise Virtualization Hosts

Hosts, also known as **hypervisors**, are the physical servers on which virtual machines run. Full virtualization is provided by using a loadable Linux kernel module called Kernel-based Virtual Machine (KVM).
KVM can concurrently host multiple virtual machines running either Windows or Linux operating systems. Virtual machines run as individual Linux processes and threads on the host machine and are managed remotely by the Red Hat Enterprise Virtualization Manager.
A Red Hat Enterprise Virtualization environment has one or more hosts attached to it.
A host is a physical 64-bit server with the Intel VT or AMD-V extensions running Red Hat Enterprise Linux 6.1 or later AMD64/Intel 64 version.
A physical host on the Red Hat Enterprise Virtualization platform:

- Must belong to only one cluster in the system.
- Must have CPUs that support the AMD-V or Intel VT hardware virtualization extensions.
- Must have CPUs that support all functionality exposed by the virtual CPU type selected upon cluster creation.
- Must have a minimum of 2 GB RAM.
- Can have an assigned system administrator with system permissions

Red Hat Enterprise Virtualization supports two types of hosts:

- Red Hat Enterprise Virtualization Hypervisor Hosts *(henceforth referred to as* RHEV Hosts *in this document)*
- Red Hat Enterprise Linux Hosts *(henceforth referred to as* RHEL Hosts *in this document)*

2.2.1 Red Hat Enterprise Virtualization Hypervisor Hosts

Red Hat Enterprise Virtualization Hypervisor hosts are installed using a special build of Red Hat Enterprise Linux with only the packages required to host virtual machines. They run stateless, not writing any changes do disk unless explicitly required.
When joining the Red Hat Enterprise Virtualization environment, the manager approves the preconfigured hosts. They are functionally equivalent to Red Hat Enterprise Linux hosts, and can coexist in the same environment.

2.2.2 Red Hat Enterprise Linux Hosts

One can use a standard Red Hat Enterprise Linux 6 installation on a capable hardware as a host. Red Hat Enterprise Virtualization supports hosts running Red Hat Enterprise Linux 6 Server AMD64/Intel 64 version.

2.2.3 Virtual Machines

A virtual machine is a software implementation of a computer. The Red Hat Enterprise Virtualization environment enables creation of virtual desktops and virtual servers.

Red Hat Enterprise Virtualization 3.1 presently supports virtualization of the following guest operating systems:

- Red Hat Enterprise Linux 3 (32 bit and 64 bit)
- Red Hat Enterprise Linux 4 (32 bit and 64 bit)
- Red Hat Enterprise Linux 5 (32 bit and 64 bit)
- Red Hat Enterprise Linux 6 (32 bit and 64 bit)
- Windows XP Service Pack 3 and newer (32 bit only)
- Windows 7 (32 bit and 64 bit)
- Windows Server 2003 Service Pack 2 and newer (32 bit and 64 bit)
- Windows Server 2008 (32 bit and 64 bit)
- Windows Server 2008 R2 (64 bit only)

3 Symantec NetBackup

The Symantec NetBackup platform simplifies the protection of one's information-driven enterprise by automating advanced technologies and standardizing operations across applications, platforms, and virtual environments. That means being able to protect completely, store efficiently, recover anywhere, and manage centrally across heterogeneous operating systems and storage hardware including tape and disk. Integrated deduplication, replication, and virtual machine protection helps customers improve storage efficiency, infrastructure use, and recovery times. A single console offers multi-site monitoring, analytics, and reporting, which allows customers to standardize operations and risk management. Used by companies around the world, Symantec NetBackup easily scales to protect the largest UNIX, Windows, and Linux environments.

The NetBackup Platform consists of the following Symantec products:

- NetBackup
- NetBackup Appliances
- NetBackup RealTime
- OpsCenter Analytics
- Enterprise Vault

Features of Symantec NetBackup include:

- Heterogeneous data protection - Protection across heterogeneous operating systems, applications and hypervisors for both disk and tape architectures.
- Centralized management - Increased efficiency by managing all data protection technologies using multiple NetBackup servers and domains from one location.
- Source and target data deduplication - Easily deploy and manage deduplication wherever needed, from remote offices to the data center
- Turnkey solution - NetBackup appliances for quickly deploying NetBackup backup and deduplication technologies
- Deep integration with storage appliances - The NetBackup OpenStorage API enables centralized management of deduplication and replication
- Effective disaster recovery - Fully automated and integrated system recovery with NetBackup Bare Metal Restore, built-in replication, and offsite tape management
- Highly scalable - Benefit from a flexible, three-tiered architecture that scales with the needs of today's growing data center
- Comprehensive data security - Flexible encryption technologies for maximum data security while in transit or in media
- Fast, granular recovery of data from applications - Quickly restore files, emails and other granular items

NetBackup Capabilities:

Figure 3-1: NetBackup Capabilities

NetBackup Scalability is enhanced by Three Tier Architecture. However, it can be configured with Two Tier Architecture by combining Master and Media Serves as a single unit.

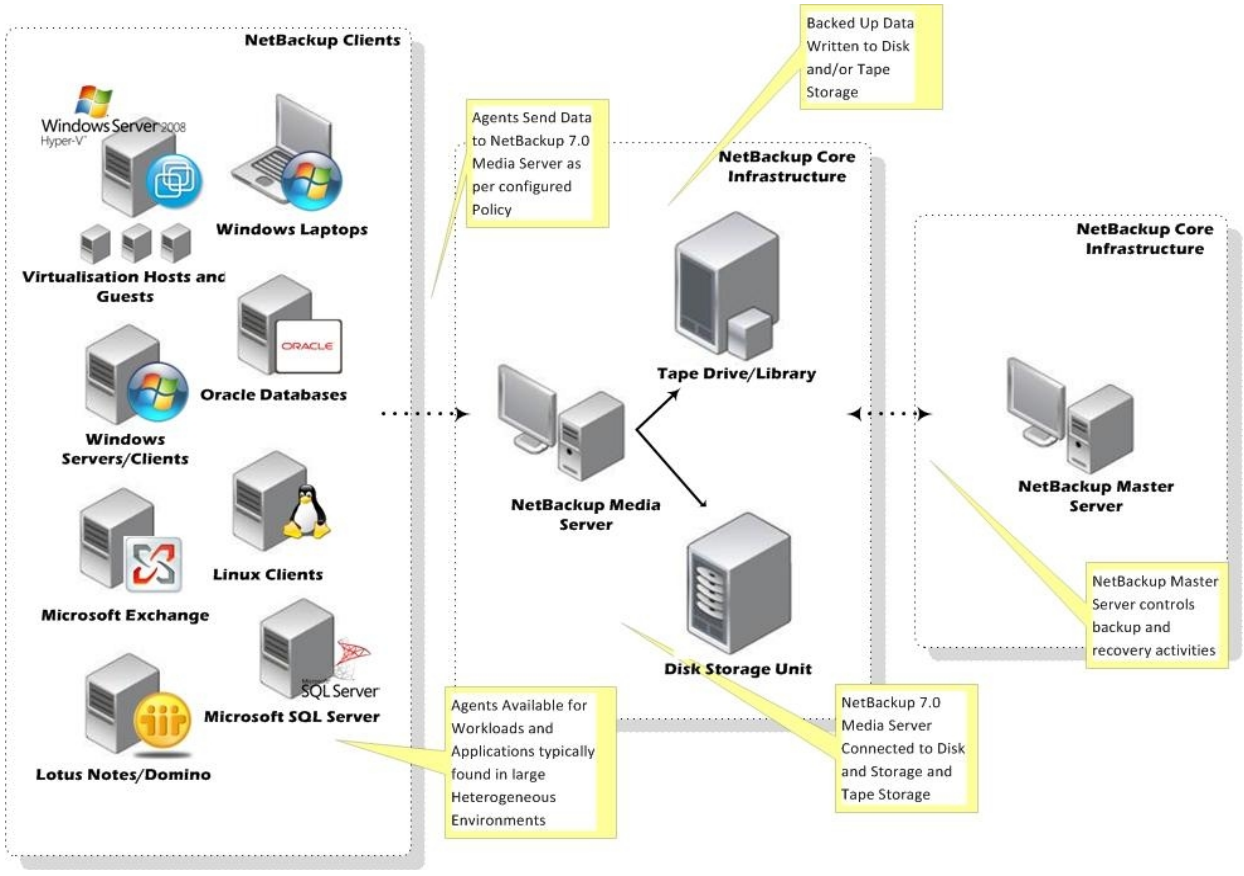

Figure 3-2: Symantec NetBackup Three-Tier Architecture

(Clients,Media and Master Server)

Main Components

- **Master Server** - This is the controller of the data protection solution. It allocates resources to the media server, keeps track of the media being used, catalogs client information, schedules and reports backup activity. (NetBackup Management)

- **EMM Server** - The Enterprise Media Manager; manages and allocates required resources. (Resource Management)

- **Media Server** - This is the data mover in the infrastructure. It has some form of storage attached to it -either directly or through the network. Media servers manage the writing and reading of data to and from media and is very I/O intensive. There can be multiple media servers classified by use case and environment connected to the same master server. (Data Management)

- **Clients**- Are hosts that have data to be backup up or restored to.

- **NetBackup OpsCenter** - Netbackup OpsCenter is a Web-based software application that can manage multiple NetBackup environments and is bundled with the NetBackup 7 distribution. This is a convergence of NetBackup Operations Manager (NOM) and Veritas Backup Reporter (VBR). This is an optional component and has to be configured separately. This is not in scope for this paper.

- **BMR Boot Server** – Bare Metal Restore Boot servers provide the environment that is required to rebuild a protected client, including resources such as shared resource trees (SRT). Boot servers also provide the resources that are used to boot the client system when it performs a network boot before restore.

4 Reference Architecture Configuration

4.1 Environment

This reference architecture environment consists of NetBackup and RHEV infrastructures. The NetBackup infrastructure comprises of a single NetBackup server that functions as a Master, Media, EMM and BMR Boot Server (NetBackup two tier architecture). The RHEV infrastructure comprises of a RHEV Manager, a RHEV Host and a RHEL Host. There are three Virtual Machines (Linux and Windows based) that reside the RHEV and RHEL Hosts.

4.1.1 Software Configuration

Operating Systems with versions used as referenced in Table 4.1.1-1: Operating System Revisions.

Hostname	Role	Software	Version
bu-netbackup	Netbackup Master	RHEL 6.3	2.6.32-279.11.1.el6
bu-rhevm	RedHat Enterprise Virtualization Manager (RHEV-M)	RHEL 6.3	2.6.32-279.11.1.el6
bu-rhevhyp1	RHEV based Host (Hypervisor)	RHEL 6.3	2.6.32-279.11.1.el6
bu-rhelhyp2	RHEL based Host (Hypervisor)	RHEL 6.3	2.6.32-279.11.1.el6
bu-vm1	Virtual Machine	RHEL6.3	2.6.32-279.11.1.el6
bu-vm2	Virtual Machine	MS-Windows 2008 R2	6.1.7600
bu-vm3	Virtual Machine	RHEL6.3	2.6.32-279.11.1.el6

Table 4.1.1-1: Operating System Revisions

Hostname	Disk	Memory	CPU
bu-vm1	50 GB	2 GB	1 Virtual core
bu-vm2	45 GB	4GB	2 Virtual core
bu-vm3	50 GB	4GB	1 Virtual core

Table 4.1.1-2: Virtual Machine Sizing

4.1.2 Hardware Configuration

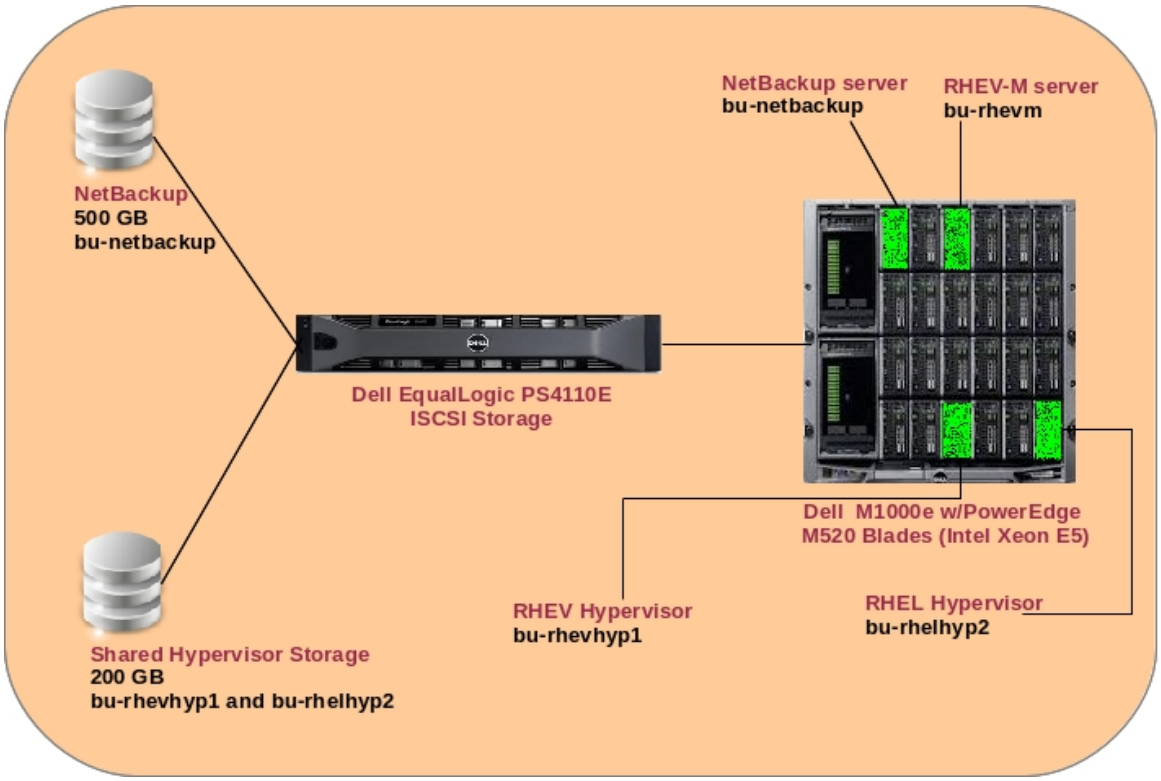

Figure 4.1.2-1: Hardware Configuration

4.1.2.1 Server Details

Hardware System	Specifications
NetBackup Host (bu-netbackup) [1 x Dell M1000e- M520 Blade]	2 Socket, 8 Core (16 cores) Intel(R) Xeon(R) E5-2450 -2.10GHz, 36 GB RAM
	2 x 146 GB SAS internal disk drive (mirrored)
	1Gb Gigabit Ethernet for public network
	10 GbE KR Ethernet for private network
RHEV-M Server (bu-rhevm) [1 x Dell M1000e- M520 Blade]	2 Socket, 8 Core (16 cores) Intel(R) Xeon(R) CPU E5-2450 -2.10GHz, 36 GB RAM
	2 x 146 GB SAS internal disk drive (mirrored)
	1Gb Gigabit Ethernet for public network
	10 GbE KR Ethernet for private network
RHEV-H Hypervisor (bu-rhevhyp1) [1 x Dell M1000e- M520 Blade]	2 Socket, 8 Core (16 cores) Intel(R) Xeon(R) CPU E5-2450 -2.10GHz, 36 GB RAM
	2 x 146 GB SAS internal disk drive (mirrored)
	1Gb Gigabit Ethernet for public network
	10 GbE KR Ethernet for private network
RHEL-H Hypervisor (bu-rhelhyp2) [1 x Dell M1000e- M520 Blade]	2 Socket, 8 Core (16 cores) Intel(R) Xeon(R) CPU E5-2450 -2.10GHz, 36 GB RAM
	2 x 146 GB SAS internal disk drive (mirrored)
	1Gb Gigabit Ethernet for public network
	10 GbE KR Ethernet for private network

Table 4.1.2.1-1: Server Hardware

4.1.2.2 Storage Hardware

The storage hardware used in this reference environment are non local provided by: EqualLogic PS4110E 10GbE iSCSI Array with (NL-SAS) 3.5” drives.

System	Disk Usage
RHEV and RHEL Hypervisors (shared disk)	200 GB
NetBackup Master Server	500 GB

Table 4.1.2.2-1: Storage

4.1.2.3 Logical Network

The following section details the Logical Network configuration used in this paper.

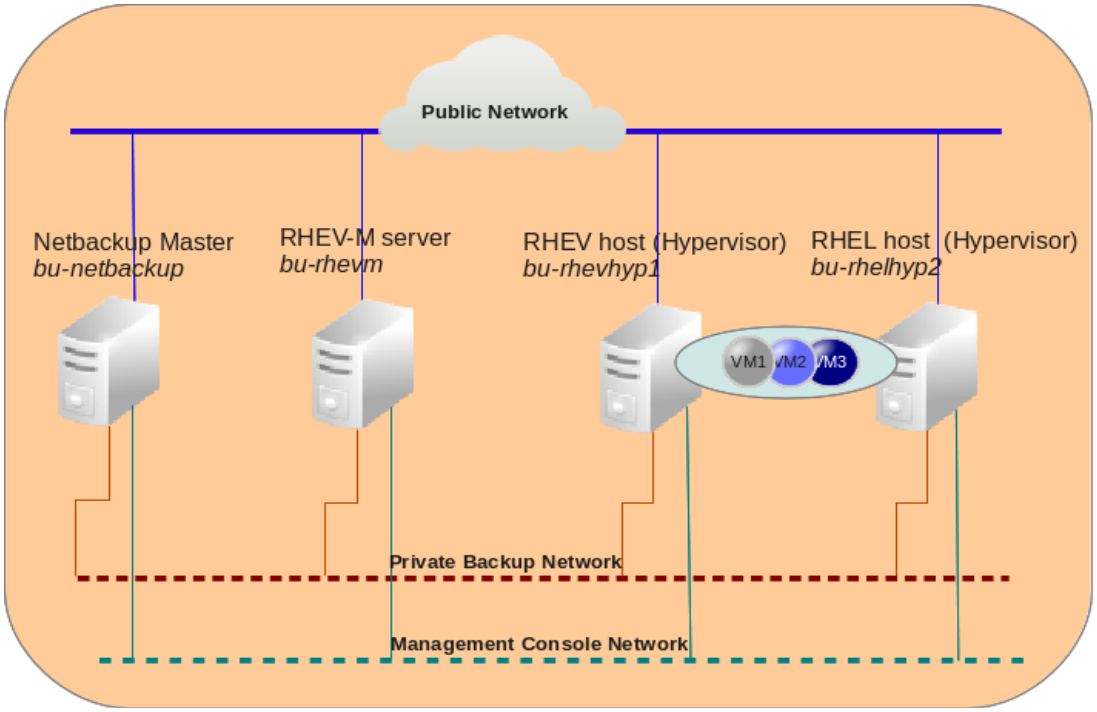

Figure 4.1.2.3-1: Logical Network Diagram

4.1.3 Security

SELinux

SELinux was enabled and enforced on all Red Hat machines in this reference environment unless otherwise noted.

4.1.4 Application Tools and Packages

Applications, tools and package revisions used referenced in

Red Hat Enterprise Virtualization Manager (RHEV-M)	3.1.0-26.el6ev
Symantec NetBackup	7.5.0.5

Table 4.1.4-1: Application Tools and Packages Revisions

5 Deploying the infrastructure

5.1 NetBackup Infrastructure Deployment

This section describes installation and configuration of Master and Media server.
The NetBackup environment in this paper runs on version NetBackup 7.5.0.5. There is a single primary server that has been setup as Master, Media, EMM and BMR Master server. The utilized configuration performs all the functional activities and was not a bottleneck in the paper's small scale configuration. While this configuration is supported by Symantec, in enterprise production environment, these are most likely to be on multiple dedicated servers to satisfy business requirements.
When performing over-the-network backup and recovery, it is recommended to segregate this traffic to a dedicated network to prevent saturation of the public network. To achieve this configuration, a private backup network was configured and all clients and hosts were attached. In this environment, the naming convention, for example the master network is as follows:

public network interface: bu-netbackup.*domain*
backup network interface: bu-netbackup-bkp.*domain*

Note- A bug has been identified where BMR Master server could not import Linux BMR clients on NetBackup 7.5 with Red Hat Enterprise Virtualization 6.3. This has been addressed in NetBackup 7.5.0.5, hence upgrade to 7.5.0.5 is a requirement if BMR functionality is expected.

5.1.1 Master Server

NetBackup Master server is configured on a physical Red Hat Enterprise Virtualization 6.3 server running Symantec NetBackup 7.5.0.5. Additional interface names for each configured network added to the configuration. Figure 5.1.1-1: NetBackup Master Server illustrates inclusion of the backup interface as an additional host pointing to the same master server. For detailed software and hardware descriptions, please refer to 4Reference Architecture Configuration**.**

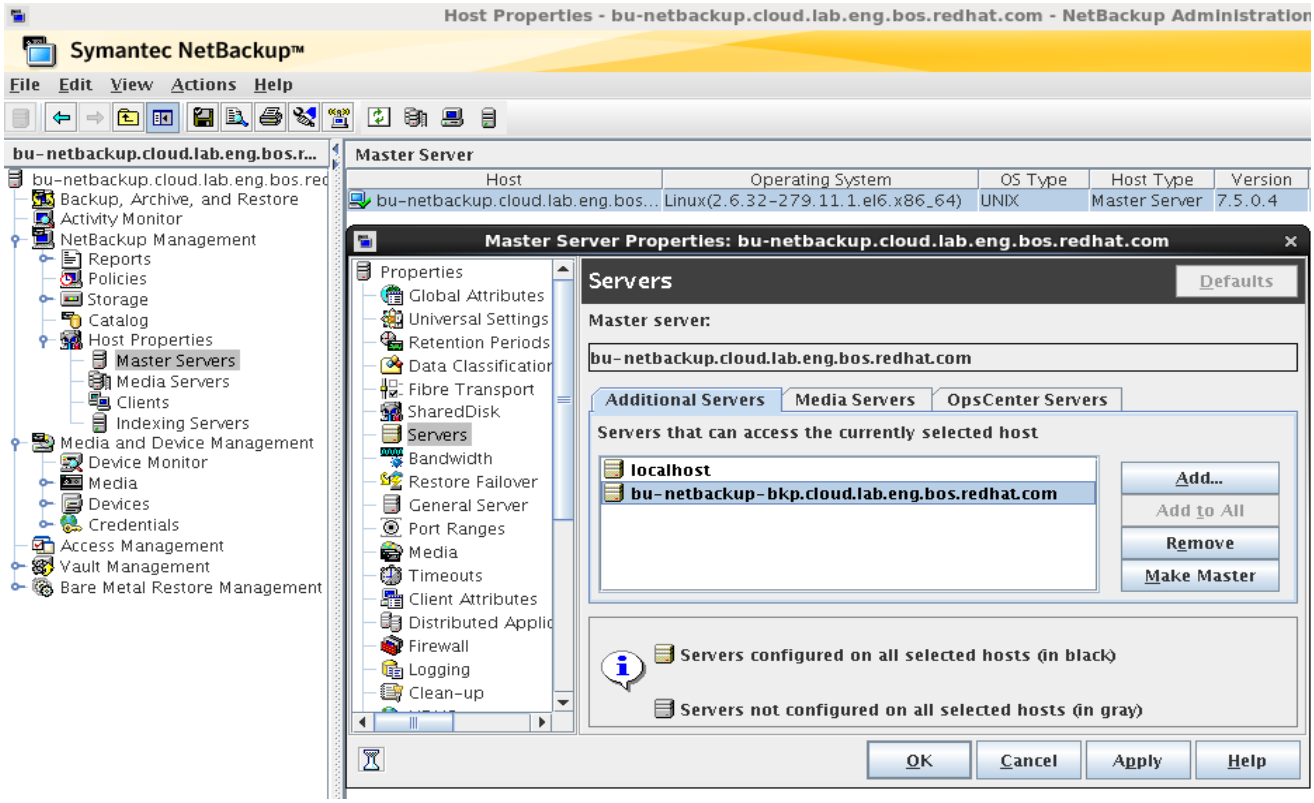

Figure 5.1.1-1: NetBackup Master Server

Installation of NetBackup Master software

This section describes installation of the NetBackup software on ***bu-netbackup*** to set it up as a master, media manager and EMM server.

```
# ls -l NetBackup_7.5_LinuxR_x86_64
total 120
drwxr-xr-x. 2 13258 13258  4096 Feb  8  2012 Doc
-r-xr-xr-x. 1 13258 13258 59354 Feb  8  2012 install
-r--r--r--. 1 13258 13258 49589 Feb  8  2012 LICENSE
dr-xr-xr-x. 4 13258 13258  4096 Feb  8  2012 linuxR_x86

#cd  NetBackup_7.5_LinuxR_x86_64
#./install
```

```
Symantec Installation Script^M
Copyright 1993 - 2012 Symantec Corporation, All Rights Reserved.
        Installing NetBackup Server Software


Do you wish to continue? [y,n] (y) y

Participate in the NetBackup Product Improvement Program? [y,n] (y) n

The NetBackup and Media Manager software is built for use on LINUX_RH_X86
hardware.
Do you want to install NetBackup and Media Manager files? [y,n] (y) y

NetBackup and Media Manager are normally installed in /usr/openv.
Is it OK to install in /usr/openv? [y,n] (y) y

A NetBackup Server or Enterprise Server license key is needed
for installation to continue.

Enter license key: XXXX-XXXX-XXXX-XXXX-XXXX-XXXX

All additional keys should be added at this time.
Do you want to add additional license keys now? [y,n] (y) n

Use /usr/openv/netbackup/bin/admincmd/get_license_key
to add, delete or list license keys at a later time.

Installing NetBackup Enterprise Server version: 7.5

Would you like to use "bu-rhelhyp3.cloud.lab.eng.bos.redhat.com" as the
configured NetBackup server name of this machine? [y,n] (y) y

Is bu-rhelhyp3.cloud.lab.eng.bos.redhat.com the master server? [y,n] (y) y

Do you want to add any media servers now? [y,n] (n) n

Enter the Enterprise Media Manager server
(default: bu-netbackup.cloud.lab.eng.bos.redhat.com):
bu-netbackup.cloud.lab.eng.bos.redhat.com

If an OpsCenter server already exists in your environment
or you plan to install one, enter the real hostname of that
OpsCenter server here.  Do not use a virtual name.  If you
do not want this local machine to be an OpsCenter server,
enter NONE.

Enter the OpsCenter server (default: NONE): NONE


NetBackup server installation complete.

File /usr/openv/tmp/install_trace.7630 contains a trace of this install.
That file can be deleted after you are sure the install was successful.

            ...output abbreviated to display only input values...
```

5.1.2 NetBackup upgrade

This step was performed to upgrade NetBackup 7.5 to NetBackup 7.5.0.5. For a RedHat Enterprise Virtualization 6.3 Client, BMR requires NetBackup 7.5.0.5. Please refer to Appendix C NetBackup Upgrade Procedure for details.

5.1.3 Media

Backup media dedicated for this reference environment consists of a 500 GB disk mounted as a filesystem to ***/backup/backup1*** on the master server ***bu-netbackup*** and is configured as the destination for backup jobs within the backup policies. **Figure 5.1.3-1: Backup to Disk Media** displays the storage unit configured for backup to disk. While this minimal setup meets the needs of this paper, an enterprise deployment would be much more extensive.

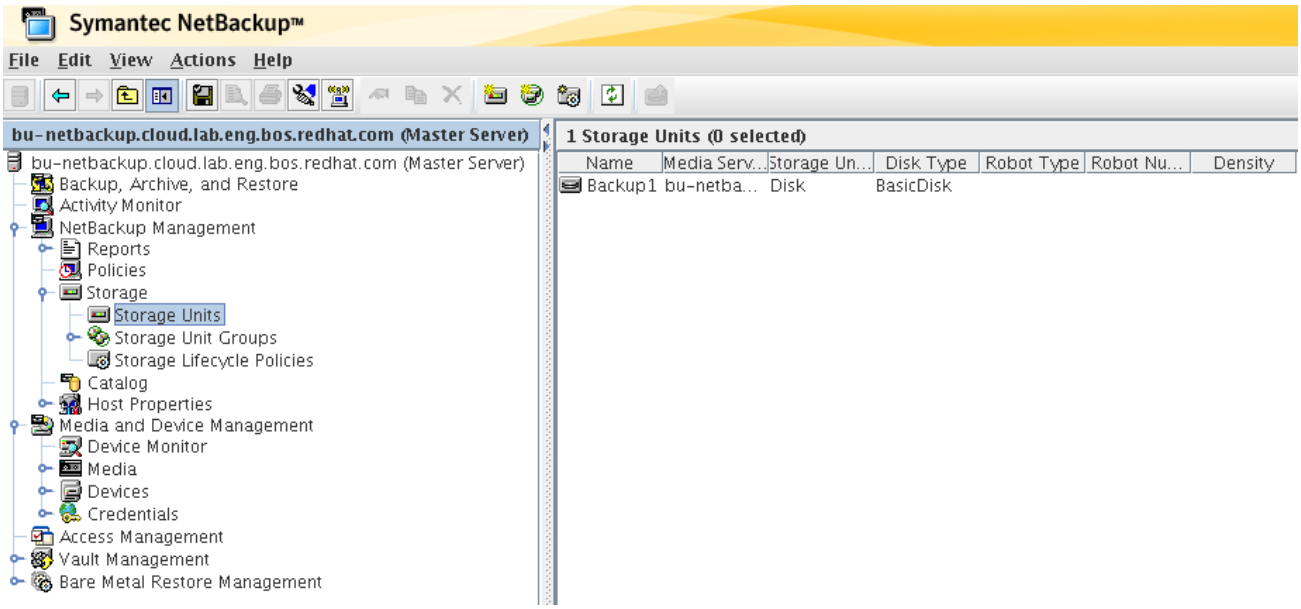

Figure 5.1.3-1: Backup to Disk Media

5.1.4 NetBackup Licensing

Although the purpose of this guide is not focused on licensing, it is important to have several NetBackup features licensed in order to complete the necessary backup operations within a Red Hat Enterprise Virtualization environment. To verify NetBackup is licensed properly to perform the necessary backup operations, choose *Help* and *License Keys* within the NetBackup console as shown below.

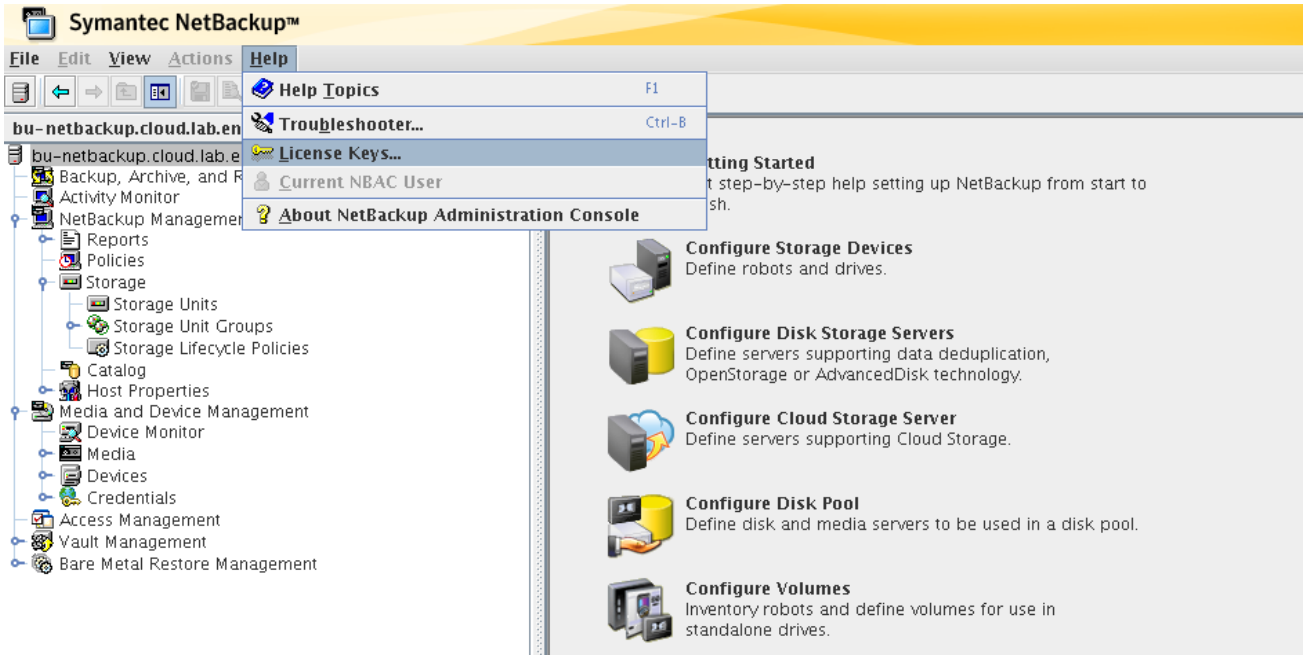

Figure 5.1.4-1: Display NetBackup Registered Keys

Command line option:

```
# /usr/openv/netbackup/bin/admincmd/get_license_key

      License Key Utility
      -------------------
      A) Add a License Key
      D) Delete a License Key
      F) List Active License Keys
      L) List Registered License Keys
      H) Help
      q) Quit License Key Utility

Enter a letter: F

Enter the name of the host (default is bu-
netbackup.cloud.lab.eng.bos.redhat.com):


Active NetBackup Features
```

```
========================
License Key:     XXXX-XXXX-XXXX-XXXX-XXXX-XXXX-XXXX-XXXX-XXXX-XXXX
Date Added:Nov 7, 2012 at 07:53:43
Host:      bu-netbackup.cloud.lab.eng.bos.redhat.com
Product:   NetBackup Enterprise Server
Client count:    See license certificate.
Expires:   Aug 17, 2013 at 01:00:00
Feature:   Base NetBackup
Feature:   Additional clients
Feature:   NDMP
Feature:   Shared Storage Option
Feature:   MS Exchange extension
Feature:   MS SQL Server extension
Feature:   DB2 extension
Feature:   Lotus Notes extension
Feature:   Oracle extension
Feature:   Informix extension
Feature:   Sybase extension
Feature:   SAP extension
Feature:   Encryption (Legacy DES 40-bit)
Feature:   Encryption (Legacy DES 56-bit)
Feature:   DataStore
Feature:   Library Based Tape Drives
Feature:   Vault
Feature:   Inline Tape Copy
Feature:   MS SharePoint Agent
Feature:   Snapshot Client
Feature:   StorageTek ACS Robotic Libraries
Feature:   Fujitsu LMF Robotic Libraries
Feature:   IBM ATL Robotic Libraries
Feature:   ADIC DAS/SDLC Robotic Libraries
Feature:   Microsoft RSM Robotic Libraries
Feature:   Remote Media Server Support
Feature:   Robotic Library Sharing Support
Feature:   Remote Client Support
Feature:   Open File Backup
Feature:   Encryption
Feature:   Bare Metal Restore
Feature:   Virtual Tape Option
Feature:   OpenStorage Disk Option
Feature:   Flexible Disk Option
Feature:   PureDisk MS SQL Server Agent
Feature:   Enterprise Vault Agent
Feature:   PureDisk MS Exchange Agent
Feature:   SAN Client
Feature:   PureDisk Option
Feature:   PureDisk Remote Office
Feature:   Accelerator
Feature:   Replication Director
========================
```

Note**:** Most of the licenses have been applied and listed here. This license list may vary according to specific environment and configuration.

5.1.5 NetBackup Policy

NetBackup policies define the rules that NetBackup follows when backing up clients. A policy can contain one or more clients, and every client must belong to at least one policy. Usually clients are grouped together by common backup and archiving requirements. Policy creation brings up questions like - **who** (client), **what** (files/data), **where** (storing media), **how** (backup characteristics) and **when** (schedule).

Table 4.2.5-1: Backup Policies describes the backup policies used in this reference environment and are grouped by similar operating systems.

Policy	Clients	Type	Schedule and Retention
linux-pol	bu-rhevm-bkp bu-vm1-bkp bu-vm3-bkp	Standard	Differential – every 8 hours Full – Every day Retention – 2 Weeks
windows-pol	bu-vm2	Standard	Differential – every 8 hours Full – Every day Retention – 2 Weeks
catalog-pol	bu-netbackup	NBU-Catalog	Differential – every 8 hours Full – Every day Retention – 2 Weeks

Table 5.1.5-1: Backup Policies

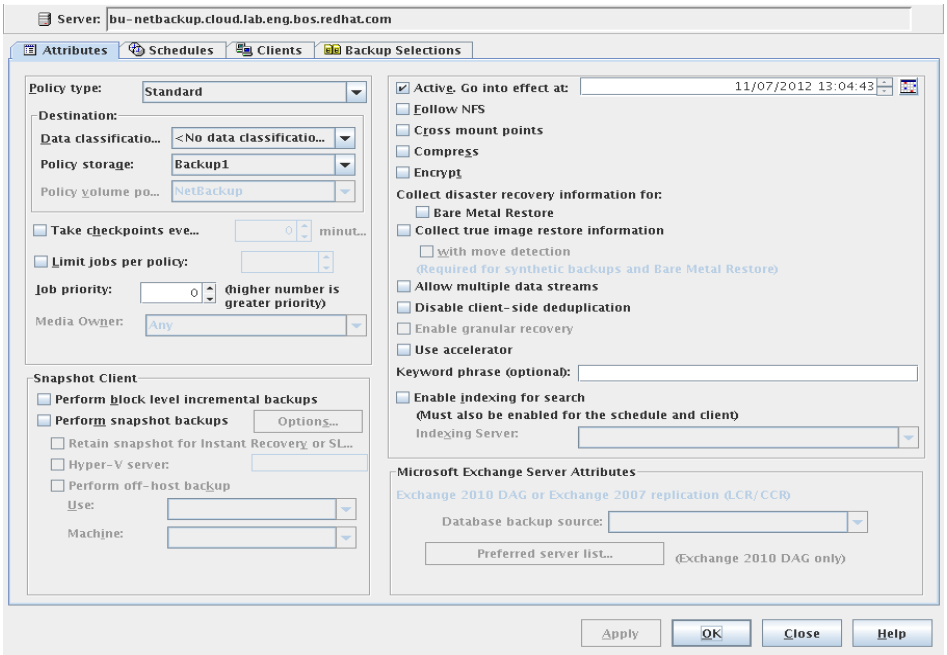

Figure 5.1.5-1: NetBackup Policy

5.2 Red Hat Enterprise Linux Client Installation

5.2.1 Network selection

The configuration described in this paper mimics a typical backup environment with a dedicated backup network, along with additional public and/or management networks on the same system. In order to ensure the backup and restore traffic passes through the dedicated backup network and does not impact other network's bandwidth, it is recommended to use the interface name that resolves to backup network when adding or installing a client. For example, to ensure the backup traffic runs through the dedicated backup network, during client installation, the client hostname used was '***bu-rhevm-bkp***' and not '***bu-rhevm***'

There are other options to identify the network to use in a multi-home setup using Preferred Network settings with '**Match**, **Prohibit** and **Only**' selections. However this option is more suited for customers with remote data centers and WAN connections. This feature is beyond the scope of this document and has not been used in this configuration. For further details on this, please refer to **Symantec NetBackup 7.5 Administrator's Guide for Unix and Linux** - http://www.symantec.com/business/support/index?page=content&id=DOC5157 .

5.2.2 Generate and Copy Shared SSH Keys

Remote installation of NetBackup agent on Linux clients from the master server can be made possible by configuring SSH with password less access from master to the the clients. On the master server, SSH Key can be generated and distributed to the clients by using ***'ssh-copy-id'*** command.

On master: As root on the master server, issue the following command:

```
 # ssh-keygen
Generating public/private rsa key pair.
Enter file in which to save the key (/root/.ssh/id_rsa):
Enter passphrase (empty for no passphrase):
Enter same passphrase again:
Your identification has been saved in /root/.ssh/id_rsa.

                              ...output abbreviated...
```

This generates a private key ***id_rsa***. The public key ***id_rsa.pub*** must be shared to the client machines to allow access. The private key must be protected and not shared. If this is compromised, the master server and connected clients are at risk of being exposed to unauthorized root access.

On clients: Verify if the root directory under the clients have the ***.ssh*** directory.

```
# ls -all -l /root | grep ssh
#
```

If no such directory exists, connect to a different server from client via ssh and this will generate the ***.ssh*** directory.. Verify by rerunning the **'ls'** command.

```
# ls -all -l /root | grep ssh
drwx------.  2 root root  4096 Nov 29 16:01 .ssh
```

On master: Run the following command to distribute the public key to the client.

```
# ssh-copy-id -i /root/.ssh/id_rsa.pub bu-vm3
Warning: Permanently added'bu-vm3,10.16.136.37'(RSA)to the list of known
hosts.
root@bu-vm3's password: <password>
Now try logging into the machine, with "ssh 'bu-vm3'", and check in:
  .ssh/authorized_keys

to make sure we haven't added extra keys that you weren't expecting.
```

On master: Verify if the password less SSH access has been enabled from master to client.

```
# ssh bu-vm3
Warning: Permanently added 'bu-vm3,10.16.136.37' (RSA) to the list of known
hosts.
Last login: Sat Dec  1 17:50:51 2012 from bu-netbackup
Kickstarted on 2012-11-26
```

5.2.3 Remote Client Install using NetBackup commands

Remote client install can be performed from the master server to a single client or multiple clients.

Prerequisites for remote client install are:

- SSH password less authorization is enabled for communication to the clients from the master server as described in Section **5.1.1 Generate and Copy Shared SSH Keys**
- The clients being installed have been already added to a policy
- Master server has NetBackup Client software installed on the master server. If not installed, it can be installed from the installation media. Select all client types during the installation pertaining to the environment
- The client names successfully resolve by name resolution

5.2.3.1 Remote install to a single client

This command installs the agent to a single client specified. For the install_client_files to work, the client must have an entry in a policy.

```
# /usr/openv/netbackup/bin/install_client_files ssh bu-rhevm-bkp
Client bu-rhevm-bkp -- Linux hardware running RedHat2.6.18
Installing NetBackup software on bu-rhevm-bkp
Warning: Permanently added 'bu-rhevm-bkp,10.16.143.89' (RSA) to the list of
known hosts.
installpbx
100%   53KB  53.0KB/s   00:00                              PBX.tar.gz
100% 4649KB   4.5MB/s   00:00                              pdinstall
100%   69KB  68.9KB/s   00:00                              pddeagent.tar.gz
```

```
100%   37MB  37.3MB/s   00:01                                  extract_java
100%   18KB  18.0KB/s   00:00                                  JRE.tar.gz
100%   30MB  29.9MB/s   00:00                                  .sizes_JRE
100%    6     0.0KB/s   00:00

                        ...output abbreviated...

Successfully updated the session cache parameters.
Starting vnetd...
Starting bpcd...
Starting nbftclnt...
Starting bmrbd...
Checking LiveUpdate registration for the following products: CLT
This may take a few minutes.
Product CLT is installed and will be registered.
Updating LiveUpdate registration now...this may take some time.

Client install complete.
       bu-rhevm-bkp install complete
```

5.2.3.2 Remote install to all clients at once

This command refers to the list of Linux clients already added to a policy in the environment and installs NetBackup client software to them.

```
# /usr/openv/netbackup/bin/install_client_files ssh ALL
```

The ***ALL*** option specifies that all clients that are configured in any backup policy on the server.

5.2.3.3 Installing client software with the sftp method

Alternatively, the client software can be copied over to the client's ***/tmp*** directory and a NetBackup command can be executed locally at the client site.
At the master server

```
#/usr/openv/netbackup/bin/install_client_files sftp bu-vm3 root
       bu-vm3 ...
Client bu-vm3 -- Linux hardware running RedHat2.6.18
Installing NetBackup software on bu-vm3 as user root
Connecting to bu-vm3...
Warning: Permanently added 'bu-vm3,10.16.136.37' (RSA) to the list of known
hosts.

sftp completed successfully.
```

Or

```
#/usr/openv/netbackup/bin/install_client_files sftp ALL user
```

At the client server, identify if the file copy was successful:

```
# ls -all | grep bp
drwxr-xr-x.  3 root root  4096 Dec  2 09:26 bp.16871
#sh /tmp/bp.<pid>/client_config  (bp.16871 in this case)
```

Execute the copied script to complete the installation.

```
#sh /tmp/bp.<pid>/client_config
Blocksize = 20 records
./
./bp_servers
./bp_client_name

Installing PBX...


                          ...output abbreviated...


Client install complete.
```

5.2.4 Remote client install using NetBackup Administration Console

Alternatively, the NetBackup GUI on the master can be used to install the NetBackup client software remotely to the clients. This option is applicable only to Linux/Unix clients with Linux/Unix master server.

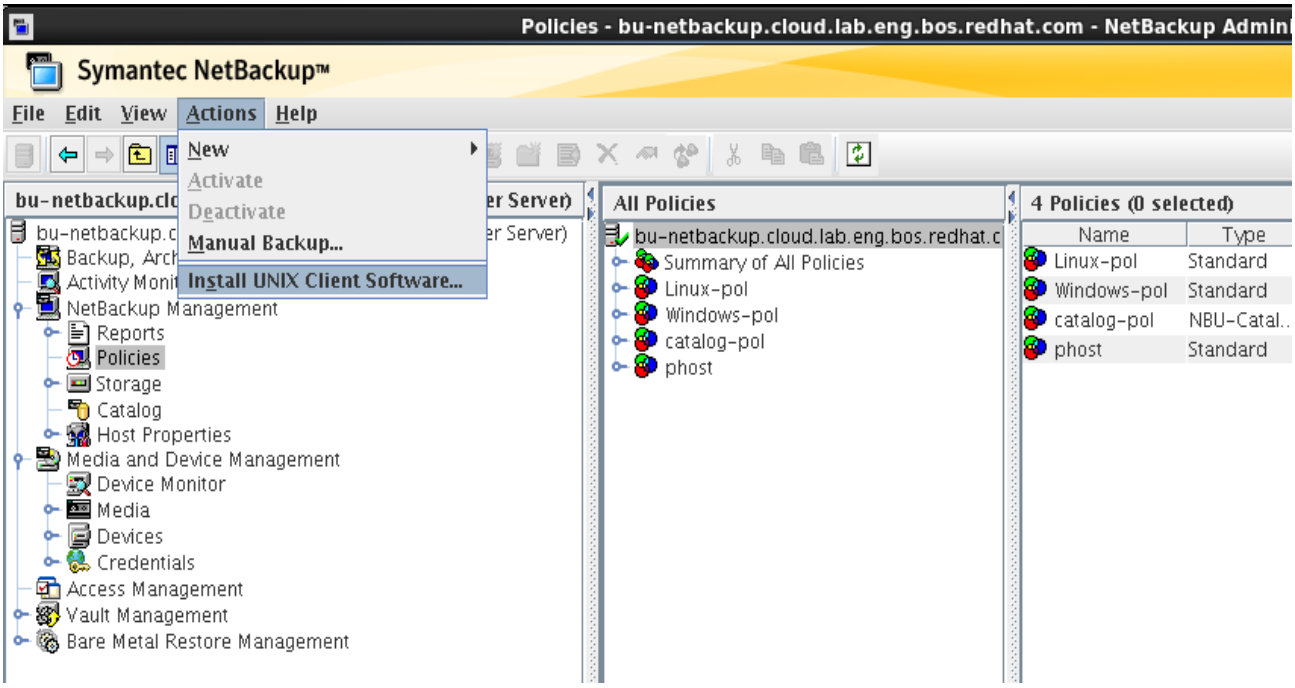

Figure 5.2.4-1: NetBackup Client Installation using Admin Console

5.2.5 Local Client Install

Local client install on Linux is similar to master server install. Please refer to topic **Installation of NetBackup Master software** for details.

5.3 Microsoft Windows Client Installation

5.3.1 Network Selection

The choice of network for NetBackup traffic between NetBackup infrastructure and the client is the same as that of Linux Client Environment. Please refer to **5.2.1Network selection** and Symantec NetBackup 7.5 Administrator's Guide for Windows http://www.symantec.com/business/support/index?page=content&id=DOC5159 for details

5.3.2 NetBackup Client Installation on Windows

There are several ways to install the NetBackup agent onto a machine running a Windows operating system. Both methods require access to the client executable from a client machine. Upon installation of the NetBackup client, the option to add additional servers is provided which enables the ability to install the client to multiple target machines at the same time as shown in below steps:

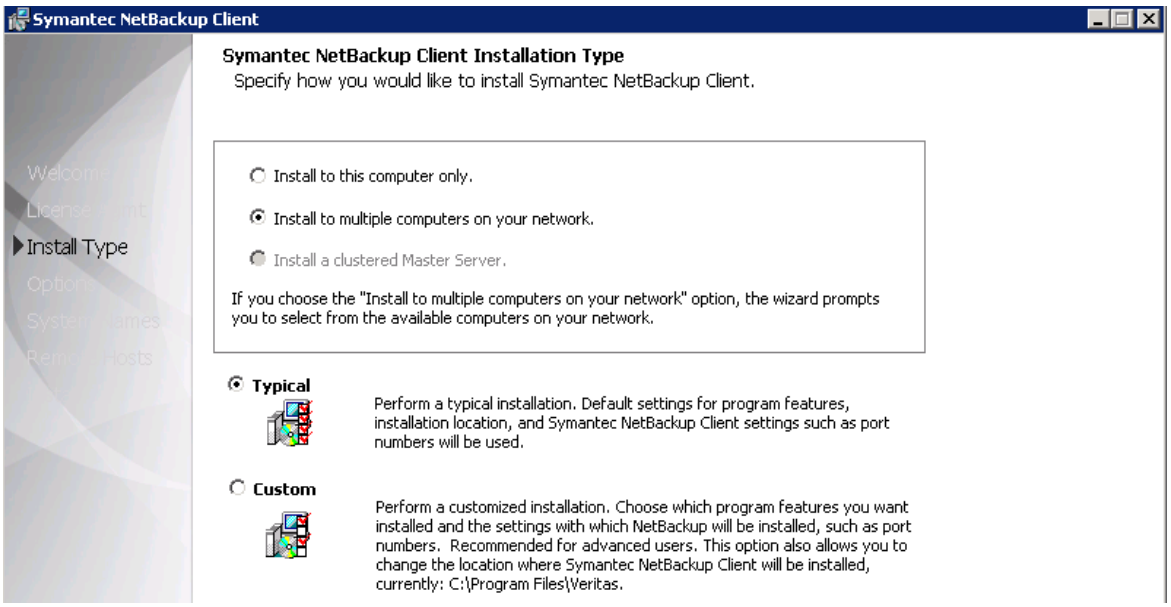

Figure 5.3.2-1: NetBackup Client Installation on Windows-step1

and

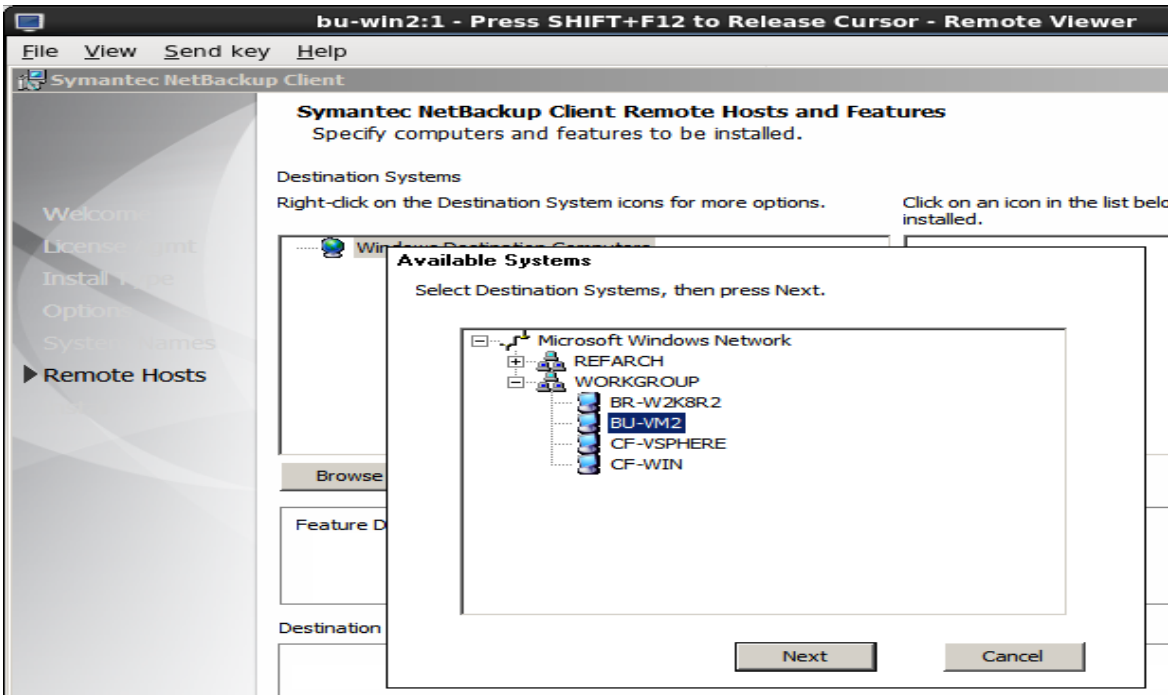

Figure 1: NetBackup Client Installation on Windows-step2

5.4 Red Hat Enterprise Virtualization Infrastructure Deployment

Red Hat Enterprise Virtualization requires installation and configuration of several components to create a functioning virtualization environment. One must install and configure each component in the order shown in this checklist **:**

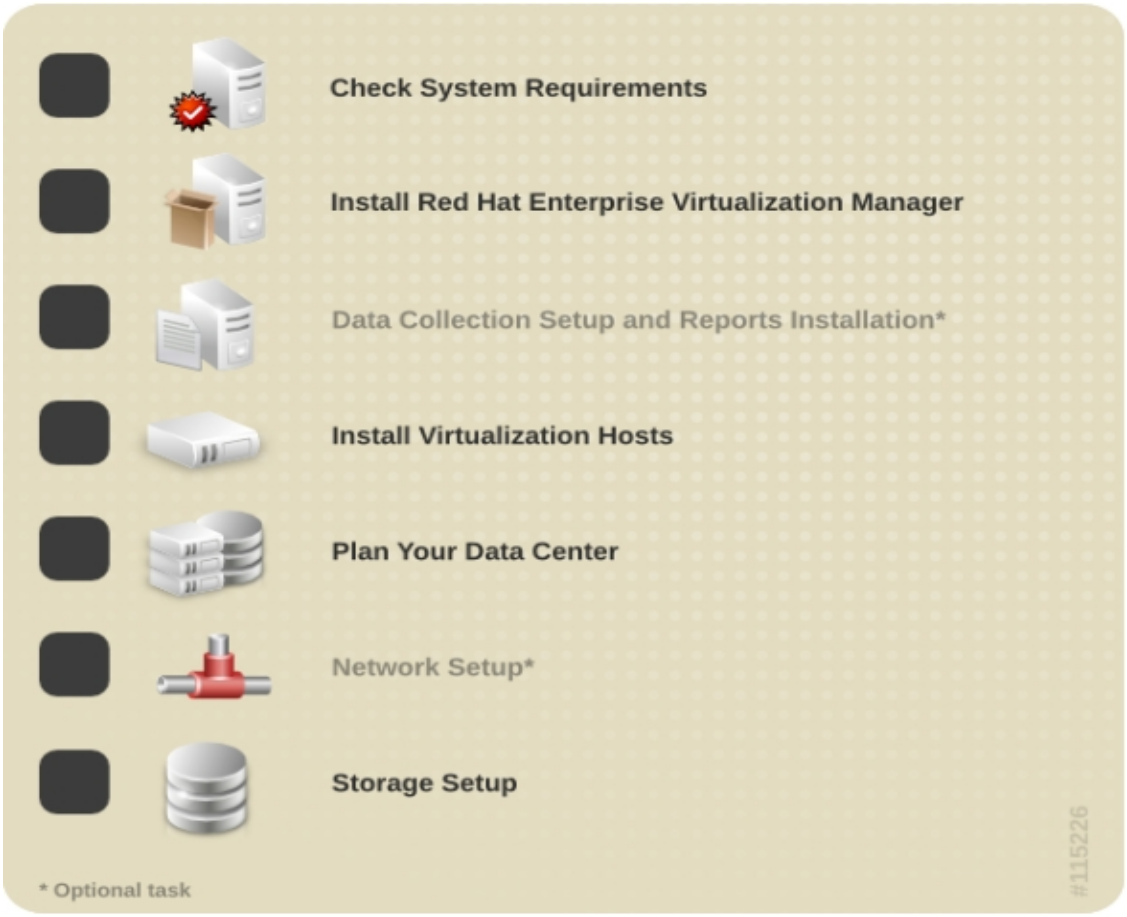

Figure 5.4-1: RHEV Environment installation steps

Note: The above installation steps are high level to describe the critical settings and sequencing. For detailed steps and descriptions, please refer to RedHat Enterprise Virtualization 3.1 Installation guide:https://access.redhat.com/knowledge/docs/enUS/Red_Hat_Enterprise_Virtualization/3.1/html/Installation_Guide/index.html

The following steps were followed to implement this infrastructure:

- **System Requirements**
  Please refer to Appendix D: Red Hat Enterprise Virtualization 3.1 Requirements

- **Install and Configure RHEV-M**
  Please refer to Appendix E: Red Hat Enterprise Virtualization Manager Installation & Configuration

- **Install Virtualization Hosts**
  Please refer to Appendix F: Hypervisor (Virtualization Host) Installation

- **Storage Setup**
  Please refer to Appendix G: Storage Configuration

6 Backup and Restore of Virtual Machines

6.1 Virtual Machine Backup

With the NetBackup client installed, policies created/defined, the next step is to execute backups. Backups can be automated periodically via a schedule or performed manually on a demand basis.

6.1.1 Scheduled backup

The schedule for automated backups is configured in the policy. Some of the basic settings include frequency (hours/days/weeks), time and date (when), level of backup (full or incremental) and retention period (weeks/months/years).

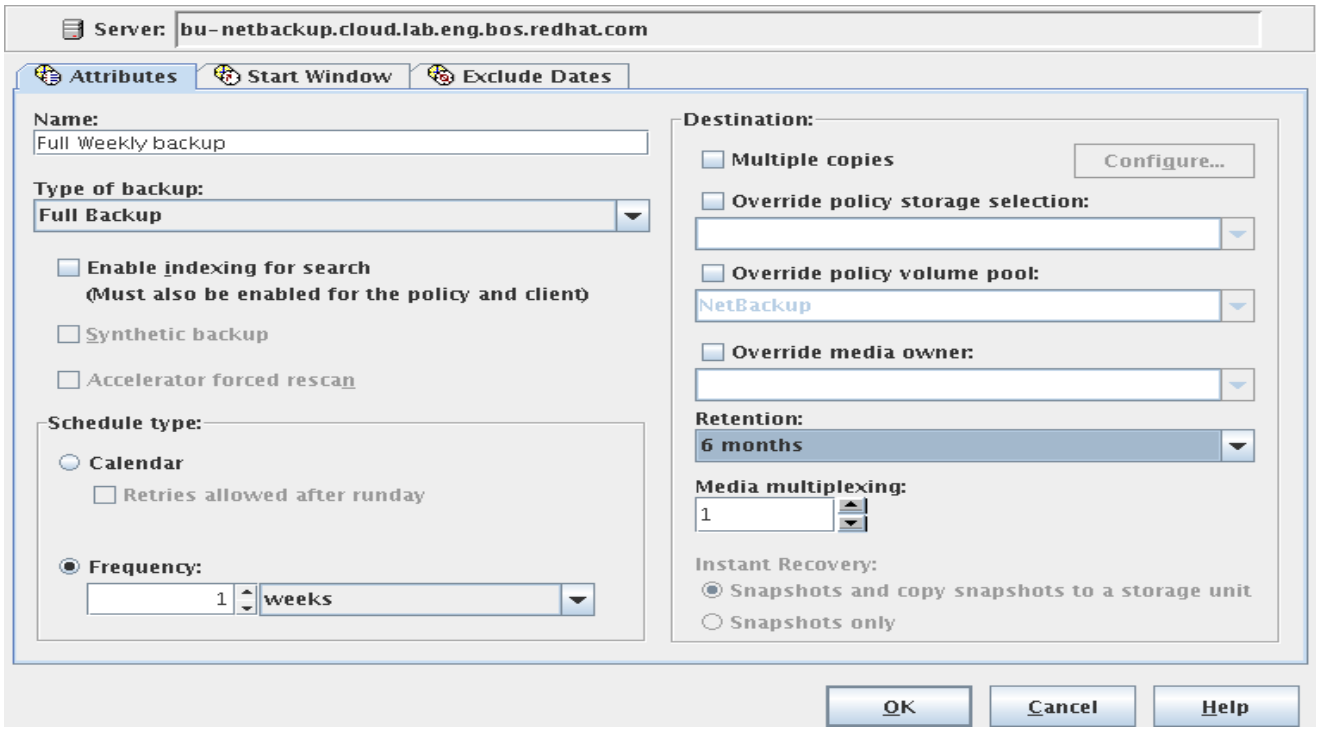

Figure 6.1.1-1: NetBackup - Backup Schedule

6.1.2 Manual Backup

It is possible to initiate manual backups at any required time. Manual backups are ideal prior to any maintenance activity. For Example -

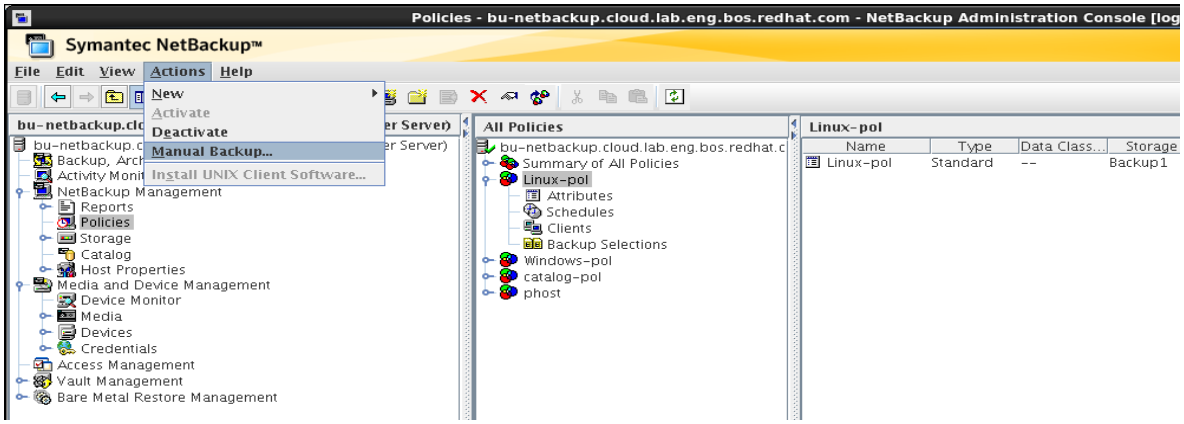

Figure 6.1.2-1: NetBackup - Manual Backup

Select backup specific information- Full/Incremental and the client:

Figure 6.1.2-2: NetBackup – Backup Specifics

Once the backup specifics like ‘Full or Differential-Incremental’ and the client for backup have been identified and OK was selected to proceed, the job status can be viewed in the Activity Monitor.

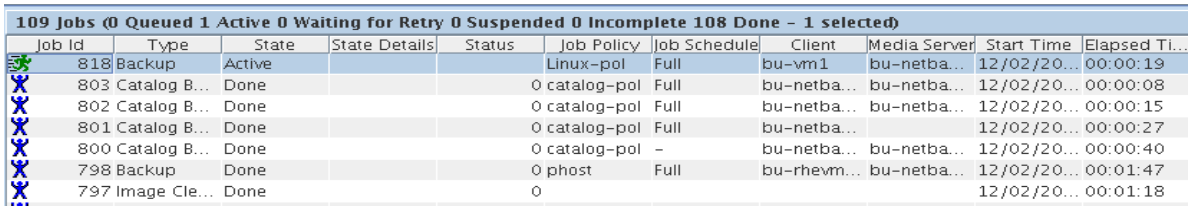

Figure 6.1.2-3: NetBackup – Backup Activity Monitor

Double clicking on the selected job provides detailed status of the job as follows:

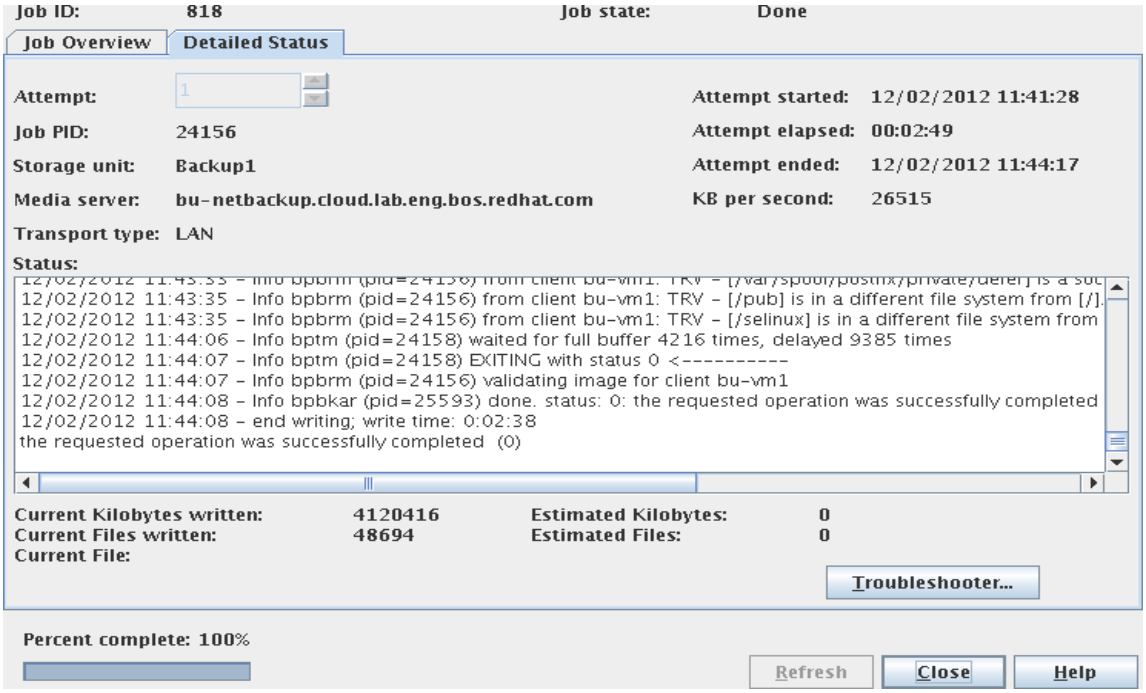

Figure 6.1.2-4: NetBackup – Backup Job Status

6.2 Virtual Machine Recovery

6.2.1 Linux Virtual Machine - file level recovery

This section describes steps to restore one or more files or directories from a previous backup. For this scenario, an Apache webserver application running on a client experiences an outage. Someone accidentally deleted key files under directory '***/var/www/webcontent***' which houses the '***index.html*** 'file and other files that contribute to the website.

The backup administrator initiates the following steps to restore the application:

1) Select a recovery point for the affected application.
2) Choose the files and their location that are to be restored.
3) Initiate file recovery.
4) Verify completion and bring the application back on line.

Source and Destination of Restore

The restore is initiated by the following menu:
Backup, Archive, and Restore --> Actions --> Specify NetBackup Machines and Policy Type

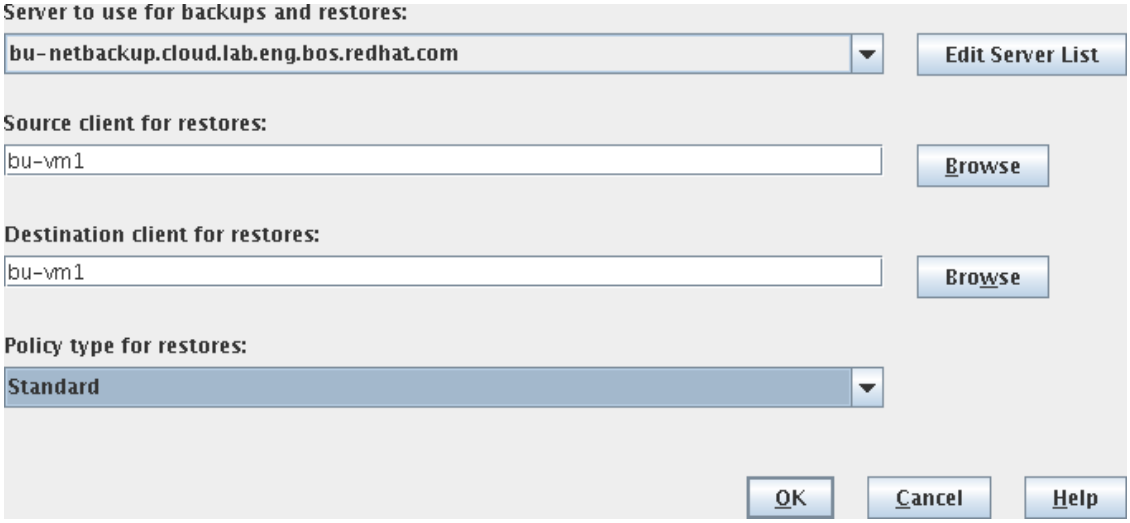

Figure 6.2.1-1: NetBackup – Restore Source & Destination

Recovery point

This menu helps select the recovery point by looking at the backup history and selecting the appropriate backup to restore from.

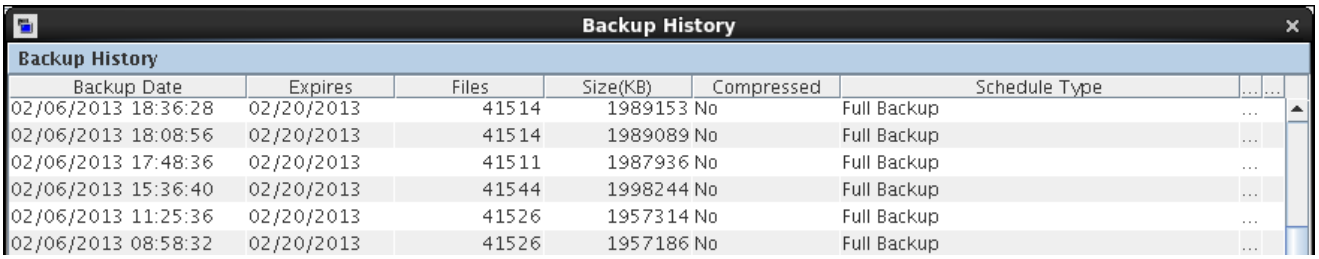

Figure 6.2.1-2: Selecting Recovery Point from History

File selection

This helps select the right file(s), directory or filesystem to restore

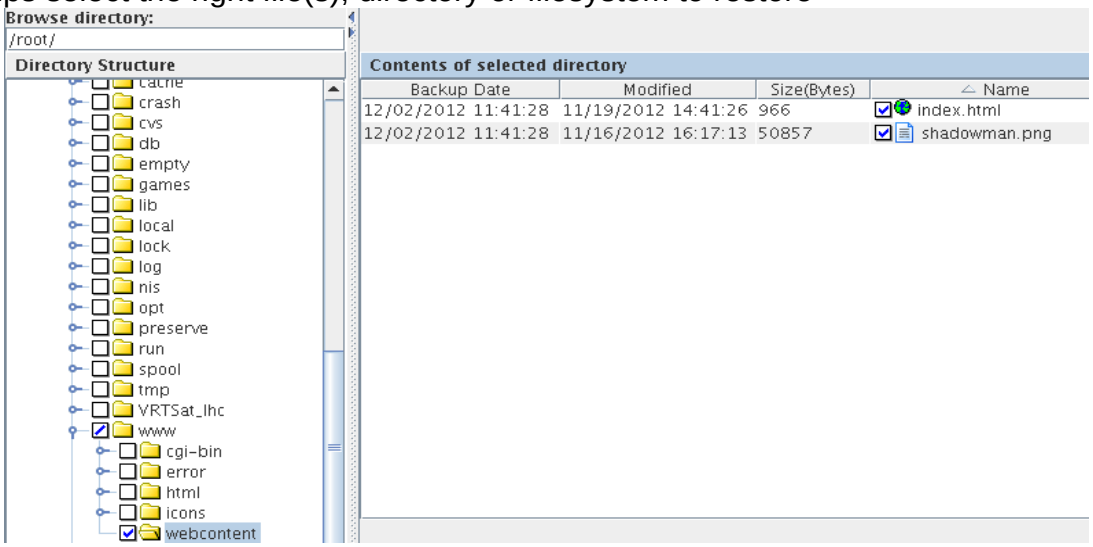

Figure 6.2.1-3: NetBackup – Restore File(s) selection

File Recovery

Restore options specified to indicate the file needs to be restored to the original location. If a file already exist, Overwrite existing files can be selected.

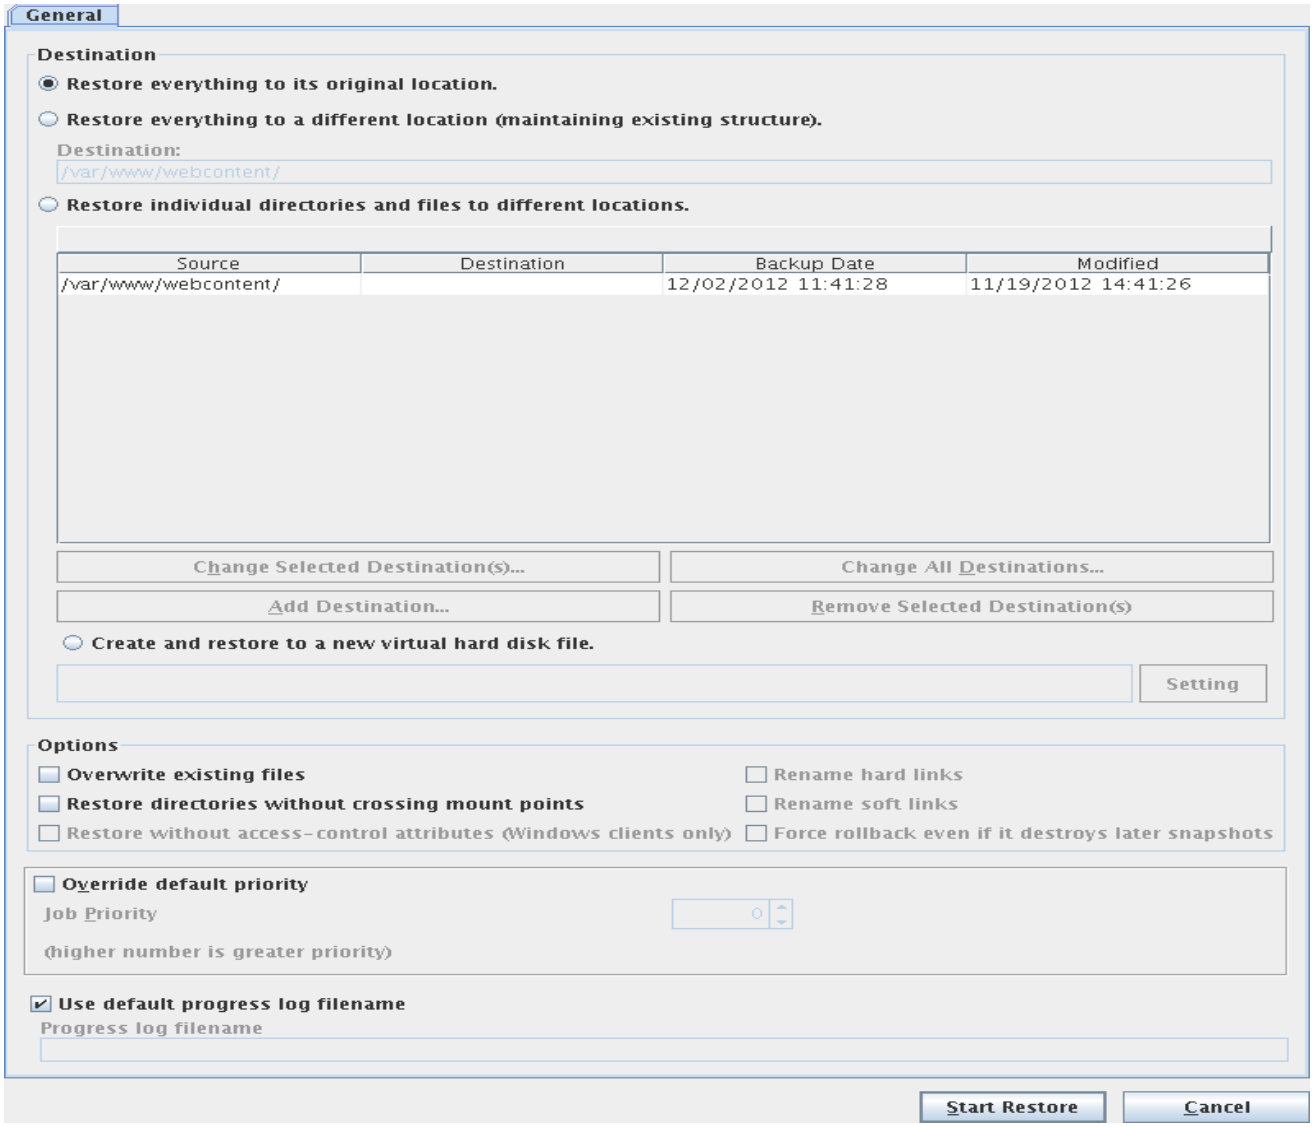

Figure 6.2.1-4: NetBackup – Restore specifics

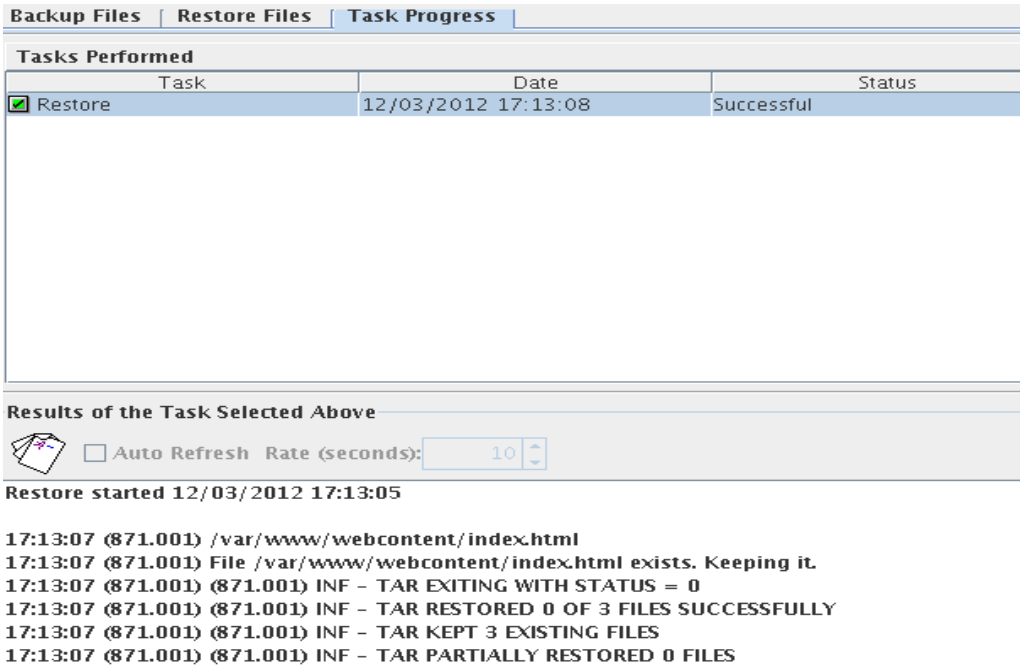

Figure 6.2.1-5: NetBackup – Restore status

Recovery Completion

The Task progress indicated that the restore was successful.

Verify Functionality and restart the application

Before restore the application (Website in this case) was not available.

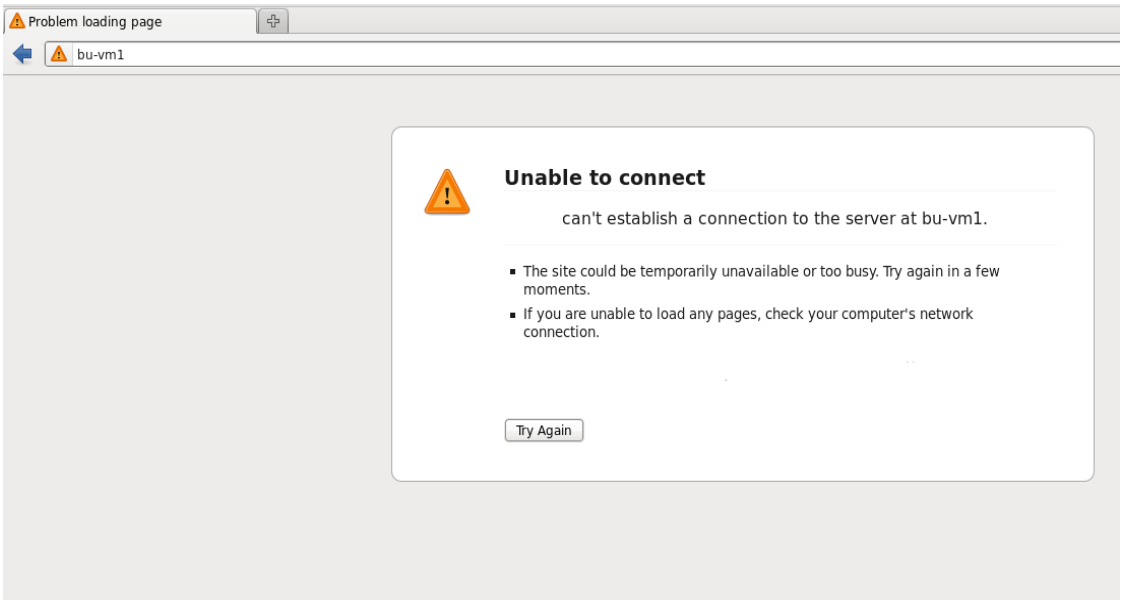

Figure 6.2.1-6: NetBackup – Application unavailable before restore

After restore was completed, the file was restored to the original location. Restarting apache rendered the web page to be available again:

```
# ls -all /var/www/webcontent
total 64
drwxr-xr-x. 2 root root  4096 Nov 19 14:41 .
drwxr-xr-x. 7 root root  4096 Nov 16 16:08 ..
-rw-r--r--. 1 root root   966 Nov 19 14:41 index.html
-rw-rw-r--. 1 root root 50857 Nov 16 16:17 shadowman.png
#
# service httpd stop
Stopping httpd:                                            [  OK  ]
#
# service httpd start
Starting httpd:                                            [  OK  ]
```

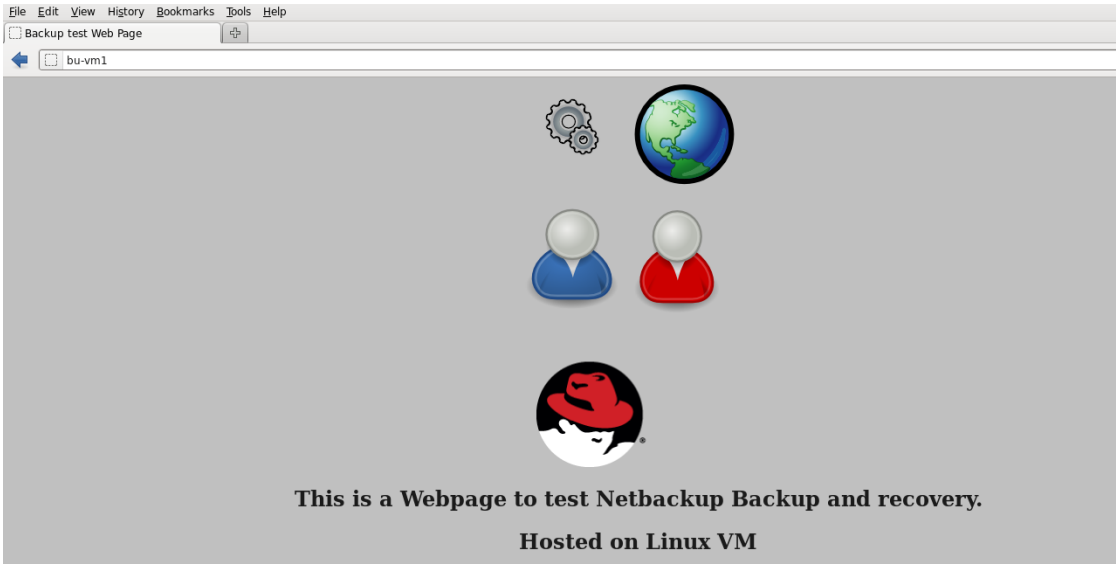

Figure 6.2.1-7: NetBackup - Application recovered after restore

6.2.2 Linux Virtual Machine - full VM recovery

This situation is best described by a disaster event where the virtual machine bu-vm2 is not responding and inaccessible via console. In order to restore the system, following steps were followed:

1) Create a VM template.
2) Gather configuration details: (previously collected when the VM was operational)
   - VM details.
   - Disk and swap partition details.
3) Create a destination VM based on this template with the gathered configuration
4) Boot up the destination VM and ensure network is functional.
5) Add a new disk to the VM with the same disk size as the original VM.
6) Partition the disk matching the original configuration
7) Create a volume group that contains root and possibly swap logical volumes
8) Mount root volume under ***/restore*** directory
9) Mount boot partition under ***/restore/boot***
10) Restore the original VM image using NetBackup
11) Edit ***grub.conf*** on the recovery disk to include the new disk boot parameters.
12) create swap area using **mkswap** the new swap partition
13) Identify UUID of the new boot partition
14) Update ***/restore/etc/fstab*** with the new UUID of boot partition and swap partition (If swap is a disk partition)
15) Boot off the destination disk
16) Update device map and boot loader.
17) Update ***/restore/boot/grub/grub.conf*** file to refer to disk hd1 instead of hd0.
18) Set destination disk bootable in the RHEV admin portal
19) Verification
20) Cleanup

The following terms were used for reference:
VM that got corrupted and needed recovery – **bu-vm2** referred to as original VM
VM created to recover the server – **bu-dest** referred to as Destination VM
Temporary boot disk referred to as recovery disk
Final boot disk referred to as destination disk

Create a VM template

This helps in provisioning a VM with preconfigured settings. This is a one time effort that can be reused for all future requirements. This will require creating a VM preferably with a unique Volume Group for the boot disk, “nb_recovery_vg” in this case. This ensures that it is different from the Volume Group name used in the original VM. Also it would help if NetBackup client is installed on it to save the effort of installing NetBackup client on new Vms created with the template. After this VM is ready, a template can be created as described in this link:
https://access.redhat.com/knowledge/docs/en-US/Red_Hat_Enterprise_Virtualization/3.0/html/Evaluation_Guide/Evaluation_Guide-Create_RHEL_Template.html
The template created in this paper was “nb_restore”

Gather Configuration details

As a pre-requisite the following details were collected on the original VM on a periodical basis prior to the event and saved at a remote location:

- VM Resource details- Virtual CPUs,Memory, disk size, and mac address of the failed VM

- Boot Volume Group name. This is to ensure the same volume group is retained if the VM is rebuilt

```
# vgs
  VG   #PV #LV #SN Attr   VSize VFree
  myvg   1   1   0 wz--n- 8.81g    0
```

- Boot disk.

```
# pvs
  PV         VG   Fmt  Attr PSize PFree
  /dev/vda3  myvg lvm2 a--  8.81g    0
```

- Boot disk partition layout

```
# fdisk /dev/vda

WARNING: DOS-compatible mode is deprecated. It's strongly recommended to
         switch off the mode (command 'c') and change display units to
         sectors (command 'u').

Command (m for help): p
Disk /dev/vda: 10.7 GB, 10737418240 bytes
16 heads, 63 sectors/track, 20805 cylinders
Units = cylinders of 1008 * 512 = 516096 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 512 bytes / 512 bytes
```

```
Disk identifier: 0x00083a3e

   Device Boot      Start         End      Blocks   Id  System
/dev/vda1   *           3         409      204800   83  Linux
Partition 1 does not end on cylinder boundary.
/dev/vda2             409        2441     1024000   82  Linux swap /
Solaris
Partition 2 does not end on cylinder boundary.
/dev/vda3            2441       20806     9255936   8e  Linux LVM
Partition 3 does not end on cylinder boundary.
```

Create destination VM

A new destination VM was created using “nb_restore” template and original VM configurations.

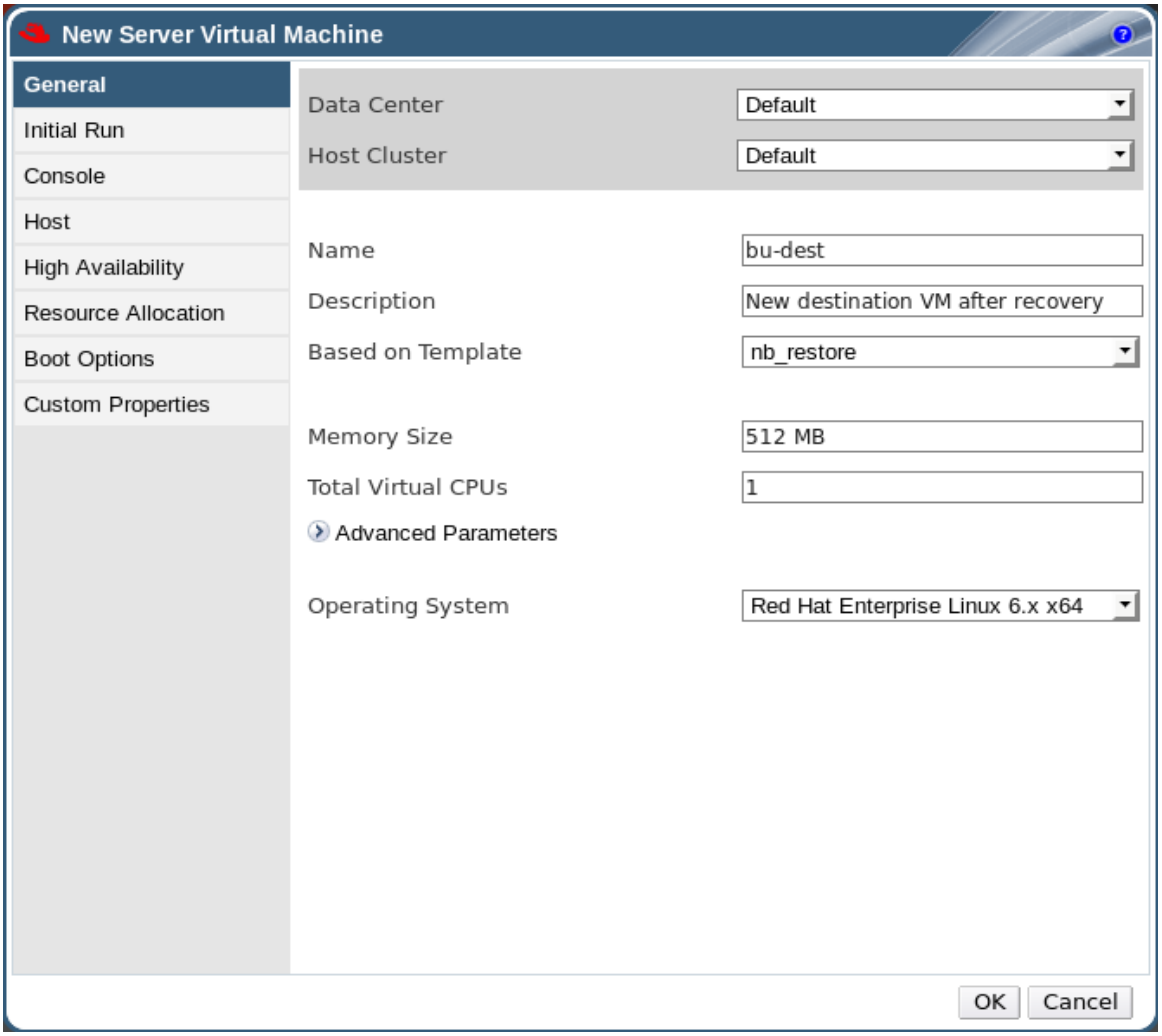

Figure 6.2.2-1: Create New Destination VM

Change the mac address on the original VM to an unused address. This step is required to reuse its mac address on the new VM. Edit the Mac address on the new VM to add the original Mac address.

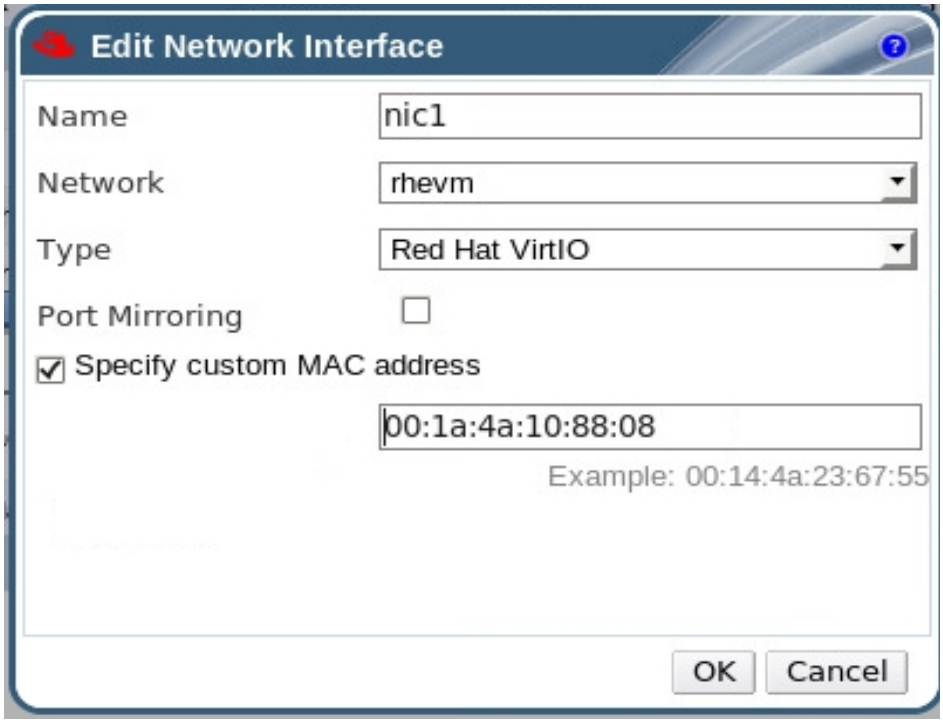

Figure 6.2.2-2: Specify Mac Address

Boot the Destination VM

Boot up the destination VM ***bu-dest***. Verify network connectivity. It is possible that the network device may have changed from eth0 to eth1. The network configuration must be changed to gain connectivity.

Verify the currently boot volume group is "***nb_recovery***"

```
# vgs
  VG             #PV #LV #SN Attr   VSize VFree
  vg_nb_recovery   1   2   0 wz--n- 9.51g    0

# lvs
  LV      VG             Attr     LSize Pool Origin Data%  Move Log Copy%
Convert
  lv_root vg_nb_recovery -wi-ao-- 7.54g
  lv_swap vg_nb_recovery -wi-ao-- 1.97g
```

Add a new disk to the VM

Add a new disk to match with boot disk on "***bu-vm2"*** (10GB):

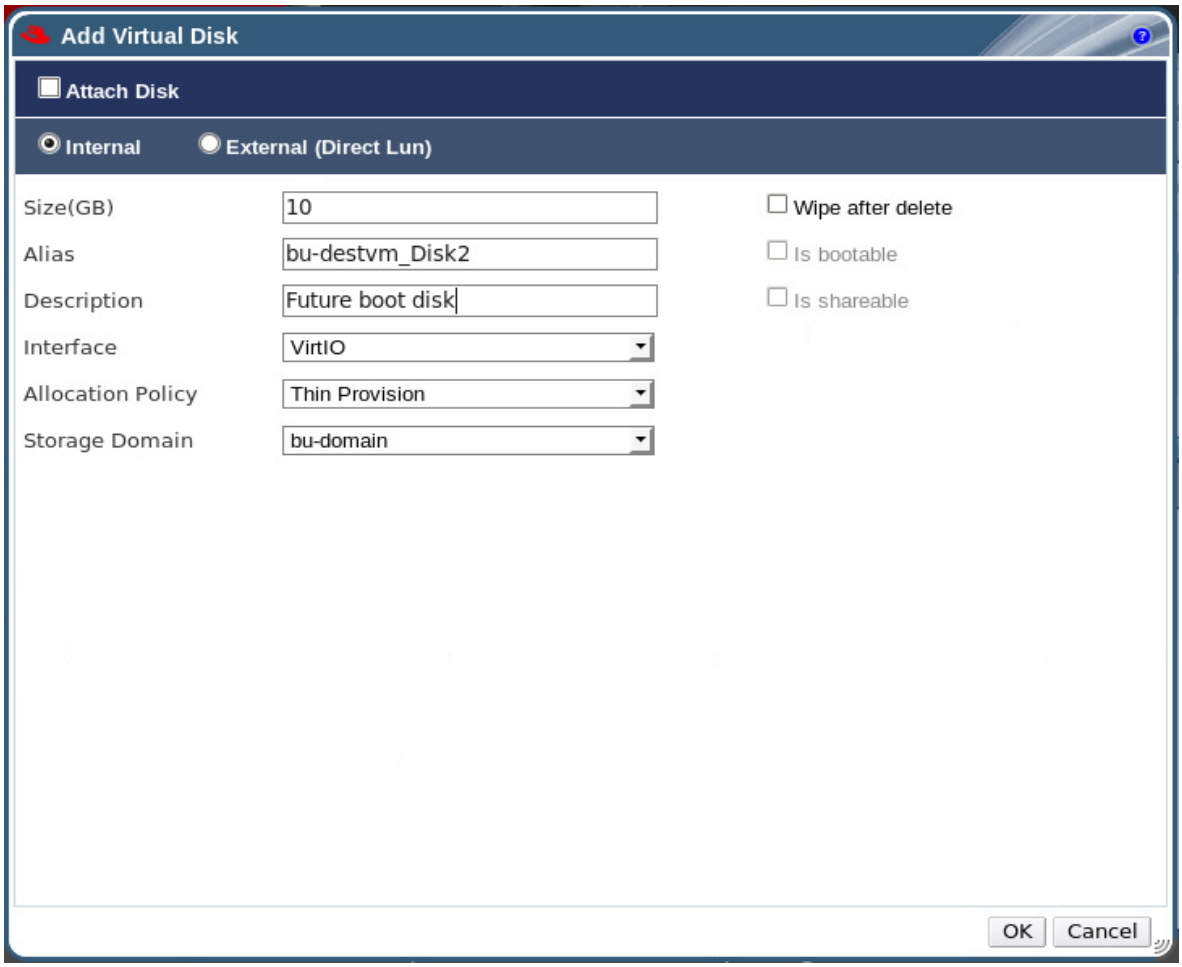

Figure 6.2.2-3: Add New Disk

Activate it:

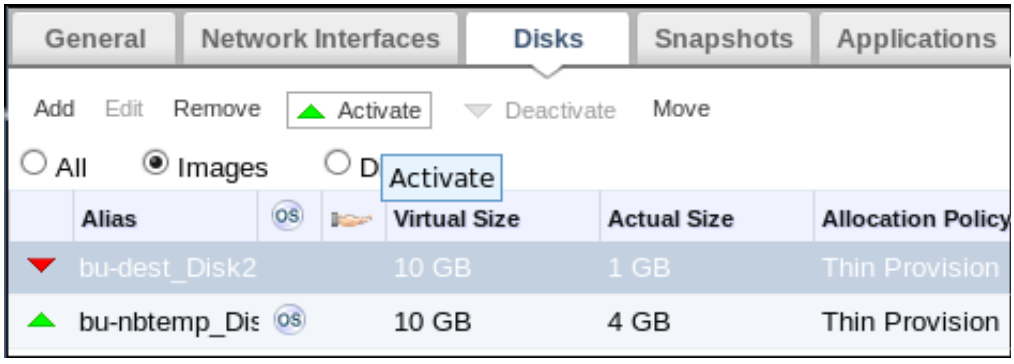

Figure 6.2.2-4: Disk Activation

Discover the disk in the Virtual Machine:

```
# fdisk -l

Disk /dev/vda: 10.7 GB, 10737418240 bytes
16 heads, 63 sectors/track, 20805 cylinders
Units = cylinders of 1008 * 512 = 516096 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 512 bytes / 512 bytes
Disk identifier: 0x000c07fe

   Device Boot      Start         End      Blocks   Id  System
/dev/vda1   *           3        1018      512000   83  Linux
Partition 1 does not end on cylinder boundary.
/dev/vda2            1018       20806     9972736   8e  Linux LVM
Partition 2 does not end on cylinder boundary.

Disk /dev/mapper/VolGroup-lv_root: 9168 MB, 9168748544 bytes
255 heads, 63 sectors/track, 1114 cylinders
Units = cylinders of 16065 * 512 = 8225280 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 512 bytes / 512 bytes
Disk identifier: 0x00000000


Disk /dev/mapper/VolGroup-lv_swap: 1040 MB, 1040187392 bytes
255 heads, 63 sectors/track, 126 cylinders
Units = cylinders of 16065 * 512 = 8225280 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 512 bytes / 512 bytes
Disk identifier: 0x00000000


Disk /dev/vdb: 10.7 GB, 10737418240 bytes
16 heads, 63 sectors/track, 20805 cylinders
Units = cylinders of 1008 * 512 = 516096 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 512 bytes / 512 bytes
Disk identifier: 0x00000000
```

Partition the new disk

Partition the new disk "***/dev/vdb***" matching the original configuration using **fdisk** utility:.

```
# fdisk /dev/vdb
Device contains neither a valid DOS partition table, nor Sun, SGI or OSF
disklabel
Building a new DOS disklabel with disk identifier 0x2d6ec9a2.
Changes will remain in memory only, until you decide to write them.
After that, of course, the previous content won't be recoverable.

Warning: invalid flag 0x0000 of partition table 4 will be corrected by
w(rite)
WARNING: DOS-compatible mode is deprecated. It's strongly recommended to
```

```
         switch off the mode (command 'c') and change display units to
         sectors (command 'u').

Command (m for help): p

Disk /dev/vdb: 10.7 GB, 10737418240 bytes
16 heads, 63 sectors/track, 20805 cylinders
Units = cylinders of 1008 * 512 = 516096 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 512 bytes / 512 bytes
Disk identifier: 0x2d6ec9a2

   Device Boot      Start         End      Blocks   Id  System

Command (m for help): m
Command action
   a   toggle a bootable flag
   b   edit bsd disklabel
   c   toggle the dos compatibility flag
   d   delete a partition
   l   list known partition types
   m   print this menu
   n   add a new partition
   o   create a new empty DOS partition table
   p   print the partition table
   q   quit without saving changes
   s   create a new empty Sun disklabel
   t   change a partition's system id
   u   change display/entry units
   v   verify the partition table
   w   write table to disk and exit
   x   extra functionality (experts only)

Command (m for help): n
Command action
   e   extended
   p   primary partition (1-4) p
Partition number (1-4): 1
First cylinder (1-20805, default 1): 3
Last cylinder, +cylinders or +size{K,M,G} (3-20805, default 20805): 409

Command (m for help): n
Command action
   e   extended
   p   primary partition (1-4) p
Partition number (1-4): 2
First cylinder (1-20805, default 1): 409
Sector 411264 is already allocated
First cylinder (410-20805, default 410): 410
Last cylinder, +cylinders or +size{K,M,G} (410-20805, default 20805): 2442

Command (m for help): n
Command action
   e   extended
   p   primary partition (1-4) p
```

```
Partition number (1-4): 3
First cylinder (1-20805, default 1): 2443
Last cylinder, +cylinders or +size{K,M,G} (2443-20805, default 20805):
100%FREE
Value out of range.
Last cylinder, +cylinders or +size{K,M,G} (2443-20805, default 20805): Last
Cylinder
Last cylinder, +cylinders or +size{K,M,G} (2443-20805, default 20805):
20805

Command (m for help): p

Disk /dev/vdb: 10.7 GB, 10737418240 bytes
16 heads, 63 sectors/track, 20805 cylinders
Units = cylinders of 1008 * 512 = 516096 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 512 bytes / 512 bytes
Disk identifier: 0x2d6ec9a2

   Device Boot      Start         End      Blocks   Id  System
/dev/vdb1               3         409      205128   83  Linux
/dev/vdb2             410        2442     1024632   83  Linux
/dev/vdb3            2443       20805     9254952   83  Linux

Command (m for help): a
Partition number (1-4): 1

Command (m for help): p

Disk /dev/vdb: 10.7 GB, 10737418240 bytes
16 heads, 63 sectors/track, 20805 cylinders
Units = cylinders of 1008 * 512 = 516096 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 512 bytes / 512 bytes
Disk identifier: 0x2d6ec9a2

   Device Boot      Start         End      Blocks   Id  System
/dev/vdb1   *           3         409      205128   83  Linux
/dev/vdb2             410        2442     1024632   83  Linux
/dev/vdb3            2443       20805     9254952   83  Linux

Command (m for help): t
Partition number (1-4): 1
Hex code (type L to list codes): 83

Command (m for help): p

Disk /dev/vdb: 10.7 GB, 10737418240 bytes
16 heads, 63 sectors/track, 20805 cylinders
Units = cylinders of 1008 * 512 = 516096 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 512 bytes / 512 bytes
Disk identifier: 0x2d6ec9a2

   Device Boot      Start         End      Blocks   Id  System
```

```
/dev/vdb1   *           3         409      205128   83  Linux
/dev/vdb2             410        2442     1024632   83  Linux
/dev/vdb3            2443       20805     9254952   83  Linux

Command (m for help): t
Partition number (1-4): 2
Hex code (type L to list codes): 82
Changed system type of partition 2 to 82 (Linux swap / Solaris)

Command (m for help): t
Partition number (1-4): 3
Hex code (type L to list codes): 8e
Changed system type of partition 3 to 8e (Linux LVM)

Command (m for help): p

Disk /dev/vdb: 10.7 GB, 10737418240 bytes
16 heads, 63 sectors/track, 20805 cylinders
Units = cylinders of 1008 * 512 = 516096 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 512 bytes / 512 bytes
Disk identifier: 0x2d6ec9a2

   Device Boot      Start         End      Blocks   Id  System
/dev/vdb1   *           3         409      205128   83  Linux
/dev/vdb2             410        2442     1024632   82  Linux swap /
Solaris
/dev/vdb3            2443       20805     9254952   8e  Linux LVM

Command (m for help): w
The partition table has been altered!

Calling ioctl() to re-read partition table.
Syncing disks.
[root@bu-vm2 ~]# fdisk /dev/vdb

WARNING: DOS-compatible mode is deprecated. It's strongly recommended to
         switch off the mode (command 'c') and change display units to
         sectors (command 'u').

Command (m for help): p

Disk /dev/vdb: 10.7 GB, 10737418240 bytes
16 heads, 63 sectors/track, 20805 cylinders
Units = cylinders of 1008 * 512 = 516096 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 512 bytes / 512 bytes
Disk identifier: 0x2d6ec9a2

   Device Boot      Start         End      Blocks   Id  System
/dev/vdb1   *           3         409      205128   83  Linux
/dev/vdb2             410        2442     1024632   82  Linux swap /
Solaris
/dev/vdb3            2443       20805     9254952   8e  Linux LVM
Command (m for help): q
```

Create a volume group

Perform the following steps to create the volume group, root logical volume and the swap logical volume if required.

Initialize the disk partition prior to volume group creation using "***pvcreate***" command:

```
# pvcreate /dev/vdb3
  Writing physical volume data to disk "/dev/vdb3"
  Physical volume "/dev/vdb3" successfully created
```

Create volume group using "vgcreate" command

```
# vgcreate myvg /dev/vdb3
```

Mount root filesystem

Create logical volume for root:

```
# lvcreate -n rootvol -l 100%FREE myvg
  Logical volume "rootvol" created
```

```
# lvs
  LV      VG   Attr     LSize Pool Origin Data%  Move Log Copy%  Convert
  lv_root myvg -wi-ao-- 8.51g
```

Format the logical volume with filesystem type matching with the original system's- "***ext4***" in this case:

```
# mkfs -t ext4 /dev/myvg/rootvol
mke2fs 1.41.12 (17-May-2010)
Filesystem label=
OS type: Linux
Block size=4096 (log=2)
Fragment size=4096 (log=2)
Stride=0 blocks, Stripe width=0 blocks
579360 inodes, 2313216 blocks
115660 blocks (5.00%) reserved for the super user
First data block=0
Maximum filesystem blocks=2369781760
71 block groups
32768 blocks per group, 32768 fragments per group
8160 inodes per group
Superblock backups stored on blocks:
      32768, 98304, 163840, 229376, 294912, 819200, 884736, 1605632
Writing inode tables: done
Creating journal (32768 blocks): done
Writing superblocks and filesystem accounting information: done

This filesystem will be automatically checked every 28 mounts or
180 days, whichever comes first.  Use tune2fs -c or -i to override.
```

Create /restore mountpoint:

```
# mkdir /restore
```

Mount root filesystem under "**/restore**" directory:

```
# mount /dev/myvg/lv_root /restore
```

Mount boot partition

Format the boot partition:

```
# mkfs -t ext3 /dev/vdb1
mke2fs 1.41.12 (17-May-2010)
warning: 327 blocks unused.
Filesystem label=
OS type: Linux
Block size=1024 (log=0)
Fragment size=1024 (log=0)
Stride=0 blocks, Stripe width=0 blocks
51400 inodes, 204801 blocks
10256 blocks (5.01%) reserved for the super user
First data block=1
Maximum filesystem blocks=67371008
25 block groups
8192 blocks per group, 8192 fragments per group
2056 inodes per group
Superblock backups stored on blocks:
      8193, 24577, 40961, 57345, 73729
Writing inode tables: done
Creating journal (4096 blocks): done
Writing superblocks and filesystem accounting information: done

This filesystem will be automatically checked every 25 mounts or
180 days, whichever comes first.  Use tune2fs -c or -i to override.
```

Create the mount point:

```
# mkdir /restore/boot
```

Mount the filesystem:

```
# mount /dev/vdb1 /restore/boot
```

Verify the newly mounted filesystems:

```
# df -h
Filesystem            Size  Used Avail Use% Mounted on
/dev/mapper/VolGroup-lv_root
                      8.5G  1.6G  6.5G  20% /
tmpfs                 246M     0  246M   0% /dev/shm
/dev/vda1             485M   32M  429M   7% /boot
/dev/mapper/myvg-rootvol
                      8.7G  148M  8.1G   2% /restore
/dev/vdb1             194M  5.6M  179M   4% /restore/boot
```

Restore the original VM image using NetBackup

Ensure NetBackup connectivity has been established between master and this VM client. Restore the backup data of the VM onto /restore partition of the recovery disk.

Restore the original VM image onto "**/restore"**directory using a previous backup of "***bu-vm2***" matching with the agreed RPO.

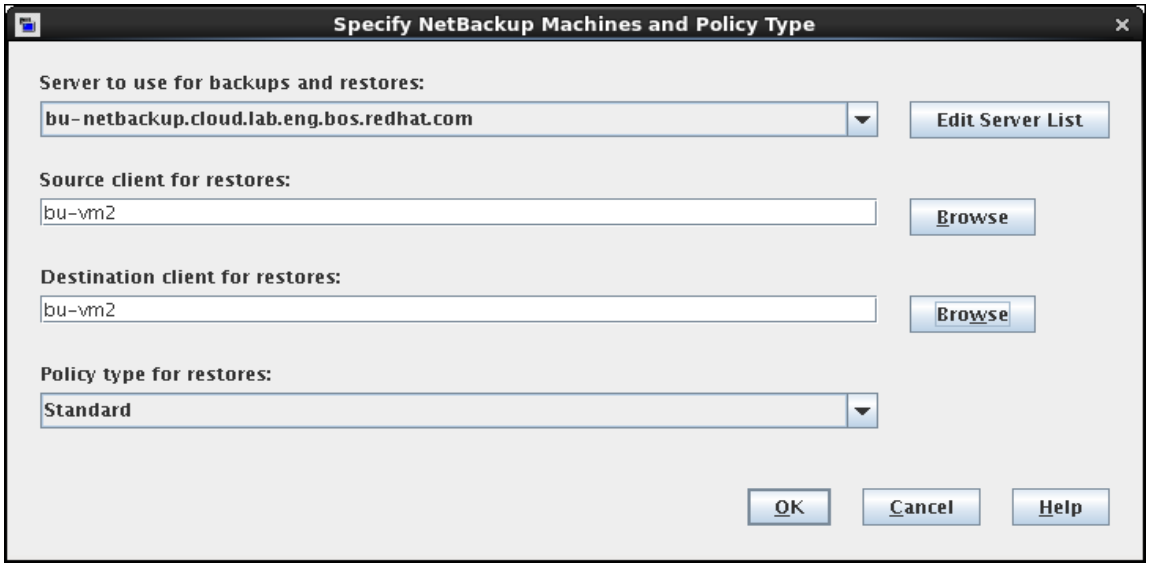

Figure 6.2.2-5: Data Restore- Source, Destination and Policy Type

Select entire root filesystem to be restored.

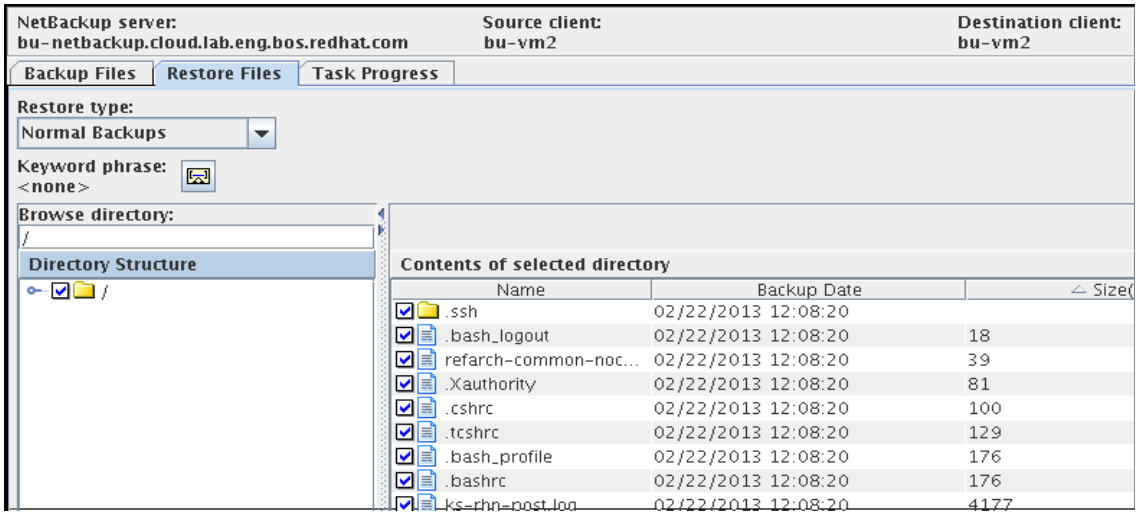

Figure 6.2.2-6: Data Restore - Selecting the contents to restore

Ensure the destination is /restore filesystem and not the original location:

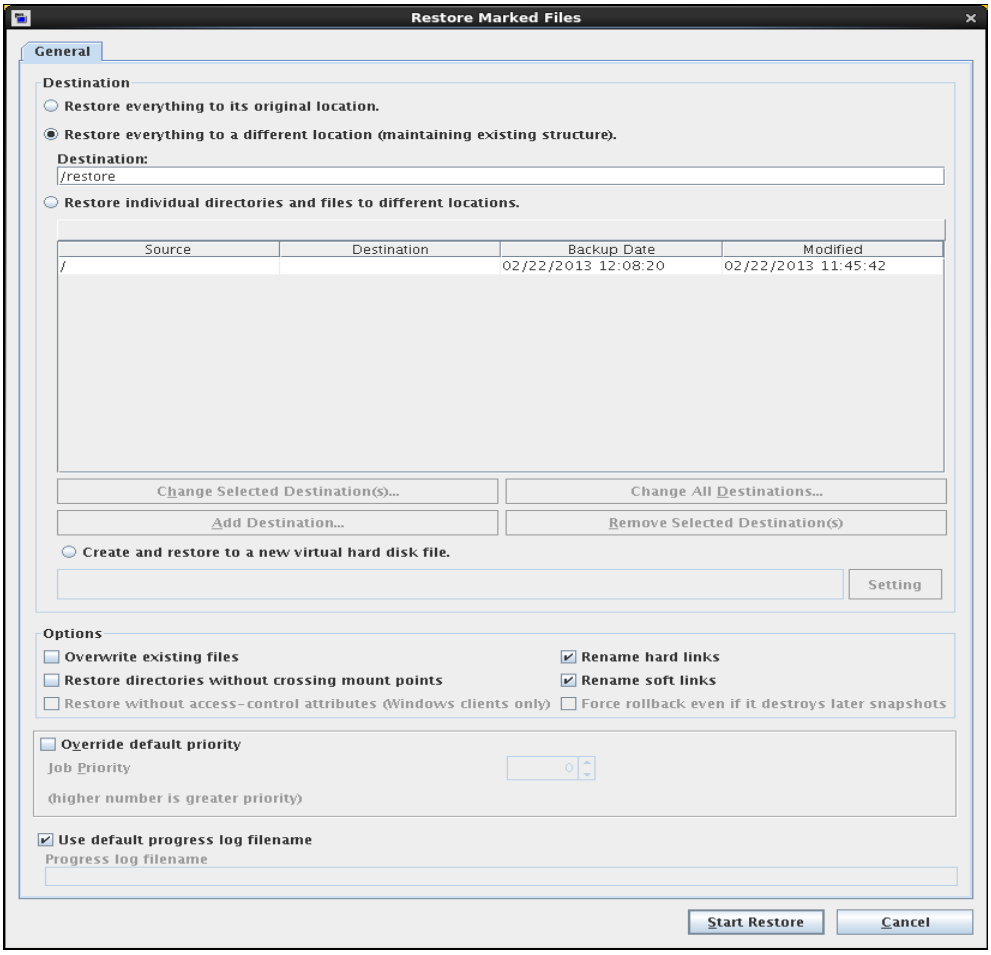

Figure 6.2.2-7: Restore to a Different Location

Verify the restore was successful:

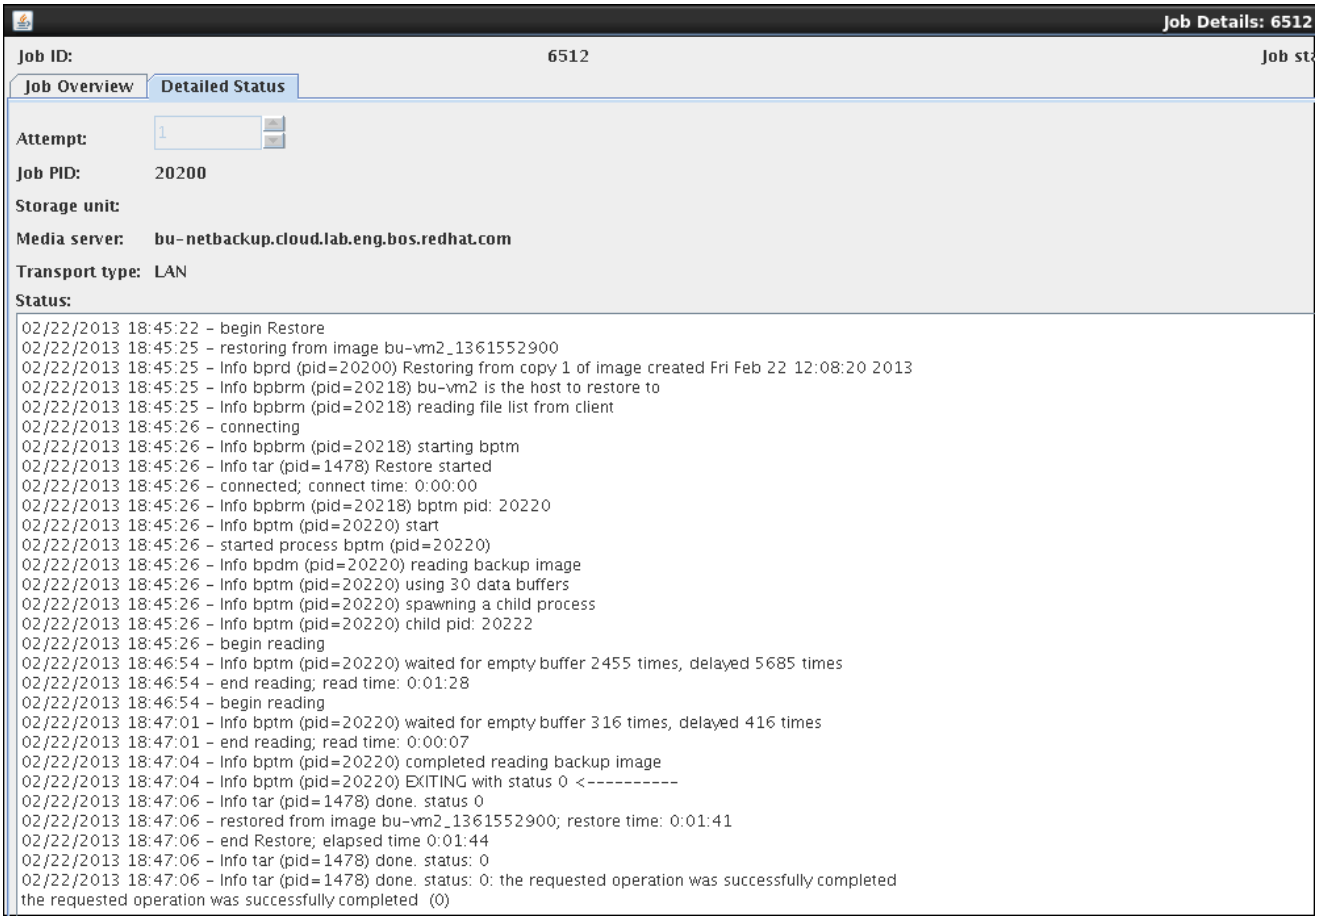

Figure 6.2.2-8: Data Restore- Job Completion Confirmation

Update grub.conf file in the recovery disk

Add the lines in the "***/boot/grub/grub.conf***" as highlighted below in bold. This update will help use the boot loader of the recovery disk but make it possible to boot the new disk during system boot.

```
# grub.conf generated by anaconda
#
# Note that you do not have to rerun grub after making changes to this file
# NOTICE:  You have a /boot partition.  This means that
#          all kernel and initrd paths are relative to /boot/, eg.
#          root (hd0,0)
#          kernel /vmlinuz-version ro root=/dev/mapper/vg_nb_recovery-
lv_root
#          initrd /initrd-[generic-]version.img
#boot=/dev/vda
default=0
timeout=5
```

```
splashimage=(hd0,0)/grub/splash.xpm.gz
hiddenmenu
title Red Hat Enterprise Linux (2.6.32-279.el6.x86_64)
        root (hd0,0)
        kernel /vmlinuz-2.6.32-279.el6.x86_64 ro
root=/dev/mapper/vg_nb_recovery-lv_root rd_NO_LUKS LANG=en_US.UTF-8
rd_LVM_LV=vg_nb_recovery/lv_root rd_NO_MD SYSFONT=latarcyrheb-sun16
crashkernel=auto rd_LVM_LV=vg_nb_recovery/lv_swap  KEYBOARDTYPE=pc
KEYTABLE=us rd_NO_DM rhgb quiet
        initrd /initramfs-2.6.32-279.el6.x86_64.img
title Red Hat Enterprise Linux - Destination Boot Disk
        root (hd1,0)
        kernel /vmlinuz-2.6.32-279.el6.x86_64 ro root=/dev/mapper/myvg-
lv_root rd_NO_LUKS LANG=en_US.UTF-8 rd_LVM_LV=myvg/lv_root rd_NO_MD
SYSFONT=latarcyrheb-sun16 crashkernel=auto rd_LVM_LV=myvg/lv_swap
KEYBOARDTYPE=pc KEYTABLE=us rd_NO_DM rhgb quiet selinux=0 2
        initrd /initramfs-2.6.32-279.el6.x86_64.img
```

Note: Please note that SELINUX is truned off with the setting "***selinux=0***". This is required to successfully boot the destination disk. However the "***2***" entry in the end, renders the system to be booted to run level 2 and does not bring up network. This protects the environment from a security point of view. SELINUX will be turned on when the destination disk is self bootable (if the original VM had SELINUX enabled)

Creating swap area on the swap device

This command will create swap area on the below device. If the swap partition is LVM based, then "***/dev/vdb2"*** will have to to replaced with "***/dev/myvg/lv_swap***".

```
# mkswap /dev/vdb2
Setting up swapspace version 1, size = 1024628 KiB
no label, UUID=a4beec84-94cc-4b65-a1cd-818c2aa0c0ec
```

Identity UUID of boot partition

Identify UUID of the boot partition "***/dev/vdb1***" and update the new value in the "***/restore/etc/fstab***" as highlighted below.

```
# blkid | grep vdb1
/dev/vdb1: UUID="05373287-d90c-429c-92a5-324ce7c684fb" TYPE="ext3"
```

Update /restore/etc/*fstab* file

Update the **/restore//etc/fstab** with UUID of boot and swap partitions gathered above:

```
# /etc/fstab
# Created by anaconda on Tue Feb 26 20:37:58 2013
# Accessible filesystems, by reference, are maintained under '/dev/disk'
# See man pages fstab(5), findfs(8), mount(8) and/or blkid(8) for more info
#
/dev/mapper/myvg-lv_root /                       ext4    defaults        1 1
UUID= 05373287-d90c-429c-92a5-324ce7c684fb /boot  ext4  defaults        1 2
UUID= a4beec84-94cc-4b65-a1cd-818c2aa0c0ec swap   swap  defaults        0 0
tmpfs                   /dev/shm                tmpfs   defaults        0 0
devpts                  /dev/pts                devpts  gid=5,mode=620  0 0
sysfs                   /sys                    sysfs   defaults        0 0
proc                    /proc                   proc    defaults        0 0
```

Booting off the destination disk

Shutdown the VM, interrupt the boot process and select the second boot disk and hit enter.

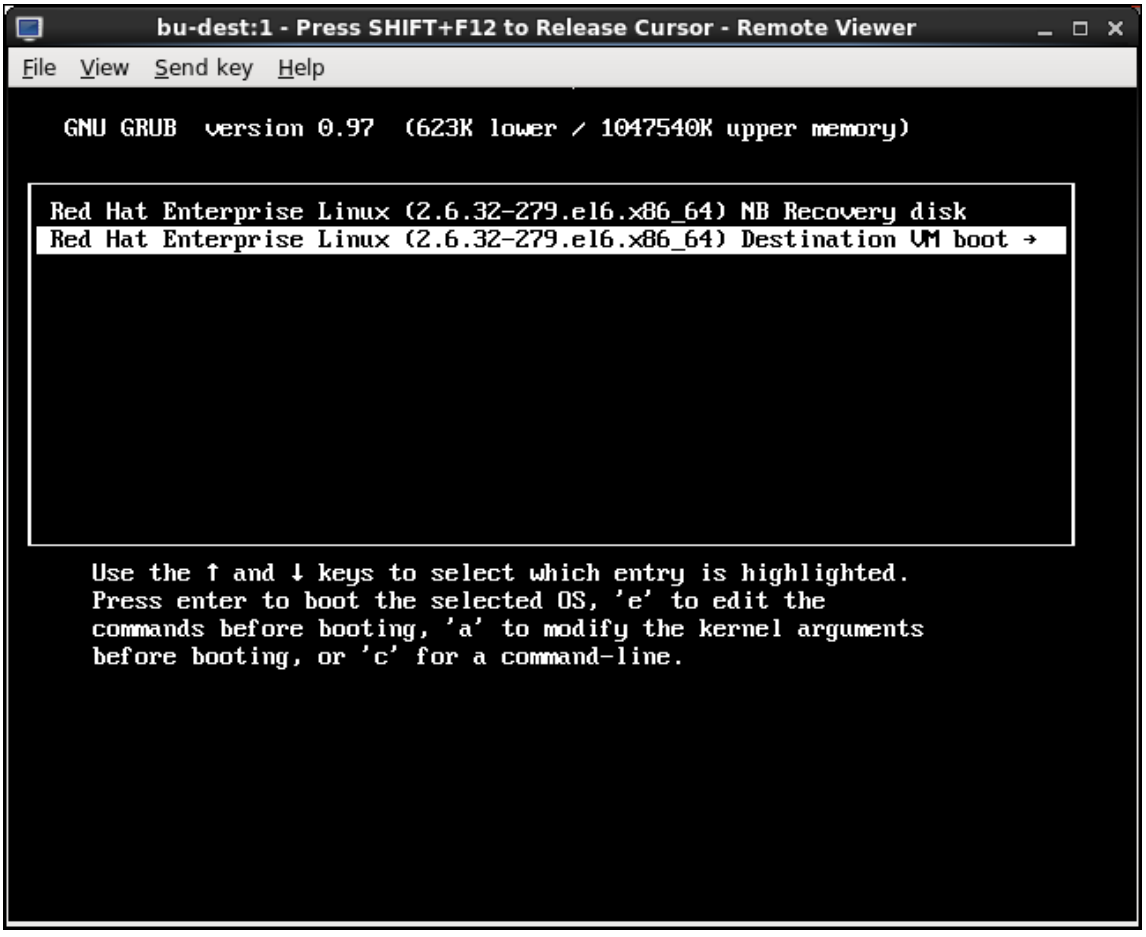

Figure 6.2.2-9: Selecting The Destination Disk for Booting

This brings up the system in run level 2. Hence the next three steps have to be performed in a console connection until the server is shutdown again to boot off the destination disk with the help of the recovery disk.

Updating device map and bootloader

Once the system successfully boots up using the destination disk, update the device map by appending the below entry to the "***/boot/grub/device.map***" file with the new device as mentioned below:

```
# echo "(hd1)     /dev/vdb" >> /boot/grub/device.map
```

/boot/grub/device.map

```
# this device map was generated by anaconda
(hd0)     /dev/vda
(hd1)     /dev/vdb
```

Update the bootloader on the destination disk using "***grub-install***" command.

```
# grub-install hd1
Installation finished. No error reported.
This is the contents of the device map /boot/grub/device.map.
Check if this is correct or not. If any of the lines is incorrect,
fix it and re-run the script `grub-install'.

# this device map was generated by anaconda
(hd0)     /dev/vda
(hd1)     /dev/vdb
```

Update grub.conf

Update the "***/boot/grub/grub.conf***" file in the destination disk, to reflect the new boot disk reference (disk hd1 instead of hd0)

```
# grub.conf generated by anaconda
#
# Note that you do not have to rerun grub after making changes to this file
# NOTICE:  You have a /boot partition.  This means that
#          all kernel and initrd paths are relative to /boot/, eg.
#          root (hd0,0)
#          kernel /vmlinuz-version ro root=/dev/mapper/myvg-lv_root
#          initrd /initrd-[generic-]version.img
#boot=/dev/vda
default=0
timeout=5
splashimage=(hd0,0)/grub/splash.xpm.gz
hiddenmenu
title Red Hat Enterprise Linux (2.6.32-279.el6.x86_64)
        root (hd1,0)
        kernel /vmlinuz-2.6.32-279.el6.x86_64 ro root=/dev/mapper/myvg-
lv_root rd_NO_LUKS LANG=en_US.UTF-8 rd_NO_MD SYSFONT=latarcyrheb-sun16
crashkernel=auto rd_LVM_LV=myvg/lv_swap  KEYBOARDTYPE=pc KEYTABLE=us
rd_NO_DM rd_LVM_LV=myvg/lv_root rhgb quiet
        initrd /initramfs-2.6.32-279.el6.x86_64.img
```

Set destination disk bootable

Shutdown the VM and switch bootable disk flag from recovery disk to the new disk using RHEV admin portal:

```
# shutdown -hy now

Broadcast message from root@bu-vm2
      (/dev/pts/0) at 9:37 ...

The system is going down for halt NOW!
```

Once the VM is down, switch the bootable parameter in the disks. This has to be performed in two steps. First the recovery disk must be edited to uncheck the bootable flag followed by checking the flag on the Destination disk.

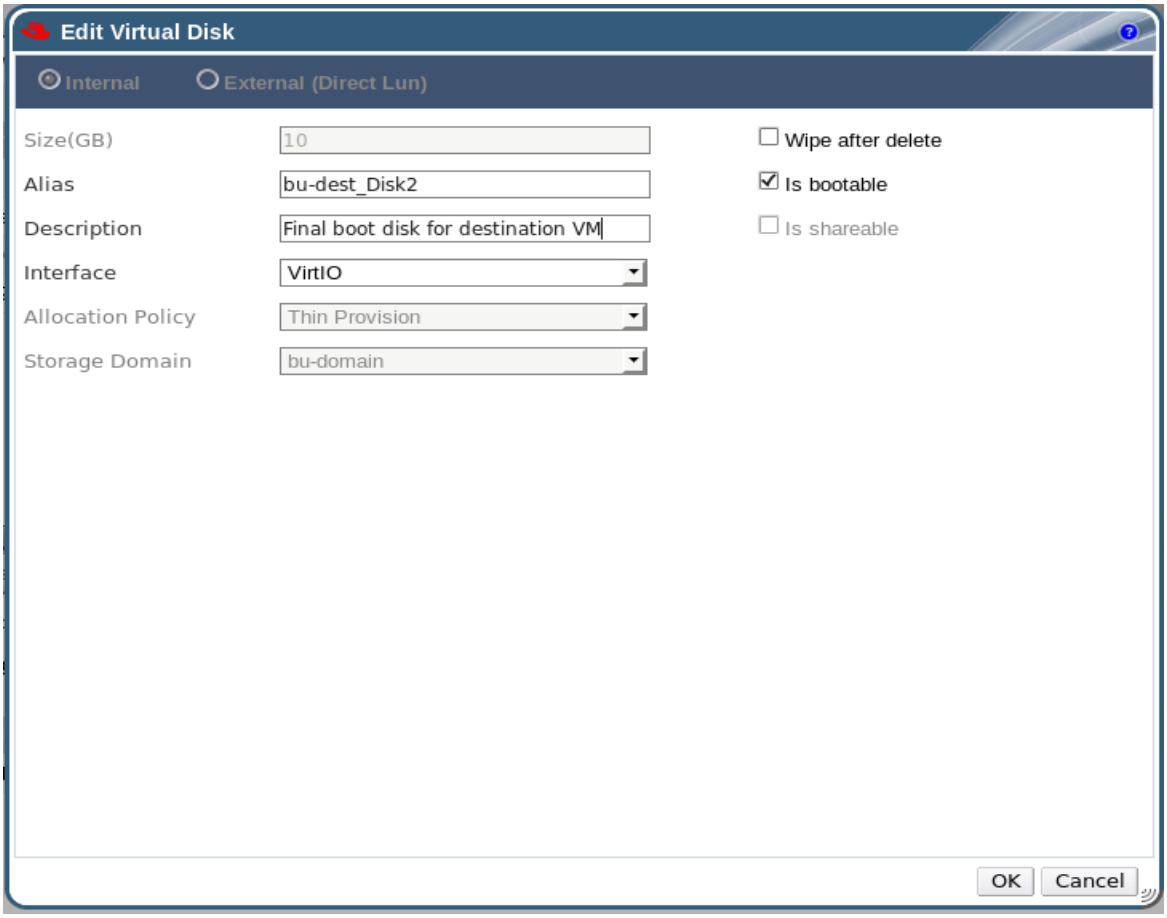

Figure 6.2.2-10: Setting the Destination Disk with Bootable Flag

Verification

Start the VM to verify if the system boots up independently with autoboot off the new disk.

Cleanup

The cleanup involves, removal of the recovery disk, changing the grub.conf disk reference, removal of the failed VM and updating the restored VM with the original VM Name.

Grub.conf update: After confirming that the server is fully functional, edit the grub.conf file and change the root disk reference from “(hd1,0) to (hd0,0).

Recovery disk removal: Shutdown the system and remove the recovery disk using the RHEV admin portal. The shutdown is necessary to ensure that the disk removal is clean. After disk removal, the server was started to run as the restored VM.
Updating the VM name: As a final step, remove the original failed VM, edit the new VM and update the Name to the original VM Name bu-vm2.

> Note: When the recovery disk is removed, the disk reference changes from “hd1” to “hd0” when the recovery disk is removed. This necessitates grub.conf update as mentioned above. Also the logical disk device name changes from /dev/vdb to /dev/vda. However LVM automatically updates the new disk device name in the volume group vgs and no action is necessary.

6.2.3 Microsoft Windows Virtual Machine - file level recovery

This section describes steps to restore one or more files or directories from a previous backup.In this scenario a user accidentally deletes the putty.exe file from the desktop location.

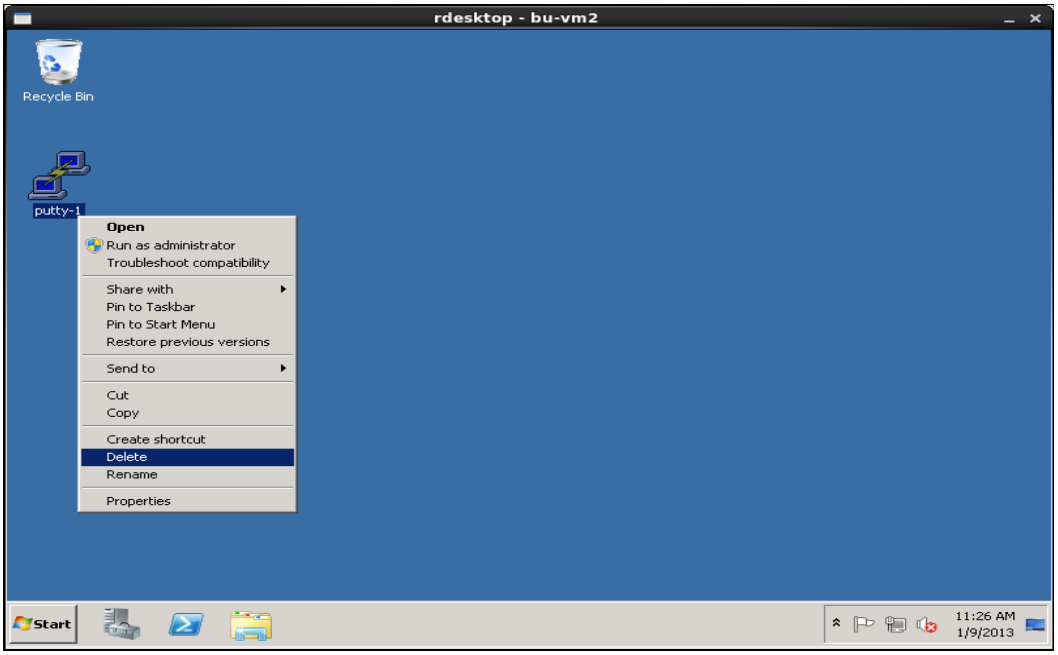

Figure 6.2.3-1: Windows File Level Restore 1

There is request from the user to restore the file.

The backup administrator initiates the following steps to restore the application:

1) Select a recovery point for the affected application.
2) Choose the files and their location that need to be restored.
3) Initiate file recovery.
4) Verify completion and confirmation of the restored file.

Source and Destination of Restore

The restore is initiated by the following menu:
Backup, Archive, and Restore --> Actions --> Specify NetBackup Machines and Policy Type

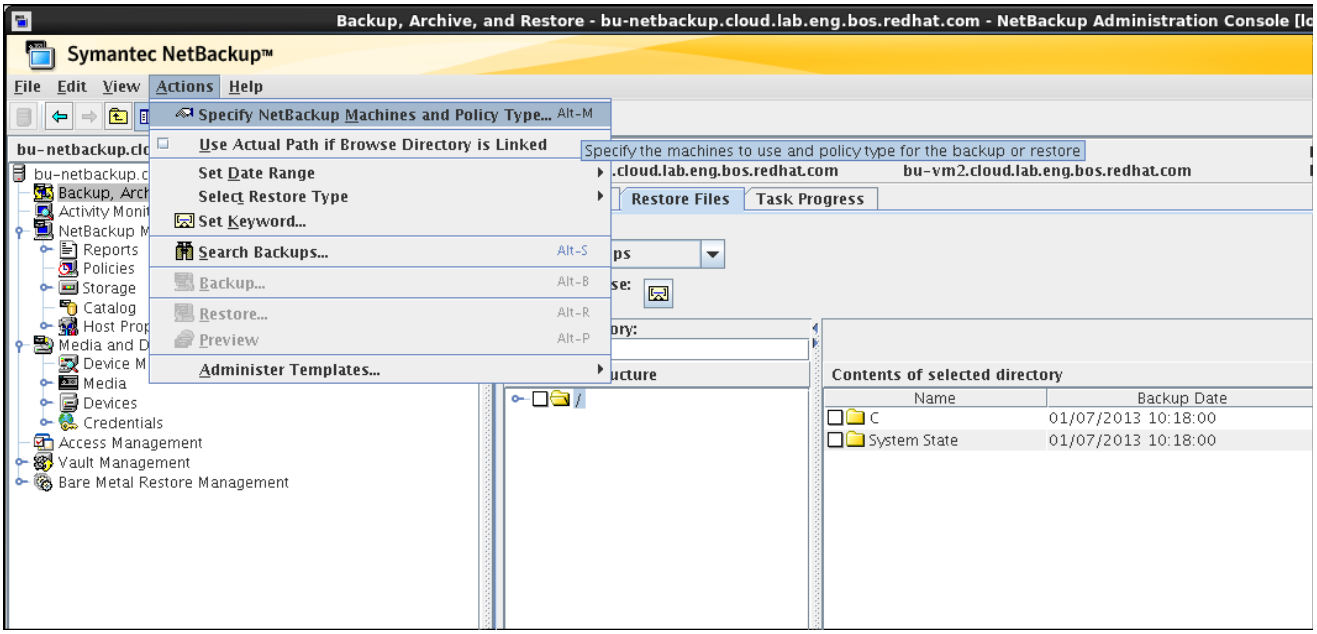

Figure 6.2.3-2: Windows File Level Restore 2

NetBackup Source, Destination and Policy Types are specified

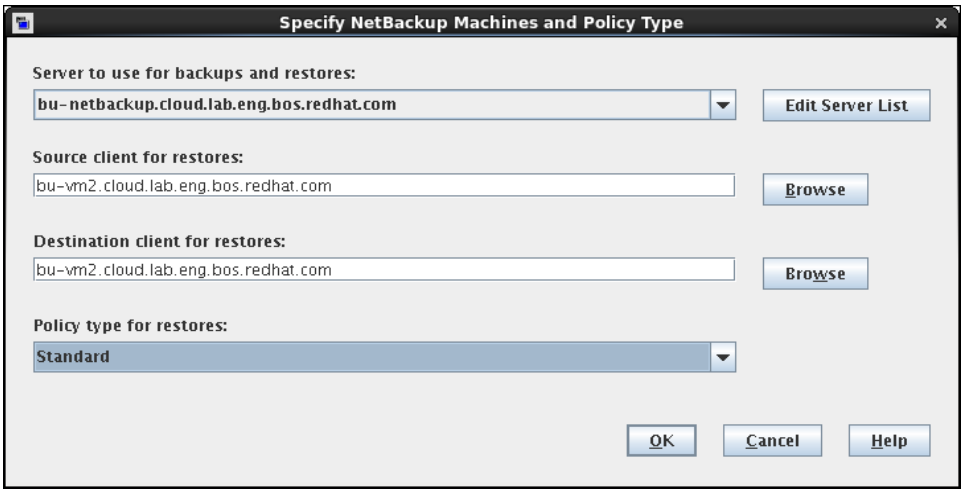

Figure 6.2.3-3: Windows File Level Restore-Source & Destination

Based on the information on when the file was deleted, the backup administrator selects the recovery point.

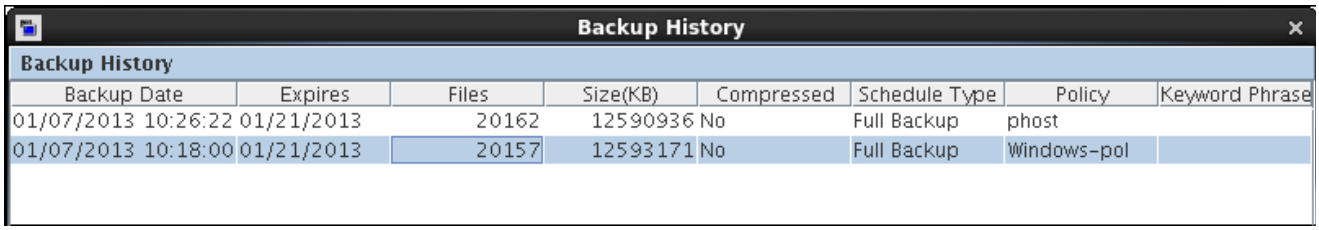

Figure 6.2.3-4: Windows File Level Restore – Recovery point

Selects the file '***putty.exe***'and begins the restore.

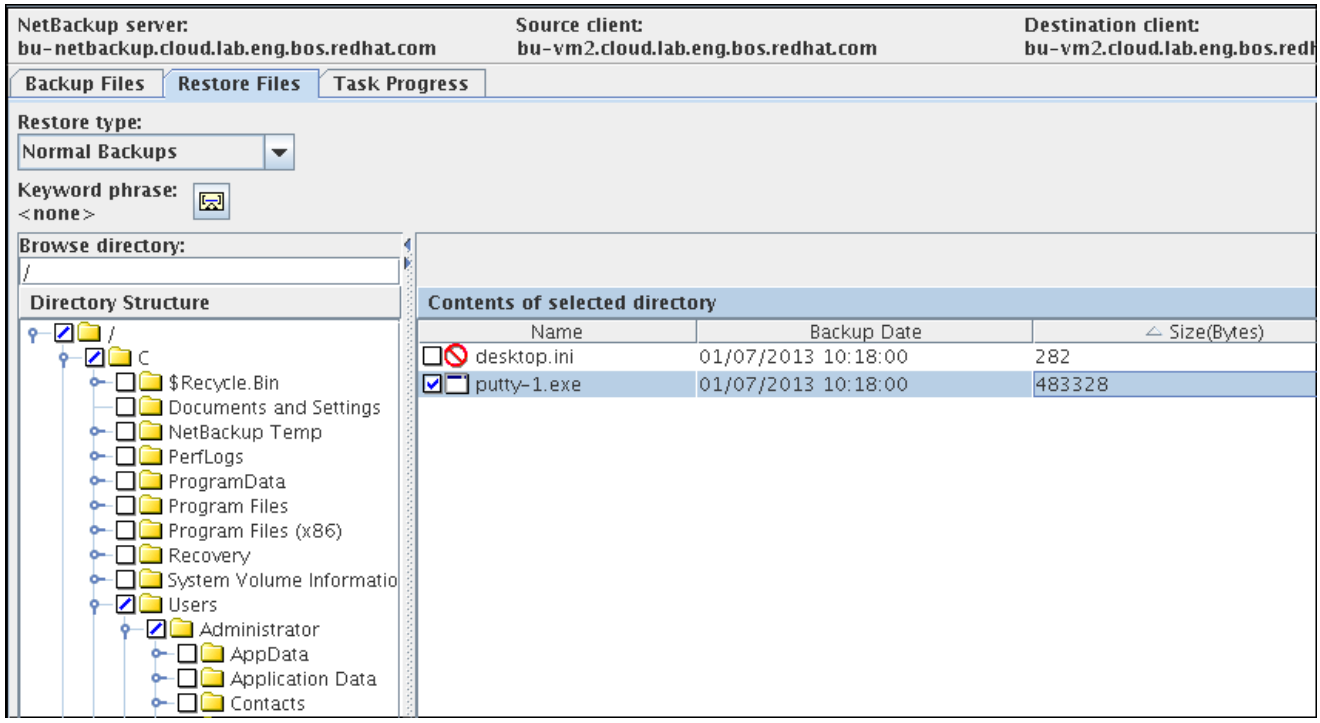

Figure 6.2.3-5: Windows File Level Restore – File Selection

Select restore options- '**Restore everything to the original location**' and '**Use default progress log filename**'

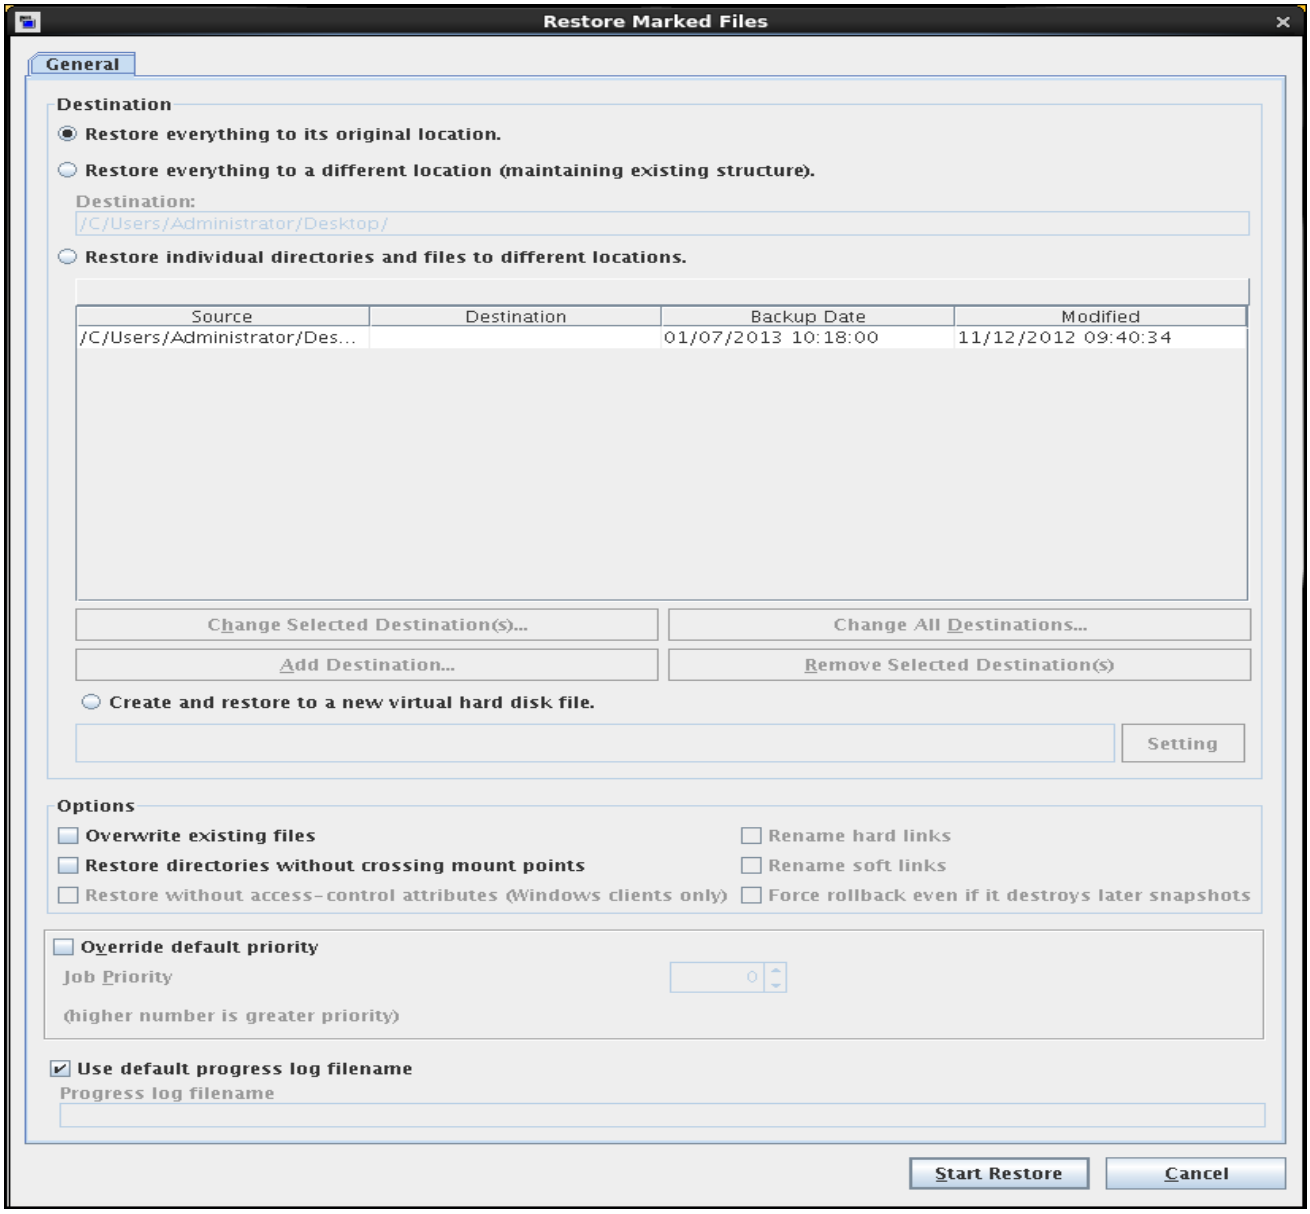

Figure 6.2.3-6: Windows File Level Restore – Restore Options

The job status shows progress of the restore job and indicates outcome of the restore job.

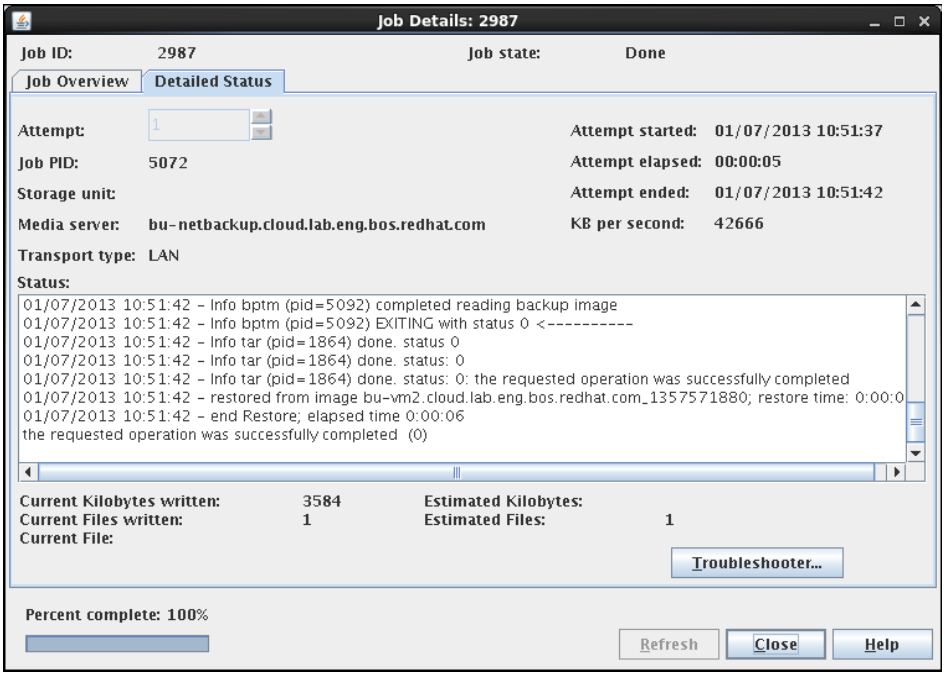

Figure 6.2.3-7: Windows File Level Restore – Job Status

Upon successful completion of the job, Figure 6.2.3-8: Windows File Level Restore – End statusdisplays that the file has been restored to prior incident level.

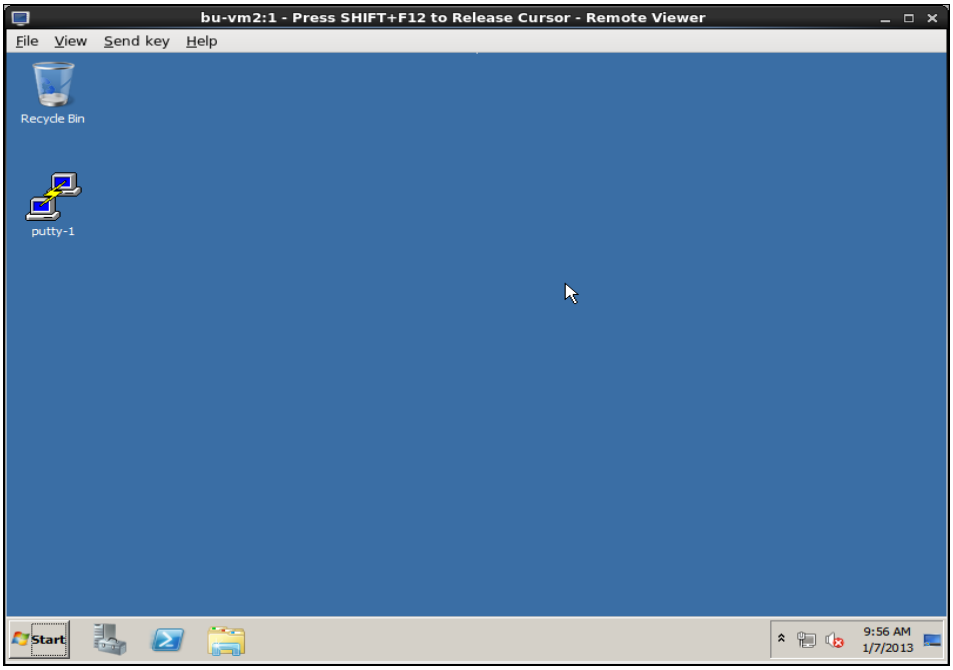

Figure 6.2.3-8: Windows File Level Restore – End status

6.2.4 Microsoft Windows Virtual Machine - full VM recovery

This situation is best described by a disaster event where the virtual machine bu-win2k8 (with hostname **bu-win1**) running Windows 2008 R2 is not responding. It cannot be accessed via console and will no longer boot. In order to restore the system, following steps were followed:

1. Provision a new virtual machine
2. Configure network settings
3. Install the NetBackup agent
4. Select files to restore
5. Verify functionality

Provision Virtual Machine

The first step in recovery is to provision a new virtual machine running the same operating system as the failed machine. This can be accomplished by deploying from a template within RHEV or creating a new virtual machine from scratch and installing the operating system through automated deployment methods or from a virtual device. For this reference environment, a template was used to deploy a base operating system. When using this method, ensure the number and size of virtual disks used, along with the network configuration, match the failed VM configuration.

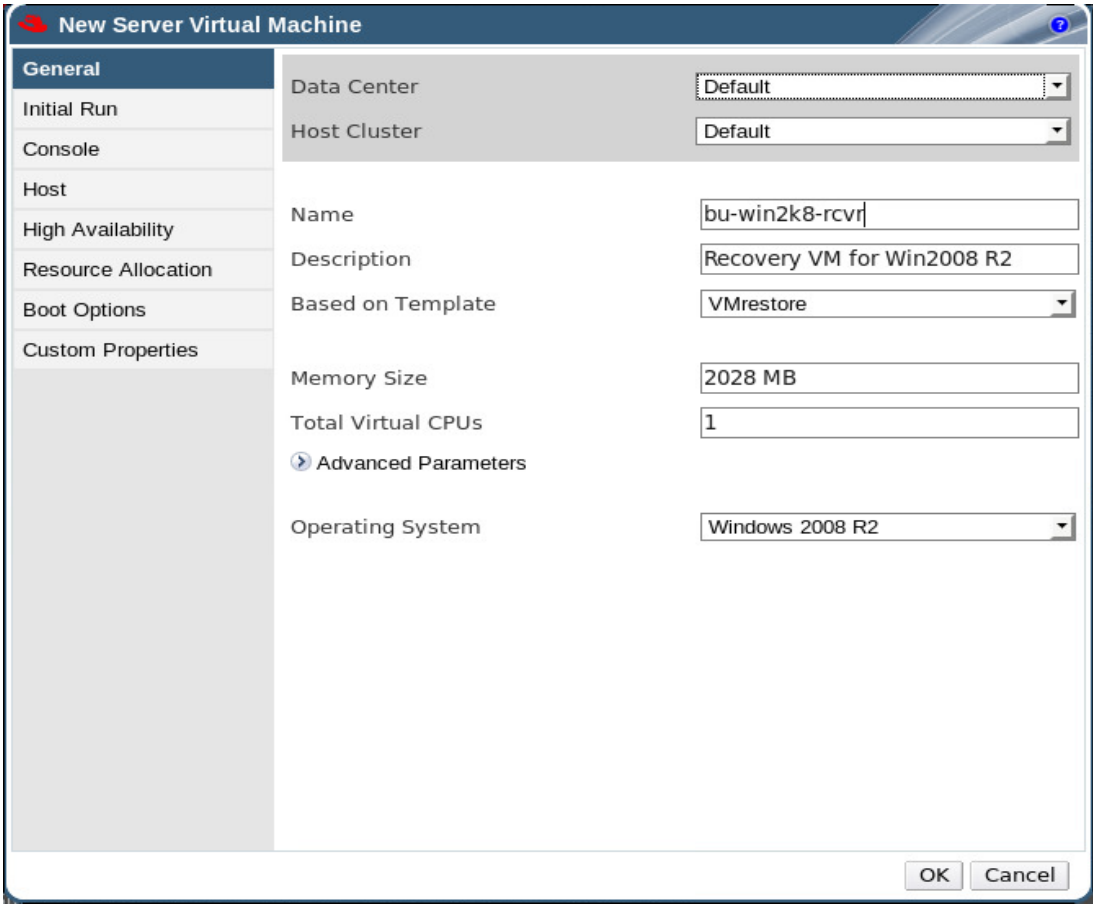

Figure 6.2.4-1: Deploy Windows VM Using RHEV Template

VM Network Configuration:

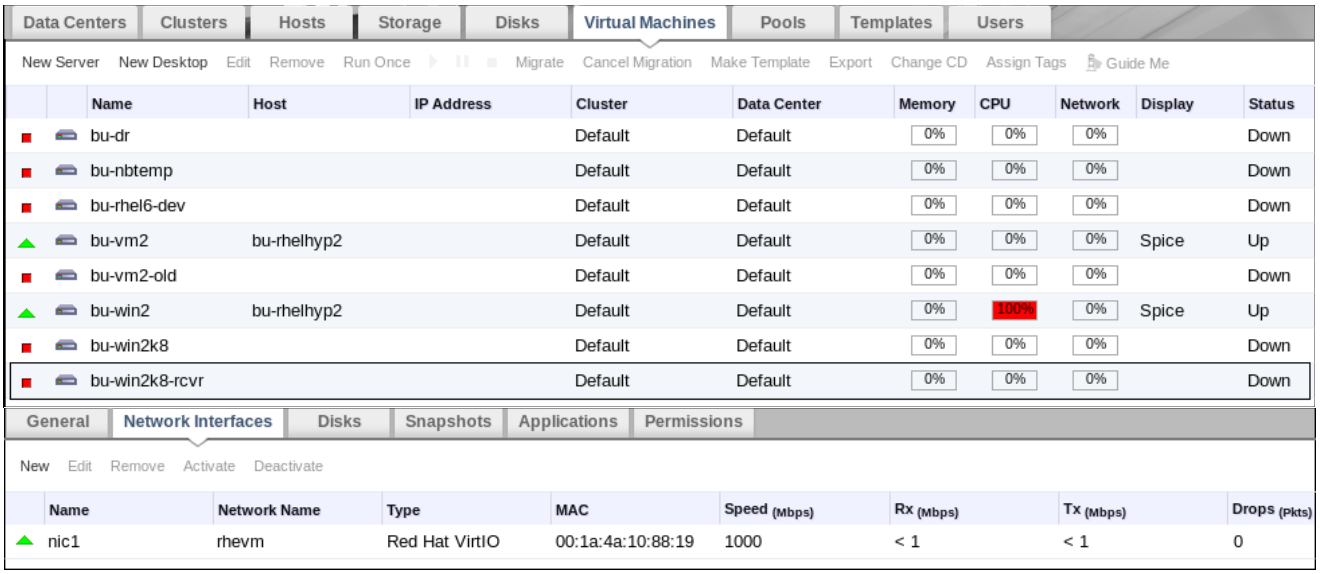

Figure 6.2.4-2: Recovery Windows VM Network Configuration

Edit and specify MAC Address of the original VM to this new recovery VM as described in Figure 6.2.2-2: Specify Mac Address. This will help boot up the new VM with **bu-win1** hostname (under bu-win2k8-rcvr VM name).

Install NetBackup Client

In this setup, the VM template **VMrestore** was created on a system that had NetBackup client Installed. Hence NetBackup client came in pre-provisioned. Otherwise if the template did not bring in the client software, it can be manually installed. For detailed specifics on installing the client, please refer to section 5.3.2 NetBackup Client Installation on Windows

File recovery

With the NetBackup agent installed and configured, file recovery can commence. When performing a system restore using the NetBackup agent running on a Windows operating system, the following is a defined process that must be followed:

1. Launch the NetBackup BAR (Backup, Archive & Restore) GUI application from the client.
2. Select the images that contain the Full and Incremental (if applicable) backups of the system drive first(s). Enable the overwrite option. **Do not elect to restore the System State/Shadow Copy Components in this step.**.
   Note: DO NOT REBOOT after the system drive(s) restore is completed.

System drive selection as shown below:

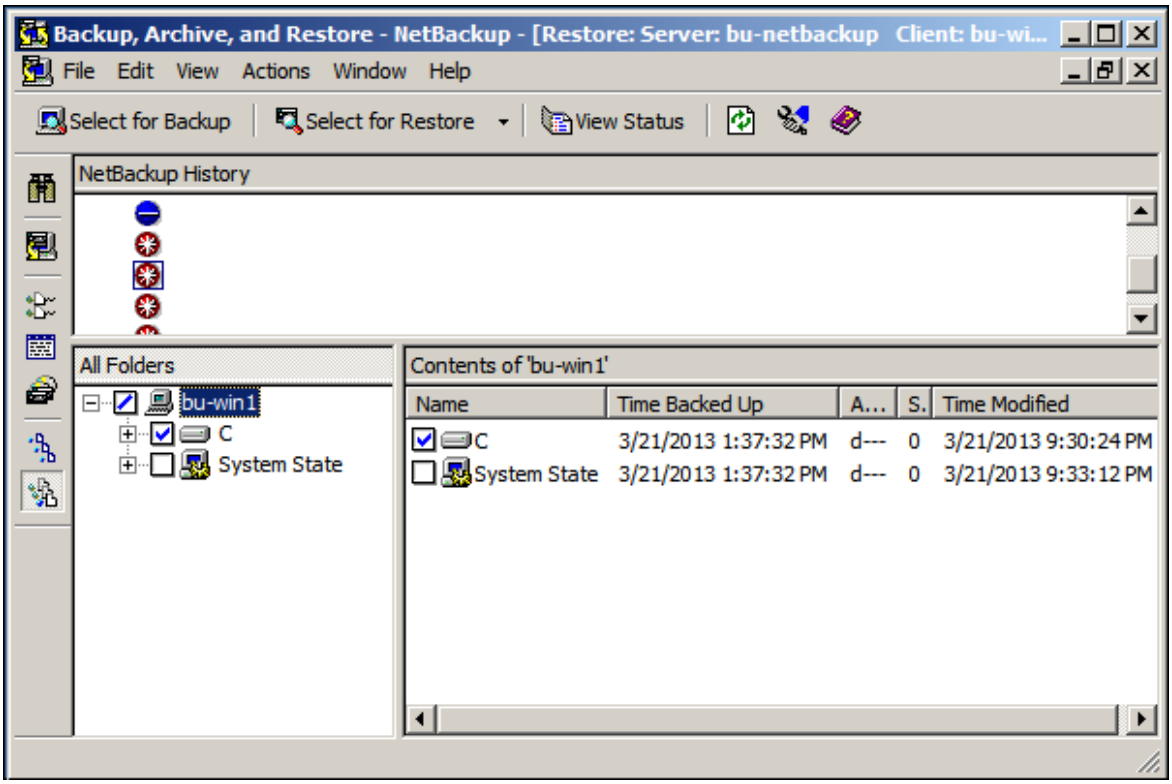

Figure 6.2.4-3: System Drive Restore Selection

3. When the restore is complete, check the client's tar log (found in *C:\Program Files>\Veritas\NetBackup\logs\tar*) for confirmation of a successful restore. Also the 'View Status' button on the NetBackup screen provides real time restore status. If the restore had errors, this log will provide more detail and any issues found should be resolved before continuing.

4. System_State/Shadow Copy Components

   CAUTION: This is the most critical part of the restore that could result in a bootable system or non-bootable system.

   Select the images that contain the Full and Incrementals (if applicable) backups and start a restore of the System State/Shadow Copy Components with the overwrite option enabled.

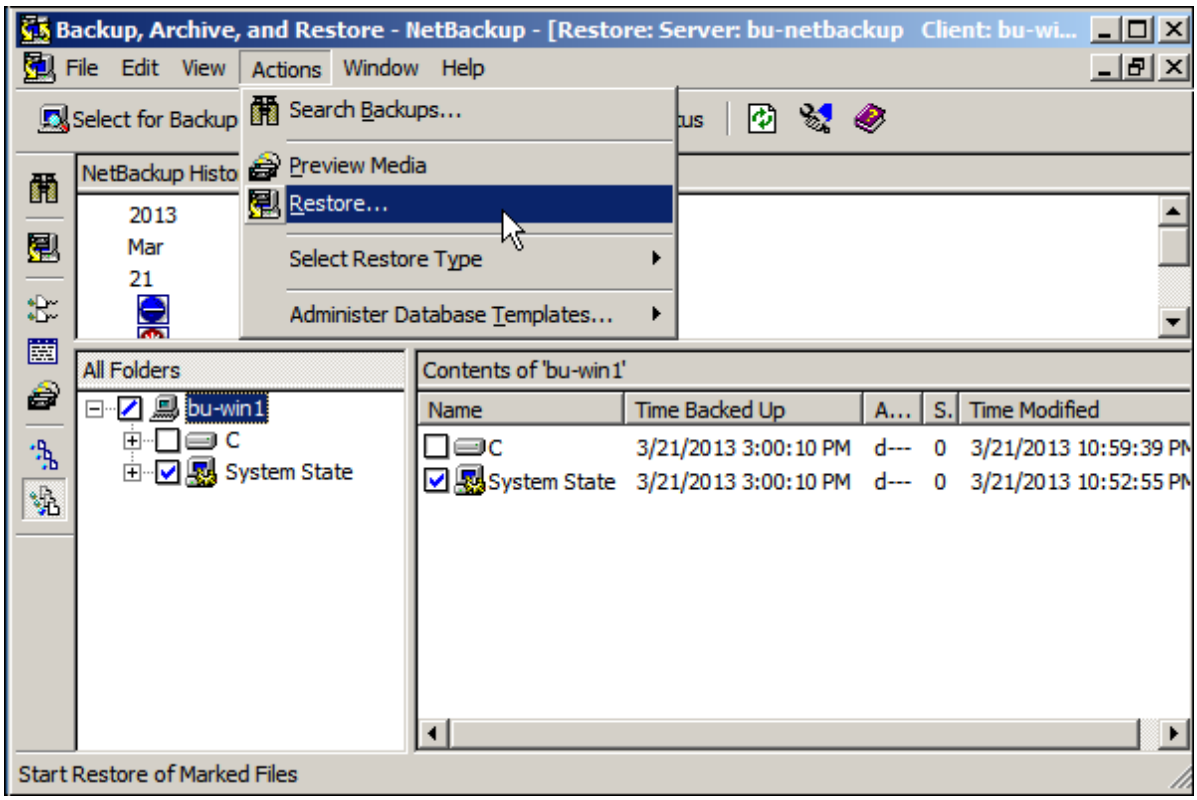

Figure 6.2.4-4: System State/Shadow Copy Component Selection

If the restore being attempted is to the ORIGINAL hardware and it is desired to restore
this information, use the W2koption.exe utility supplied with the NetBackup Windows Client. **THIS SHOULD NOT BE USED IF RESTORE IS NOT TO ORIGINAL HARDWARE**. Running W2koption.exe:

A. Restore the System_State or Shadow Copy Components. **DO NOT REBOOT**.

B. Upon completion of the restore, check the tar log again for any problems. Once again, **DO NOT REBOOT**.

Restore Status:

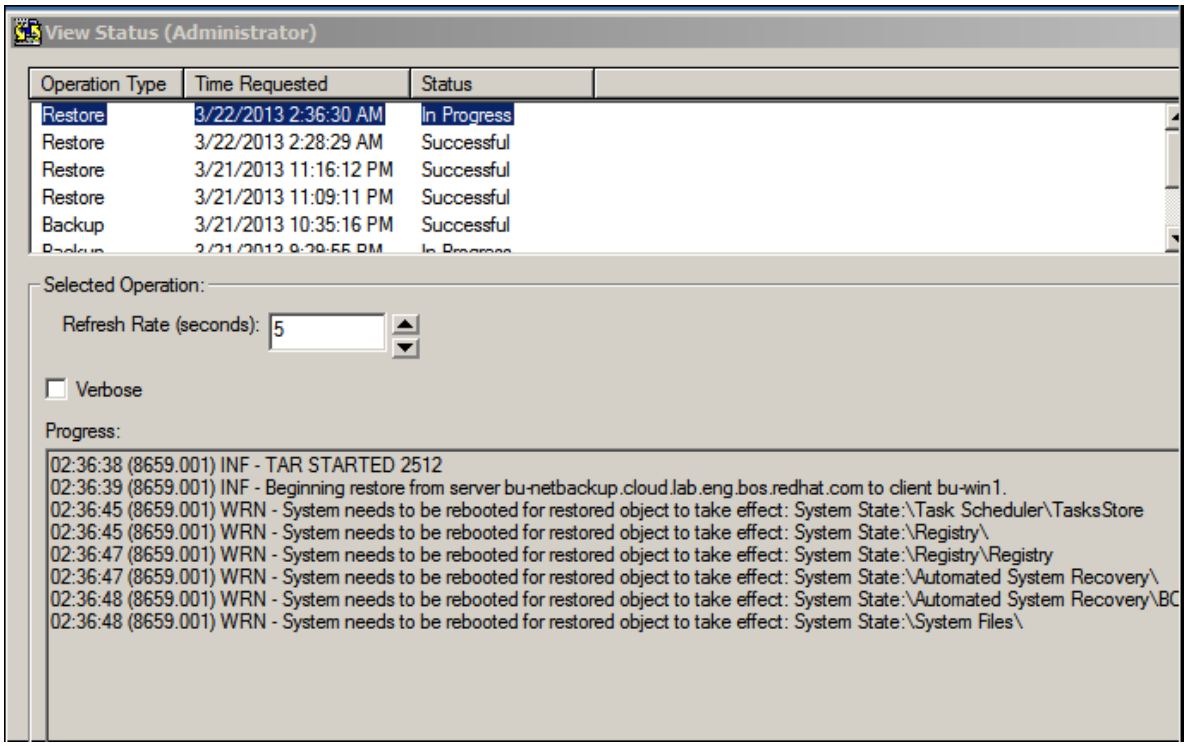

Figure 6.2.4-5: View Status Output

5. **[Optional]** Restore other data
   Perform restores of other drive letters (non-system Drives) before rebooting. (Alternatively, you could perform this step after the reboot.)

6. Stop NetBackup Client Service(bpinetd) and Reboot

   It is now prudent to reboot the server however the NetBackup client should be stopped prior to the reboot to ensure the registry information is pushed.

```
net stop "NetBackup Client Service"
```

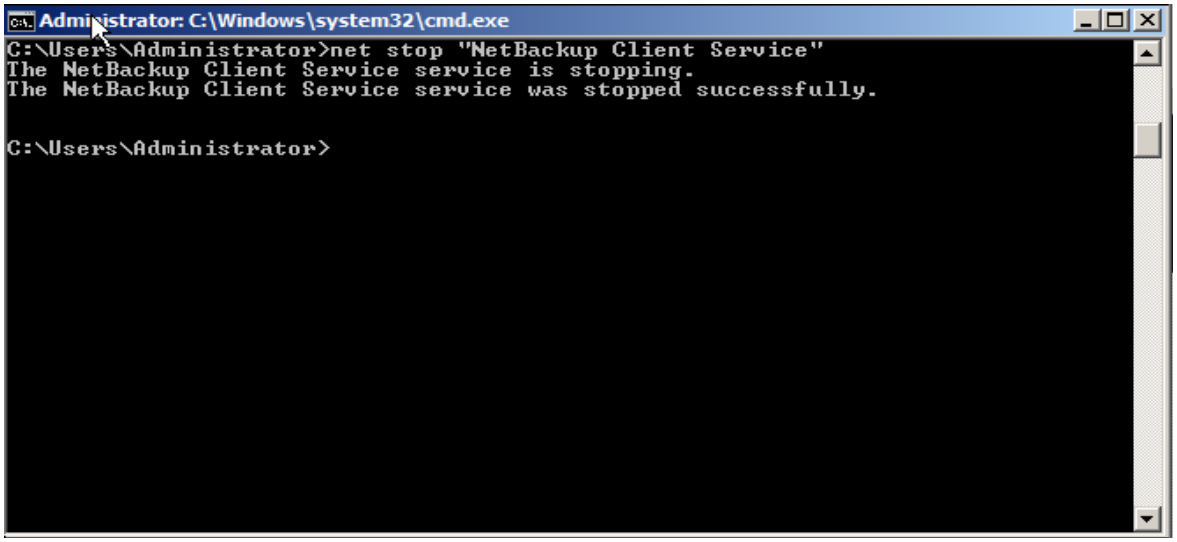

Figure 6.2.4-6: Stopping NetBackup Client

Verify Functionality

Once the server completes the reboot after restoring, verify if the VM has been restored similar to the original VM.

Cleanup

As a final step, remove the original failed VM, edit the new VM and update the Name, to the original VM Name bu-win2k8.

7 Backup and Restore of RedHat Enterprise Virtualization Manager

7.1 NetBackup Client Installation

Like any other Linux server being backed up, the RHEV-M server must be attached to a policy based on the backup requirements and the NetBackup client software has to be installed. Please refer to Section **5.1.2 Remote Client Install using NetBackup commands** for client installation.

7.2 Backup

In addition to performing full or incremental backups, it is necessary to ensure that a copy of the database is dumped as an sql file within the server that gets backed up during the regular backup process.

Backing up the RHEV Databases using the backup.sh Script

There are three databases under Red hat Enterprise Virtualization Manager that could be backed up for recovery purposes:

1) **engine** database.
2) **ovirt_engine_history** which tracks engine database over time.
3) **`rhevmreports`** for Red Hat Enterprise Virtualization Manager Reports functionality.

The Red Hat Enterprise Virtualization Manager includes a script ***`/usr/share/ovirt-engine/dbscripts/backup.sh`*** to automate these database backups.

Detailed description of the ***backup.sh*** script is mentioned under the link:
https://access.redhat.com/knowledge/docs/en-US/Red_Hat_Enterprise_Virtualization/3.1/html/Administration_Guide/appe-Backups.html

This script can be further included in the ***/usr/openv/netbackup/bin/bpstart_notify*** to automatically backup RHEV-M database every time a backup is executed using NetBackup.

/usr/openv/netbackup/bin/bpstart_notify script:

```
#!/bin/bash
service ovirt-engine stop
cd /usr/share/ovirt-engine/dbscripts
#
###backup Engine Database
/usr/share/ovirt-engine/dbscripts/backup.sh -s localhost -p 5432 -d engine
-u postgres -l /usr/tmp/rhevm-backups/
#
###backup rhevmreport database
/usr/share/ovirt-engine/dbscripts/backup.sh -s localhost -p 5432 -d
rhevmreport -u postgres -l /usr/tmp/rhevm-backups/
#
###backup ovirt_engine_history (dwh) database
/usr/share/ovirt-engine/dbscripts/backup.sh -s localhost -p 5432 -d
ovirt_engine_history -u postgres -l /usr/tmp/rhevm-backups/
#
service ovirt-engine start
```

Note:

1. In this scenario, there was no necessity for a post backup script. However if need arises, it can be made possible by creating a ***bpend_notify*** script in ***/usr/openv/netbackup/bin*** directory for any required automation.
2. The ***backup.sh*** command will generate a new sql file with date and time stamp, each time the backup or this ***bpstart_notify*** *script is executed.* Hence old backups have to be manually moved or cleaned up using a log rotate script.
3. While **engine** database is a requirement, **ovirt_engine_history** and **rhevmreports** are optional. If not implemented, they can be removed from the above mentioned script.
4. Installation and configuration of History and Reports are detailed in Appendix E.2Install and Configure History and Reports Database

The following job details screen confirms the starting and finishing of ***bpstart_notify script*** during a regular backup process.

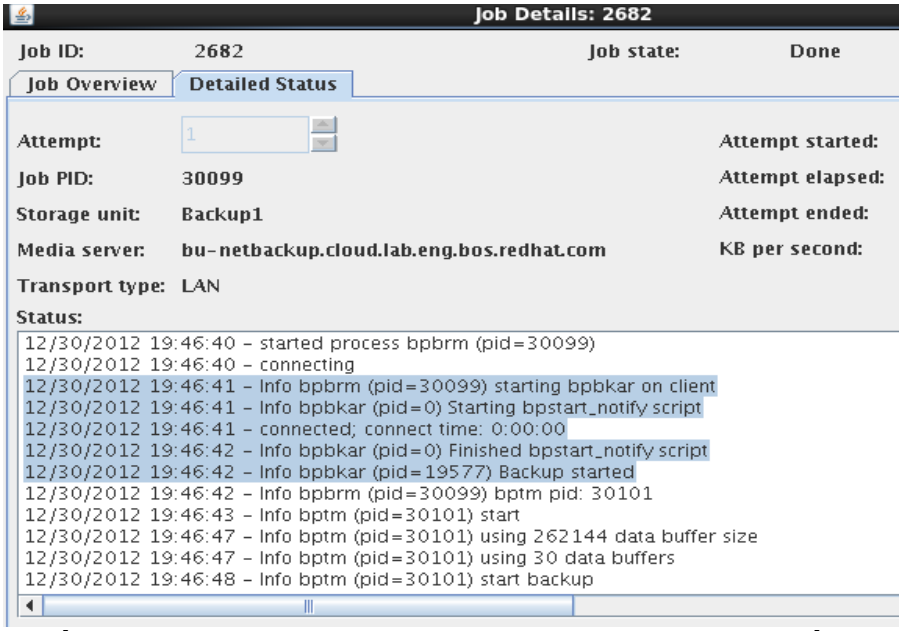

Figure 7.2-1: NetBackup – Pre-Backup Script

7.3 Recovery

7.3.1 Red Hat Enterprise Virtualization Manager (RHEV-M) Application Failure

This section is practically the same as a file recovery step under **Section 5.4.1 File level restore** for a Linux client.

7.3.2 Red Hat Enterprise Virtualization Manager (RHEV-M) Database Corruption

This scenario involves having a corrupted Red Hat Enterprise Virtualization Manager database. Upon login to the RHEV-M Portal, a login error for the *admin* account is presented indicating '**action failure due to database connection failure**' as displayed below:

Figure 7.3.2-1: Database Corruption – RHEV-M

Restoring the Engine Database Using the restore.sh Script

The Red Hat Enterprise Virtualization Manager includes a script ***/usr/share/ovirt-engine/dbscripts/restore.sh*** to automate database restoration. Using this script on the Manager server, the corrupted database can be recovered.

```
# service ovirt-engine stop
Stopping engine-service:                                   [  OK  ]

#/usr/share/ovirt-engine/dbscripts/restore.sh -p 5432 -d engine -u postgres
-f /usr/tmp/rhevm-backups/engine_Wed_Feb_13_19:40:10_EDT_2013.sql  -r
ar_13_19:40:10_EDT_2013.sql  -r
Restore of database engine from /usr/tmp/rhevm-
backups/engine_Wed_Mar_13_19:40:10_EDT_2013.sql started...
SET
SET
                    ...output abbreviated...
ALTER TABLE
```

```
ALTER TABLE
REVOKE
REVOKE
GRANT
GRANT
Restore of database engine from /usr/tmp/rhevm-
backups/engine_Wed_Mar_13_19:40:10_EDT_2013.sql completed.

#Starting engine-service:                                  [  OK  ]
```

Verify functionality

Verify if the RHEV-M functionality has been restored by a successful login to the RHEV-M admin portal and all details are accessible.

Detailed description of the ***restore.sh*** script is mentioned in the link https://access.redhat.com/knowledge/docs/en-US/Red_Hat_Enterprise_Virtualization/3.1/html/Administration_Guide/Restoring_the_engine_database_using_the_restore.sh_script.html.

> Note: The History and Reports databases can be restored the same way by replacing **engine** database with **ovirt_engine_history** and **rhevmreports** respectively to the above command with appropriate backup files.

7.3.3 Full Red Hat Enterprise Virtualization Manager (RHEV-M) Machine Crash

This scenario involves a total system crash and the hardware could not be rendered usable at that point. A previously -backed up system image and its contents have to be restored on to a similar hardware for business to continue as usual. While it is possible to install a fresh OS image and restore the system at file level, Symantec provides an automated solution to restore the entire server as a Bare Metal Restore. Please refer to Appendix H: Bare Metal Restore Essentials for more details on BMR.

Bare Metal Restore

In this scenario, a different hardware exactly identical to the original RHEV-M server has been used for BMR. If there are variations in the hardware, additional configuration may be needed. For details, please refer to the administration guide - http://www.symantec.com/business/support/index?page=content&id=DOC5163
While BMR can be accomplished across network, in this guide, boot media in the form of virtual CD has been used to boot the new server before data restore can be begin.

Bare Metal Restore steps:

1) Selection of BMR option in the policy
2) Verification of inclusion of BMR configuration during regular backups
3) Verification that the client is listed in the BMR client list
4) Prepare to restore when incident happens
5) Mounting the iso image before the new hardware is booted up
6) BMR initiation and providing RHEV-M server and NetBackup Master server details
7) Verification after BMR has been successfully completed and the server has been rebooted.

Note: The first three steps mentioned above are performed while the RHEV-M server is functional. The last four are performed after a system failure.

Verify the selection of BMR option in the policy

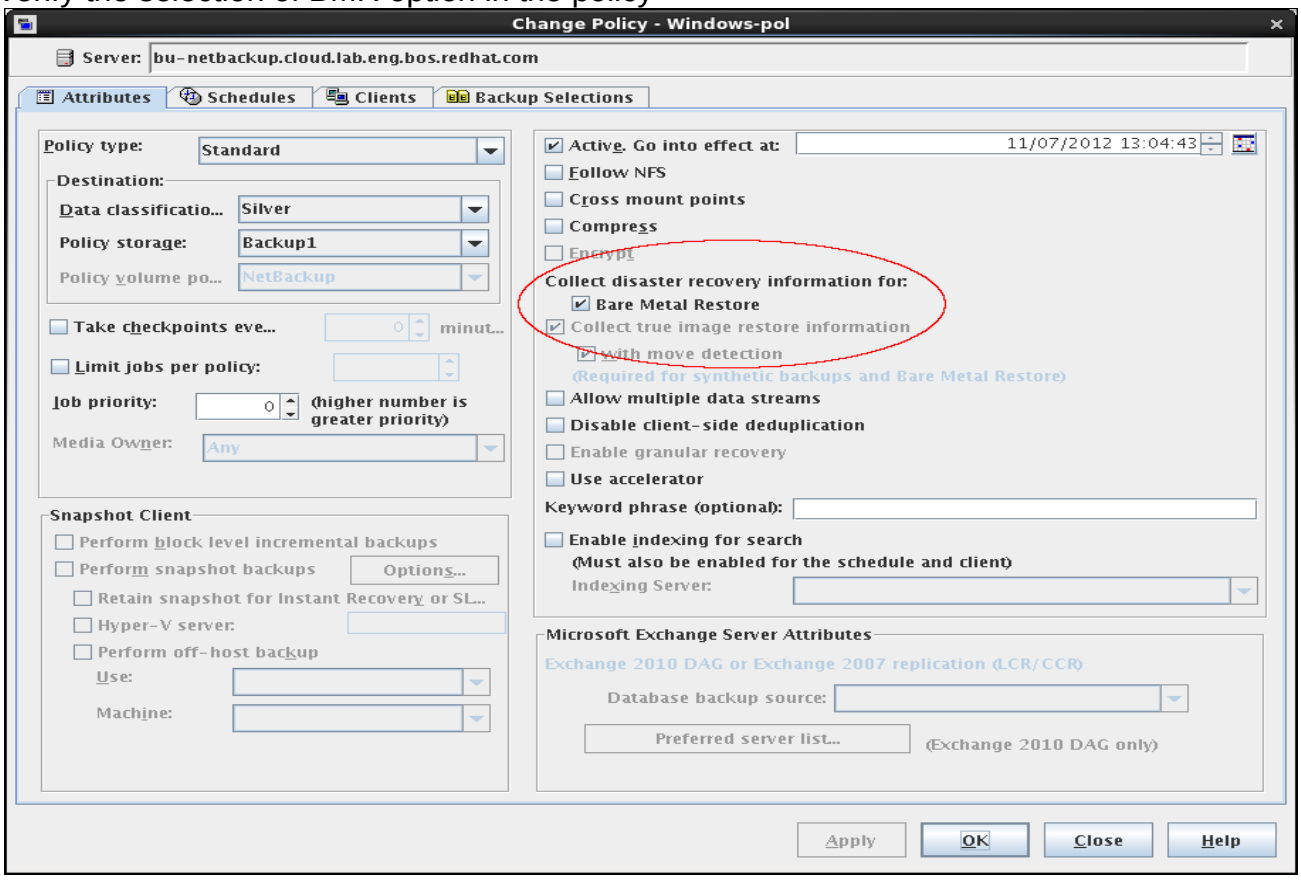

Figure 7.3.3-1: BMR option in Policy

Verify that the client is listed in the BMR client list.
Bare Metal Restore Management--> **Hosts**--> **Bare Metal Restore Clients**

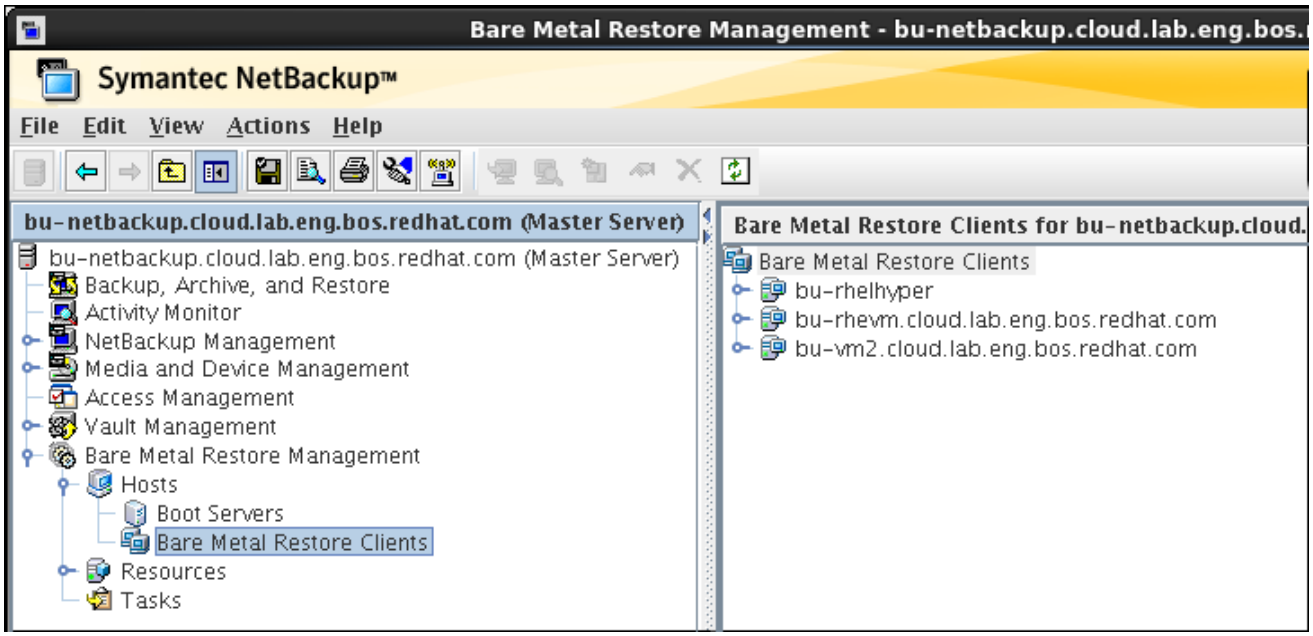

Figure 7.3.3-2: BMR Client List

Right click on the client and select **Prepare to Restore** after incident happens:

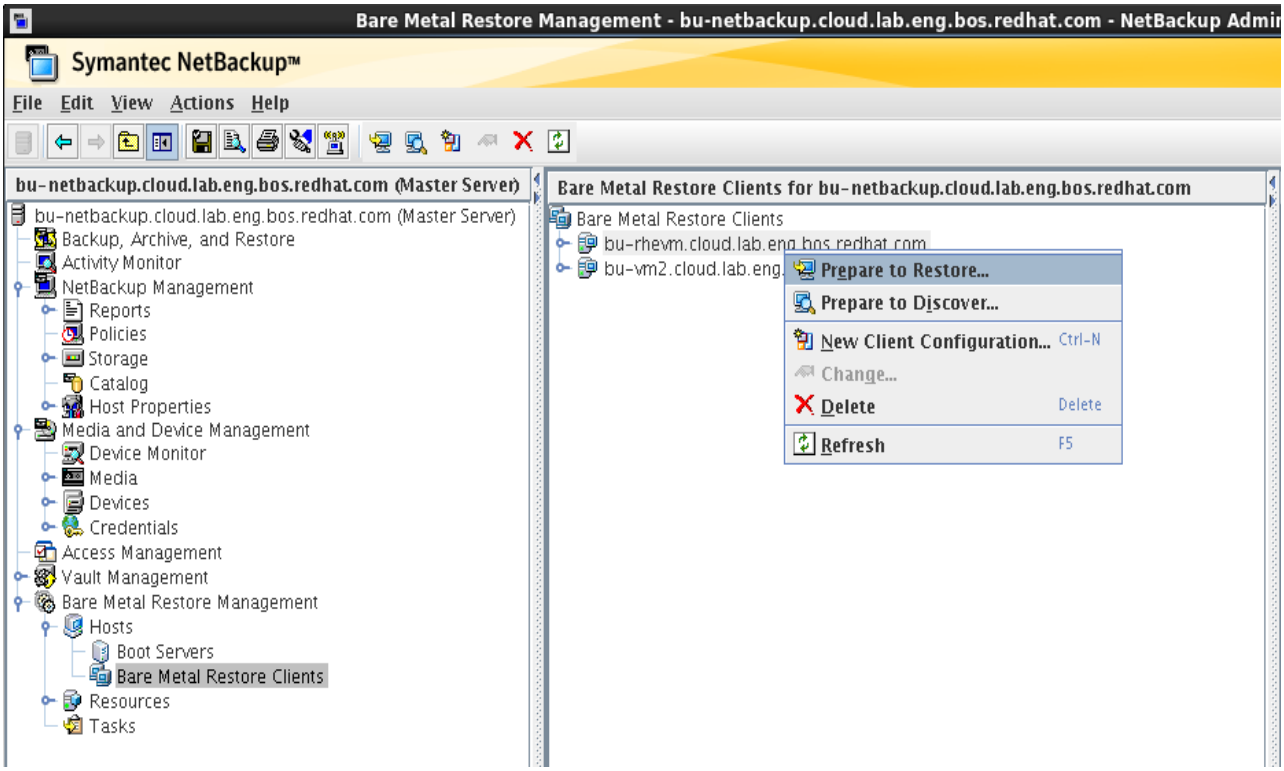

Figure 7.3.3-3: BMR Prepare to Restore

The Prepare to Restore brings up a window to select configuration and SRT
Verify the chosen options and click “OK”” to proceed. The warning about OS level mismatch can be ignored for minor OS variations.

Figure 7.3.3-5: BMR Prepare to Restore confirmation

Mount the iso image before the new hardware is booted up. This is similar to mounting a bootable CD or a virtual CD/DVD. In this case, the hardware is a Dell Blade. The virtual CD/DVD can be mounted on the IDRAC console:

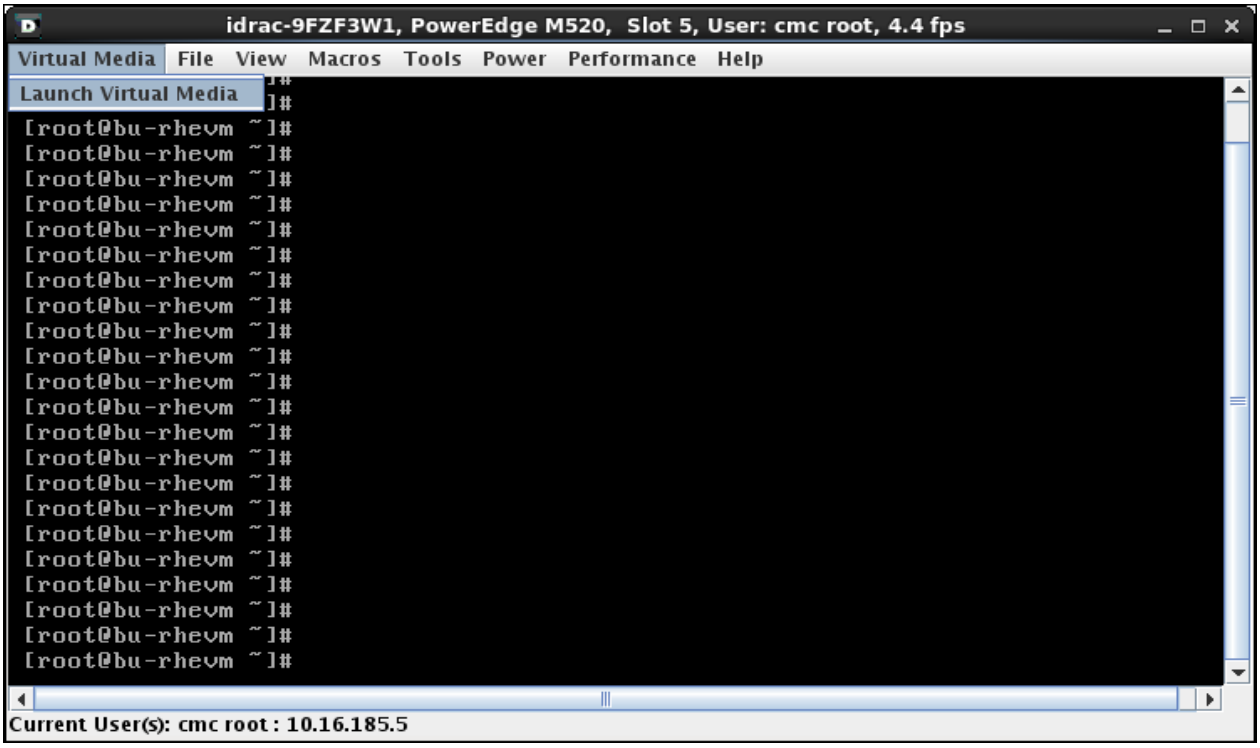

Figure 7.3.3-6: IDRAC Console launching Virtual Media

In the virtual console an iso image is loaded by selecting **Add Image button** and selecting the iso image file. The check for the right image has to be selected as circled in red.

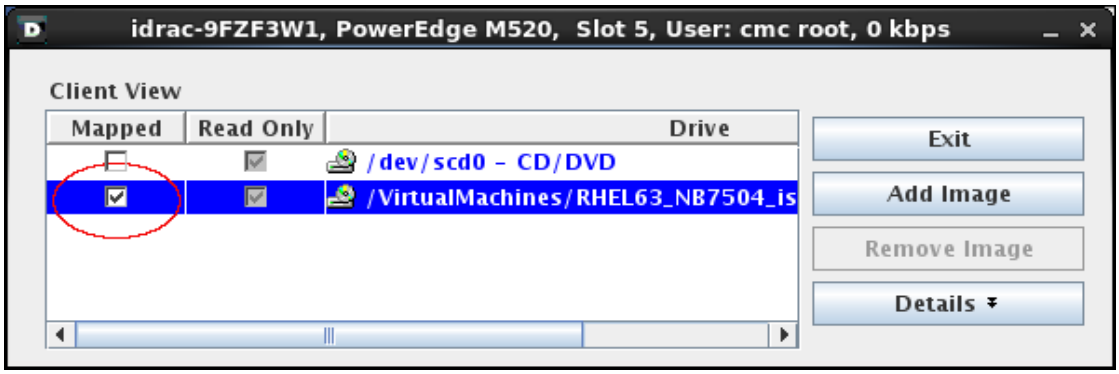

Figure 7.3.3-7: Mounting the SRT ISO image

BMR process starts.

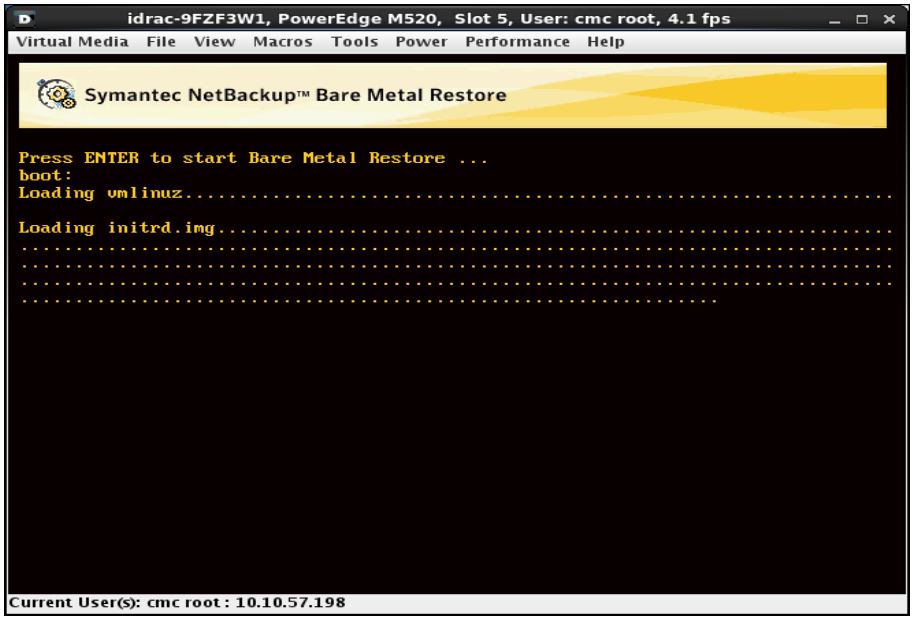

Figure 7.3.3-8: BMR Initiation

Provide RHEV-M server, Master server and network details as prompted in the screen below.

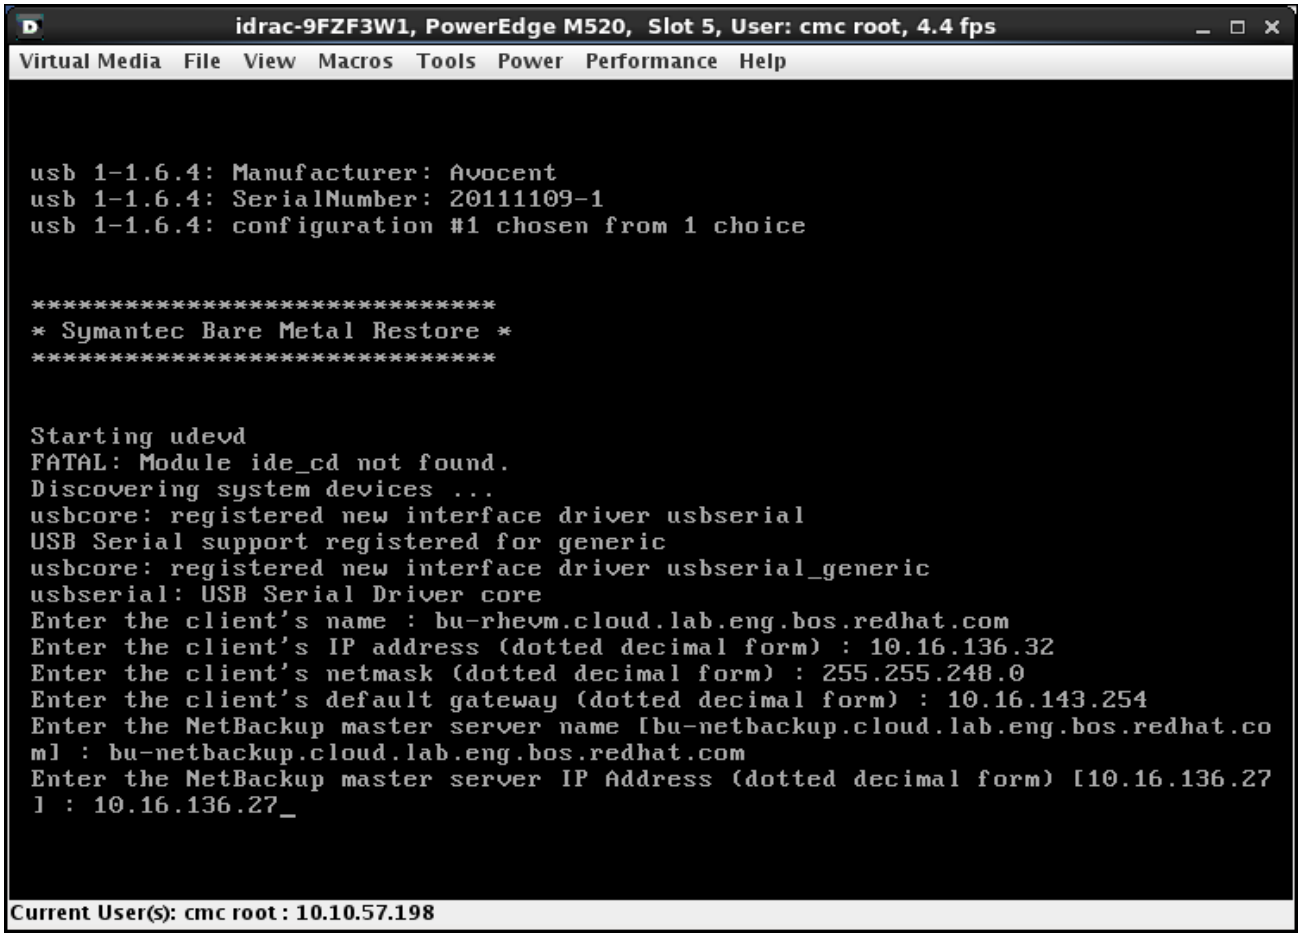

Figure 7.3.3-9: BMR Client & Master Details input

After this input, BMR will orchestrate the restore with all the necessary drivers, disk and network configurations provided by the master. Once the restore has been completed, the machine will prompt for a reboot. Verify machine functionality post-reboot.

> Note- If the RHEV-M server uses DHCP, then the RHEV-M server's IP must be associated with this new MAC Address unique to this hardware.

8 Conclusion

This paper demonstrated the ability to backup and recover a Red Hat Enterprise Virtualization environment using Symantec NetBackup to maintain business continuity. The following steps were successfully performed:

- Symantec NetBackup setup and configuration within a Red Hat Enterprise Virtualization environment
- Virtual Machine file level backup and recovery
- Virtual Machine operating system backup and recovery
- Red Hat Enterprise Virtualization Manager application backup and recovery
- Red Hat Enterprise Virtualization Manager database backup and recovery
- Red Hat Enterprise Virtualization full system backup and recovery

There are almost endless possibilities with maintaining business continuity. When planning for backup and recovery, many considerations need to be taken into account and will vary based on specific business needs. This drives the ultimate list of requirements for IT departments to plan for. Some requirements may include:

- Data deduplication
- Off-site storage
- Data retention policies
- Data encryption
- Reporting
- Scalability

In conclusion, business continuity is of utmost importance for business operations. It is critical to maintain and operate successful backup and recovery procedures and plans regardless of the tools used. Requirements for each environment and businesses will vary however the need to recover from unplanned disasters never changes.

Appendix A: Revision History

Revision 2.0	Friday March 22, 2013	Balaji Jayavelu

- Updated for RHEV and NetBackup releases to include increased functionality and workflows.

Appendix B: NetBackup Environment Requirements

The following requirements were met to proceed with normal functioning of the NetBackup 7.5.0.5 environment running on RHEL 6.3 server. However further settings may be required suiting to different configuration or environment needs. Please refer to Symantec NetBackup Installation Guide - http://www.symantec.com/business/support/index?page=content&id=DOC5154&key=15143 for further details.

B.1 Name Resolution

NetBackup has a requirement that all the active interfaces used for communication among peers (master, media and clients etc) must have name resolution. This can be achieved by using Domain Name Service (DNS), where the public and private network interfaces are to be registered and should resolve to the '**nslookup**' command. It should be noted that in RedHat Enterprise Linux, **dig, host and nslookup** utilities are designed to obtain names from name servers only. Hence maintaining entries in ***/etc/hosts*** file is not going to help in communication among NetBackup Tiers.

B.2 NetBackup Firewall Requirements

The following port(s) were opened bidirectional on the master/media and clients:

Service	Port	Protocol
VERITAS_PBX*	1556	TCP
bprd	13720	TCP
bpdbm**	13721	TCP
vnetd	13724	TCP
bpcd	13782	TCP

Table B.2-1: Services and Ports - NetBackup

*Not necessary in this case since Master and EMM servers are the same.
** Not necessary in this case since Master and Media servers are the same.

Note: Additional ports may need to be open depending upon specific environment and configuration.

Contents of */etc/sysconfig/iptables*

```
# Firewall configuration written by system-config-firewall
# Manual customization of this file is not recommended.
*filter
:INPUT ACCEPT [0:0]
:FORWARD ACCEPT [0:0]
:OUTPUT ACCEPT [0:0]
-A INPUT -m state --state ESTABLISHED,RELATED -j ACCEPT
-A INPUT -p icmp -j ACCEPT
-A INPUT -i lo -j ACCEPT
-A INPUT -m state --state NEW -m tcp -p tcp --dport 22 -j ACCEPT
-A INPUT -m state --state NEW -m tcp -p tcp --dport 443 -j ACCEPT
-A INPUT -m state --state NEW -m tcp -p tcp --dport 13724 -j ACCEPT
-A INPUT -m state --state NEW -m tcp -p tcp --dport 13720 -j ACCEPT
-A INPUT -m state --state NEW -m tcp -p tcp --dport 13782 -j ACCEPT
-A INPUT -j REJECT --reject-with icmp-host-prohibited
-A FORWARD -j REJECT --reject-with icmp-host-prohibited
COMMIT
```

Appendix C: NetBackup Upgrade Procedure

Bare Metal Restore on NetBackup 7.5 version has a compatibility issue with RedHat Enterprise Linux 6.3 clients, where BMR fails to import client config file. This issue has been fixed in NetBackup version 7.5.0.5 and is a requirement if BMR functionality is expected.

Upgrading NetBackup 7.5 to 7.5.0.5 maintenance release

For proper functioning of BMR for RedHat Enterprise Linux 6.3 clients, NetBackup was updated to the latest maintenance release. This is an optional step to utilize BMR feature of NetBackup.
The 7.5.0.5 maintenance release can be downloaded from Symantec website:
http://www.symantec.com/business/support/index?page=content&id=TECH199269

After downloading and un-tarring the software the NB_update.install utility will update to the required version.

Close the NetBackup user interfaces after confirming that there are no active jobs running.

```
# /usr/openv/netbackup/bin/bp.stop_all
```

Install NB_7.5.0.5 and NB_CLT_7.5.0.5 Maintenance Release binaries.

```
# cd /tmp
# /bin/sh NB_update.install
```

Restart NetBackup processes.

```
# /usr/openv/netbackup/bin/bp.start_all
```

NOTE: Selecting the server Maintenance Release automatically installs the client Maintenance Release if the client (CLT) .gz file and the README exist in the installation directory. The server install fails if the (CLT) .gz file and the README are not present and the CLT update has not been previously installed. The client Maintenance Release is NOT installed automatically during a reinstall of the server Maintenance Release.

Appendix D: Red Hat Enterprise Virtualization 3.1 Requirements

D.1 Red Hat Enterprise Virtualization Manager Requirements

D.1.1 Hardware Requirements

The minimum and recommended hardware requirements outlined here are based on a typical small to medium sized installation. The exact requirements vary between deployments based on sizing and load. Please use these recommendations as a guide only.

Minimum:

- A dual core CPU.
- 4 GB of available system RAM that is not being consumed by existing processes.
- 25 GB of locally accessible, writeable, disk space.
- 1 Network Interface Card (NIC) with bandwidth of at least 1 Gbps.

Recommended:

- A quad core CPU or multiple dual core CPUs.
- 16 GB of system RAM.
- 50 GB of locally accessible, writeable, disk space.
- 1 Network Interface Card (NIC) with bandwidth of at least 1 Gbps.

The Red Hat Enterprise Virtualization Manager runs on Red Hat Enterprise Linux. To confirm whether or not specific hardware items are certified for use with Red Hat Enterprise Linux refer to http://www.redhat.com/rhel/compatibility/hardware/.

D.1.2 Software Requirements

- Red Hat Enterprise Virtualization Manager requires Red Hat Enterprise Linux 6.3 Server or later.
- Complete successful installation of the operating system prior to commencing installation of the Red Hat Enterprise Virtualization Manager.

D.1.3 Required Channels

Host	Channels	Purpose
bu-rhevm	rhel-x86_64-server-6	RHEL Base
	rhel-x86_64-server-supplementary-6	Java Runtime Env (JRE)
	rhel-x86_64-server-6-rhevm-3.1	RHEV Manager
	jbappplatform-6-x86_64-server-6-rpm	Jboss Application Platform

Table D.1.3-1: Required Channels

D.1.4 Firewall Requirements

Ports	Protocol	Purpose
-	ICMP	ICMP Ping request when the Hypervisor registers with RHEVM
22	TCP	SSH (Optional)
8080,8443	TCP	HTTP and HTTPS access to the manager

Table D.1.4-1: Firewall port details

Additionally, ports 111 and 2049 may need to be opened up if NFS storage is used.
The firewall settings can be automated while performing *rhevm-setup* step.
However if there is a need to perform manual firewall configuration, the configured rules need to be equivalent to those found in */etc/ovirt-engine/iptables.example*

```
# Generated by ovirt-engine installer
#filtering rules
*filter
:INPUT ACCEPT [0:0]
:FORWARD ACCEPT [0:0]
:OUTPUT ACCEPT [52:9697]
:RH-Firewall-1-INPUT - [0:0]
-A INPUT -j RH-Firewall-1-INPUT
-A FORWARD -j RH-Firewall-1-INPUT
-A RH-Firewall-1-INPUT -i lo -j ACCEPT
-A RH-Firewall-1-INPUT -p icmp -m icmp --icmp-type any -j ACCEPT
-A RH-Firewall-1-INPUT -m state --state RELATED,ESTABLISHED -j ACCEPT
-A RH-Firewall-1-INPUT -p tcp -m state --state NEW -m tcp --dport 22 -j
ACCEPT
-A RH-Firewall-1-INPUT -m state --state NEW -p tcp --dport 8080 -j ACCEPT
-A RH-Firewall-1-INPUT -m state --state NEW -p tcp --dport 8443 -j ACCEPT
-A RH-Firewall-1-INPUT -m state --state NEW -p udp --dport 111 -j ACCEPT
-A RH-Firewall-1-INPUT -m state --state NEW -p tcp --dport 111 -j ACCEPT
-A RH-Firewall-1-INPUT -m state --state NEW -p udp --dport 892 -j ACCEPT
-A RH-Firewall-1-INPUT -m state --state NEW -p tcp --dport 892 -j ACCEPT
-A RH-Firewall-1-INPUT -m state --state NEW -p udp --dport 875 -j ACCEPT
-A RH-Firewall-1-INPUT -m state --state NEW -p tcp --dport 875 -j ACCEPT
-A RH-Firewall-1-INPUT -m state --state NEW -p udp --dport 662 -j ACCEPT
-A RH-Firewall-1-INPUT -m state --state NEW -p tcp --dport 662 -j ACCEPT
-A RH-Firewall-1-INPUT -m state --state NEW -p tcp --dport 2049 -j ACCEPT
-A RH-Firewall-1-INPUT -m state --state NEW -p tcp --dport 32803 -j ACCEPT
-A RH-Firewall-1-INPUT -m state --state NEW -p udp --dport 32769 -j ACCEPT
#drop all rule
-A RH-Firewall-1-INPUT -j REJECT --reject-with icmp-host-prohibited
COMMIT
```

D.2 Virtualization Host (Hypervisor) Requirements

D.2.1 Hardware requirements

Processor:

Virtualization hosts must have at least one CPU. Red Hat Enterprise Virtualization supports the use of these CPU models in virtualization hosts:

- AMD Opteron G1/G2/G3/G4
- Intel Conroe/Penryn/Nehalem/Westmere/Sandybridge

All CPUs must have support for the Intel® 64 or AMD64 CPU extensions, and the AMD-V™ or Intel VT® hardware virtualization extensions enabled. Support for the `No eXecute` flag (NX) is also required. To check that your processor supports the required flags, and that they are enabled:

- At the Red Hat Enterprise Linux or Red Hat Enterprise Virtualization Hypervisor boot screen press any key and select the **Boot** or **Boot with serial console** entry from the list. Press **Tab** to edit the kernel parameters for the selected option. After the last kernel parameter listed ensure there is a **Space** and append the `rescue` parameter.
- Press **Enter** to boot into rescue mode.
- At the prompt which appears, determine that the processor has the required extensions and that they are enabled by running this command:
  ```
  # grep -E 'svm|vmx' /proc/cpuinfo | grep nx
  ```
  If any output is shown, the processor is hardware virtualization capable. If no output is shown it is still possible that the processor supports hardware virtualization.
- As an additional check, verify that the `kvm` modules are loaded in the kernel:
  ```
  # lsmod | grep kvm
  ```
  If the output includes `kvm_intel` or `kvm_amd` then the `kvm` hardware virtualization modules are loaded and the system meets requirements.

RAM:

It is recommended that virtualization hosts have at least 2 GB of RAM. The amount of RAM required varies depending on:

- guest operating system requirements
- guest application requirements
- memory activity and usage of guests

A maximum of 1 TB of RAM per virtualization host is currently supported.

Storage:

Virtualization hosts require local storage to store configuration, logs, kernel dumps, and for use as swap space.
It is recommended that each virtualization host has at least 2 GB of internal storage. The minimum supported internal storage for each Hypervisor is the total of that required to

provision the following partitions:

- The root partitions require at least 512 MB of storage.
- The configuration partition requires at least 8 MB of storage.
- The recommended minimum size of the logging partition is 2048 MB.
- The data partition requires at least 256 MB of storage. By default all disk space remaining after allocation of swap space will be allocated to the data partition.
- The swap partition size recommended as follows:
    - 2 GB of swap space for systems with 4 GB of RAM or less, or
    - 4 GB of swap space for systems with between 4 GB and 16 GB of RAM, or
    - 8 GB of swap space for systems with between 16 GB and 64 GB of RAM, or
    - 16 GB of swap space for systems with between 64 GB and 256 GB of RAM.

PCI Device:

Virtualization hosts must have at least one network interface with a minimum bandwidth of 1 Gbps. It is recommended that each virtualization host have two network interfaces with a minimum bandwidth of 1 Gbps to support network intensive activity, including virtual machine migration.

D.2.2 Software Requirements

- Virtualization hosts must run version 6.3, or later
- Complete successful installation of the operating system prior to commencing installation in case of RHEL Hypervisor.

D.2.3 Firewall Requirements

Firewall settings on these hosts are automatically configured if the check box is selected while adding the host to the RHEV-M.

Ports	Protocol	Purpose
-	ICMP	ICMP Ping request when Hypervisor registers with RHEV-M
22	TCP	SSH (Optional)
8080,8443	TCP	Connect from browser
5634 to 6166	TCP	Guest Console connections
16514	TCP	libvirt VM migration traffic
49152 to 49216	TCP	VDSM VM migration traffic
54321	TCP	VDSM

Table D.2.3-1: Hypervisor Firewall Requirement

Configuring Virtualization Host Firewall:

Red Hat Enterprise Virtualization requires that a number of network ports be open to support virtual machines and remote management of the hypervisor from RHEVM. Remove any existing firewall rules using the `--flush` parameter to the `iptables` command.

```
# iptables --flush
```

Adding new firewall rules to configuration

```
#iptables --append INPUT -m state --state ESTABLISHED,RELATED -j ACCEPT
#iptables --append INPUT -p icmp -j ACCEPT
#iptables --append INPUT -i lo -j ACCEPT
#iptables --append INPUT -p tcp --dport 22 -j ACCEPT
#iptables --append INPUT -p tcp --dport 16514 -j ACCEPT
#iptables --append INPUT -p tcp --dport 54321 -j ACCEPT
#iptables --append INPUT -p tcp -m multiport --dports 5634:6166 -j ACCEPT
#iptables --append INPUT -p tcp -m multiport --dports 49152:49216 -j ACCEPT
#iptables --append INPUT -j REJECT --reject-with icmp-host-prohibited
#iptables --append FORWARD -m physdev ! --physdev-is-bridged -j REJECT \
--reject-with icmp-host-prohibited
```

Save the updated firewall configuration.

```
# service iptables save
```

Enable iptables service

```
# chkconfig iptables on
# service iptables restart
```

Virtualization Hypervisor firewall table: /etc/sysconfig/iptables

```
# Generated by iptables-save v1.4.7 on Fri Dec  14 12:49:12 2012
*filter
:INPUT ACCEPT [0:0]
:FORWARD ACCEPT [0:0]
:OUTPUT ACCEPT [104:14240]
-A INPUT -m state --state RELATED,ESTABLISHED -j ACCEPT
-A INPUT -p icmp -j ACCEPT
-A INPUT -i lo -j ACCEPT
-A INPUT -p tcp -m state --state NEW -m tcp --dport 22 -j ACCEPT
-A INPUT -j REJECT --reject-with icmp-host-prohibited
-A INPUT -m state --state RELATED,ESTABLISHED -j ACCEPT
-A INPUT -p icmp -j ACCEPT
-A INPUT -i lo -j ACCEPT
-A INPUT -p tcp -m tcp --dport 22 -j ACCEPT
-A INPUT -p tcp -m tcp --dport 16514 -j ACCEPT
-A INPUT -p tcp -m tcp --dport 54321 -j ACCEPT
-A INPUT -p tcp -m multiport --dports 5634:6166 -j ACCEPT
-A INPUT -p tcp -m multiport --dports 49152:49216 -j ACCEPT
-A INPUT -j REJECT --reject-with icmp-host-prohibited
-A FORWARD -j REJECT --reject-with icmp-host-prohibited
-A FORWARD -m physdev ! --physdev-is-bridged -j REJECT --reject-with icmp-
host-prohibited
COMMIT
# Completed on Fri Dec  14 12:49:12 2012
```

D.3 RHEV-M Client Requirements

The Administration Portal, and the User Portal of RHEV-M can be accessed from a client with a supported web browser. The portals support the following clients and browsers:

- Mozilla Firefox 10, and later, on Red Hat Enterprise Linux is required to access both portals.
- Internet Explorer 8, and later, on Microsoft Windows is required to access the User Portal.
- Internet Explorer 9, and later, on Microsoft Windows is required to access the Administration Portal.

To access virtual machine consoles it is a requirement to install a supported SPICE client.

When accessing the portal(s) using Mozilla Firefox, the SPICE client is provided by the spice-xpi package, which needs to be manually installed using `yum`.
When accessing the portal(s) using Internet Explorer the SPICE ActiveX control will automatically be downloaded and installed.

Appendix E: Red Hat Enterprise Virtualization Manager Installation & Configuration

The process of deploying a RHEV-M server requires the following:

1) Base operating system installation (With SELinux turned on by default in this case)
2) Opening relevant Firewall ports
3) Sourcing and installation of the right RHEV packages on the server
4) RHEV-M configuration using 'rhevm-setup' tool
5) Verification of the RHEV-M environment

Operating System

Please refer to Table 4.1.1-1: Operating System Revisionsfor OS details.

Channel subscription

To install Red Hat Enterprise Virtualization Manager one must first register the target system to Red Hat Network and subscribe to the required software channels.

```
# rhn-channel -add –channel=rhel-x86_64-server-6
# rhn-channel -add –channel=rhel-x86_64-server-supplementary-6
# rhn-channel -add –channel=rhel-x86_64-server-6-rhevm-3.1
# rhn-channel -add --channel=jbappplatform-6-x86_64-server-6-rpm
```

Alternatively, the required channels can be subscribed using Subscription Manager as detailed in section E.3Subscribing to Channels Using Subscription Manager.

E.1 Install and Configure Engine Database

1. Ensuring most up to date versions of all installed packages are in use.

```
# yum upgrade
```

2. Initiation of installation of the rhevm package and all dependencies. This needs to be executed as root

```
# yum install rhevm
```

Configuration

Once package installation is complete the Red Hat Enterprise Virtualization Manager must be configured. The **rhevm-setup** command is provided to assist with this task. Once all required values have been provided, the updated configuration is applied and the Red Hat Enterprise Virtualization Manager services are started.

As root the rhevm-setup needs to be executed.

```
# rhevm-setup

In order to proceed the installer must stop the JBoss service
```

```
Would you like to stop the JBoss service? (yes|no): yes

RHEV Manager uses httpd to proxy requests to the application server.
It looks like the httpd installed locally is being actively used.
The installer can override current configuration .
Alternatively you can use JBoss directly (on ports higher than 1024)
Do you wish to override current httpd configuration and restart the
service?
['yes'| 'no'] [yes] : yes

HTTP Port [80] : 8080

HTTPS Port [443] : 8443

Host fully qualified domain name. Note: this name should be fully
resolvable
[rhevm31.demo.redhat.com] : bu-rhevm.cloud.lab.eng.bos.redhat.com

Password for Administrator (admin@internal) : suitable_password

Organization Name for the Certificate
[demo.redhat.com] :cloud.lab.eng.bos.redhat.com

The default storage type you will be using ['NFS'|'FC'|'ISCSI'] [NFS]:
ISCSI

Enter DB type for installation ['remote'| 'local'] [local] : local

Database password (required for secure authentication with the locally
created database) : suitable_password

Enter DB type for installation ['remote'| 'local'] [local] : local
Enter the host IP or host name where DB is running: 10.16.136.32

Enter DB port number [5432] : 5432

Enter DB admin user name [postgres] : postgres

Remote database password : suitable_password
Confirm password : suitable_password

Configure NFS share on this server to be used as an ISO Domain? ['yes'|
'no'] [yes] : Yes

Local ISO domain path [/usr/local/exports/iso] : /var/lib/exports/iso

Firewall ports need to be opened.
The installer can configure iptables automatically overriding the current
configuration. The old configuration will be backed up.
Alternately you can configure the firewall later using an example iptables
file found under /etc/ovirt-engine/iptables.example
Configure iptables ? ['yes'| 'no']: Yes
```

The entered values are displayed for confirmation. If ‘yes’ is selected, the configuration will

proceed.

```
RHEV Manager will be installed using the following configuration:
==================================================================
override-httpd-config:           yes
http-port:                       8080
https-port:                      8443
host-fqdn:                       rhevm.cloud.lab.eng.bos.redhat.com
auth-pass:                       ********
org-name:                        cloud.lab.eng.bos.redhat.com
default-dc-type:                      ISCSI
db-remote-install:               local
db-local-pass:                   ********
nfs-mp:                          /var/lib/exports/iso
config-nfs:                           yes
override-iptables:               yes
Proceed with the configuration listed above? (yes|no): yes
```

This will complete the RHEV-M installation and point the URL to access RHEV-M admin portal:

https://bu-rhevm.cloud.lab.eng.bos.redhat.com:8443/

```
Additional information displayed in the output to assist with login and
usage:

* A default ISO share has been created on this host.
If IP based access restrictions are required, please edit
/var/lib/exports/iso entry in /etc/exports
* The installation log file is available at: /var/log/ovirt-engine/engine-
setup_2012_08_29_16_38_10.log
* Please use the user "admin" and password specified in order to login
into RHEV Manager
* To configure additional users, first configure authentication domains
using the 'rhevm-manage-domains' utility
```

E.2 Install and Configure History and Reports Database

The Red Hat Enterprise Virtualization Manager optionally includes a comprehensive management history database, which can be utilized by any application to extract a range of information at the data center, cluster, and host levels.

In addition to the history database Red Hat Enterprise Virtualization Manager Reports functionality is also available as an optional component. Red Hat Enterprise Virtualization Manager Reports provides a customized implementation of JasperServer, and JasperReports.

Install Required Packages

Using yum installer as root, both datawarehouse and reports packages can be installed in the RHEV-M server as follows:

```
# yum install rhevm-reports
```

Configure History Database

```
# rhevm-dwh-setup
Would you like to stop the JBoss service? (yes|no):Yes
```

Configure Red Hat Enterprise Virtualization Manager Reports

```
# rhevm-reports-setup
In order to proceed the installer must stop the JBoss service
Would you like to stop the JBoss service? (yes|no):yes
Please choose a password for the admin users (rhevm-admin and superuser):
password
Re-type password: password
```

Once these steps have been completed, the reporting portal can be invoked to view the RHEV-M reports as mentioned below in this case:

https://bu-rhevm.cloud.lab.eng.bos.redhat.com:8443/rhevm-reports/login.html

Also the three databases - **engine**, **ovirt_engine_history** and **rhevmreports** can be viewed on the postgres table output as below:

```
# psql -U postgres
psql (8.4.13)
Type "help" for help.

postgres=# \l
                                  List of databases
         Name         |  Owner   | Encoding |  Collation  |    Ctype    |
Access privileges
----------------------+----------+----------+-------------+-------------
+-----------------------
 engine               | engine   | UTF8     | en_US.UTF-8 | en_US.UTF-8 |
 ovirt_engine_history | postgres | UTF8     | en_US.UTF-8 | en_US.UTF-8 |
 postgres             | postgres | UTF8     | en_US.UTF-8 | en_US.UTF-8 |
 rhevmreports         | postgres | UTF8     | en_US.UTF-8 | en_US.UTF-8 |
 template0            | postgres | UTF8     | en_US.UTF-8 | en_US.UTF-8 |
=c/postgres
                                                                          :
postgres=CTc/postgres
 template1            | postgres | UTF8     | en_US.UTF-8 | en_US.UTF-8 |
=c/postgres
                                                                          :
postgres=CTc/postgres
(6 rows)

postgres=# \q
#
```

E.3 Subscribing to Channels Using Subscription Manager

To install packages signed by Red Hat one must register the target system to Red Hat Network. Then the entitlement from the entitlement pool can be used to subscribe the system to channels.

1. Run the `subscription-manager` ***`register`*** command to register the system with Red Hat Network. To complete registration successfully Red Hat Network **Username** and **Password** must be provided when prompted.

```
# subscription-manager register
```

2. Identify available entitlement pools

   To subscribe the system to channels, the identifiers for the relevant entitlement pools must be located using the ***`list`*** action of the ***`subscription-manager`***.

   For example, following is the command to identify available subscription pools for `Red Hat Enterprise Virtualization`:

```
# subscription-manager list --available | grep -A8 "Red Hat Enterprise
Virtualization"
```

3. Subscribe system to entitlement pools

   Using the pool identifiers located in the previous step, subscribe the system to the required entitlements. When a system is subscribed to an entitlement pool, the system is automatically subscribed to the channels in the entitlement. The main channel is automatically enabled, other channels in the entitlement must be enabled manually:

```
# subscription-manager subscribe –pool=POOLID
```

4. Enable additional subscription channels

   When a system is subscribed to an entitlement with a main channel and some additional channel, only the main channel is enabled by default. Other channels are available, but disabled. The additional channels must be enabled using the `yum-config-manager` command as the **`root`** user:

```
# yum-config-manager --enable CHANNEL
```

The system is now registered with Red Hat Network and subscribed to the channels required.

Appendix F: Hypervisor (Virtualization Host) Installation

Red Hat Enterprise Virtualization supports both virtualization hosts which run the Red Hat Enterprise Virtualization Hypervisor (reffered to as **RHEV Hypervisor**), and those which run Red Hat Enterprise Linux (referred to as **RHEL Hypervisor**). Both types of virtualization hosts can coexist in the same Red Hat Enterprise Virtualization environment. Prior to installing virtualization hosts, it is necessary to ensure that a RHEV-M has been configured, available and all the hardware requirements for the hosts have been met.

F.1 Installing Red Hat Enterprise Virtualization Host (Hypervisor)

This installation contains the following steps:

1) Setting up RHEL server using iso image
2) Configuring the server with RHEV-M information
3) Approve the newly added host to the RHEV-M
4) Network and storage settings
5) ISCSI configuration (if applicable) on the host.

Setting up the RHEL server

The Red Hat Enterprise Virtualization Hypervisor (v.6 x86_64) RHN channel contains the Hypervisor packages. The Hypervisor iso image is contained in the rhev-hypervisor package.

On the RHEV-M server, query if the iso file of the hypervisor exists. If so this iso image file can be copied/used to install the RHEV hypervisor server.

```
# rpm -ql rhev-hypervisor6.noarch | grep 6_3.iso
/usr/share/rhev-hypervisor/rhevh-6.3-20121107.0.el6_3.iso
```

This iso file was loaded as a virtual media on the blade console and the physical host (blade) was booted off this image.

Figure F.1-1: RHEV Hypervisor installation media selection

The server was rebooted after successful installation of this image:

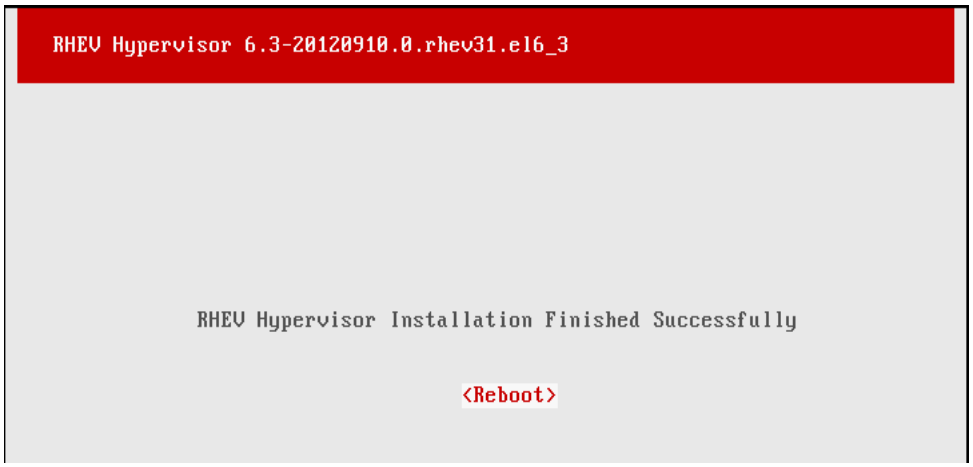

Figure F.1-2: Reboot after Installation

Appropriate options need to be selected such as Language, Disk to use and Administrator Password etc before selecting install. Once the installation has been completed, the reboot button needs to be selected. **Figure 4.3.5-2: RHEV Hypervisor Installation**
After reboot, login as 'admin' with the previously set password.

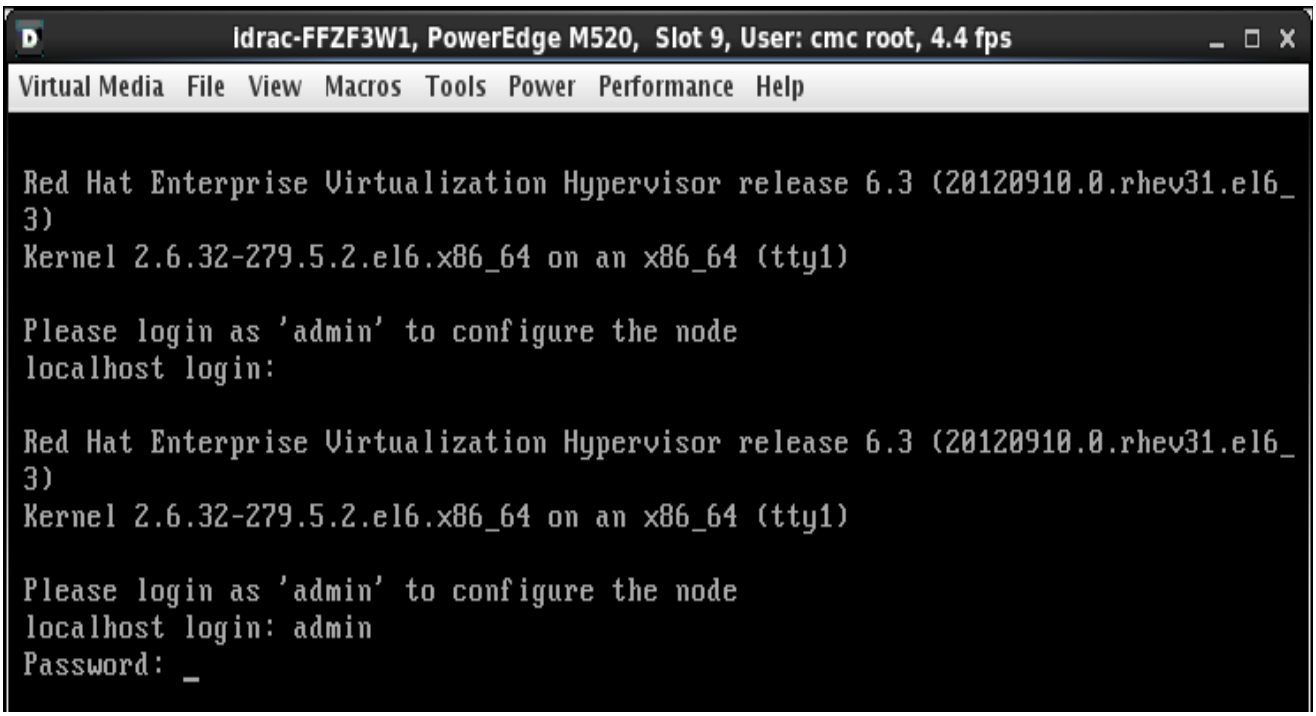

Figure F.1-3: RHEV Hypervisor Configuration login

Upon login, the RHEV Hypervisor tool gets initiated to configure the following settings:

Network Settings

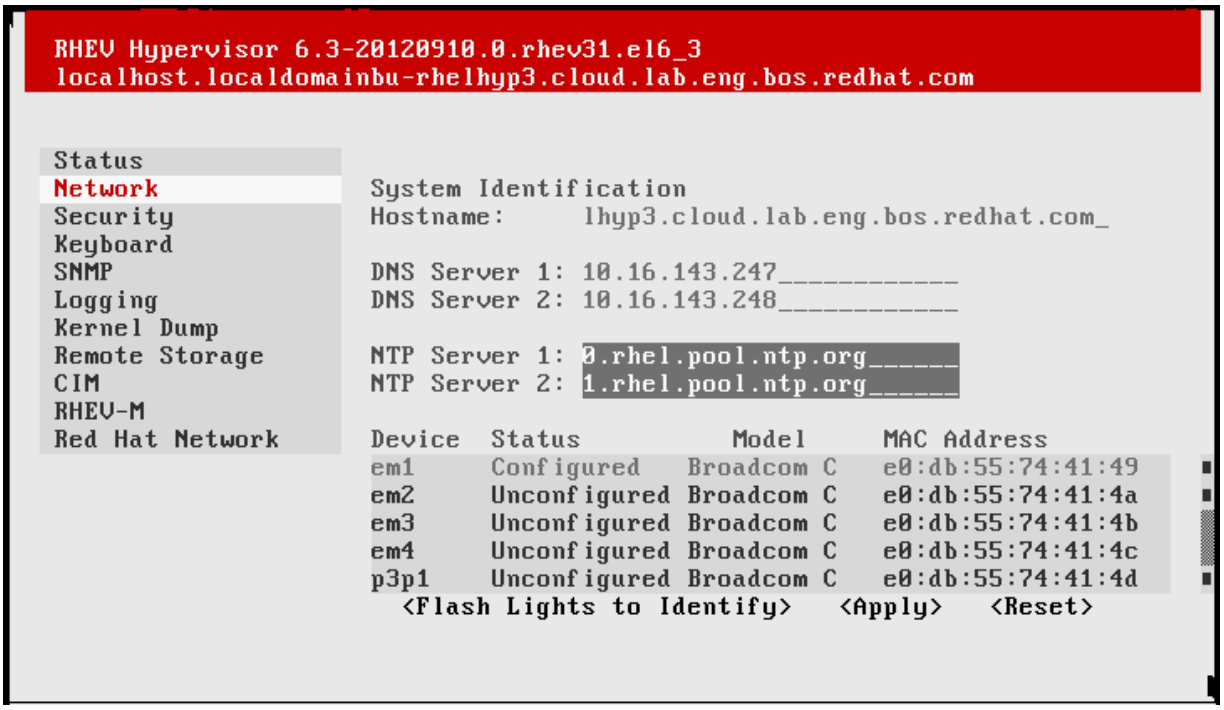

Figure F.1-4: RHEV Hypervisor Configuration -Network

Security

Enable ssh password authentication and provide root password as mentioned below:

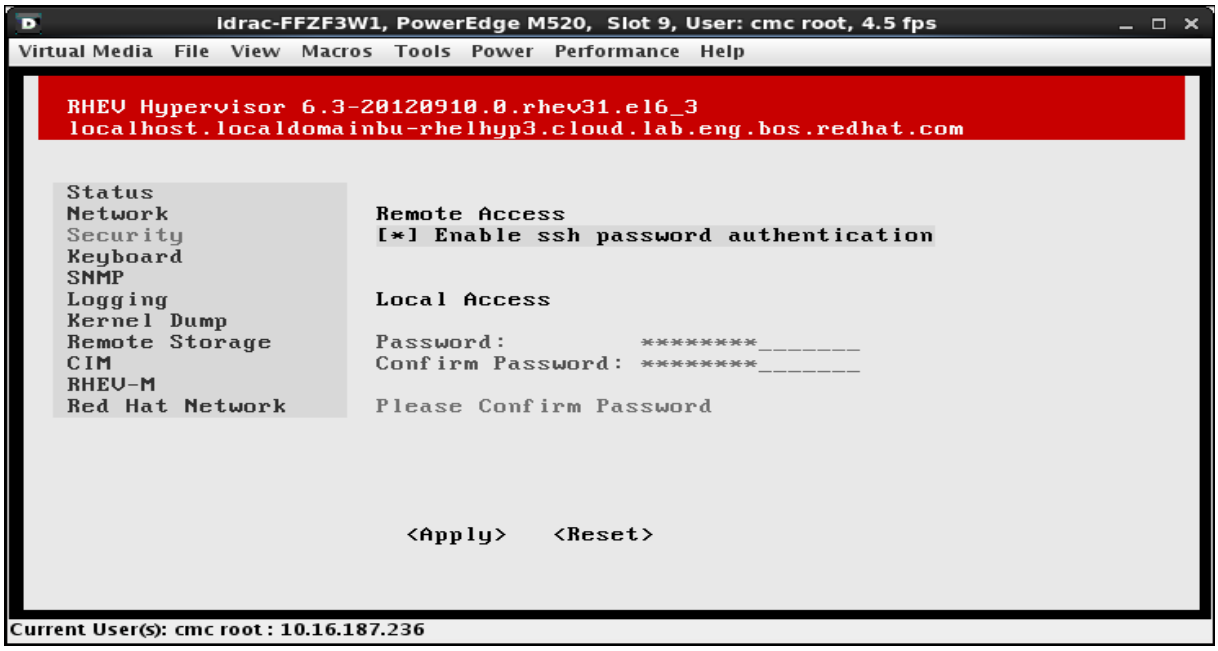

Figure F.1-5: RHEV Hypervisor Configuration -Security

RHEVM

Configure the server with RHEV-M information.

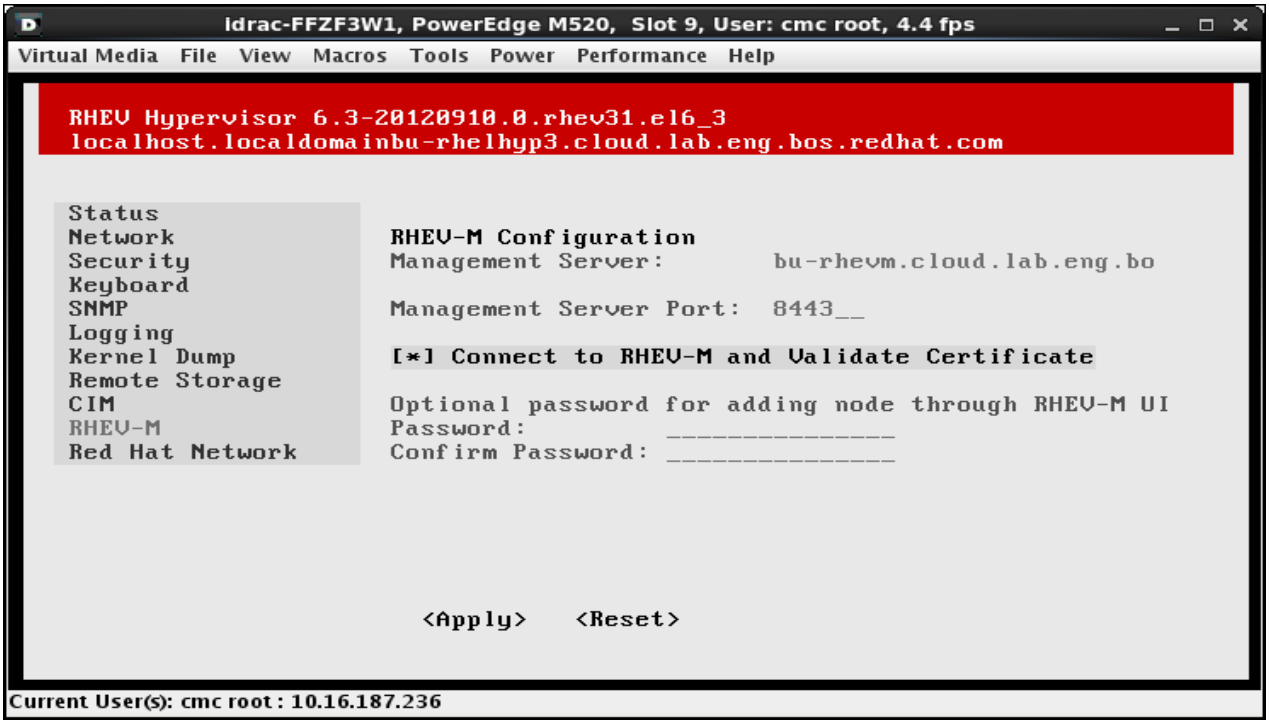

Figure F.1-6: RHEV Hypervisor Configuration - RHEVM

Approve Certificate

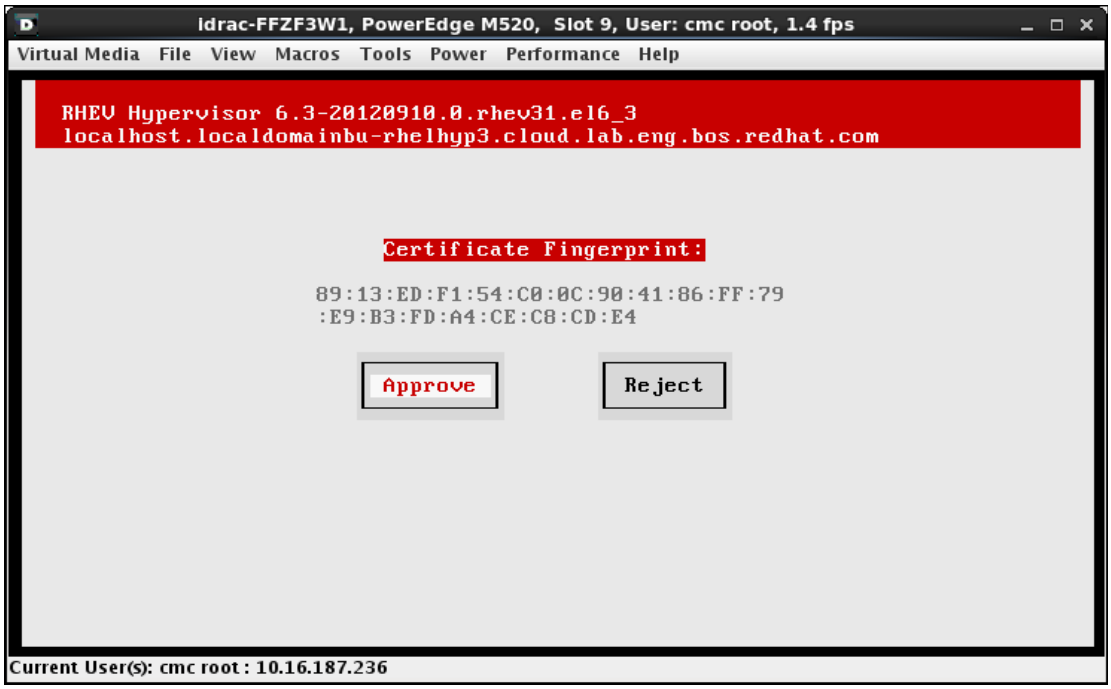

Figure F.1-7: RHEV Hypervisor Configuration - Approval

Confirm selection and restart the server:

Figure F.1-8: RHEV Hypervisor Configuration - Restart

Approving the new host at the RHEV-M portal

After the restart, login to RHEVM admin portal. The new hypervisor (host) will be listed in the host tab waiting for approval.

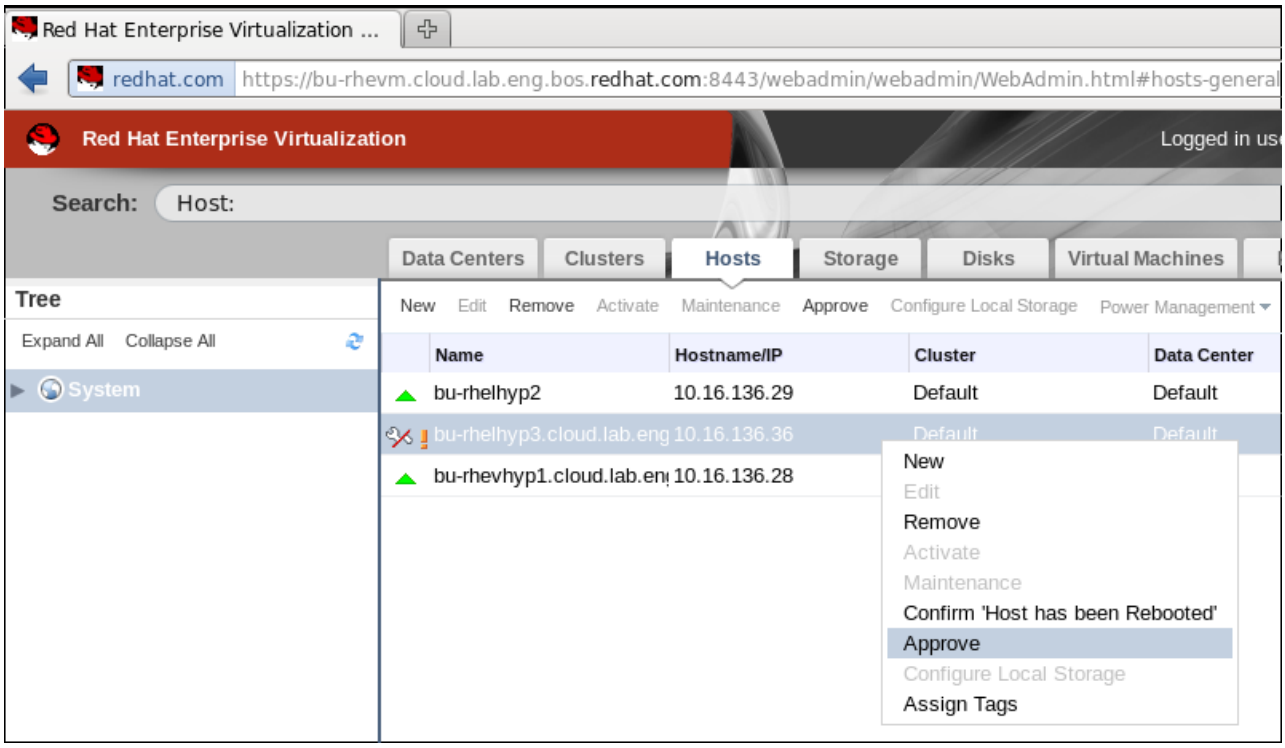

Figure F.1-9: RHEV Hypervisor -Adding to RHEV-M

Configure Power Management

Once the approve button is selected, power management window will pop up.

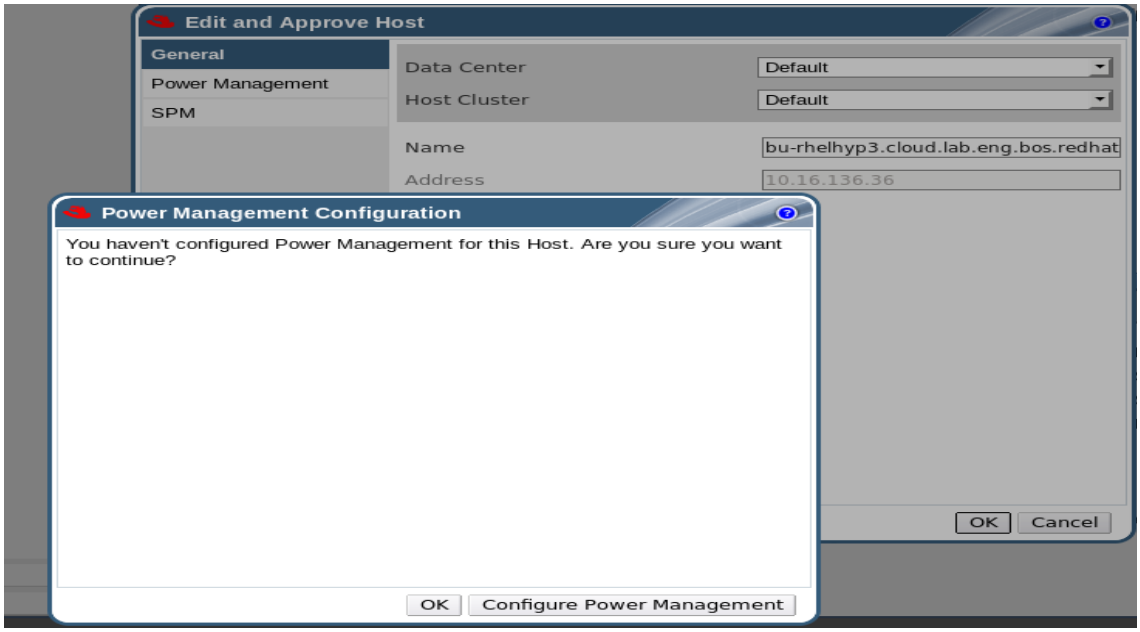

Figure F.1-10: RHEV Hypervisor Power Management

Power Management Settings

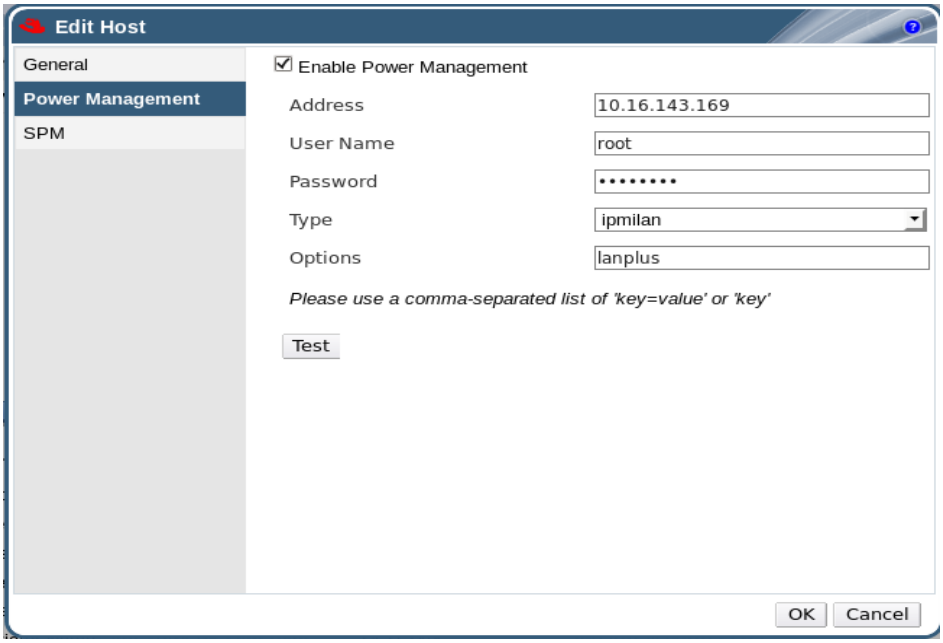

Figure F.1-11: RHEV Hypervisor Power Management Settings

The power management settings are provided with management console information. In this case, since this is a Dell hardware, this value is the same as DRAC console access settings.

Network and storage settings

The next step is to configure network settings for the host.

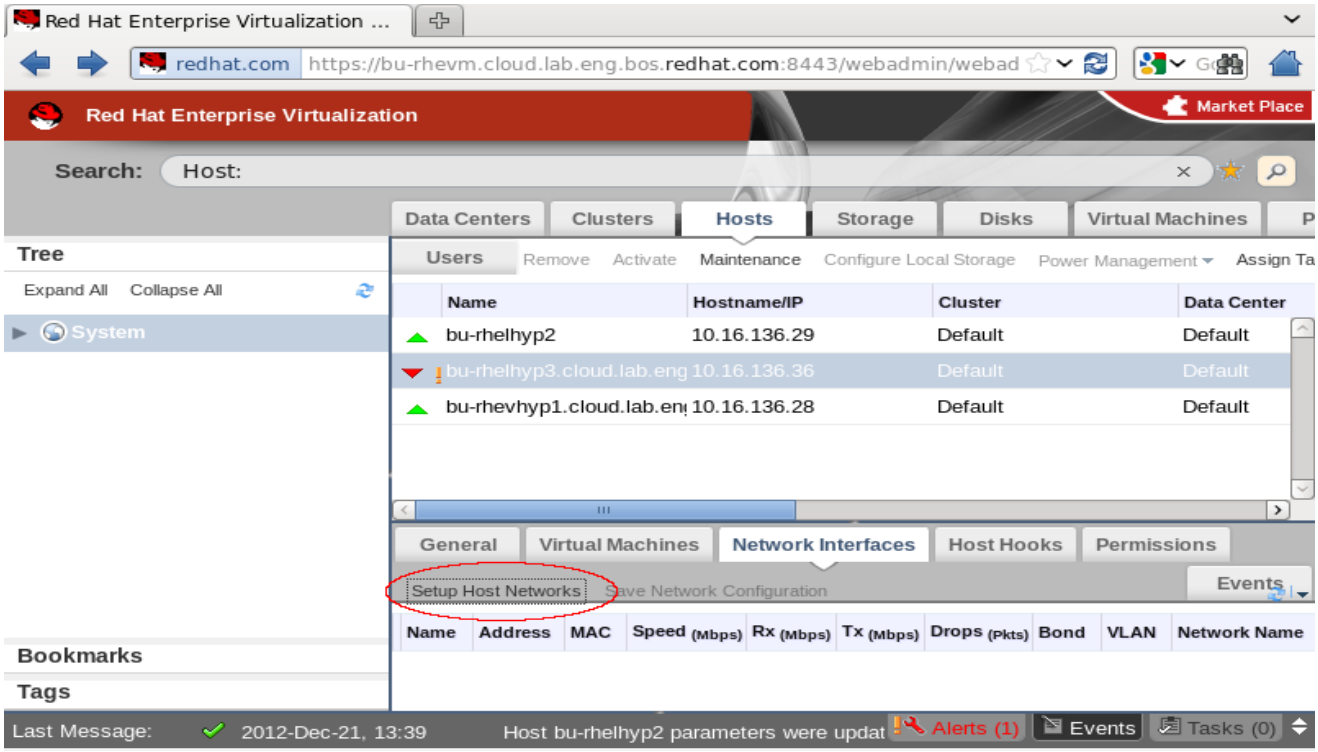

Figure F.1-12: RHEV Hypervisor -Network Settings 1

Unassigned Logical Networks:

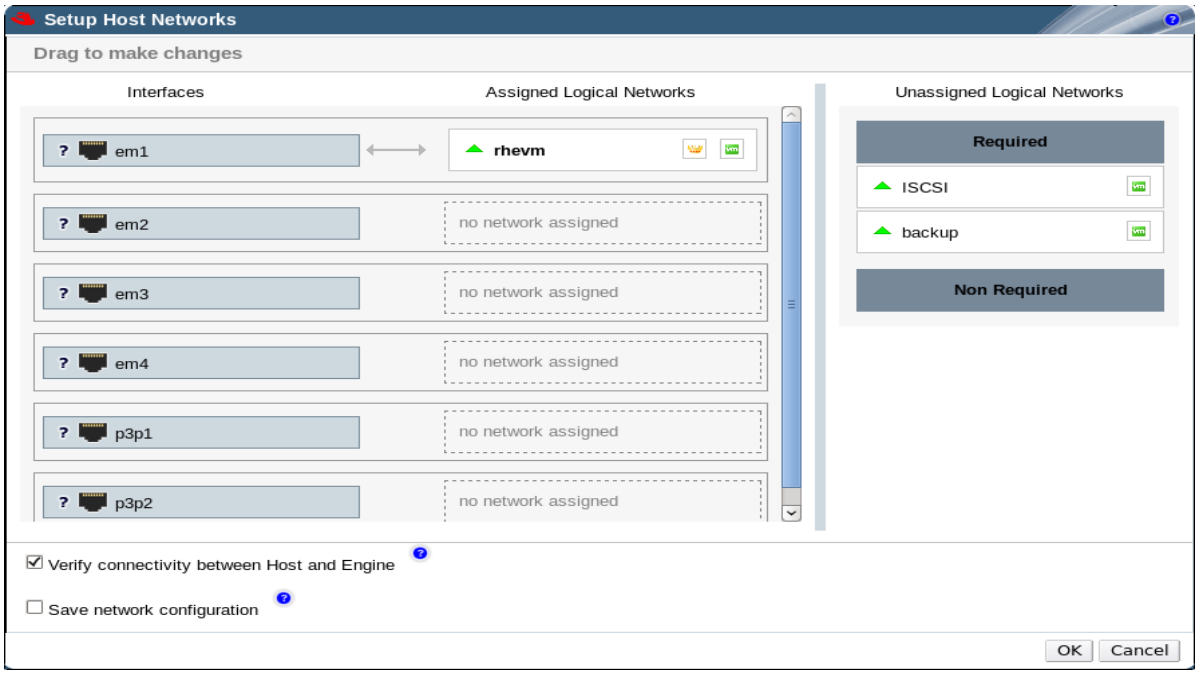

Figure F.1-13: RHEV Hypervisor -Network Settings 2

Please note that the public interface is already configured by default (rhevm interface).

ISCSI and *backup* logical networks are added by dragging the boxes on the right to the appropriate interfaces as detailed below.

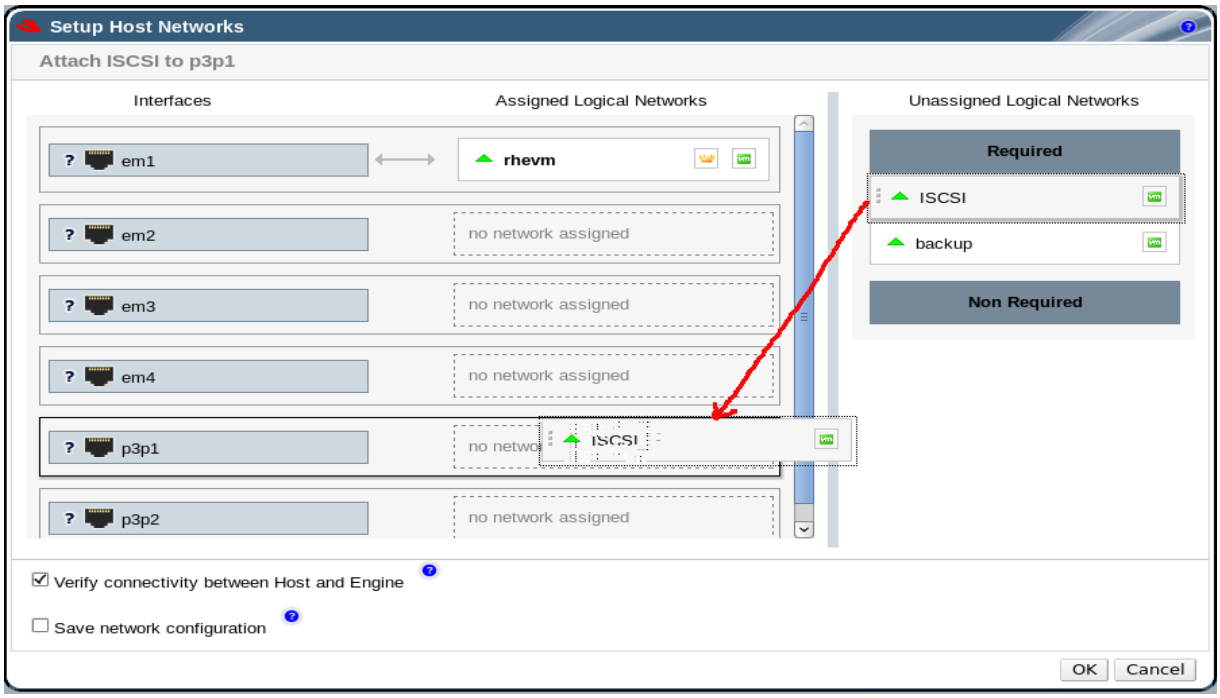

Figure F.1-14: RHEV Hypervisor -Network Settings 3

After assigning the logical networks to the right interface, the Setup Host Networks Window appears as follows:

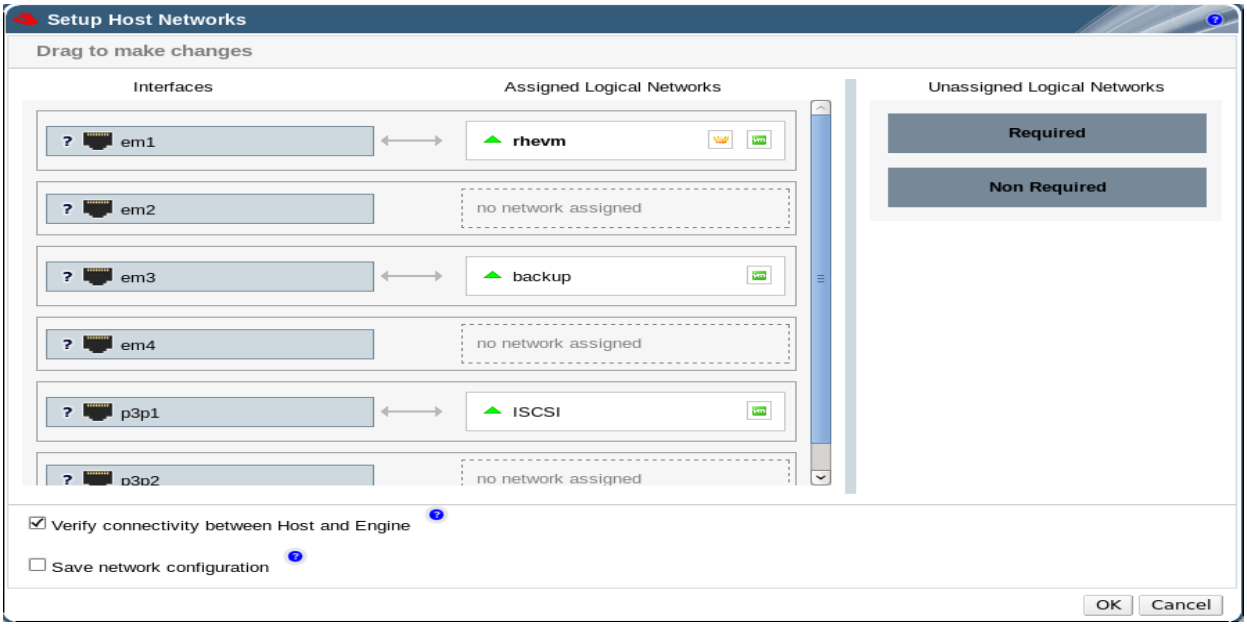

Figure F.1-15: RHEV Hypervisor -Network Settings 4

Select the pencil button by placing the mouse on top of the box as displayed below:

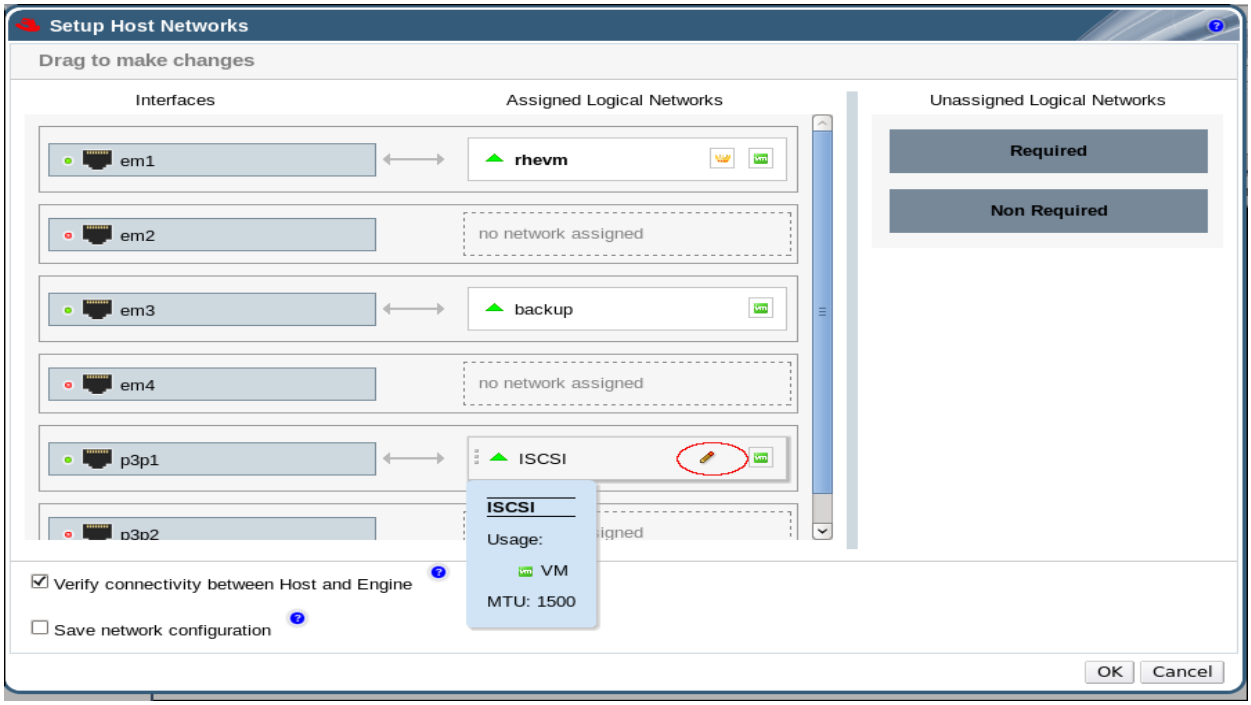

Figure F.1-16: RHEV Hypervisor - Network Settings 5

Enter network settings pertaining to the interface:

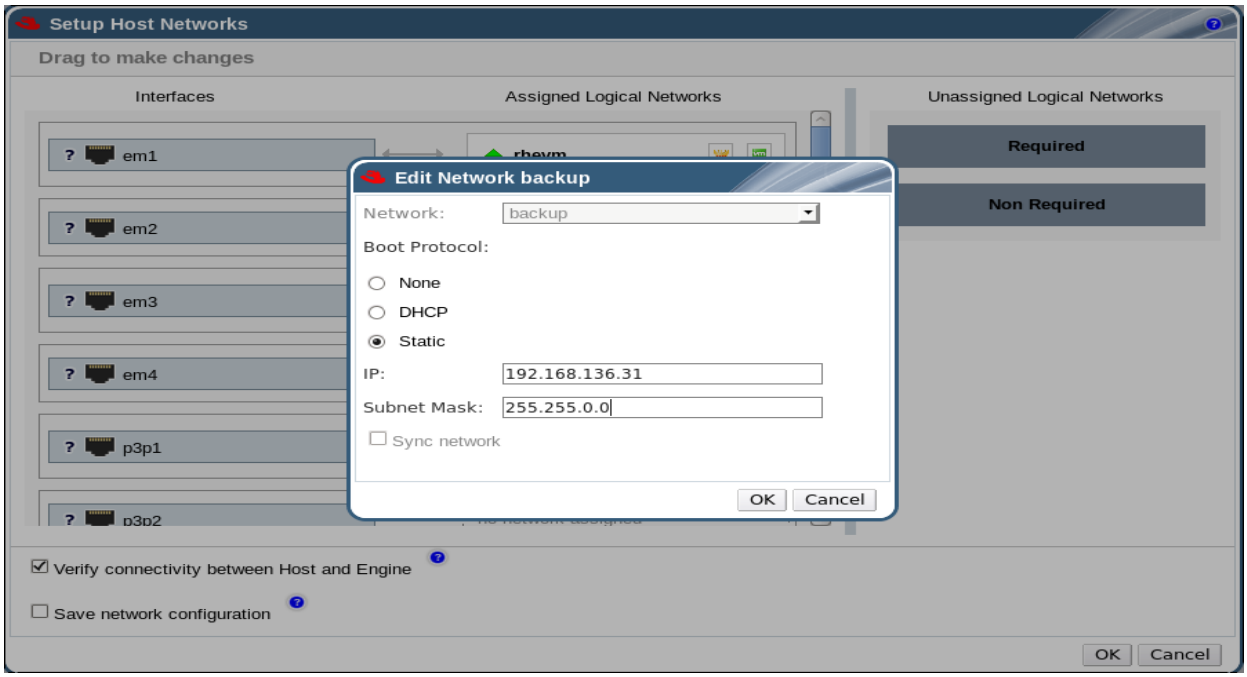

Figure F.1-17: RHEV Hypervisor - Network Settings 6

Once all the setting have been entered, check the *Save network configuration* box to complete the step.

Once all the network settings have been completed, the active interfaces are displayed as follows:

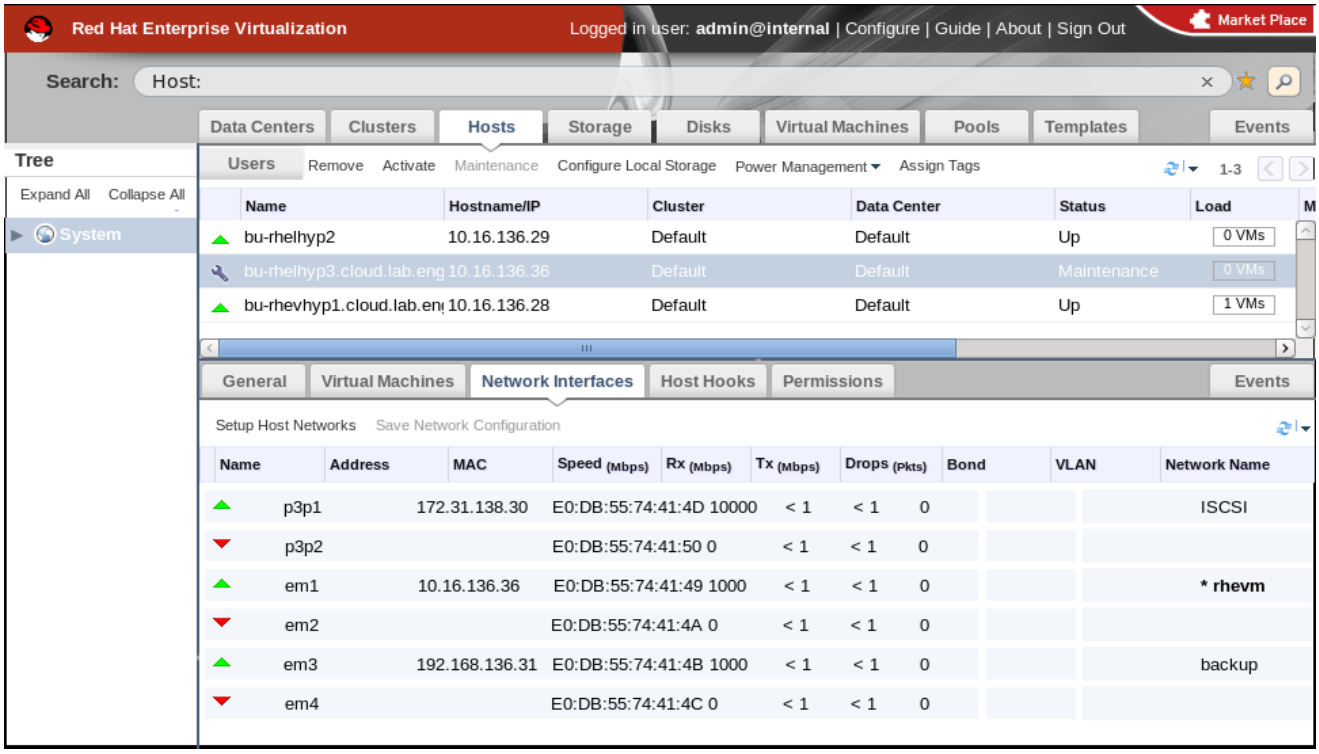

Figure F.1-18: RHEV Hypervisor - Network Settings 7

With all the network interfaces active (primary interface: *rhevm*, backup interface: *backup* & ISCSI interface: *ISCSI*), the host was activated to display a status of 'Up' with green indicator.

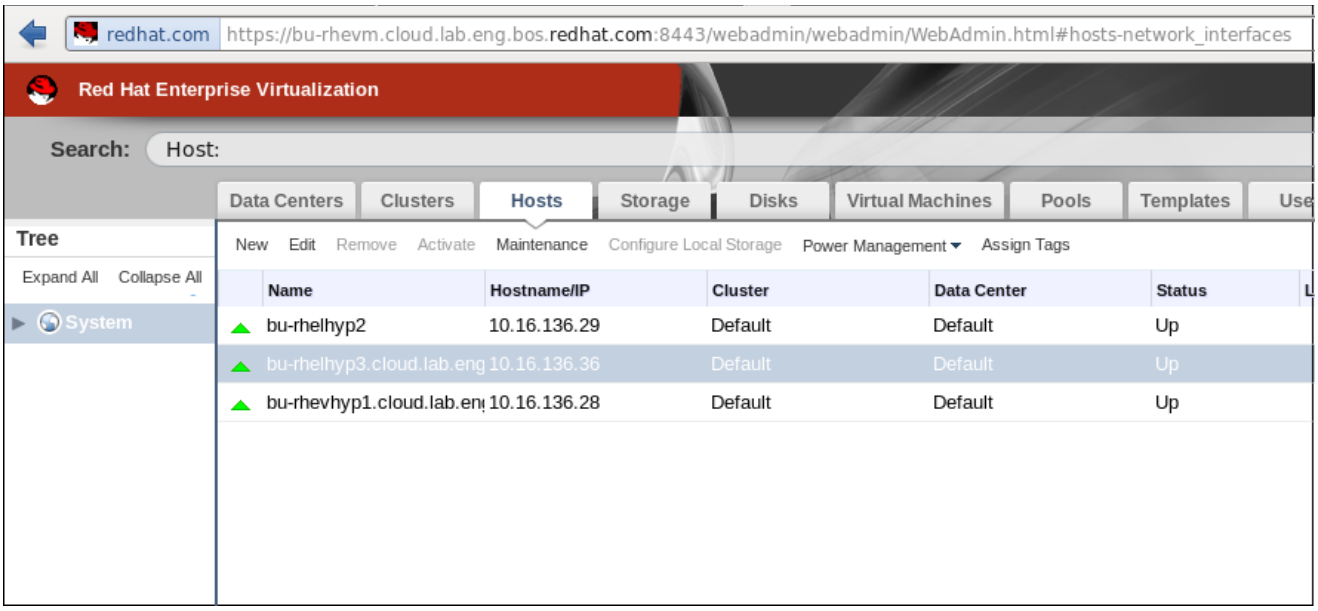

Figure F.1-19: RHEV Active Hosts

F.2 Red Hat Enterprise Linux Hypervisor Configuration

The process of deploying a RHEL Hypervisor requires the following:

1. Base operating system installation (With SELinux turned on by default in this case).
2. Adding RHN Channels relevant to RHEV Management Agents.
3. ISCSI settings to access the shared storage for the bu-domain in the RHEV environment.
4. Verification of the RHEL Hypervisor environment.

Registration and Subscription

The server needs to be registered and the required channels needs to be subscribed.
As root the server was registered providing RedHat Network username and password.

```
# rhn_register
```

The required channels can be subscribed or added. In this case, *rhn-channel* command was used to add the channels.

```
# rhn-channel --add --channel=rhel-x86_64-server-6
# rhn-channel --add –channel=rhel-x86_64-rhev-mgmt-agent-6
```

Adding a Red Hat Enterprise Linux Host

1. Click the **Hosts** resource tab to list the hosts in the results list.
2. Click **New** to open the **New Host** window.
3. Use the drop-down menus to select the **Data Center** and **Host Cluster** for the new host.
    1. Enter the **Name**, **Address**, and **Root Password** of the new host.
4. **Select Automatically configure host firewall** check box.
5. Configure the **Power Management** and **SPM** using the applicable tabs
6. Click **OK** to add the host

Alternatively step 5 can be skipped and firewall settings can be manually setup as mentioned in Appendix section D.2.3Firewall Requirements.

Appendix G: Storage Configuration

This environment details use of ISCSI storage. Once the ISCSI configuration has been completed at the ISCSI Storage Array, the following two steps are followed to utilize the newly carved storage in the RHEV environment:

1) ISCSI configuration at the Hypervisor
2) Creating a storage domain with the new storage

G.1 ISCSI Configuration at the Hypervisor

In this setup, a 200GB lun was created at the ISCSI Storage Array to be used as a shared storage for the VMs in the RHEV environment.

On the Hypervisor the initiatorname.iscsi and iscsid.conf files were configured for authentication:
An unique initiator name for the target has to be created for each host the target/volume is being shared. Edit the /etc/iscsi/*initiatorname.iscsi* file to add this unique initiator name created at the ISCSI storage array:

```
#
InitiatorName=iqn.1994-05.com.redhat:bu-hyp3
```

Edit the /etc/iscsi/*iscsid.conf* file to add authentication information (bold entries):

```
#
# Open-iSCSI default configuration.
# Could be located at /etc/iscsi/iscsid.conf or ~/.iscsid.conf
#
# Note: To set any of these values for a specific node/session run
# the iscsiadm --mode node --op command for the value. See the README
# and man page for iscsiadm for details on the --op command.
#
#####################
# iscsid daemon config
#####################
# If you want iscsid to start the first time a iscsi tool
# needs to access it, instead of starting it when the init
# scripts run, set the iscsid startup command here. This
# should normally only need to be done by distro package
# maintainers.
#
# Default for Fedora and RHEL. (uncomment to activate).
iscsid.startup = /etc/rc.d/init.d/iscsid force-start
#
# Default for upstream open-iscsi scripts (uncomment to activate).
# iscsid.startup = /sbin/iscsid

############################
# NIC/HBA and driver settings
############################
# open-iscsi can create a session and bind it to a NIC/HBA.
```

```
# To set this up see the example iface config file.

#*****************
# Startup settings
#*****************

# To request that the iscsi initd scripts startup a session set to
"automatic".
# node.startup = automatic
#
# To manually startup the session set to "manual". The default is
automatic.
node.startup = automatic

# For "automatic" startup nodes, setting this to "Yes" will try logins on
each
# available iface until one succeeds, and then stop.  The default "No" will
try
# logins on all available ifaces simultaneously.
node.leading_login = No

# *************
# CHAP Settings
# *************

# To enable CHAP authentication set node.session.auth.authmethod
# to CHAP. The default is None.
#node.session.auth.authmethod = CHAP
node.session.auth.authmethod = CHAP

# To set a CHAP username and password for initiator
# authentication by the target(s), uncomment the following lines:
#node.session.auth.username = username
#node.session.auth.password = password
node.session.auth.username = initiatorCHAP
node.session.auth.password = password_string

# To set a CHAP username and password for target(s)
# authentication by the initiator, uncomment the following lines:
#node.session.auth.username_in = username_in
#node.session.auth.password_in = password_in

# To enable CHAP authentication for a discovery session to the target
# set discovery.sendtargets.auth.authmethod to CHAP. The default is None.
#discovery.sendtargets.auth.authmethod = CHAP
# To set a discovery session CHAP username and password for the initiator
# authentication by the target(s), uncomment the following lines:
#discovery.sendtargets.auth.username = username
#discovery.sendtargets.auth.password = password

# To set a discovery session CHAP username and password for

                              ...output abbreviated...
```

Enable iscsi and iscsid services:

```
# service iscsi stop
                                                           [  OK  ]
# service iscsid stop
                                                           [  OK  ]
# service iscsi start
                                                           [  OK  ]
# service iscsid start
                                                           [  OK  ]
```

Verify the iscsi status:

```
# service iscsi status
iSCSI Transport Class version 2.0-870
version 2.0-872.41.el6
Target: iqn.2001-05.com.equallogic:0-1cb196-c6260c201-55d0000000d50951-bu-
hyp
      Current Portal: 172.31.143.201:3260,1
      Persistent Portal: 172.31.143.200:3260,1
            **********
            Interface:
            **********
            Iface Name: default
            Iface Transport: tcp
            Iface Initiatorname: iqn.1994-05.com.redhat:bu-hyp3
            Iface IPaddress: 172.31.138.30

            *****
            CHAP:
            *****
            username: initiatorCHAP
            password: ********
            username_in: <empty>
            password_in: ********
            ************************
            Negotiated iSCSI params:
            ************************
                        ...output abbreviated...

            **********
            Interface:
            **********
            Iface Name: default
            Iface Transport: tcp
            Iface Initiatorname: iqn.1994-05.com.redhat:bu-hyp3
            Iface IPaddress: 172.31.138.30
            Iface HWaddress: <empty>
            Iface Netdev: <empty>
            SID: 38
            iSCSI Connection State: LOGGED IN
            iSCSI Session State: LOGGED_IN
```

```
            Internal iscsid Session State: NO CHANGE
            *********
            Timeouts:
            *********
            Recovery Timeout: 120
            Target Reset Timeout: 30
            LUN Reset Timeout: 30
            Abort Timeout: 15

            ************************
            Attached SCSI devices:
            ************************
            Host Number: 55  State: running
            scsi55 Channel 00 Id 0 Lun: 0
                  Attached scsi disk sdc      State: running
```

ISCSI Target discovery

Identify if the host can see the target disk at the array:

```
# iscsiadm -m discovery -t st -p 172.31.143.200:3260
172.31.143.200:3260,1 iqn.2001-05.com.equallogic:0-1cb196-c6260c201-
55d0000000d50951-bu-hyp
```

Login to the disk:

```
# iscsiadm -m node --targetname iqn.2001-05.com.equallogic:0-1cb196-
c6260c201-55d0000000d50951-bu-hyp -p 172.31.143.200:3260--login
# BEGIN RECORD 2.0-872.41.el6
node.name = iqn.2001-05.com.equallogic:0-1cb196-c6260c201-55d0000000d50951-
bu-hyp
node.tpgt = 1
node.startup = automatic
node.leading_login = No
iface.hwaddress = <empty>
iface.ipaddress = <empty>
                          ...output abbreviated...
```

This ensures that the storage target/volume created at the Storage array has been presented to the hypervisor host. This step has to be repeated for each host that shares the target/volume.

G.2 Adding a New Storage Domain

This step is to create a storage domain using RHEV-M admin console.
Red Hat Enterprise Virtualization platform supports iSCSI storage by creating a storage domain from a volume group made of pre-existing LUNs. Neither volume groups nor LUNs can be attached to more than one storage domain at a time.

Adding iSCSI Storage :

1. Click the *Storage* resource tab to list the existing storage domains in the results list
2. Click the *New Domain* button to open the New Domain window
3. Enter the Name of the new storage domain

4. Select the datacenter (Default in this case)
5. Select am active Use host. This must a Hypervisor already attached to the RHEV-M

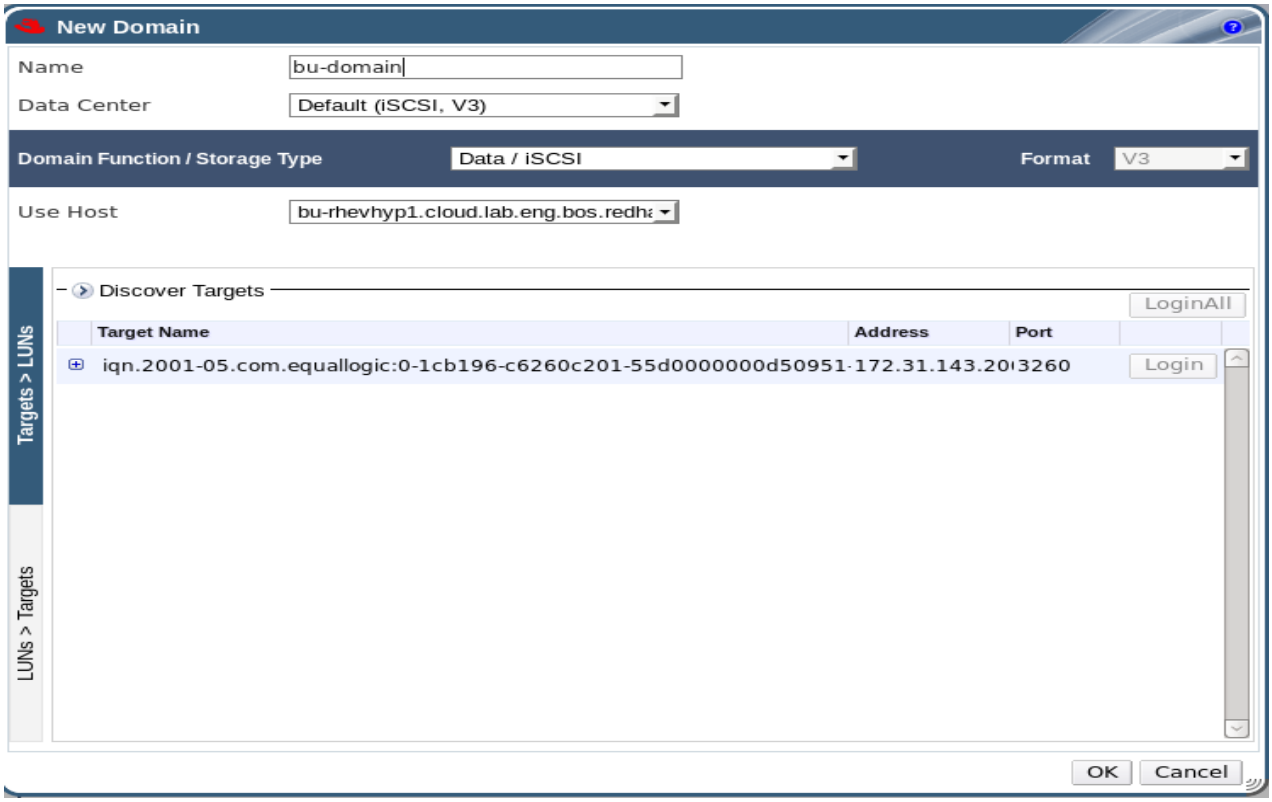

Figure G.2-1: ISCSI Storage Configuration

Appendix H: Bare Metal Restore Essentials

Bare Metal Restore (BMR) is the Server Recovery option of NetBackup. BMR is an automated server recovery process that helps avoid manual re-install of operating systems or configure hardware. Using this, complete server restores can be accomplished either on the same hardware or onto a similar hardware should the original server become unusable. Bare Metal Restore requires additional Symantec licensing and could be either configured during initial NetBackup setup or can be added upon at a later time.

Bare Metal Restore works with normal NetBackup backups. Clients are backed up to NetBackup servers as always. However, an additional procedure automatically runs before every scheduled backup to record the state of the machine configuration, including disk layouts and TCP/IP configuration. BMR automatically captures and records any changes to a machine's configuration at the next scheduled backup. For more details on BMR please refer to the administration guide - http://www.symantec.com/business/support/index?page=content&id=DOC5163

The BMR restore process begins by booting the client . It can be rendered by network boot or media boot (DVD/CD etc). This document focuses on using boot media only.

BMR requirements:

- Master Server that controls the BMR process. Bundled with NetBackup master.
- Boot servers that manage and provide the resources used to rebuild systems.
- Client software that is installed when the NetBackup client software is installed. No special installation or configuration is required. This can be installed on the master server or on a client.
- Shared Resource Tree SRT. This is a collection of software and configurations whose purpose is to bring the protected system to a state from which the original files can be restored.

H.1 BMR Master Server

Bare Metal Restore components get installed when NetBackup is installed. `compat-libstdc++` needs to be installed on the server as a pre-requisite:

```
# rpm -aq | grep compat-libstdc
compat-libstdc++-33-3.2.3-61.x86_64
```

After the master has been installled, BMR needs to be activated with a license and BMR database needs to be created. After activation, BMR master server daemon needs to be running. If not, it needs to be started:

```
#/usr/openv/netbackup/bin/rc.bmrd start
```

The following command, creates the BMR database:

```
#/usr/openv/netbackup/bin/bmrsetupmaster
```

H.2 Boot Server Registration

```
# ./bmrsetupboot -register
[Info] V-127-76 Setting up BMR boot server completed successfully.
```

H.3 Shared Resource Tree

A shared resource tree (SRT) is a collection of the following:

- Operating system files
- NetBackup client software
- Programs that format drives, create partitions, rebuild file systems, and restore the original files using the NetBackup client software
- Resources needed to boot the client system and begin the restore process.

In this is performed in a two step process.

1) Create a generic SRT
2) Create a CD image version (iso file in this case) based on the generic SRT

Also the version of NetBackup used here is NetBackup 7.5.0.5. Hence the SRT is created with NetBackup Client 7.5 followed by adding a NetBackup Maintenance Pack as described below to upgrade it to version 7.5.0.5.

Software needed for the SRT:
RHEL 6.3 software: */data/6.3/iso/rhel-server-6.3-x86_64-dvd.iso*
NetBackup Client software 7.5: */data/netbackup7.5/netbackup_7.5_clients.iso*
NetBackup Client Software 7.5.0.5 /data/netbackup7.5.0.5/NB_CLT_7.5.0.5 (directory)

The NetBackup Client 7.5.0.5 enhancement to NetBackup 7.5 was downloaded from Symantec website: http://www.symantec.com/business/support/index?page=content&id=TECH194138

Creating a generic SRT

```
# cd /usr/openv/netbackup/bin
# ./bmrsrtadm

Select one of the following options:
    1.  Create a new Shared Resource Tree.
    2.  Create a new CD image based Shared Resource Tree.
    3.  Copy an existing Shared Resource Tree to a new location.
    4.  Import a Shared Resource Tree.
    5.  Modify an existing Shared Resource Tree.
    6.  Delete an existing Shared Resource Tree.
    7.  List Shared Resource Trees available on this server.
    8.  Quit.
```

```
Enter your selection (1-8) [1] : 1
Enter the name of the SRT to create : rhel6.3_main_srt
Enter the description of the new SRT :
Main SRT for Rhel6.3 on NB 7.5.0.5 reference
Enter the desired RedHat level (3/3.0 or 4/4.0 or 5/5.0 or 6/6.0) [6.0] :
Enter the desired architecture (i686 or x86_64) [x86_64] :
Enter the directory in which to place the new SRT [/srt] :
Creating repository to stow files from distribution media.

The following media is required:

    Red Hat Enterprise Linux Server release 6.0 (x86_64) - disc 1 of 5

Please load the media now.

Load media from (? for help) [/dev/cdrom] :
/data/6.3/iso/rhel-server-6.3-x86_64-dvd.iso
Mounting media ... ok.
Extracting
files ..........................................................................
.............................. ok.
Unmounting media ... ok.
The media repository is now complete.
Working ........................................................................
............... ok.
Working ................................................... ok.
Preparing boot files - please stand by...

The following media is required:

    NetBackup Client

Please load the media now.

Load media from (? for help) [/data/netbackup7.5/netbackup_7.5_clients.iso]
:
Mounting media ... ok.
The NetBackup Client installation script will run now.


Symantec Installation Script
Copyright 1993 - 2012 Symantec Corporation, All Rights Reserved.


     Installing NetBackup Client Software


Please review the SYMANTEC SOFTWARE LICENSE AGREEMENT located on
the installation media before proceeding.  The agreement includes
details on the NetBackup Product Improvement Program.

For NetBackup installation and upgrade information specific to your
platform and to find out if your installed EEBs or hot fixes are
contained in this release, check out the Symantec Operations Readiness
Tools (SORT) Installation and Upgrade Checklist and Hot fix and EEB
```

```
Release Auditor, respectively, at https://sort.symantec.com/netbackup.

Do you wish to continue? [y,n] (y)
Do you want to install the NetBackup client software for this client? [y,n]
(y)

This package will install Linux/RedHat2.6.18 client.

This package will install NetBackup client 7.5.

Enter the name of the NetBackup server :
bu-netbackup.cloud.lab.eng.bos.redhat.com

Would you like to use "bu-netbackup.cloud.lab.eng.bos.redhat.com" as the
configured
name of the NetBackup client? [y,n] (y)

Client binaries are located in
/mnt/NBClients/anb/Clients/usr/openv/netbackup/client/Linux/RedHat2.6.18.


Installing PBX...
Please wait while installation is in progress...
Installation completed Successfully
Installation log located here: /var/tmp/installpbx-6354-020113145150.log


Unpacking SYMCnbclt package.
Checking for pre-existing SYMCnbclt package.
Installing SYMCnbclt package.
Installation of SYMCnbclt was successful.
ore details regarding SYMCnbclt can be found in file
/tmp/install_cltpkg_trace.6273 on bu-
netbackup.cloud.lab.eng.bos.redhat.com.
Checking network connections.
DNS Lookup failed for host bu-netbackup.cloud.lab.eng.bos.redhat.com
error:-3!
bp.conf: IP_ADDRESS_FAMILY = AF_INET: default value, no update needed

No [x]inetd process found.


File /usr/openv/tmp/install_trace.6059 contains a trace of this install.
That file can be deleted after you are sure the install was successful.

The NetBackup Client installation script has completed.
Unmounting media ... ok.
[Info] V-125-668 SRT "rhel6.3_main_srt" has been initialized successfully.

------------------------------------------------------------------------
---
SRT name:      rhel6.3_main_srt
Location:      /srt/rhel6.3_main_srt
Description:   Test SRT for Rhel6.3 on NB 7.5.0.5
Exclusive use: (none)
```

```
--------------------------------------------------------------------------------

You may make modifications to this SRT.
Select one of the following options:

    1.  Install Symantec NetBackup Maintenance Pack.
    2.  Install Veritas Volume Manager and Veritas File System.
    3.  Update Veritas Volume Manager and Veritas File System.
    4.  Install Veritas Security Services.
    5.  Install additional patches/drivers.
    6.  Change SRT description.
    7.  Change client exclusive use of this SRT.
    8.  Quit.

Enter your selection (1-8) [8] : 1

The following media is required:

    NetBackup Client Maintenance Pack

Please load the media now.

Load media from (? for help, 'none' to return to the menu)
[/data/netbackup7.5.0.5/NB_CLT_7.5.0.5] :
Mounting media ... mount: warning: /srt/rhel6.3_main_srt/mnt seems to be
mounted read-write.
ok.
The NetBackup Client installation script will run now.

There are 2 packs available in /mnt:
(* denotes installed pack)

        NB_7.5.0.5
        NB_CLT_7.5.0.5

Enter pack name (or q) [q]: NB_CLT_7.5.0.5


                      ... output abbreviated ...



The NetBackup Client installation script has completed.
Unmounting media ... ok.

You may make modifications to this SRT.
Select one of the following options:

    1.  Install Symantec NetBackup Maintenance Pack.
    2.  Install Veritas Volume Manager and Veritas File System.
    3.  Update Veritas Volume Manager and Veritas File System.
    4.  Install Veritas Security Services.
    5.  Install additional patches/drivers.
    6.  Change SRT description.
    7.  Change client exclusive use of this SRT.
```

```
    8.  Quit.

Enter your selection (1-8) [8] :
[Info] V-125-669 SRT "rhel6.3_main_srt" has been created successfully.
```

As the second step, create a CD image based SRT

```
# ./bmrsrtadm

Select one of the following options:

    1.  Create a new Shared Resource Tree.
    2.  Create a new CD image based Shared Resource Tree.
    3.  Copy an existing Shared Resource Tree to a new location.
    4.  Import a Shared Resource Tree.
    5.  Modify an existing Shared Resource Tree.
    6.  Delete an existing Shared Resource Tree.
    7.  List Shared Resource Trees available on this server.
    8.  Quit.

Enter your selection (1-8) [1] : 2
Enter the name of an existing SRT : rhel6.3_main_srt
Enter the name of the new SRT to create : RHEL63_NB7504_iso_srt
Enter the description of the new SRT [Test SRT for Rhel6.3 on NB 7.5.0.5] :
CD image SRT for RHEL63 on NB7505
Enter the directory in which to place the new SRT CD image [/srt] :
[Info] V-125-718 Media image has been successfully created.
```

Verify the new srt iso file:

```
# ll /srt | grep RHEL63_NB7505
-rw-r--r--.  1 root root 2072152064 Feb  1 14:58 RHEL63_NB7505_iso_srt.iso
```

This image file can be used as the boot media (uploaded as virtual CD in this case during BMR restore)