Technical Report

Red Hat Enterprise Virtualization 3.0 and NetApp Storage Deployment Guide

Jon Benedict, NetApp
March 2012 | TR-3940

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

1 Deploying Red Hat Enterprise Virtualization on NetApp

1.1 Overview

NetApp technology enables data centers to extend their virtual infrastructures to include the benefits of advanced storage virtualization. NetApp unified storage platforms offer industry-leading technologies in the areas of storage efficiencies, instantaneous virtual machine (VM) and datastore cloning for virtual servers, and virtual data center backup and business continuance solutions. Red Hat Enterprise Linux® (RHEL) also offers the benefits of flexible deployment:

- It can be deployed as a bare-metal operating system, as a hypervisor, or as a virtual guest operating system.
- From a storage perspective, RHEL KVM supports both storage area network (SAN: iSCSI, Fibre Channel [FC], Fibre Channel over Ethernet [FCoE]) and network-attached storage (NAS: Network File System [NFS]) for shared virtual machine storage.

NetApp and Red Hat maintain a long-term strategic alliance that includes end-to-end solution testing between Red Hat products and NetApp storage. As a result of this testing, NetApp has developed operational guidelines and best practices for storage arrays running the NetApp Data ONTAP® operating system in support of Red Hat Enterprise Linux. These guidelines have been extended to include RHEL-based KVM virtualization.

The following sections provide the steps necessary to deploy Red Hat Enterprise Virtualization 3.0 and NetApp storage together in accordance with TR-3914: Red Hat Enterprise Virtualization 3 and NetApp Storage: Best Practices Guide.

1.2 Intended Audience

This document is intended for system architects, system administrators, and storage administrators who are deploying Red Hat Enterprise Virtualization on NetApp storage.

1.3 Scope

The steps provided are specific to where Red Hat and NetApp technologies intersect. That is to say that because Red Hat and NetApp provide excellent product documentation, this document cover topics that are different or not covered by existing documents.

1.4 Use in Conjunction with Existing Documents

To avoid unnecessary duplication in the areas of RHEL 6 KVM hosts and virtual machines, some sections of this deployment guide have been truncated. In these cases, there are references made to existing documents such as:

- TR-3848: RHEL 6, KVM, and NetApp Storage: Best Practices Guide
- TR-4034: Red Hat Enterprise Linux 6, KVM, and NetApp Storage Deployment Guide
- TR-3914: Red Hat Enterprise Virtualization 3 and NetApp Storage: Best Practices Guide
- Red Hat Enterprise Linux 6 Product Guides
- Red Hat Enterprise Virtualization Product Guides

These documents should be reviewed thoroughly prior to planning or deploying RHEV 3.0 with NetApp storage. Additionally, this document does not attempt to recreate Red Hat's product documentation. Instead, this document attempts to highlight where Red Hat Enterprise Virtualization and NetApp storage intersect or where deployment steps for one technology affect the other.

1.5 Topics Out of Scope

Deployment steps associated with IP and Fibre Channel networks are not covered in this document. However, a good understanding of these topics is necessary for properly configuring items such as VLANs, switched fabrics, and other related technologies.

1.6 Best Practices and Supportability

Following the steps in this document is fully supported by both NetApp and Red Hat for their respective products. However, the ways and means described here cannot possibly cover every scenario or business need. Rather, this deployment guide should serve as a starting point. Should your environment require topics not covered in this document, technical resources from Red Hat and NetApp are available to provide guidance that is compliant with support.

1.7 Thick and Thin Hypervisors

As described in TR-3914: Red Hat Enterprise Virtualization 3 and NetApp Storage: Best Practices Guide, Red Hat supports both thick and thin hypervisors. This document focuses mostly on the thin or RHEV-H hypervisor available with Red Hat Enterprise Virtualization. There are use cases for using both thick and thin hypervisors in an RHEV environment.

Figure 1) Thick and thin hypervisors.

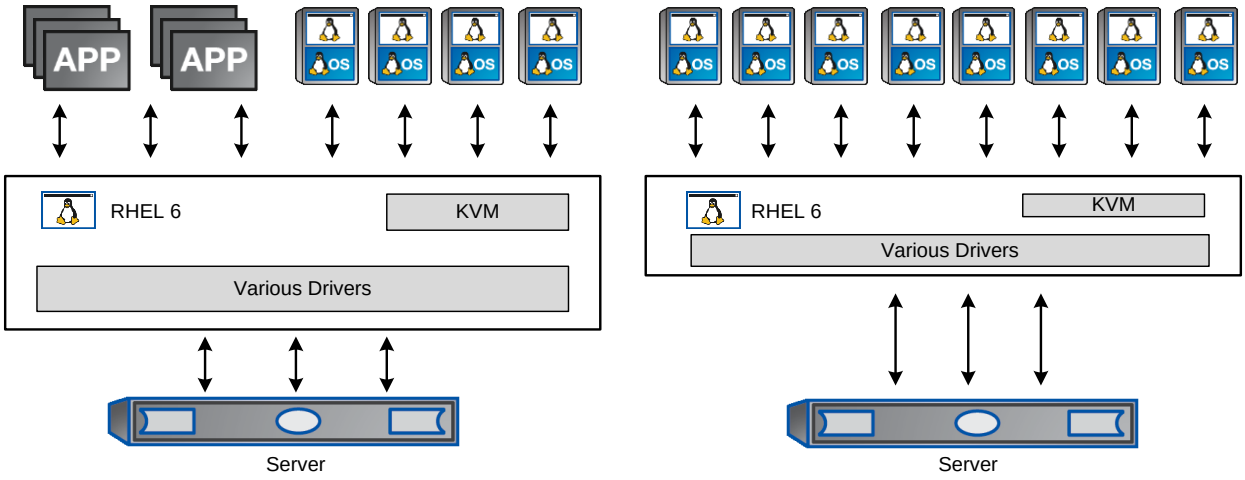

2 Deploying NetApp Storage Best Practices for RHEV 3

It is assumed that a NetApp controller is already configured, at least from a base standpoint. That is to say that a NetApp FAS controller running Data ONTAP 7.3.5 or higher or 8.0.1 7-Mode is already deployed and accessible. Data ONTAP Cluster-Mode is not covered in this technical report.

2.1 Prerequisites

The following items should already be configured:

- Data ONTAP installed on its own root FlexVol® volume and aggregate

- Basic networking to include a primary interface, management interface, and host name that is fully resolvable in DNS
- A second NetApp FAS controller configured in active-active configuration (or MetroCluster™ configuration)
- At least one large disk aggregate that is separate from the root aggregate to be used for data storage

2.2 NetApp Storage Security Best Practices

1. Run the following commands to enhance security on the storage controller:

```
options rsh.access none
options rsh.enable off
options webdav.enable off
options security.passwd.rules.everyone on
options security.passwd.rules.minimum 8
options security.passwd.rules.maximum 14
options security.passwd.rules.minimum.alphabetic 2
options security.passwd.rules.minimum.digit 1
options security.passwd.rules.minimum.symbol 1
options security.passwd.rules.history 6
options security.passwd.lockout.numtries 6
options security.passwd.firstlogin.enable on
options telnet.enable off
options autologout.console.timeout 5
options autologout.console.enable on
options autologout.telnet.timeout 5
options autologout.telnet.enable on
options httpd.enable off
options httpd.timeout 300
options ssh.idle.timeout 60
```

2. Create a new administrator account to replace the root login:

```
useradmin user add newadmin -g administrators
```

3. Log in with the new administrator account to verify that it works.
4. Disable the root account from the new administrator account:

```
options security.passwd.rootaccess.enable off
```

5. Run the following commands per existing user:

```
useradmin user modify <acct> -g <group> -M 90
useradmin user modify <acct> -g <group> -m 1
```

6. Secure Shell (SSH) protocol must be configured and enabled. Use the following one-time command to generate host keys:

```
secureadmin setup -q ssh 768 512 768
```

7. Use the following options to configure and enable SSH:

```
options ssh.access *
options ssh.enable on
options ssh.idle.timeout 60
options ssh.passwd_auth.enable on
options ssh.pubkey_auth.enable on
options ssh1.enable off
options ssh2.enable on
```

2.3 Licenses for Storage Features

If this is a new NetApp FAS controller, specific licenses must be obtained and installed prior to deploying storage. The most relevant licenses to this deployment guide include, but are not limited to:

- FCP (includes FCoE)
- iSCSI
- NFS
- Deduplication
- SnapRestore®
- SnapMirror®
- SyncMirror® local (if using MetroCluster)
- Cluster remote (if using MetroCluster)

One or more licenses can be installed as follows from the NetApp command line:

```
license add <first_license> <second_license> <third_license>
```

2.4 Enabling Storage Protocols

NFSv3

1. Add a license for NFS:

```
license add <LICENSE_KEY>
```

2. Set the following recommended options that enable NFS version 3:

```
options nfs.tcp.enable on
options nfs.udp.enable off
options nfs.v3.enable on
```

3. Enable NFS:

```
nfs on
```

Note: If this is part of an active-active pair, run these steps against both controllers.

FCP and FCoE

1. License FCP by entering the following command:

```
license add <LICENSE_KEY>
```

2. Start the FCP service by entering the following command:

```
fcp start
```

3. Record the WWPN or FC port name for later use by entering the following command:

```
fcp show adapters
```

4. Check whether the Fibre Channel ports are targets or initiators by entering the following command:

```
fcadmin config
```

5. Make a Fibre Channel port into a target.

   Note: Only Fibre Channel ports that are configured as targets can be used to connect to initiator hosts on the SAN.

   For example, make a port called `0b` into a target port run by entering the following command:

```
fcadmin config -t target 0b
```

Note: If an initiator port is made into a target port, a reboot will be required. NetApp recommends rebooting after completing the entire configuration, because other configuration steps might also require a reboot.

Note: If this is part of an active-active pair, run these steps against both controllers.

iSCSI

The following steps configure the iSCSI service on a storage system. These steps do not include any host-specific configuration tasks.

1. From the storage controller CLI, license the iSCSI protocol:

```
license add <LICENSE_KEY>
```

Note: The following commands can all be run in the vFiler® context.

2. Start the iSCSI service:

```
iscsi start
```

3. Enable or disable the appropriate storage system interfaces for use with iSCSI.

   This example enables the iSCSI protocol on interfaces `e0a` and `e0c` and disables iSCSI on interfaces `e0b`, `e0d`, and `e0M`.

   Note: By default, all storage system interfaces are enabled for iSCSI when the service is started.

```
iscsi interface enable e0a e0c
iscsi interface disable e0b e0d e0M
```

4. Set the default iSCSI authentication method to **`deny`**:

```
iscsi security default –s deny
```

Note: If this is part of an active-active pair, run these steps against both controllers.

2.5 Provisioning NetApp Storage for RHEV

Prior to creating NFS exports or LUNs, one or more FlexVol volumes must be created. This in turn requires that at least one disk aggregate be created based on NetApp best practices.

Creating a FlexVol Volume

1. The following information is required to create a flexible volume: the volume's name and size and the aggregate on which it will exist. For example, to create a volume called `rhev_vol` on aggregate `aggr1` with a size of 10GB, run the following commands:

```
vol create rhev_vol01 aggr1 10g
vol options rhev_vol01 create_ucode on
vol options rhev_vol01 convert_ucode on
vol options rhev_vol01 nosnap on
```

Creating a LUN

LUNs are created within a FlexVol volume.

1. To create a 2GB LUN in volume rhev_vol called rhev_lun, run the following command:

```
lun create –s 2g –t linux /vol/rhev_vol01/rhev_lun01
```

2. Create an igroup:

For FCP or FCoE, an igroup named rhev_igroup01 is created:

```
igroup create -f -t linux rhev_igroup01 20:00:00:e0:8b:9d:a1:4d
```

Note: The `-t` option specifies the OS type, and `-f` specifies FC.

Note: Change the WWPN to reflect your environment.

Note: For iSCSI, an igroup named rhev_igroup02 is created:

```
igroup create -i -t linux rhev_igroup02 iqn.2010-04.com:fujitsu.host02
```

Note: The `-t` option specifies the OS type, and `-i` specifies iSCSI.

Note: Change the IQN to reflect your environment. The IQN can be found on the HBA BIOS (if using a hardware-based initiator) or in the `/etc/iscsi/initiatorname.iscsi` file (if using the native RHEL software-based initiator).

3. Map the LUN used to boot the igroup (in this example to the iSCSI igroup):

```
lun map /vol/rhev_vol01/rhev_lun01 rhev_igroup02
```

4. Mount and format the LUN if using for data, or install the operating system if it is a boot LUN.

Note: Depending on the operating system, the concept of file system alignment might need to be addressed. Refer to TR-3747: Best Practices for File System Alignment in Virtual Environments.

Creating a Thin Provisioned NFSv3 Export

1. Create a new FlexVol volume, such as `rhev_vol02`.
2. Run the following commands to modify the volume options and the volume Snapshot™ options. These commands set the default FlexVol volume to be thin provisioned:

```
vol options rhev_vol02 guarantee none
vol options rhev_vol02 fractional_reserve 0
vol autosize rhev_vol02 on
vol autosize rhev_vol02 -m <<2*SIZE>>g
vol options rhev_vol02 try_first volume_grow
snap reserve rhev_vol02 0
snap autodelete rhev_vol02 on
```

Note: The autodelete setting differs from the one presented in RA-0007: Storage Efficiency Every Day: How to Achieve and Manage Best-in-Class Storage Use.http://media.netapp.com/documents/RA-0007.pdf In RA-0007, autodelete is set to `off`. Here it is set to `on` to allow the storage to be imported into Provisioning Manager.

Note: Setting the autosize option to two times the original size of the volume allows the volume to double its original size by using more space from the aggregate. Be certain to understand the consequences of this setting and to monitor free space in all aggregates.

3. Export a new volume read/write to a host called 192.168.27.40 and put the entry permanently into the `/etc/exports` file:

```
exportfs -p rw=192.168.27.40 /vol/rhev_vol02
```

The volume should now be accessible from the host.

4. The `/etc/exports` file can also be modified manually. When this modification is complete, run the following command to export all existing entries in the `/etc/exports` file:

```
exportfs -a
```

Note: The `exportfs` command has many options. For more information, refer to Data ONTAP 7.3.3 Documentation.

Thin Provisioning SAN

1. Run the following commands to modify the volume options and the volume Snapshot options. These commands set the default FlexVol volume and LUN to be thin provisioned:

```
vol options rhev_vol01 guarantee none
vol options rhev_vol01 fractional_reserve 0
vol autosize rhev_vol01 on
vol autosize rhev_vol01 -m <<2*SIZE>>g
vol options rhev_vol01 try_first volume_grow
snap reserve rhev_vol01 0
snap autodelete rhev_vol01 on
lun set reservation /vol/rhev_vol01/rhev_lun01 disable
```

Note: The autodelete setting differs from the one presented in RA-0007: Storage Efficiency Every Day: How to Achieve and Manage Best-in-Class Storage Use. In RA-0007, autodelete is set to `off`. Here it is set to `on` to allow the storage to be imported into Provisioning Manager.

Note: Setting the autosize option to two times the original size of the volume allows the volume to double its original size by using more space from the aggregate. Be certain to understand the consequences of this setting and to monitor free space in all aggregates.

Enabling Deduplication

1. Enable deduplication on volume /vol/rhev_vol01 by running:

```
sis on /vol/rhev_vol01
```

2. Start the initial scan (which should be run only once).

   Note: The initial scan of the data is the most resource intensive, so make certain this is an acceptable time to run it.

```
sis start -f -s /vol/rhev_vol01
```

3. Answer `yes` to the question.
4. Change the default deduplication schedule by running:

```
sis config -s SCHEDULE /vol/rhev_vol01
```

   Where `SCHEDULE` can be in one of these formats:

   - `[day_list][@hour_list]`
   - `[hour_list][@day_list]`
   - `auto`

Note: The `auto` setting runs only when 20% of the data in the volume has changed.

3 Scaling Out and Separating the Environment

In this section, MultiStore® and vFiler are used to provide scalable and mobile virtual storage controllers. The following procedure creates an IP space and assigns a designated interface to that IP space. All commands in this procedure are run from the Data ONTAP CLI.

3.1 Creating an IP Space and vFiler Unit

1. Install the MultiStore license.
2. Create an IP space for a vFiler unit or multiple vFiler units:

```
ipspace create var_ipspace01
```

3. Assign an interface that will be used for a vFiler unit to the newly created IP space:

```
ipspace assign var_ipspace01 var_interface01
```

The following procedure provisions one nondefault vFiler unit with a dedicated IP space, and one additional data volume is also assigned to the vFiler unit.

4. Create a primary storage unit for the vFiler unit:

```
vol create var_vfiler01_rootvol var_aggr01 256m
```

5. Create a vFiler unit and specify the root volume and IP space that will be used.

   Note: The IP address used by the vFiler unit must not be configured when the vFiler unit is created.

   Note: Quotas must be turned off before assigning a qtree or volume to a vFiler unit; they may be turned back on after the resource is assigned.

6. After assigning an IP space to a vFiler unit, the IP space cannot be changed without destroying the vFiler unit.

```
vfiler create var_vfiler01 -n -s var_ipspace01 -i var_interface01 /vol/var_vfiler01_rootvol
```

7. Disallow rsh and any other protocols not needed by the newly created vFiler unit by using the vfiler disallow command.

   The following command disables rsh, cifs, iscsi, http, and ftp on the vFiler unit:

```
vfiler disallow var_vfiler01 proto=rsh proto=cifs proto=iscsi proto=http proto=ftp
```

8. Use the ifconfig command on the hosting storage system to configure the interface as Up with the IP address specified during the creation of the vFiler unit.

   Note: The IP address used here must be the same address used in the vfiler create command:

```
ifconfig var_interface01 var_ipaddress01 netmask var_netmask01 partner var_partner_interface01
```

9. Modify the routing table of the IP space the vFiler unit is using with the vfiler run command. This command adds a default route for the newly created IP space that the Vfiler unit is using.

```
vfiler run var_vfiler01 route add default var_dgateway01 1
```

10. Add an additional data volume to the new vFiler unit:

```
vfiler add var_vfiler01 /vol/var_vfiler01_datavol
```

11. Apply the security best practices to the vFiler unit.

4 Storage Network Best Practices for Red Hat KVM

4.1 Storage Architecture Concepts

The following deployment examples assume that redundant switches are set up and that Link Aggregate Control Protocol (LACP) has been enabled and configured on the relevant switch ports.

4.2 IFGRP LACP

Since this type of interface group requires two or more Ethernet interfaces and a switch that supports LACP, make sure that the switch is configured properly. The commands are the same for Gigabit Ethernet (GbE) or 10 Gigabit Ethernet (10GbE).

Run the following command on the command line and also add it to the `/etc/rc` file, so it is activated upon boot. This example assumes that there are two network interfaces called `e0a` and `e0b` and that an interface group called `vif01` is being created:

```
ifgrp create lacp vif01 -b ip e0a e0b
wrfile -a /etc/rc "ifgrp create lacp vif01 -b ip e0a e0b"
```

Note: All interfaces must be in `down` status before being added to an interface group.

4.3 VLAN

1. Run the following commands to create tagged VLANs. This example assumes that there is a network ifgrp called `vif01` and that VLANs are being created that are tagged with numbers 10, 20, and 30. For failover to work properly, the commands must also be run on the other controller in an active-active pair.

```
vlan create vif01 10 20 30
wrfile -a /etc/rc "vlan create e0a 10 20 30"
```

Note: Although this example uses an ifgrp network interface, a non-ifgrp (e0a, e0b) can be used instead.

4.4 IP Config

1. Run the following commands on the command line.

Note: This example assumes that the user has a network interface called vif01-10 with IP address 192.168.10.10 and default router 192.168.10.1. If routing is not needed, skip the commands starting with `route add`.

```
ifconfig vif01-10 192.168.10.10 netmask 255.255.255.0 mtusize 1500 flowcontrol none
wrfile -a /etc/rc "ifconfig vif01-10 192.168.10.10 netmask 255.255.255.0 mtusize 1500 flowcontrol
none"
route add default 192.168.10.1 1
wrfile -a /etc/rc "route add default 192.168.10.1 1"
routed on
wrfile -a /etc/rc "routed on"
```

Note: The interface can be either physical (for example, e0a), an ifgrp (VIF), or a VLAN.

Run this routine once for every physical or virtual interface for which an IP address is needed.

If MultiStore is used, routing is common to every vFiler unit in an IP space.

If a 10Gb per second interface is not used, remove the `flowcontrol none` option.

If jumbo frames are needed, change the `mtusize` option to 9000.

5 Deploying an Infrastructure Cluster for RHEV

In following best practices for deploying RHEV and NetApp storage, a subset of hypervisors and virtual machines is used to host most, if not all, of the infrastructure-level services. All of the hypervisors are of the thick type using RHEL 6.

1. Deploy at least two RHEL 6 KVM hosts, using best practices in TR-3848: Best Practices for RHEL 6, KVM, and NetApp Storage.
2. The only required infrastructure virtual machine is an RHEL 6 guest that hosts the RHEV-M application, which is mentioned in the following section:
3. Deploy optional virtual machines to host applications such as NetApp Operations Manager, Red Hat Network Satellite, a Kickstart server, or other infrastructure-level applications.
4. Deploy a third RHEL 6 server to host Snap Creator™ for backup management. Although it can be a virtual machine, it should not be hosted on the hypervisor that it backs up. (If Snap Creator pauses the hypervisor on which it is hosted, it will not be able to resume itself.)

5.1 Deploying RHEV-M with NetApp

The steps to deploy RHEV-M do not differ from Red Hat's product guides, with one exception. It should be deployed on a virtual machine within an infrastructure cluster, as noted earlier. The following steps describe the overall work flow for deploying RHEV-M. The following steps assume that NetApp storage has been provisioned for an infrastructure cluster as well as storage for ISO images and virtual machine storage.

1. Deploy an infrastructure cluster of at least two RHEL 6 KVM hosts, as described in TR-3848: RHEL 6, KVM, and NetApp Storage: Best Practices.
2. Deploy an RHEL 6 virtual machine on the infrastructure cluster. It must have a fully resolvable host name. See the appendix for the ports that should be allowed through the firewall.
3. Register the virtual machine to Red Hat Network (RHN) and subscribe it to the RHEV-M-specific software channels. For example:

```
rhn_register
```

4. Follow the prompts to complete the registration, then subscribe to the required channels:

```
rhn-channel --add --channel=rhel-x86_64-server-6-rhevm-3
rhn-channel --add --channel=jbappplatform-5-x86_64-server-6-rpm
rhn-channel --add --channel=rhel-x86_64-server-supplementary-6
```

5. Update the system, install the RHEV-M package, and configure RHEV-M:

```
yum –y upgrade
yum –y install rhevm
rhevm-setup
```

Note: Accept all of the defaults for ports. Do not set up a default storage domain. Do not allow the configuration tool to rewrite the IPtables firewall.

6. Configure the firewall on the RHEV-M virtual machine. The list of required ports is included in the appendix.
7. After the basic install and configuration are complete, at least one RHEV data center will need to be created. This will include the following items:
   - At least one RHEV cluster
   - At least one logical network (with VLAN tagging)
   - At least one hypervisor
   - At least one storage domain

6 Deploying Hypervisors in RHEV

6.1 Deploying RHEV-H

1. Download the RHEV-H ISO image from Red Hat Network (RHN). Depending on how it is to be deployed, it can be burned to CD, written to a USB drive, or converted for used with a PXE network.
2. Configure the physical server to boot from the NetApp controller using FC, FCoE, or iSCSI.
3. Boot to the RHEV-H image and follow the prompts (if installing interactively) or allow PXE to automatically deploy the image. If deploying interactively, the installer will also prompt for the boot device.

Note: When the initial install is complete, the RHEV-H host will reboot. Log in to the host with the user admin and the password that was chosen during the installation. No root user is available.

If deploying interactively, the RHEV-H configuration will require at least a host name, network interface, DNS server, NTP server, and host name for RHEV-M. These items are configured through a text user interface (TUI).

When the configuration is complete, the RHEV-H host will show up in RHEV-M with a status of Pending Approval.

6.2 Adding an RHEL 6 Hypervisor to RHEV-M

These steps assume that the RHEL 6 KVM host is fully installed and configured per TR-3848: RHEL 6, KVM, and NetApp Storage: Best Practices. This includes but is not limited to security configuration, storage configuration, and package selection. In addition, the RHEL 6 KVM host must have a fully resolvable host name. It is also assumed that the RHEL 6 KVM host boots from a SAN device on the NetApp controller.

1. Register the RHEL 6 KVM host to Red Hat Network.
2. Subscribe the system to the Red Hat Enterprise Virt Management software channel.
3. Make sure that there is a complete host name listing in the `/etc/hosts` file.
4. Configure IPtables to allow required ports as listed in the appendix.

   Note: Configuring virtual networks or SSH host keys is not necessary (RHEV-M will take care of this).
5. Log in to RHEV-M, and click the Hosts tab.
6. Click the New button. Provide the host name, IP address, and root password, then click OK.
7. RHEV-M will then install a number of packages and reboot the host.

7 Configuring RHEV-M

7.1 Configuring an RHEV Data Center

1. Log in to the RHEV-M Web portal and click the New button on the Data Centers tab.
2. Enter a name and description of the data center to be created.
3. Select the storage type of the data center. Select the storage protocol to be used for storing virtual machines in the data center from the drop-down menu to your data center.
4. Select the compatibility level of the data center, 2.2 or 3.0, and click OK.
5. When the Guide Me dialog box is displayed, click Configure Later.

7.2 Configuring an RHEV Cluster

1. Log in to the RHEV-M Web portal and select the Clusters tab, then click the New button.
2. When the New Cluster dialog box is displayed, select the General tab.
3. Select an existing data center from the drop-down menu and then enter a cluster name and description.
4. Select the CPU name for hosts in this cluster and then select the compatibility level.
5. Select the memory optimization to be utilized.
6. Select the Resilience Policy tab to define if high availability is to be a consideration for any virtual guest in the RHEV cluster.
7. Click OK to create the cluster.
8. When the Guide Me dialog box is displayed, click Configure Later.

7.3 Configuring RHEV Logical Networks

1. Log in to the RHEV-M Web portal and select the Data Centers tab, then the Logical Networks tab.
2. Click the New button. Enter a name and description. Check the “Enable VLAN tagging” box and also enter a VLAN number. This should match the VLAN configured on the network.
3. Select the RHEV cluster(s) that will have access to this VLAN, then click OK.
4. The remaining steps are completed while configuring network interfaces on the hypervisors.
   a. Select the Hosts tab, then the hypervisor that will have an interface and VLAN configured, then the Network Interfaces tab.
   b. Select the physical interface to be configured, then click Add/Edit.
   c. Select the network from the drop-down menu and then select either DHCP or Static for the configuration. If Static is selected, an IP address and network mask are also required.
   d. Check the “Save network configuration” box, then click OK.

      These steps can be repeated several times (especially on 10GbE devices) to maintain and operate multiple logical networks and VLANs on the same physical device.

7.4 Configuring RHEV Storage Domains

Although there are three types of storage domains in RHEV, this guide will only discuss the use of a data domain and an ISO domain. The Red Hat Enterprise Virtualization product guides should be followed for specific details. The following sections describe the workflow required.

Configuring a Data Domain

1. Log in to the RHEV-M Web portal and select the Storage tab, then click the New Domain button.
2. Enter a name for the new storage domain and then select the RHEV data center that will use this new storage domain. The type of storage protocol for the storage domain is determined at RHEV data center creation time.

FC or FCoE Data Domain

1. Select Data/FCP.
2. Select the hypervisor node that will attach to the storage initially. This will trigger a scan of the SCSI bus.
3. Select each LUN that is to be configured and used.
4. Click OK.

iSCSI Data Domain

1. Select Data/iSCSI.
2. Select the hypervisor node that will attach to the storage initially.
3. Enter the NetApp host or IP. If CHAP is used, enter the CHAP username and password.
4. Click Discover. Select each LUN that is to be configured and used.
5. Click OK.

NFS Data Domain

1. Select Data/NFS.
2. Select the hypervisor node that will attach to the storage initially.
3. Enter the export path from the NetApp controller.

Note: Prior to performing any actions in RHEV-M, the NFS export must be mounted manually to configure proper ownership for RHEV. Failure to perform this initial step will result in the inability to mount NFS storage:

```
mount /path/to/export /mnt; cd /mnt
chown -R 36.36 *
cd; umount /mnt
```

Configuring an ISO Domain

1. Select ISO (it is only available through NFS).
2. Select the hypervisor node that will attach to the storage initially.
3. Enter the export path from the NetApp controller.

Note: Prior to performing any actions in RHEV-M, the NFS export must be mounted manually to configure proper ownership for RHEV. Failure to perform this initial step will result in the inability to mount NFS storage:

```
mount /path/to/export /mnt; cd /mnt
chown -R 36.36 *
cd; umount /mnt
```

Populating ISO Domain

The ISO domain is used to store ISO images for operating system installation, VirtIO drivers, and virtual floppy disks.

1. Log in to the RHEV-M host and open a console.
2. Type the following command to upload ISO and VFD images to an existing ISO domain:

```
rhevm-iso-uploader --iso-domain=<name_of_iso-domain> upload rhel6.iso virtio-driver.iso virtio-driver.vfd
```

8 RHEV Guest Configuration

8.1 Provisioning Virtual Machines in RHEV 3.0

This section describes how to provision virtual machines within RHEV using the RHEV-M portal.

Note: The following steps assume that there is a kickstart server available on the network and that there are installation ISO images available to boot from in the ISO storage domain.

The configuration of virtual machines is covered in detail in TR-3848: RHEL 6, KVM, and NetApp Storage: Best Practices and TR-4034: Red Hat Enterprise Linux 6, KVM, and NetApp Storage: Deployment Guide.

Procedure for provisioning virtual machines within RHEV using the RHEV-M portal:

1. Log in to the RHEV-M Web portal and select the Virtual Machines tab.
2. Click the New Server button.
3. In the New Server Virtual Machine dialog box, select Data Center and Cluster where the virtual machine is operating, and then enter a name and description. Adjust the memory size, number of CPU cores, and number of CPU sockets, and select the specific operating system from the drop-down menu.
4. Select the Console tab and select VNC from the drop-down menu. Click OK.
5. Click the Configure Network Interfaces button. Select the logical network to which the network interface will attach. Click OK.

6. Click Configure Virtual Disks. Enter the size of the disk and check the “Wipe after delete” option. Click OK.
7. Right-click the newly created virtual machine and select Run Once. Check the Attach CD option and select the installation ISO image from the drop-down menu. Click OK.
8. Right-click the newly created virtual machine and select Console to interact with the installation.

9 RHEV 3.0 and NetApp Storage Management Tools

Any management servers in this section can all be virtualized on a separate group of infrastructure hosts. By virtualizing the management servers, they gain the same benefits as the production virtual machines such as mobility, availability, and centralized data management on the NetApp controller(s).

9.1 RHEV-M

To launch the RHEV-M Web portal, open Internet Explorer and go to the URL https://hostname.domain.com:8080, where host name and domain are replaced with the fully qualified domain name for your RHEV-M server. Select Administrator Portal.

RESTful API

To view the RESTful API guide, open Internet Explorer and go to the URL https://hostname.domain.com:8080, where host name and domain are replaced with the fully qualified domain name for your RHEV-M server. Select REST API Guide.

9.2 RHN and RHN Satellite

Red Hat Network Satellite can be deployed as a virtual machine within the RHEV infrastructure cluster.

For detailed information on Red Hat Network, see Red Hat Network Documentation.

For detailed information, including deployment procedures, for Red Hat Network Satellite server, see the documents available at Red Hat Network Satellite.

9.3 NetApp Operations Manager

NetApp Operations Manager can be deployed as a virtual machine within the RHEV infrastructure cluster. For detailed deployment instructions and best practices, see the Operation Manager, Provisioning Manager, and Protection Manager documents referenced in the appendix.

9.4 Kickstart Server

The Kickstart server can be deployed as a virtual machine within the RHEV infrastructure cluster. See the Red Hat Enterprise Linux 6 Deployment Guide for information on setting up and using Kickstart.

10 Deploying a Backup Framework for RHEV

10.1 Snapshot Data ONTAP

The deployment and configuration instructions for Data ONTAP Snapshot are identical to those in TR-4034: RHEL 6, KVM, and NetApp Storage: Deployment Guide.

10.2 Snap Creator 3.x Preinstallation

The deployment and configuration instructions for preinstallation of Snap Creator are identical to those in TR-4034: RHEL 6, KVM, and NetApp Storage: Deployment Guide.

10.3 Snap Creator 3.x Server

The deployment and configuration instructions for Snap Creator server are identical to those in TR-4034: RHEL 6, KVM, and NetApp Storage: Deployment Guide.

10.4 Snap Creator 3.x Agent

The Snap Creator agent can only be deployed on virtual guests and thick hypervisors. The Snap Creator agent cannot be installed on a thin RHEV-H hypervisor.

The deployment and configuration instructions for Snap Creator agent and Snap Creator profiles are identical to those in TR-4034: RHEL 6, KVM, and NetApp Storage: Deployment Guide.

11 Deploying and Configuring Disaster Recovery for RHEV 3

NetApp SnapMirror is used in conjunction with NetApp Snap Creator for disaster recovery. The deployment and configuration guidelines of NetApp SnapMirror for RHEV 3 are identical to those in TR-4034: RHEL 6, KVM, and NetApp Storage: Deployment Guide.

12 Site Failover for RHEL 6 KVM and NetApp Storage

12.1 Deploying SnapMirror Async for Site Failover

The following procedures assume that the primary site has already been set up, the secondary NetApp FAS controller has been set up, and that a SnapMirror relationship (DR) has also be configured between the two NetApp FAS controllers. This configuration can be done at the same time as the primary site or easily added on at a later date.

Procedures for Deploying SnapMirror for Site Failover

1. Configure the network in such a manner that both sites have the same subnets, IP spaces, and VLANs. If possible, the sites should have full Layer 2 and Layer 3 connectivity, but it is only a requirement for the SnapMirror relationship.
2. When configuring RHEV clusters in a site failover scenario, place all hypervisors in the same RHEV data center. However, the different sites will be in different RHEV clusters.
3. Configure the RHEV hypervisors at the secondary site. The virtual bridges, VLANs, subnets, and IP spaces must be identical to the primary site.
4. Configure failover targets (IP network). This means that the secondary controller must have the means to host the same IP addresses as the primary controller to support the same storage targets. This can typically be done by way of creating aliases on the existing ifgrps. These failover targets will remain down until they are needed in a site failover. After establishing the SnapMirror relationship, be sure to copy over NFS export permissions. Like the failover IPs, these exports should not be active except in the case of a site failover.

Note: If the secondary controller carries its own production workload in addition to the role of secondary site RHEV storage, then additional NICs might need to be added to the NetApp storage controller. Otherwise, configure aliases on the existing NICs to handle the IP traffic in a failover situation. In either case, these failover IPs will remain dormant until a failover situation arises.

Figure 2) Example of the secondary site layout.

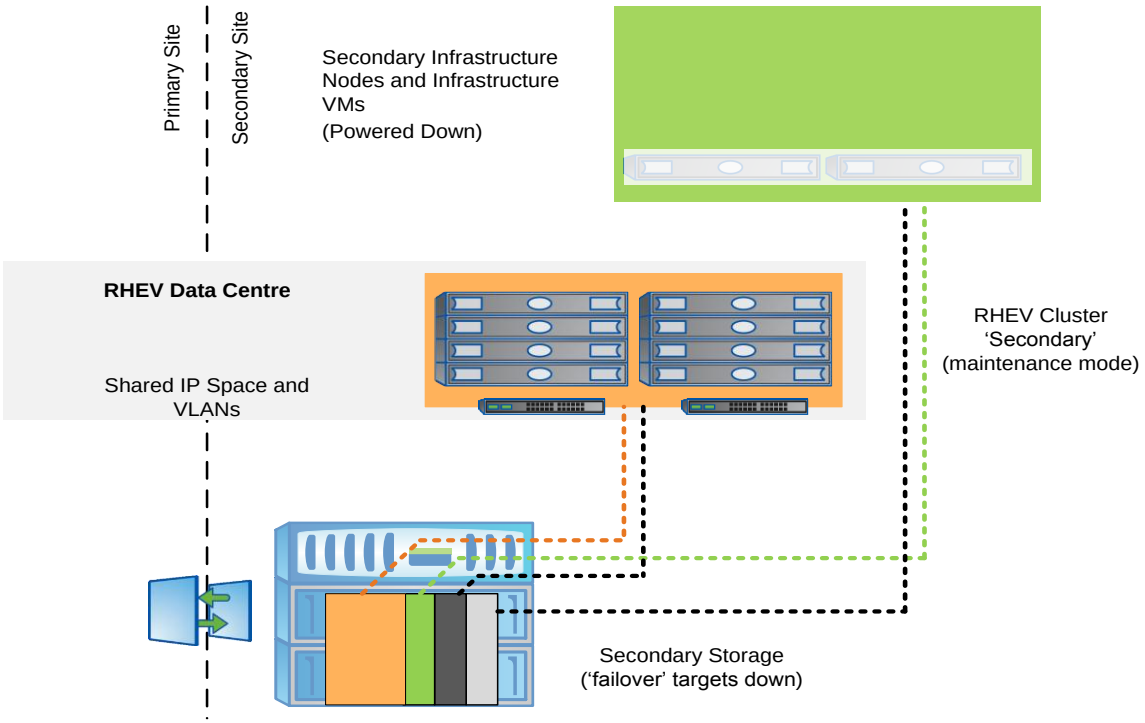

5. Configure zoning and igroups. Unlike IP networking, FC zoning is not shared between the sites. Hypervisors at the primary site are to remain zoned for the primary NetApp FAS controller. Hypervisors at the secondary site should only be zoned for the secondary NetApp FAS controller. The same rules apply to igroups for FC and iSCSI. However, be sure to properly map the secondary igroups to the right LUNs.
6. Query and log all LUN serial numbers on the primary and secondary controllers. When a volume is copied with SnapMirror, the LUNs at the secondary site will have a different serial number. In a site failover situation, they will need to be changed. Prior to giveback, the original serial numbers will need to be returned. Run the following command to query the serial number for a given LUN:

```
lun serial /vol/InfraVMVol/vmlun
```

   Note: This is a critical step in the process. Log the serial numbers for all LUNs and keep copies at the secondary site.

7. The secondary hypervisor nodes must have the same IP connectivity (bridges, VLANs, and so on) as the ones at the primary site.

Procedures for Site Failover with SnapMirror Async

1. Primary site is declared down.
2. Isolate IP traffic from primary site in case it comes back online while the DR site is running.
3. Break the SnapMirror relationship on the secondary NetApp FAS controller to make the volumes read-write:

```
snapmirror break name_of_flexvol
```

4. If using LUNs, bring the LUN(s) offline and make note of the original serial number. Change serial numbers on LUNs to the serial numbers from the LUNs at the primary site. Finally, bring it back online:

```
lun offline /vol/InfraVMVol/vmlun
lun serial /vol/InfraVMVol/vmlun
              Serial#: P3OoG4WVLFoa
lun serial /vol/InfraVMVol/vmvol <new_serial_#>
lun online /vol/InfraVMVol/vmlun
```

5. Bring up failover ifgrps on the secondary NetApp controller.
6. Bring up the RHEL 6 KVM hosts (infrastructure cluster) at the secondary site, if they are not running.
7. Start the virtual machine that is running RHEV-M, and then verify that RHEV-M is accessible by logging into the RHEV-M Web portal.
8. Bring up the hypervisors at the secondary site and bring them out of maintenance mode. One of the hypervisors should automatically attach to the storage. When RHEV-M comes up, the secondary hypervisor nodes will not be able to take over storage (attain SPM) until the primary hypervisor nodes are confirmed shut down or rebooted within RHEV-M. Confirm this, and then activate the hypervisors.
9. Move the critical virtual machines from the primary cluster to the secondary cluster (this can be scripted using the RESTful API).
10. Start the critical virtual machines in RHEV.
11. Failover is complete.

Procedures for Site Giveback with SnapMirror Async

1. When the primary site is back up and ready to operate again, be sure that it is still isolated.
2. Gracefully bring down the virtual machines in RHEV.
3. Place the hypervisors at the secondary site in maintenance mode.
4. Bring down failover IPs and ifgrps on the NetApp controller at the secondary site.
5. If using LUNs, change serial numbers on LUNs back to original:

```
lun offline /vol/InfraVMVol/vmlun
lun serial /vol/InfraVMVol/vmvol <orig_serial_#>
lun online /vol/InfraVMVol/vmlun
```

6. Next, resync the volumes between sites. From the NetApp controller at the primary site, run the following command from the primary controller:

```
snapmirror resync -S ice3170-3b:InfraVMVol ice3170-3a:InfraVMVol
```

From the NetApp controller at the secondary site, run the following command:

```
snapmirror resync -S ice3170-3a:InfraVMVol ice3170-3b:InfraVMVol
```

7. Put the full network back online.
8. Start up the RHEL 6 KVM hosts (infrastructure cluster) at the primary site.
9. Move the virtual machines back to the primary cluster.
10. Start the virtual machine that is running RHEV-M, then verify that RHEV-M is accessible by logging into the RHEV-M Web portal.
11. Start the virtual machines in RHEV.

    Giveback is complete.

12.2 Deploying MetroCluster for Site Failover

MetroCluster is fairly straightforward to deploy when an environment is first being set up, but can be challenging to deploy as part of an existing environment.

Configuring MetroCluster

1. Configure the network. All Layer 2 and Layer 3 networking between the sites must be shared. That is to say that IP spaces, VLANs, and VIFs should be accessible between the sites.
2. Configure the hardware. MetroCluster has specific hardware requirements and specific hardware configuration that are outside of the scope of this document. It is likely that a NetApp professional services consultant will implement the hardware portion of the configuration.
3. Configure partner IPs on each NetApp FAS controller for each network interface, VLAN, and ifgrp.
4. Configure zoning and igroups. Zoning and igroups do not need special configuration. Storage zoning and NetApp igroups should be site specific.
5. Create a mirrored aggregate:

```
aggr create aggr1 -m 12
```

This will create a mirrored aggregate named aggr1 with 12 disks. Six disks will be on the primary controller, and six disks will be on the secondary controller.

6. Create FlexVol volumes, NFS exports, and LUNs using standard NetApp best practices.
7. Disable the changing of LUN serial numbers on forced takeover:

```
options cf.takeover.change_fsid off
```

   Note: This will alleviate the need to change the LUN serial numbers during failover or failback.

8. Create RHEL 6 KVM hosts as prescribed in this document.
9. Configure RHEV-M as a VM, as prescribed in this document. When creating the RHEV cluster for the secondary hypervisor nodes, be sure that the cluster is part of the same RHEV data center as the primary hypervisor nodes. This is a critical step in the process.
10. When setting up the secondary hypervisors, configure them to be a separate RHEV cluster in the same RHEV data center as the primary hypervisors. Then place them in maintenance mode or power them down.

    Note: Placing the secondary hypervisor nodes in the same RHEV data center, but different RHEV cluster, is a crucial part of successfully bringing RHEV up in the secondary site.

Figure 3) Example of an RHEV 3 layout with MetroCluster.

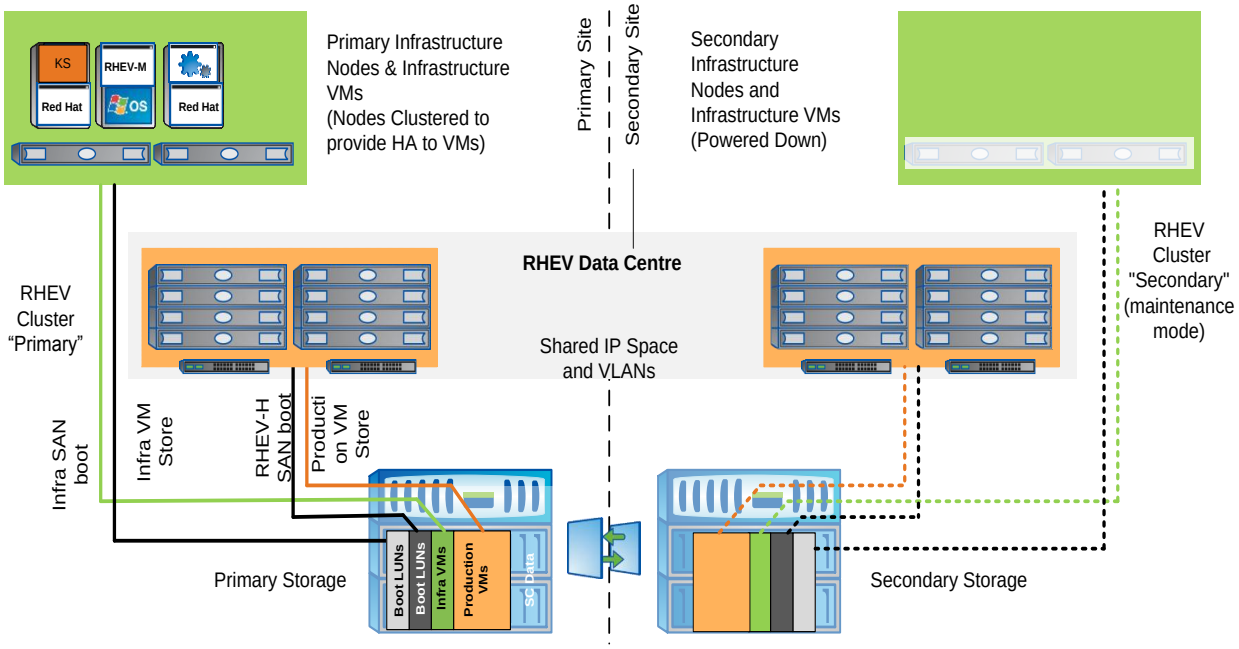

Procedures for Site Failover with MetroCluster

1. The primary site fails (total loss of site power, catastrophe, and so on), and the secondary FAS controller recognizes that it no longer has connectivity with any part of the primary storage.
2. Isolate the primary site controller from the secondary site in case it comes back up. The easiest way to do this is to turn off the power to the controller (leaving power on to the shelves). Otherwise, disconnect the cluster interconnect at the secondary site.

Caution

This is a critical step that has serious consequences if not followed.

Note: If the primary storage comes back up, it will be required to service storage requests at the same time as the secondary controller.

3. Initiate a forced (manual) takeover on the secondary storage:

```
cf forcetakeover -d
```

Note: If the RHEL 6 KVM hosts are still operational, then they can still reach the secondary storage. However, if the distance is great, then latency might be an issue. For this reason, the following procedures assume that everything is being failed over.

4. Bring up the secondary infrastructure nodes and start the infrastructure VMs (including RHEV-M). When RHEV-M comes up, the secondary hypervisor nodes will not be able to take over storage (attain SPM) until the primary hypervisor nodes are confirmed shut down or rebooted within RHEV-M. Confirm this, and then activate the hypervisors.
5. Move the critical virtual machines from the primary cluster to the secondary cluster.
6. Start the critical virtual machines.

   The failover is complete.

Procedures for Site Giveback with MetroCluster

Caution

When the primary site is brought back online, do not bring the primary NetApp controller up yet, only the disk shelves.

1. Save the state of any network changes within RHEV-M and bring down the production VMs gracefully and place the secondary hypervisor nodes back in maintenance mode.
2. Bring down the DR infrastructure VMs gracefully, followed by the secondary infrastructure nodes.
3. Bring up the primary disk shelves and properly reestablish aggregate mirrors and make sure that they are synced. Only after this has been confirmed can the primary controller be brought back up.
   a. Confirm that the remote storage can be seen from the secondary controller:

```
aggr status -r
```

   b. Go into partner mode on the secondary controller:

```
partner
```

   c. From within partner mode, make a determination as to which aggregates are at what site and which aggregates are out of sync:

```
aggr status -r
```

   d. Before resyncing aggregates, take the secondary site aggregates offline:

```
aggr offline aggr1
```

   e. Resync the aggregates:

```
aggr mirror aggr1 -v aggr1
```

   Note: Before continuing, be sure *all* aggregates are resynced.

4. The primary controller will pause its own boot procedure until a giveback is initiated from the secondary controller.

Caution

Do not initiate the giveback until the following command is run from the secondary NetApp FAS controller: `cf giveback`.

Note: The secondary NetApp FAS controller will reboot.

5. Bring up the infrastructure cluster at the primary site and start the virtual machine that is hosting RHEV-M.
6. Move the virtual machines back to the primary cluster.
7. Start the RHEL 6 KVM guests.

   Giveback is complete.

Appendix: Ports to Allow Through Firewall

Table 1) Ports to allow through firewall.

Port	Protocol	Description
22	TCP	SSH
80, 443	TCP	HTTP, HTTPS
111	TCP, UDP	Portmap
123	TCP	NTP
16514	TCP	libvirt
3260	TCP, UDP	iSCSI (optional)
53	TCP, UDP	DNS
5353	TCP, UDP	mDNS
54321	TCP	KVM interhost communication
5900-5910 (range can be increased)	TCP	VNC consoles (optional)
32803, 662	TCP	NFS client
49152–49216	TCP	KVM migration
67, 68	TCP, UDP	DHCP
8080	TCP	RHEV 3, Snap Creator, and Operations Manager portals
8088	TCP	NetApp Management Console
8443	TCP	Secure Operations Manager Console
8488	TCP	Secure NetApp Management Console
9090	TCP	Snap Creator agent
N/A	N/A	ICMP

Table 2 lists ports for LDAP-based authentication. Ports that are specific to LDAP based on Active Directory® are noted as such.

Table 2) Ports for LDAP-based authentication.

Port	Protocol	Description
135	TCP	MS RPC (Active Directory)
1025, 1026	TCP	Login and replication (Active Directory)
389	TCP	LDAP
636	TCP	Secure LDAP
445	TCP	Microsoft® DS (Active Directory)

Port	Protocol	Description
139	TCP	SMB (Active Directory)
137 and 138	UDP	NetBIOS (Active Directory)
88	UDP	Kerberos v5 (Active Directory)

References

- Home page for KVM
  www.linux-kvm.org
- Red Hat and Microsoft Virtualization Interoperability
  http://www.redhat.com/promo/svvp/
- KVM: Kernel-Based Virtual Machine
  www.redhat.com/f/pdf/rhev/DOC-KVM.pdf
- Red Hat Enterprise Linux 6 Virtualization Guide
  http://docs.redhat.com/docs/en-US/Red_Hat_Enterprise_Linux/6/html/Virtualization_Administration_Guide/index.html
- Red Hat Enterprise Virtualization Administration Guide
  http://docs.redhat.com/docs/en-US/Red_Hat_Enterprise_Virtualization/3.0/html/Administration_Guide/index.html
- Red Hat Enterprise Virtualization Installation Guide
  http://docs.redhat.com/docs/en-US/Red_Hat_Enterprise_Virtualization/3.0/html/Installation_Guide/index.html
- Red Hat Enterprise Linux 6 Deployment Guide
  http://docs.redhat.com/docs/en-US/Red_Hat_Enterprise_Linux/6/html/Deployment_Guide/index.html
- Red Hat Enterprise Linux 6 Installation Guide
  http://docs.redhat.com/docs/en-US/Red_Hat_Enterprise_Linux/6/html/Installation_Guide/index.html
- Red Hat Enterprise Linux 6 Security-Enhanced Linux User Guide
  http://docs.redhat.com/docs/en-US/Red_Hat_Enterprise_Linux/6/html/Security-Enhanced_Linux/index.html
- Best Practices for File System Alignment in Virtual Environments
  http://www.netapp.com/us/library/technical-reports/tr-3747.html
- Storage Best Practices and Resiliency Guide
  http://media.netapp.com/documents/tr-3437.pdf
- NetApp Deduplication for FAS and V-Series Deployment and Implementation Guide
  http://www.netapp.com/us/library/technical-reports/tr-3505.html
- SnapMirror Async Overview and Best Practices Guide
  http://www.netapp.com/us/library/technical-reports/tr-3446.html
- Operation Manager, Protection Manager, and Provisioning Manager Sizing Guide
  http://www.netapp.com/us/library/technical-reports/tr-3440.html
- Operations Manager, Provisioning Manager, and Protection Manager Best Practices Guide
  http://www.netapp.com/us/library/technical-reports/tr-3710.html
- RHEL 6, KVM, and NetApp Storage: Best Practices
  http://media.netapp.com/documents/tr-3848.pdf
- RHEL 6, KVM, and NetApp Storage: Deployment Guide
  http://media.netapp.com/documents/tr-4034.pdf
- Red Hat Enterprise Virtualization 3 and NetApp Storage: Best Practices
  http://media.netapp.com/documents/tr-3914.pdf

- RHEV 3 and NetApp Storage: Deployment Guide
  http://media.netapp.com/documents/tr-3940.pdf

Version History

Version	Date	Document Version History
Version 2.0	March 2012	Updated to include support for RHEV 3.0 and Data ONTAP 8.0.1.