



Red Hat Gluster Storage 3.3

Deploying Red Hat Gluster Storage using RHS One

Fast deployment for NAS or Content Repository

Red Hat Gluster Storage 3.3 Deploying Red Hat Gluster Storage using RHS One

Fast deployment for NAS or Content Repository

Laura Bailey
lbailey@redhat.com

Legal Notice

Copyright © 2018 Red Hat, Inc.

The text of and illustrations in this document are licensed by Red Hat under a Creative Commons Attribution–Share Alike 3.0 Unported license ("CC-BY-SA"). An explanation of CC-BY-SA is available at

<http://creativecommons.org/licenses/by-sa/3.0/>

. In accordance with CC-BY-SA, if you distribute this document or an adaptation of it, you must provide the URL for the original version.

Red Hat, as the licensor of this document, waives the right to enforce, and agrees not to assert, Section 4d of CC-BY-SA to the fullest extent permitted by applicable law.

Red Hat, Red Hat Enterprise Linux, the Shadowman logo, JBoss, OpenShift, Fedora, the Infinity logo, and RHCE are trademarks of Red Hat, Inc., registered in the United States and other countries.

Linux ® is the registered trademark of Linus Torvalds in the United States and other countries.

Java ® is a registered trademark of Oracle and/or its affiliates.

XFS ® is a trademark of Silicon Graphics International Corp. or its subsidiaries in the United States and/or other countries.

MySQL ® is a registered trademark of MySQL AB in the United States, the European Union and other countries.

Node.js ® is an official trademark of Joyent. Red Hat Software Collections is not formally related to or endorsed by the official Joyent Node.js open source or commercial project.

The OpenStack ® Word Mark and OpenStack logo are either registered trademarks/service marks or trademarks/service marks of the OpenStack Foundation, in the United States and other countries and are used with the OpenStack Foundation's permission. We are not affiliated with, endorsed or sponsored by the OpenStack Foundation, or the OpenStack community.

All other trademarks are the property of their respective owners.

Abstract

This document is a step-by-step process to setting up Network Attached Storage (NAS) or a Content Repository based on Red Hat Gluster Storage, using the RHS One fast deployment method.

Table of Contents

CHAPTER 1. OVERVIEW OF RED HAT STORAGE ONE	3
PART I. SETTING UP YOUR HARDWARE	4
CHAPTER 2. CONSIDERATIONS FOR HARDWARE SETUP	5
2.1. VERIFY THAT YOU HAVE ALL REQUIRED PARTS	5
2.2. ENSURE THE RACK IS STABLE	5
2.3. DETERMINE COMPONENT PLACEMENT	5
CHAPTER 3. INSTALLING YOUR SERVER INTO A RACK	6
CHAPTER 4. INSERTING PHYSICAL STORAGE DEVICES	9
CHAPTER 5. CONNECTING POWER, NETWORK, AND USER INTERFACE DEVICES	10
CHAPTER 6. VERIFYING YOUR HARDWARE SETUP	11
PART II. SETTING UP YOUR SOFTWARE	12
CHAPTER 7. CONSIDERATIONS FOR RED HAT STORAGE ONE COLONIZER	13
7.1. REQUIRED INFORMATION	13
7.2. DEPLOYMENT, NOT ADMINISTRATION	13
7.3. CLIENT ACCESS TO STORAGE	13
CHAPTER 8. RUNNING THE COLONIZER	15
CHAPTER 9. POST-DEPLOYMENT TASKS	20
9.1. TESTING DEPLOYMENT PERFORMANCE	20
9.2. EXPANDING YOUR STORAGE	20
9.2.1. Expanding your Network Attached Storage	20
9.2.2. Expanding your Content Repository	20
9.3. SUBSCRIBING YOUR SERVERS TO RED HAT NETWORK	21
9.4. UPDATING YOUR SERVERS	21
9.5. OTHER ADMINISTRATIVE TASKS	21
PART III. UNDERSTANDING RED HAT STORAGE ONE	22
CHAPTER 10. SUPPORTED WORKLOADS	23
10.1. NETWORK ATTACHED STORAGE CONFIGURATION OVERVIEW	23
10.2. CONTENT REPOSITORY CONFIGURATION OVERVIEW	23
CHAPTER 11. YOUR HARDWARE	25
11.1. PHYSICAL FEATURES OF YOUR SERVER	25
11.1.1. The front of your server	25
11.1.2. The back of your server	25
11.2. YOUR STORAGE DRIVES	26
CHAPTER 12. SAFETY RECOMMENDATIONS	28
12.1. HARDWARE LOCATION	28
12.2. HARDWARE TRANSPORT AND MOVEMENT	28
12.3. RACK AND HARDWARE STABILITY	28
12.4. SAFE OPERATING TEMPERATURE	29
12.5. ELECTRICAL SAFETY RECOMMENDATIONS	29
CHAPTER 13. NETWORKING REQUIREMENTS	31

CHAPTER 1. OVERVIEW OF RED HAT STORAGE ONE

Red Hat Storage One is a simple way to purchase and deploy Red Hat Gluster Storage 3.3 for [Network Attached Storage](#) or [Content Repository](#) workloads.

Red Hat Storage One hardware is purchased through original equipment manufacturers. When your hardware arrives, the deployment process is fast and simple:

1. **Set up your hardware.**
 - a. **Review all hardware considerations, network considerations, and safety and advisory material.**
 - b. **Rack mount your Red Hat Storage One servers**
 - c. [Insert all physical storage devices into each server](#)
 - d. [Connect and configure all networking, power, and user interface devices](#)
 - e. [Verify your hardware setup](#)
2. **Set up your software.**
 - a. [Review all software prerequisites](#)
 - b. **Run the configuration software**

PART I. SETTING UP YOUR HARDWARE

CHAPTER 2. CONSIDERATIONS FOR HARDWARE SETUP

2.1. VERIFY THAT YOU HAVE ALL REQUIRED PARTS

Each Red Hat Storage One server is shipped with two quick release rail assemblies, assorted screws, and 12 hot-swap 3.5 inch hard disk drives in drive carriers.

If any of these parts are missing, check your order details and contact Super Micro Computer, Inc.

2.2. ENSURE THE RACK IS STABLE

Bolt the server rack to the floor or attach stabilizers to the rack to ensure that the rack remains stable throughout this process.

2.3. DETERMINE COMPONENT PLACEMENT

Determine where you will install each component in this rack, keeping in mind the network and power requirements of your servers, as described in [Chapter 11, *Your hardware*](#).

Ensure that you install components starting from the bottom of the rack and the heaviest component, to ensure rack stability.

CHAPTER 3. INSTALLING YOUR SERVER INTO A RACK

This section shows you how to install the SC826 chassis into a rack unit using the quick release rails provided with your hardware.

Because there are a number of different rack units available, the assembly procedure may differ slightly for your rack unit. Refer to the installation instructions provided with your rack unit in addition to this documentation.



NOTE


You may appreciate the help of a second person when performing some of these steps.

1. Locate the rack mounting kit

The rack mounting kit consists of two quick release rail assemblies, which contain:

- an inner chassis rail extension, which secures to the rail that came pre-attached to your server chassis
- an outer rack rail, which secures directly to the rack

2. Separate the inner and outer rails



Separate the inner chassis rail extension from the outer rack rail by doing the following for each rail assembly:

- Extend the rail assembly to its full length.
- Press the quick release tab.


The inner chassis rail extension separates from the outer rack rail.

3. Install the inner rail extensions

For each side of the server chassis:

- a. Place the inner rail extension against the side of the chassis, inside the pre-attached inner rail.
Ensure that the extension faces outward, and that you align the holes in the inner rail extension with the hooks on the chassis.
- b. Slide the inner rail extension toward the front of the chassis so that the inner rail and inner rail extension are completely collapsed.
- c. Secure the inner rail extension to the chassis by using the screws and holes provided.

4. Install the outer rails




For each side of the rack:

- a. While ensuring that the rail remains horizontal, secure the back end of the outer rail to the rack using the screws provided.
- b. Retract the small outer rail by pressing the button at the join of the two outer rails.
- c. Secure the hooks at the front of the outer rail into the holes at the front of the rack.
Optionally, use the screws provided to secure the front of the outer rail onto the rack.

5. Install the chassis in the rack

You may need a second person to help you with this part.



- a. Fully extend the outer rails.
- b. Align the inner rails on the chassis with the outer rails on the rack.
- c. Slide the inner rails into the outer rails, ensuring that you keep the racks moving parallel with even pressure on both sides.
Push the server completely into the rack. The rail clicks when the server chassis locks into the rail correctly.
- d. Optionally, secure the front of the chassis to the rack by using the screws provided.

CHAPTER 4. INSERTING PHYSICAL STORAGE DEVICES

1. Locate the hard drives for this server

Each server is provided with 12 hard drives in drive carriers.

If you ordered multiple servers, ensure that you use the correct hard drives for each server. The server serial number is located on a decal at the rear of the server. This server serial number matches the serial number on the box containing the drives and drive carriers.


2. Determine the correct location for each drive carrier

The front of each server has 12 drive bays for your drive carriers. Your drive carriers are numbered from 0 to 11 according to the drive bay into which they should be inserted.

3. Insert the drive carriers

From the front of the server, insert each drive carrier into the drive bay that matches its number. You do not need to remove power from the system or access the inside of the server chassis to insert or remove drive carriers.

Figure 4.1. Front view of server chassis



- a. Ensure that the drive carrier handle is fully extended.
- b. Align the drive carrier with the drive bay and insert fully.
- c. Push the drive carrier handle into the drive carrier into its secured position.



CHAPTER 5. CONNECTING POWER, NETWORK, AND USER INTERFACE DEVICES

Connect your servers to power, network cables, and any user interface devices that you want to use during deployment.

- Connect RJ45 and SFP+ networking cables from the rear of the server to the appropriate networking switches.

The management and storage ports individually support 10Gb speeds. These ports are automatically bonded together to form a corresponding `mgmt_team` and `storage_team` interface at the operating system level. Red Hat recommends using at least 2 SFP+ ports (1 across each physical NIC device) for resilience.

See [Networking requirements](#) and [Physical features of your server](#) for information about setting up an appropriate network for your Red Hat Storage One deployment.

- Connect the two power supply modules to appropriate power sources. Red Hat recommends using separate power sources to provide redundancy in case of power failures.
- Connect a VGA monitor and other user interface devices if direct local interaction is required. Red Hat Storage One can be set up from any single server, so you do not require user interface devices for more than one server.

See the documentation for your server chassis for further advice on connecting power and user interface devices.

CHAPTER 6. VERIFYING YOUR HARDWARE SETUP

1. **Power on the servers**

Turn the servers on and ensure that they boot correctly to the login prompt.

2. **Verify storage detection**

Ensure that all drive bays show a solid blue indicator light; this indicates SAS drive detection.

3. **Verify login access**

Verify that you can log in to all servers using the username **root** and the password **redhat**.

4. **Verify network access**

Ensure that all servers have network access as described in [Chapter 13, *Networking requirements*](#).

PART II. SETTING UP YOUR SOFTWARE

CHAPTER 7. CONSIDERATIONS FOR RED HAT STORAGE ONE COLONIZER

The Colonizer is a command line tool that prompts you for various system details and then automatically prepares a Red Hat Gluster Storage environment for either a Content Repository or Network Attached Storage workload. Review this section carefully to ensure that you are aware of the limitations of the Colonizer before you start the software deployment process.

7.1. REQUIRED INFORMATION

You need to collect or decide the following details for all nodes before you run the Colonizer.

- hostnames
- IP addresses
- subnet mask
- default gateway (optional)
- DNS server addresses (optional)
- NTP server addresses (optional; you can use the defaults if you do not have preferred time sources)
- virtual IP addresses (required for storage access over NFS/SMB)
- For SMB with Active Directory: AD server, domain, and Admin credentials

7.2. DEPLOYMENT, NOT ADMINISTRATION

The Colonizer is used only for initial deployment, and cannot be used to modify the environment after initial deployment and configuration. It creates disks and bricks on all servers, and configures a Red Hat Gluster Storage volume that spans an initial set of servers. For Content Repository workloads, the initial volume spans the first six servers. For Network Attached Storage workloads, the initial volume spans the first four servers.

The deployment created using the Colonizer can then be modified, scaled, and updated in the same way as a manually installed and configured deployment of Red Hat Gluster Storage. For details, see the [Red Hat Gluster Storage 3.3 Documentation](#).

7.3. CLIENT ACCESS TO STORAGE

Clients are expected to connect to the configured storage volumes using one of the following protocols.

Native FUSE

See [Native Client](#) in the *Red Hat Gluster Storage 3.3 Administration Guide* for more information.

NFS Ganesha

Connecting clients using NFS Ganesha requires additional virtual IP addresses for the high availability service. The number of virtual IP addresses required is two-thirds the number of servers in your Red Hat Storage One deployment, with a minimum of 4 virtual IP addresses required for any

deployment. Deployments with the minimum 4 servers also require 4 virtual IP addresses. Configure these virtual IP addresses on the storage network. See [NFS Ganesha](#) in the *Red Hat Gluster Storage 3.3 Administration Guide* for more information.

SMB

See [SMB](#) in the *Red Hat Gluster Storage 3.3 Administration Guide* for more information. For Active Directory integration, see [Integrating Red Hat Gluster Storage with Active Directory](#).


CHAPTER 8. RUNNING THE COLONIZER

1. Log in

Log in to any one server using the username **root** and the password **redhat**.

2. Start deployment

The Colonizer automatically configures your Red Hat Gluster Storage cluster according to details in the template file provided by your vendor.



IMPORTANT

Do not modify the provided template file, as doing so invalidates Red Hat Storage One support and may prevent the Colonizer from completing successfully.

Run the following command from any of the provided servers to start the configuration process.

```
# gluster-colonizer.py -f /usr/share/gluster-colonizer/oemid/oemid-
config.yml
```

3. Agree to the installation

```
Welcome to the Red Hat Storage One deployment tool!
```

```
This node will be configured as the deployment node for your
Gluster storage pool. Before proceeding, please ensure that
all RHS One Gluster nodes are connected to the management
network infrastructure and are booted.
```

```
Do you wish to continue? [Y/n]
```

Enter **Y** to begin the configuration process.

4. Specify the total deployment size

```
How many RHS One Gluster nodes are you deploying?
```

```
Number of nodes (valid range is 4-24):
```

Enter the number of nodes that you want to configure.

- For Network Attached Storage workloads, you can enter any **even** number from 4 to 24.
- For Content Repository configurations, you can enter **6, 12, 18, or 24**.

5. Specify your client type

```
Please choose the client access method you will use for the default
storage volume.
```

- ```
1. NFS
2. Gluster Native Client (FUSE)
3. SMB
```

```
Client method? [1]
```

Enter the number associated with the type of client access to configure. The default client type is NFS.

## 6. For SMB, optionally, configure Active Directory integration

If you select SMB you are given the option to configure Active Directory integration.

For SMB, Active Directory integration can optionally be configured. The provided method will use winbind to connect the Gluster nodes to Active Directory and join the domain. This will require an Active Directory username and password for an account with rights to add systems to the domain.

Would you like to configure your RHS One Gluster nodes for Active Directory now? [Y/n]

If you do not want to configure Active Directory integration for SMB, enter **n** and skip the rest of this step.

If you want to configure Active Directory integration for SMB, enter **Y** and answer the following Active Directory integration prompts.

**Table 8.1. Colonizer prompts for Active Directory integration**

| Prompt text                                                                                                                                                                                                                                                                                                                                              | Action                                                                                  |
|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------|
| <p>The Samba HA cluster requires a single short name for entry in the Active Directory tree. This is the name by which the cluster will be referenced in DNS.</p> <p>Samba cluster short name:</p>                                                                                                                                                       | Provide a short name for the Samba cluster.                                             |
| <p>Please provide the fully-qualified Active Directory domain name and the credentials for a user with rights to add systems to the domain.</p> <p>Active Directory domain name: AD admin username [Administrator]: Active Directory admin password: Confirm admin password:</p>                                                                         | Provide the Active Directory domain name, AD admin username, and the AD admin password. |
| <p>Samba uses an identity mapping (idmap) module to map Active Directory SIDs to POSIX UIDs. You will need to select an idmap module that is appropriate for your environment.</p> <p>Which idmap module would you like to use?</p> <p>1. TDB (Samba default)</p> <p>2. AutoRID</p> <p>3. Skip selection and configure manually</p> <p>idmap module?</p> | Enter <b>1</b> to select the default module.                                            |

| Prompt text                                                                                                                                                   | Action                                      |
|---------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------|
| <p>The default idmap range is 1000000-4000000. If you would like to change this, enter a new value here with the same notation (m-n).</p> <p>idmap range?</p> | Enter the idmap range that you want to use. |

## 7. Answer server configuration prompts

Provide answers to the Colonizer prompts to complete server configuration.

The following table explains the prompts that appear, in order of appearance. Not all Colonizer output is part of this table; other informational messages are logged by the Colonizer as configuration progresses. Examples of content that is automatically populated by the Colonizer are rendered in *italicised text*.

**Table 8.2. Colonizer prompts for server-side configuration**

| Prompt text                                                                                                                                          | Action                                                                                                                                                                                                                                                                                         |
|------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <p>We will now configure your Gluster nodes for your storage network. [...]</p> <p>Do you wish to continue? [Y/n]</p>                                | Ensure that you have the information described in <a href="#">Section 7.1, “Required information”</a> prepared, then enter <b>Y</b> to continue.                                                                                                                                               |
| <p>All nodes are expected to be on the same subnet, so we will first collect all shared network information.</p> <p>Storage network domain name:</p> | Enter the storage network domain name, for example, <b>example.com</b> .                                                                                                                                                                                                                       |
| <p>Storage network and CIDR mask:</p>                                                                                                                | Enter the storage network and netmask, for example, <b>10.10.10.0/24</b> .                                                                                                                                                                                                                     |
| <p>Gateway and DNS fields may be left blank for now, if you prefer.</p> <p>Gateway IP address:</p>                                                   | Enter the IP address of the gateway, or leave blank and press <b>[Enter]</b> to continue.<br><br><b>This information is required for Active Directory integration.</b>                                                                                                                         |
| <p>DNS1 server address:</p>                                                                                                                          | Enter the IP address (not the FQDN) of a primary DNS server, or leave blank and press <b>[Enter]</b> to continue.<br><br><b>This information is required for Active Directory integration.</b> Use the IP address of the Active Directory server as the value for <b>DNS1 server address</b> . |

| Prompt text                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                | Action                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| DNS2 server address:                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       | Enter the IP address (not the FQDN) of a secondary DNS server, or leave blank and press <b>[Enter]</b> to continue.                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <p>NTP will be configured for time synchronization. You may enter up to four NTP servers below. If you would prefer to use the RHEL public NTP servers, simply press Enter at the first prompt and the default servers will be applied.</p> <p>Note: Using the default public NTP servers requires that all of the RHS One Gluster nodes have access to the Internet.</p> <p>NTP Server 1 (press Enter to accept defaults):</p>                                                                                            | <p>Enter the IP address or hostname of up to four servers that you want to use as time sources. Alternatively, press <b>[Enter]</b> to use the Red Hat Enterprise Linux defaults.</p> <p><i>Configuring this node as the deployment master...</i> is displayed.</p>                                                                                                                                                                                                                                                                                          |
| <p>The deployment tool requires DHCP service on the management network. If the service is not already available on the management network, we can start a DHCP server on this node now. This local DHCP service will only operate for the duration of the deployment process.</p> <p>Do you wish to start the DHCP service here? [y/N]</p>                                                                                                                                                                                 | <p>The default option assumes that a DHCP network already exists. Enter <b>N</b> to use an existing DHCP network. The management subnet is detected automatically.</p> <p>Enter <b>Y</b> to have the Colonizer to create a temporary DHCP network for the duration of the configuration process. At the next prompt, enter a management subnet, or press <b>[Enter]</b> to use the default.</p>                                                                                                                                                              |
| We will now begin node discovery. Do you wish to continue? [Y/n]                                                                                                                                                                                                                                                                                                                                                                                                                                                           | Enter <b>Y</b> to continue. Automatic node discovery begins. This takes up to 5 minutes to complete successfully.                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <p>You may choose to either assign hostnames and static IP addresses to your nodes manually, or the deployment tool can assign them for you automatically from the <i>10.11.12.0/24</i> network.</p> <p>Note: These hostnames and IP addresses are expected to remain static, so choose your values carefully. If you choose automatic assignment, the deployment tool assumes that there are no other devices on the storage network.</p> <p>Would you like to proceed with (m)anual or (a)utomatic assignment? [m/a]</p> | <p>If you want to manually assign hostnames and static IP addresses, enter <b>m</b>. If you have a large number servers, be aware that this requires a lot of data entry.</p> <p>If you want to use automatically assigned hostnames and IP addresses, enter <b>a</b>.</p> <p><b>WARNING:</b> In automatic mode, IP addresses are assigned sequentially and are not checked for availability. Automatic setup is therefore not recommended when there are IP addresses in the subnet other than those being used for the Red Hat Storage One deployment.</p> |
| <p>Please confirm your deployment details:</p> <p>[...]</p> <p>Do you wish to continue with this configuration? Y/n]</p>                                                                                                                                                                                                                                                                                                                                                                                                   | <p>Verify that your deployment details are listed correctly.</p> <p>Enter <b>Y</b> to continue the configuration process, or <b>n</b> to abort the process.</p>                                                                                                                                                                                                                                                                                                                                                                                              |

| Prompt text                                                                                                                                                                                                                                                                 | Action                                                                                                                                                                                                                                                                      |
|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <p>Your systems have factory SSH keys for the ansible user. These keys should not be considered secure. It is highly recommended that we replace these keys now with a newly-generated set.</p> <p>Would you like to proceed with creating a new set of SSH keys? [Y/n]</p> | <p>Enter <b>Y</b> to create a new set of SSH keys. This is required for production deployments.</p>                                                                                                                                                                         |
| <p>The default root password on your RHS One Gluster nodes must be reset. Be careful to select a secure password, and note that the password will be updated for the root user on all nodes.</p> <p>Please enter the new root password:</p>                                 | <p>Enter a new password to be used for the <b>root</b> user on all Red Hat Storage One servers, and enter the same password again to confirm.</p> <p><i>Begin RHS One Gluster validation phase is displayed, along with various other configuration check messages.</i></p> |
| <p>Next we will initiate the Gluster installation - OK? [Y/n]</p>                                                                                                                                                                                                           | <p>Enter <b>Y</b> to continue.</p>                                                                                                                                                                                                                                          |
| <p>WARNING: This step will delete any existing Gluster configurations and will wipe the LVM block devices and filesystems for drives other than the system drive.</p> <p>Are you sure you want to continue? [Y/n]</p>                                                       | <p>Enter <b>Y</b> to continue.</p>                                                                                                                                                                                                                                          |

From this point, the deployment process requires no further input. Deployment is complete when the following message is displayed:

```
Your Red Hat Storage One Gluster deployment is now complete!
```

A summary of your configuration choices is displayed on screen. You can see these at any time in the `/root/colonizer.README.txt` file.



## IMPORTANT

To enable storage access for SMB clients, you need to do some extra configuration.

Follow the instructions in [Mounting Volumes Using SMB](#) to configure a new Samba username and password. You can then use the username and password to mount your storage.

## CHAPTER 9. POST-DEPLOYMENT TASKS

### 9.1. TESTING DEPLOYMENT PERFORMANCE

After deployment using the Colonizer is complete, the Colonizer prompts you to run a series of performance tests.



#### IMPORTANT

These tests use the FUSE client protocol to check whether your systems are performing at a minimum acceptable level for your selected workload. They are not benchmarks specific to your selected client protocol.

+ Red Hat recommends running your own performance tests to aid you in tracking and improving the aspects of performance that are most relevant for your workload.

You have the option to run a series of performance tests to validate your RHS One Gluster environment. The performance tests can take an hour or longer to complete. It is recommended that you perform these tests now, but you may also choose to run them at a later time.

Would you like to start the performance tests now? [Y/n]

Enter **Y** to run the performance tests now.

Enter **n** if you do not want to run the performance tests at this time.

Performance test results are stored in **/root/g1-perf-results.out**.

### 9.2. EXPANDING YOUR STORAGE

The Red Hat Storage One process installs all required packages and sets up all required file systems and bricks, but sets up only one Gluster volume using the minimum number of servers for that workload. For Network Attached Storage workloads, the volume created at deployment uses the first four servers. For Content Repository workloads, the volume created at deployment uses the first six servers.

This lets you use your servers to create new volumes or expand the created volume as necessary.

#### 9.2.1. Expanding your Network Attached Storage

To create a new arbitrated distributed-replicated volume suitable for the Network Attached Storage workload, see [Creating an arbitrated replicated volume](#) in the *Red Hat Gluster Storage 3.3 Administration Guide*.

To expand the existing volume across the next set of servers, see [Expanding a volume](#) in the *Red Hat Gluster Storage 3.3 Administration Guide*.

#### 9.2.2. Expanding your Content Repository

To create a new distributed dispersed volume suitable for the Content Repository workload, see [Creating Distributed Dispersed Volumes](#) in the *Red Hat Gluster Storage 3.3 Administration Guide*.



To expand the existing volume across the next set of servers, see [Expanding a dispersed or distributed-dispersed volume](#) in the *Red Hat Gluster Storage 3.3 Administration Guide*.

## 9.3. SUBSCRIBING YOUR SERVERS TO RED HAT NETWORK

To gain access to software updates from Red Hat as they become available, subscribe your servers to Red Hat Network.

Registering a system to Red Hat Network requires valid Red Hat Customer Portal credentials. Either use the credentials supplied with your hardware purchase, or visit [Create a Red Hat Login](#) to create a new account.

When you have valid credentials, follow the instructions in [Subscribing to the Red Hat Gluster Storage Server Channels](#) in the *Red Hat Gluster Storage 3.3 Installation Guide* to subscribe your servers to the appropriate software repositories.

## 9.4. UPDATING YOUR SERVERS

For information about updating your servers, see [Updating Red Hat Gluster Storage from 3.3.x to 3.3.y](#) in the *Red Hat Gluster Storage 3.3 Installation Guide*.

## 9.5. OTHER ADMINISTRATIVE TASKS

See the [Product Documentation for Red Hat Gluster Storage 3.3](#) for information on any other configuration changes you want to make to your deployment.

## **PART III. UNDERSTANDING RED HAT STORAGE ONE**

# CHAPTER 10. SUPPORTED WORKLOADS

## 10.1. NETWORK ATTACHED STORAGE CONFIGURATION OVERVIEW

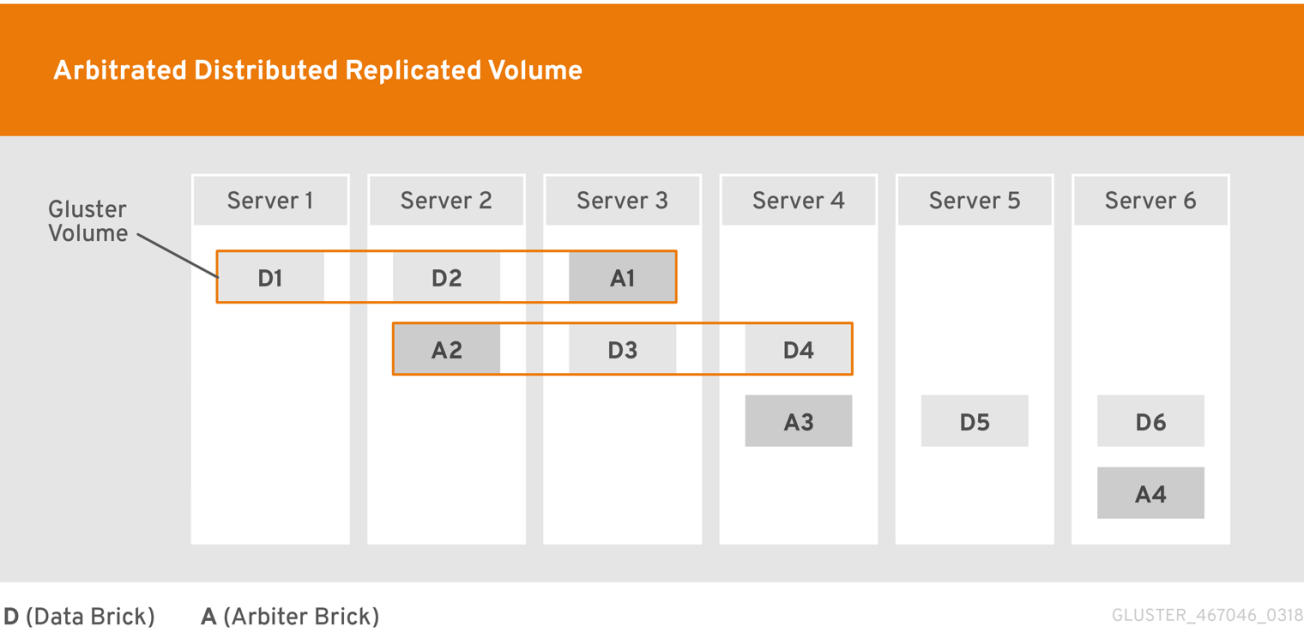
A Network Attached Storage (NAS) configuration uses between 4 and 24 servers in increments of 2 servers (4, 6, 8, 10, etc.).

Each server is provided with 12 drives, which are configured as a 10+2 RAID 6 device. For further information about this hardware, see [The front of your server](#).

The Colonizer uses these RAID devices to create 1-2 thin provisioned logical storage pools on each server. The entirety of each pool is used to create a logical volume. Logical volumes are then consumed by Red Hat Gluster Storage to create either a data brick or an arbiter brick.

The Colonizer then uses all bricks on each of the first four servers to create one arbitrated distributed replicated volume across these servers. Each sub-volume is configured in a chain so that each server has a maximum of one data brick and one arbiter brick, as shown in [Figure 10.1, “Network Attached Storage configuration”](#). Bricks that are not on the first four servers are prepared but not consumed as part of Red Hat Storage One deployment.

**Figure 10.1. Network Attached Storage configuration**



See [Creating Arbitrated Replicated Volumes](#) in the *Red Hat Gluster Storage 3.3 Administration Guide* for more information about arbitrated volumes.

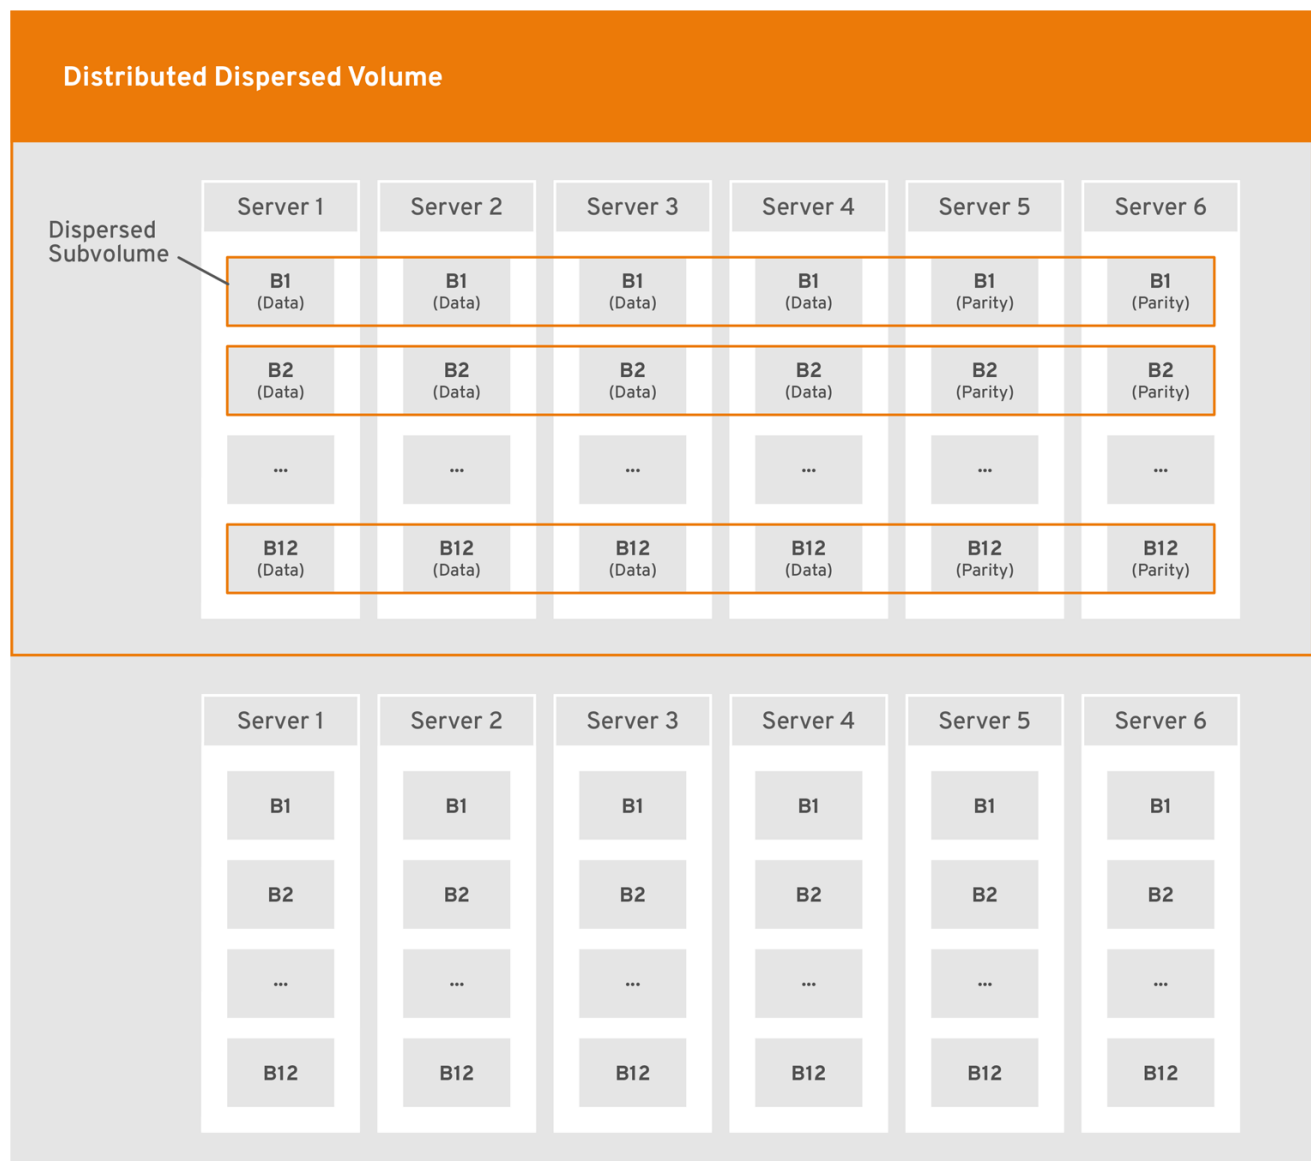
## 10.2. CONTENT REPOSITORY CONFIGURATION OVERVIEW

A Content Repository configuration uses between 6 and 24 servers in increments of 6 servers (6, 12, 18, or 24).

Each server is provided with 12 drives, which are configured as single-disk RAID 0 devices. The Colonizer uses the disks on each server to create a thin provisioned logical storage pool on each server. This pool is then divided into logical volumes. Each logical volume then provides the storage for one Gluster brick, so that there are 12 bricks per server.

The Colonizer then uses all 12 bricks on each of the first six servers to create a distributed dispersed volume in a 12x(4+2) configuration across those servers, as shown in [Figure 10.2, “Content Repository configuration”](#). Bricks that are not on the first six servers are prepared but not consumed as part of Red Hat Storage One deployment.

**Figure 10.2. Content Repository configuration**



B (Brick)

GLUSTER\_467046\_0318

See [Creating Distributed Dispersed Volumes](#) in the *Red Hat Gluster Storage 3.3 Administration Guide* for more information.




## CHAPTER 11. YOUR HARDWARE

### 11.1. PHYSICAL FEATURES OF YOUR SERVER

#### 11.1.1. The front of your server

Figure 11.1. Front view of server chassis




The front of your server chassis has the following:

- 12 drive bays (marked 0-11 according to logical drive bay number)
- control panel (marked 12)

#### 11.1.2. The back of your server

Figure 11.2. Back view of server chassis



The back of your server chassis has the following:

- 2 redundant power supply modules (A)
- 2 operating system boot drives in RAID1 configuration (B)
- Hardware RAID controller with 2GB cache, 16 drive license, and battery backup (C)
- 2 Dual-Port 10G Ethernet cards with SFP+ (D), used for the storage network
- I/O ports (1-10) as per the following table:

Table 11.1. I/O Ports

| Number | Description |
|--------|-------------|
|--------|-------------|

|    |                                                                                        |
|----|----------------------------------------------------------------------------------------|
| 1  | COM1                                                                                   |
| 2  | USB 0 (USB 3.0)                                                                        |
| 3  | USB 1 (USB 3.0)                                                                        |
| 4  | 1G RJ45 port for IPMI. Its IP address is set in the corresponding section of the BIOS. |
| 5  | USB 2 (USB 3.0)                                                                        |
| 6  | USB 3 (USB 3.0)                                                                        |
| 7  | 1G/10G RJ45 port for LAN1, used for the management network                             |
| 8  | 1G/10G RJ45 port for LAN2, used for the management network                             |
| 9  | VGA Port                                                                               |
| 10 | Unit identifier switch (UID)                                                           |

## 11.2. YOUR STORAGE DRIVES

Each hard drive carrier has two indicator lights: an activity indicator, and a status indicator.

The following table shows the meaning of each possible state of the indicator lights.

**Table 11.2. Activity Indicator (Blue LED)**

| Light state         | Drive activity                     |
|---------------------|------------------------------------|
| Off                 | SATA drive installed, OR empty bay |
| On continuously     | SAS drive installed                |
| Blinking on and off | I/O activity                       |

**Table 11.3. Status Indicator (Red LED)**

| Light state                | SAS/SATA drive status |
|----------------------------|-----------------------|
| On continuously            | Failed drive          |
| On for 5 seconds, then off | Identify drive        |

| Light state                    | SAS/SATA drive status |
|--------------------------------|-----------------------|
| Blinking once per second       | Rebuild drive         |
| Blinking twice per second      | Hot spare             |
| Blinking four times per second | Power on              |

## CHAPTER 12. SAFETY RECOMMENDATIONS

This section contains important safety and care information for your Red Hat Storage One hardware. Read this section in its entirety before you begin to install servers into server racks.



### WARNING

Failing to follow these recommendations can lead to injury or death.



### WARNING

Electrostatic discharge can damage electronic components. Avoid damage by always following these recommendations:

- Wear a grounded wrist strap when handling printed circuit boards.
- Handle printed circuit boards only by their edges.
- Store printed circuit boards in anti-static bags when not in use.

### 12.1. HARDWARE LOCATION

Do not set up your servers in an area that is affected by heat, electrical noise, or electromagnetic fields.

Set up your servers in a Restricted Access Location; that is, a dedicated equipment room, service closet, or similar. Ensure that this location is clean and free of dust and that it is well ventilated.

Leave enough space around the rack for proper air flow and service access, and so that you can completely open the door. Supermicro recommends at least 25 inches (62.5 centimeters) of clearance at the front of the rack, and at least 30 inches (75 centimeters) of clearance at the back of the rack.

### 12.2. HARDWARE TRANSPORT AND MOVEMENT

Do not use the handles on the front of the server to lift the server. These handles are designed to extend the system from a rack while supported by rack rails and are not designed to bear the full weight of the server.

### 12.3. RACK AND HARDWARE STABILITY



**WARNING**

Failing to follow these recommendations can lead to injury or death.

Ensure that the leveling jacks on the bottom of the rack are extended to the floor so that the full weight of the rack rests on them.

If you are using a single rack, attach stabilizers to the rack. If you are using multiple racks, couple the racks together for stability. Attach any stabilizers before mounting or servicing any hardware in the rack.

Install hardware starting from the bottom of the rack and working your way up, even if there is only one piece of hardware in the rack. Install heavier hardware below lighter hardware.

Ensure that hardware is placed so as to safely distribute mechanical load.

Never use rack mounted hardware as a shelf or a work space.

Always ensure that the rack is stable before extending a server or other component from the rack.

Never extend more than one server or component at a time. Extending two or more components simultaneously may destabilize the rack.

## 12.4. SAFE OPERATING TEMPERATURE

**WARNING**

Failing to follow these recommendations can lead to injury or death.

Mount hardware in racks such that each piece of hardware has the appropriate amount of space and airflow.

When you are not servicing hardware, keep the rack doors and all covers and panels closed to maintain a safe operating temperature. Additionally, ensure that the ambient temperature inside the rack does not exceed your hardware manufacturer's maximum rated ambient temperature (TMRA).

Ensure that you allow the hardware time to cool before touching them, especially when handling drives and power supply modules.

## 12.5. ELECTRICAL SAFETY RECOMMENDATIONS



## **WARNING**

Failing to follow these recommendations can lead to injury or death.

Read the electrical and general warnings online in *Standardized Warning Statements for AC Systems*:  
[http://www.supermicro.com/about/policies/safety\\_information.cfm](http://www.supermicro.com/about/policies/safety_information.cfm).

Use a regulating uninterruptible power supply (UPS) to protect hardware from power surges and voltage spikes, and to ensure systems continue to operate during a power failure.

Do not overload circuits, as this can reduce the effectiveness of surge protection and supply wiring. Review equipment nameplate ratings to determine the best way to connect your hardware to power.

Ensure that your hardware and rack are reliably grounded at all times.

## CHAPTER 13. NETWORKING REQUIREMENTS

Red Hat Storage One requires a minimum of two networks:

### Management network

The management network is a dedicated virtual LAN on the network switch ports that connect the Red Hat Gluster Storage servers. It requires multicast DNS support. During deployment, the management network requires DHCP, and can either use an existing DHCP configuration, or be automatically configured with DHCP by the Colonizer. After deployment, this network can be used for storage administration tasks.

Use at least one of the onboard Ethernet ports, LAN1 and LAN2, to create the management network. These ports are automatically teamed in active-backup mode as the **mgmt\_team** virtual interface, and do not require any special switch configuration. Red Hat recommends connecting both LAN1 and LAN2 if possible.

### Storage network

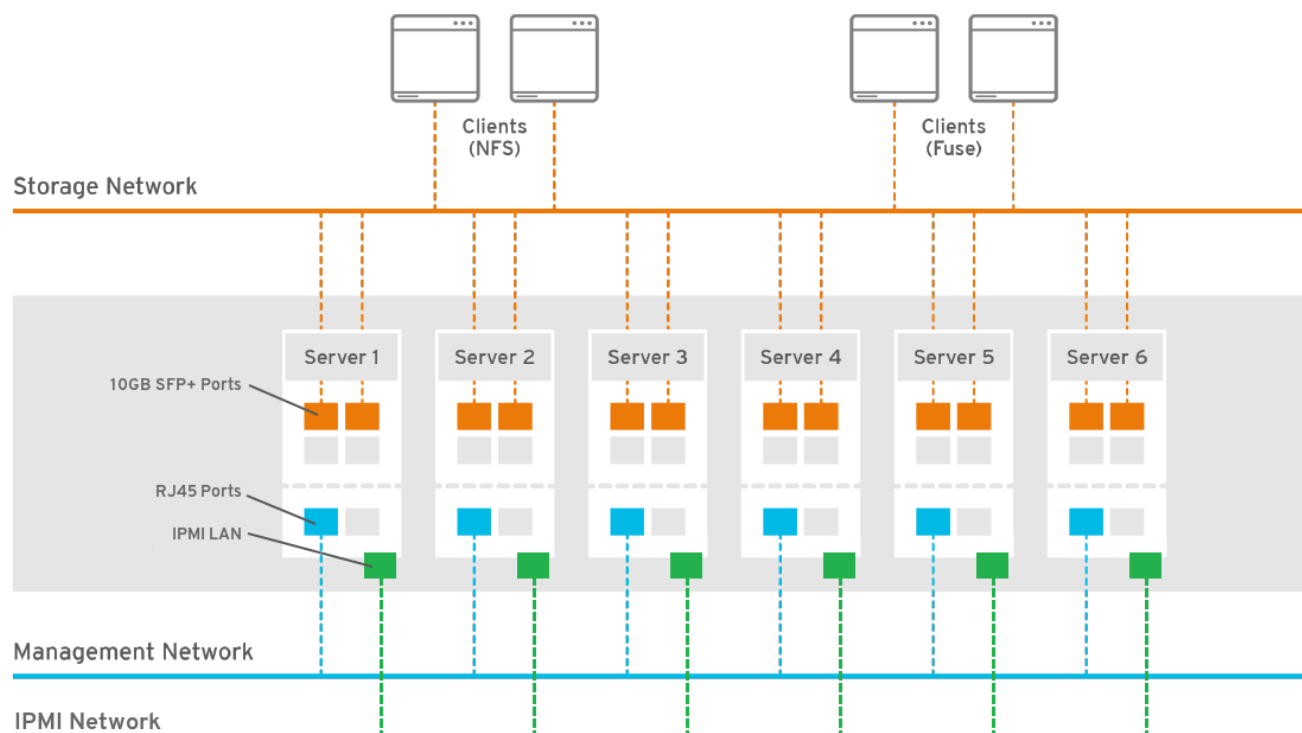
The storage network is the network used for Red Hat Gluster Storage processes and data in production.

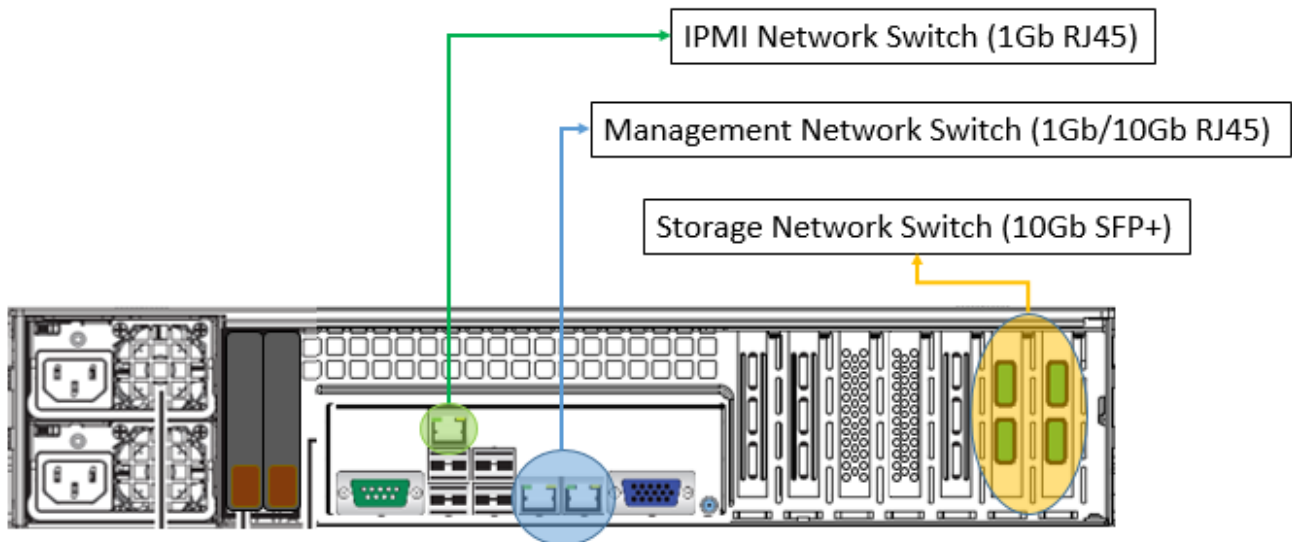
Use at least two of the PCIe Ethernet ports provided to create this network. These ports are automatically teamed in active-backup mode as the **storage\_team** virtual interface, and do not require any special switch configuration. Red Hat recommends connecting all four ports if possible.

### (Recommended) IPMI network

Your Red Hat Storage One server uses its IPMI LAN network interface to search for a DHCP server and request the next available IP address. Additional information regarding IPMI configuration, including how to set a static IP address, can be found in the, [IPMI User's Guide](#). The default login with username **ADMIN** and password **ADMIN** should be changed immediately for security purposes. See [Best practices for managing servers with IPMI features enabled](#) for more information about recommendations for IPMI.

**Figure 13.1. Network configuration diagram**



**Figure 13.2. Network ports**

You must provide correctly configured networking hardware (cabling and switches) to provide network access between your Red Hat Storage One servers.

You must also provide the following information during deployment for permanent use by the storage network in production:

- A short host name and static IP address for each server
- A storage network domain name
- A static storage network subnet

If multiple interfaces are active on these networks, and DHCP is not already running, the Colonizer configures a server-side, temporary DHCP network for the duration of the initial configuration process. This DHCP network broadcasts and assigns IP addresses to all devices on the network, but the Colonizer does not attempt to configure devices that are not part of Red Hat Storage One. You can continue to use this network after deployment is complete, to administer Red Hat Storage One on an ongoing basis.

Red Hat also recommends providing the following:

- (Optional) A default gateway
- (Optional) One or more DNS servers