
LOCKSS Node Setup Guide

LOCKSS Program, Stanford University Libraries

Last updated 2018-09-09

Table of Contents

1. Overview	1
2. Checklist	1
3. Setting Up the Machine	4
3.1. Hardware Considerations	4
3.2. Installing the Operating System	4
3.3. Configuring Storage Arrays	5
4. Installing the LOCKSS Software	7
5. Configuring the LOCKSS Software	8
6. Verifying the Installation	11
A. Installing CentOS 7	12
B. CentOS 7 Storage Array Configuration	29

1. Overview

This document explains how to set of a Linux system to install and operate the LOCKSS software.

The general outline of the setup process is as follows.

Procedure 1. Summarized setup process

1. Install the operating system and configure storage volumes. See Section 3.
2. Install the LOCKSS software. See Section 4.
3. Configure the LOCKSS software. See Section 5.
4. Verify the installation. See Section 6.

Skip Step 1 if you are using an existing machine with a Linux operating system that can install packages with **rpm**.

2. Checklist

This document uses symbolic placeholders for values that are relevant to your particular environment, such as host names and IP addresses. See Table 1 and make a note of the values corresponding to your situation. Example values are given for illustration purposes.

Table 1. Symbolic placeholders

Symbol	Meaning
	Basic network information
<code>\${NODEHOST}</code>	The fully qualified host name of the LOCKSS node.

Symbol	Meaning
	Example value: lockss.myuniversity.edu Your value:
<code>`\${NODEIP}`</code>	The IP address of the LOCKSS node. Example value: 192.168.1.123 Your value:
<code>`\${NODENETMASK}`</code>	The netmask of the LOCKSS node. Example value: 255.255.255.0 Your value:
<code>`\${NODEGATEWAY}`</code>	The IP address of the LOCKSS node's gateway. Example value: 192.168.1.1 Your value:
<code>`\${NODEDNS1}`</code>	The IP address of the LOCKSS node's primary DNS server. Example value: 8.8.8.8 Your value:
<code>`\${NODEDNS2}`</code>	Optional. The IP address of the LOCKSS node's secondary DNS server. Example value: 8.8.4.4 Your value:
Advanced network information (optional)	
<code>`\${NODENATIP}`</code>	Optional. The external IP address of the LOCKSS node, if network address translation (NAT) is in use. Example value: 172.31.255.1 Your value:
<code>`\${PROXYIP}`</code>	Optional. The IP address of a proxy the LOCKSS node is required to use for outgoing traffic. Example value: 192.168.1.100 Your value:
<code>`\${PROXYPORT}`</code>	Optional. The port number for the outgoing proxy at <code>`\${PROXYIP}`</code> . Example value: 8888 Your value:
E-mail information	
<code>`\${MAILHOST}`</code>	The host name of the mail relay the LOCKSS node uses. Example value: smtp.myuniversity.edu

Symbol	Meaning
	Your value:
<code>\${MAILUSER}</code>	The user name for the mail relay at <code>\${MAILHOST}</code> , if applicable. Example value: mailuser Your value:
<code>\${MAILPASS}</code>	The password for the mail relay at <code>\${MAILHOST}</code> , if applicable. Example value: mailpass Your value:
<code>\${ADMINADDR}</code>	The e-mail address of the LOCKSS node's administrator. Example value: jsmith@myuniversity.edu Your value:
Web user interface information	
<code>\${USERNET}</code>	User network accessing the Web user interface (UI) of the LOCKSS node, in CIDR notation. If the Web user interface is accessed from multiple user networks, this will be a list. Example value: 192.168.1.0/24 Your value(s):
<code>\${UIUSER}</code>	User name used to access the Web user interface of the LOCKSS node. Recommended value: lockss Your value:
LOCKSS network configuration information	
<code>\${PLNCONFIG}</code>	The URL of the LOCKSS network's configuration file. This value is given to you by the administrator of the network's infrastructure server. If you are joining the Global LOCKSS Network (GLN), the value is <code>http://props.lockss.org:8001/daemon/lockss.xml</code> . Example value: <code>http://infra.mybigpln.org:8001/mybigpln/lockss.xml</code> Your value:
<code>\${PLNCODE}</code>	A short code representing the name of the LOCKSS network. This value is given to you by the administrator of the network's infrastructure server. If you are joining the Global LOCKSS Network (GLN), the value is <code>prod</code> . Example value: mybigpln Your value:

3. Setting Up the Machine

3.1. Hardware Considerations

Hardware considerations vary depending on the expected level of activity incurred by the node and the expected cumulative size of the content to be preserved by the node.

We recommend a 64-bit dual-core CPU (quad-core preferred), with 8GB of memory (or more) and a boot-capable CD or DVD drive. The machine can be a dedicated server or a virtual machine.



Note

Whether you choose a physical or virtual machine, we no longer recommend a 32-bit architecture for new installations.

We recommend devoting one (preferably two) small disks (for instance solid-state disks) to the system itself (boot partition, EFI system partition, swap partition, root partition), and bundling commodity hard disks into one or more arrays (with software RAID) for the node's storage. If you cannot devote one or two disks to the system, you can alternatively dedicate a modest amount of space from the first storage array to that function. The number and cumulative capacity of the storage arrays depends on the expected size of the content in the network to be harvested and preserved by the node. In its simplest form, a minimal LOCKSS node could consist of a single array of disks, jointly housing a storage array and the system partitions, in a 1U form factor. Some LOCKSS boxes house as many as sixty hard disks in a 4U form factor.

For software RAID, we recommend RAID6 over RAID5. While RAID5 offers some protection against failed disks, the data in the array becomes vulnerable while one disk in the array has failed and is being replaced and repopulated. The data in a RAID6 array does not become vulnerable until two disks in the array have failed; if a single disk in the array fails, the data in the array is not at high risk during the window of time it might take to purchase, install and repopulate a replacement disk.

Note in the planning of your storage needs that the disks providing redundancy to the array -- two per array in RAID6, one per array in RAID5 -- do not contribute to the total usable storage capacity of the array. For instance, an array of six 4TB disks will yield 16TB of usable storage in RAID6, or 20TB in RAID5, minus file system overhead.

You can also use remote storage, for example over iSCSI or NFS. For more information about these options, contact <lockss-support@lockss.org>.

You are encouraged to review your hardware configuration and ask any questions you might have prior to installation by contacting the LOCKSS Team over e-mail at <lockss-support@lockss.org>.

3.2. Installing the Operating System

For the operating system, we recommend **CentOS 7**, a flavor of Linux based on Red Hat Enterprise Linux. See Appendix A for a CentOS 7 installation guide, with the following additional considerations for Procedure A.1:

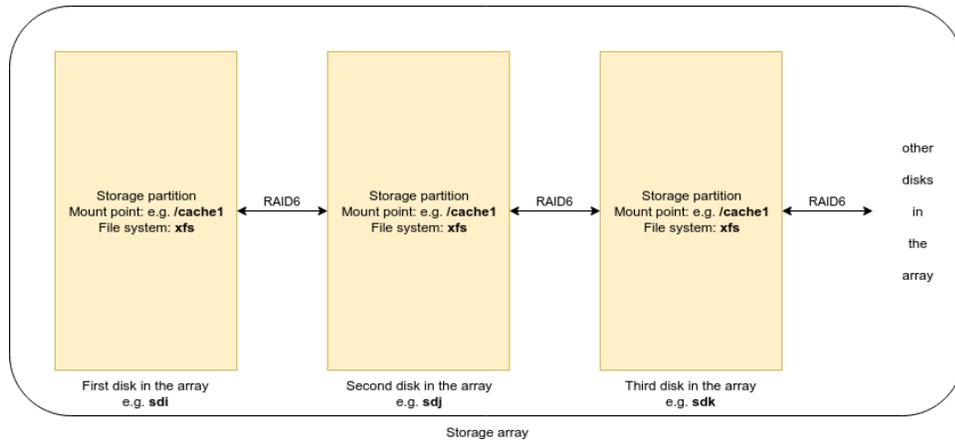
- In Step 12, you will need the networking-related values from Table 1 to set up the node's networking interface: `${NODEHOST}`, `${NODEIP}`, `${NODENETMASK}`, `${NODEGATEWAY}`, `${NODEDNS1}`, and `${NODEDNS2}`.
- In Step 18, you will need to configure storage volumes. The recommended layout and procedure is detailed in Section 3.3.

3.3. Configuring Storage Arrays

Section last updated: 2018-09-09

For the node's storage, we recommend bundling disks into arrays, preferably with **RAID6**, otherwise with RAID5. Under CentOS, we recommend the **XFS** file system for these storage arrays, which should have a single mount point each spanning the whole array. Historically in the LOCKSS system, storage mount points have been named `/cache0`, `/cache1`, `/cache2`, etc. This simple layout is illustrated in Figure 1. Only if you cannot dedicate one or two disks to the system partitions will the first storage array (`/cache0`) look different; see "Shared system array" and Figure 4 below.

Figure 1. Storage array



The system itself requires several partitions: a boot partition (mount point `/boot`), an EFI system partition (mount point `/boot/efi`), a swap partition, and a root partition (mount point `/`). The EFI system partition must use a FAT-compatible file system; CentOS recognizes the `/boot/efi` mount point and assigns it the file system type "EFI System Partition". For the other system partitions, under CentOS we recommend the XFS file system.

We recommend allocating 512MB to both the boot partition and the EFI system partition. We further recommend allocating twice as much as the machine has physical memory to the swap partition. If you are dedicating one or two disks to the system partitions, the remainder of the system disk or disks is for the root partition. If you are using a shared system array instead, you need to decide how much space to allocate to the root partition. We do not recommend less than 20GB of usable space (that is, 20GB on each of the disks involved in the RAID1 array of the root partition).

We recommend the following layout, and offer two alternatives, in order of preference:

- **Dedicated system array.** Ideally, two small disks (for example solid-state disks) are bundled together into a system array. The first disk in the array is the boot device. It has the boot partition, the EFI system partition, and the swap partition, with no counterparts on the second disk in the array. The remainder of the first disk is devoted to the root partition, mirrored using RAID1 to its counterpart on the second disk. This layout is illustrated in Figure 2.
- **Dedicated system disk.** Alternatively, a small disk (for example a solid-state disk) is designated as the system disk. This disk is the boot device, and features the boot partition, the EFI system partition, the swap partition, and the root partition. This layout is illustrated in Figure 3.
- **Shared system array.** If neither option is applicable, a modest amount of storage from the first storage array can be devoted to the system partitions. The first disk in the first storage array is the boot device. It has the boot partition, the EFI system partition, and the swap partition, with no counterparts on the other

disks in the array. Then it features the root partition, mirrored to its counterparts on each of the other disks in the array using RAID1. (Finally, the remainders of each disk in the array are bundled together into the first storage partition.) This layout is illustrated in Figure 4.

Figure 2. Dedicated system array

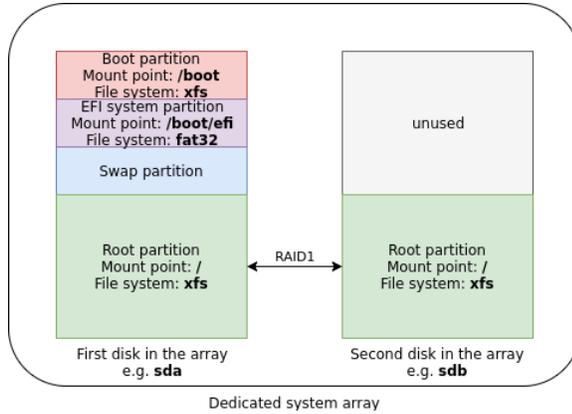


Figure 3. Dedicated system disk

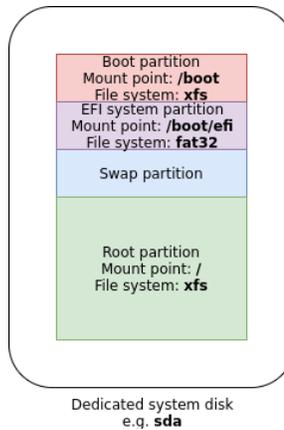
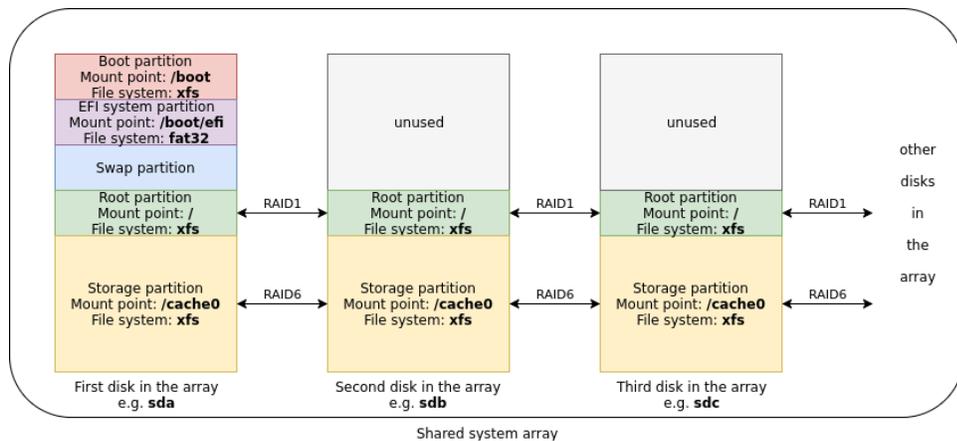


Figure 4. Shared system array



In Appendix B, we present a procedure to configure storage arrays in CentOS 7. (For other operating systems, refer to the usage manual for disk partitioning instructions.)

4. Installing the LOCKSS Software

Once the basic installation of your Linux system is complete, you are ready to install the LOCKSS software.

The LOCKSS Team offers a basic integration script to perform pre-requisite system steps on CentOS 7.

Procedure 2. LOCKSS software installation

1. Install the **wget** utility and **net-tools** as root with this command:

```
sudo yum install wget net-tools
```

Follow the on-screen prompts.

2. In a working directory, for instance `/tmp`, download the integration script from GitHub with this command:

```
wget https://github.com/lockss/admin-tools/raw/master/centos/install-lockss-centos7.sh
```

3. Run the script as root:

```
sudo sh install-lockss-centos7.sh
```

4. For each one of the storage mount points defined in Section 3.3, typically called `/cache0`, `/cache1`, `/cache2`, etc., do the following as root:

```
sudo mkdir /cache0/gamma /cache1/gamma /cache2/gamma ...
sudo chown lockss:lockss /cache0/gamma /cache1/gamma /cache2/gamma ...
sudo chmod 0750 /cache0/gamma /cache1/gamma /cache2/gamma ...
```

This has the effect of creating a single directory in each called `gamma` (and setting permissions appropriately).

The end result of this section is illustrated in Figure 5.

Figure 5. Result of integration script

```
331:443... connected.
HTTP request sent, awaiting response... 200 OK
Length: 5157 (5.0K) [text/plain]
Saving to: 'install-lockss-centos7.sh'

100%[=====>] 5,157      --.-K/s   in 0.001s

2017-06-07 23:58:32 (7.34 MB/s) - 'install-lockss-centos7.sh' saved [5157/5157]

[thib@localhost ~]$ sudo sh install-lockss-centos7.sh
 * Creating user lockss...
 * Creating user lcap...
 * Adding LOCKSS RPM repository...
 * Importing LOCKSS GPG key...
 * Installing software...
Please run /etc/lockss/hostconfig . See /etc/lockss/README for details.
 * Enabling lockss service...
lockss.service is not a native service, redirecting to /sbin/chkconfig.
Executing /sbin/chkconfig lockss on
 * Enabling yum-cron service...
 * Adjusting firewall...
success
success
 * Done.
[thib@localhost ~]$ _
```

If you encounter any difficulty during this process, contact lockss-support@lockss.org for assistance.

5. Configuring the LOCKSS Software

The next step is to configure the LOCKSS software by running the `/etc/lockss/hostconfig` script as root. The script asks a series of configuration questions. For many, a value is suggested in square brackets; you can accept it by simply hitting **Enter**.

Procedure 3. hostconfig procedure

1. Run the following command:

```
sudo /etc/lockss/hostconfig
```

You may be prompted for your password.

2. Enter the fully qualified host name of the machine, `_${NODEHOST}` and hit **Enter**. For instance, the example value in Table 1 would be `lockss.myuniversity.edu`.
3. Enter the IP address of the machine, `_${NODEIP}` and hit **Enter**. For instance, the example value in Table 1 would be `192.168.1.123`.
4. Next you will be asked if the machine is behind Network Address Translation (NAT).
 - If the machine is not behind NAT (common), simply hit **Enter** to accept the default value `n` (for "no").
 - If the machine is behind NAT (uncommon):
 - Enter `y` (for "yes") and hit **Enter**.
 - Enter the external IP address of the machine behind NAT (`_${NODENATIP}`), and hit **Enter**. For instance, the example value in Table 1 would be `172.31.255.1`.
5. Enter the initial list of subnets that should be granted access to the administrative user interface (UI) of the LOCKSS instance. The default value suggested in square brackets is the class C subnet the machine is currently on. Type in the value `_${USERNET}`, which can be a single subnet, or a list of subnets separated by semicolons (see Table 1). Hit the **Enter** key to validate your entry.

Note that this setting can be adjusted later in the administrative UI without re-running the configuration script.

6. Enter the desired LCAP port (the port over which the LOCKSS node will communicate with other LOCKSS nodes). Unless you are doing something exotic like running multiple LOCKSS instances out of a single IP address, we recommend you hit **Enter** to accept the default value `9729`.
7. Enter the desired proxy port (the port on which the LOCKSS node can run a Web proxy for certain client IP addresses). We recommend you simply hit **Enter** to accept the default value `8080`.
8. Enter the desired administrative user interface (UI) port (the port on which the LOCKSS node runs its Web-based UI). We recommend you simply hit **Enter** to accept the default value `8081`.
9. Enter the host name of the mail relay for the machine, `_${MAILHOST}` and hit **Enter**. Using the example value in Table 1 you would enter `smtp.myuniversity.edu`.
10. Next you will be asked if the mail relay requires a username and password.
 - If the mail relay does not require a username and password (common), hit **Enter** to accept the default value `n` (for "no").

- If the mail relay requires a username and password (uncommon):
 - Enter **y** (for "yes") and hit **Enter**.
 - Enter the username for the mail relay, `#{MAILUSER}` and hit **Enter**. Using the example value in Table 1 you would enter `mailuser`.
 - Enter the password for the mail relay, `#{MAILPASS}` and hit **Enter**. Using the example value in Table 1 you would enter `mailpass`.
- 11. Enter the e-mail address of the administrator of the LOCKSS instance, `#{ADMINADDR}` and hit **Enter**. Using the example value in Table 1 you would enter `jsmith@myuniversity.edu`.
- 12. Enter the path to the Java executable (**java**). We recommend hitting **Enter** to accept the suggested value derived by the system.
- 13. Enter any command-line switches you wish to pass to the Java executable. By default, this is not necessary and you should simply leave the value empty and simply hit **Enter**.
- 14. Enter the URL of your LOCKSS network's configuration file, `#{PLNCONFIG}`. Only accept the default value if you are joining the Global LOCKSS Network, otherwise enter the value supplied by your LOCKSS network administrator. Using the example value in Table 1 you would enter `http://infra.mybigpln.org:8001/mybigpln/lockss.xml`.
- 15. Next you will be asked if you require a proxy to access your LOCKSS network's configuration file.
 - If you do not require a proxy to access the configuration file (common), hit **Enter** to accept the default value `NONE`.
 - If you require a proxy to access the configuration file (uncommon), enter `#{PROXYIP}`:
`#{PROXYPORT}` and hit **Enter**.
- 16. You will then be asked if you would like to enable the configuration failover feature.
 - If you wish to enable the failover feature (recommended):
 - Hit **Enter** to accept the default answer **y** (for "yes").
 - You will then be asked the maximum age of the configuration failover file. Accept the default value by hitting **Enter**.
 - If you do not wish to enable the failover feature, enter **n** (for "no"), then hit **Enter**.
- 17. Enter the code name of your LOCKSS network, `#{PLNCODE}`. (The configuration script refers to it as the preservation group name.) Enter the value supplied by your LOCKSS network administrator, or if you are joining the Global LOCKSS Network (GLN), accept the default value `prod`. Using the example value in Table 1, you would enter `mybigpln`.
- 18. Enter the list of `gamma` directories created in Procedure 2 Step 4, separated by semicolons, then hit **Enter** to validate your entry. In the example machine used throughout this document, there are two storage partitions, `/cache0` and `/cache1`, so the corresponding `gamma` directories are `/cache0/gamma` and `/cache1/gamma`, meaning you would enter the value `/cache0/gamma;/cache1/gamma`.
- 19. Enter the path of a temporary storage area for use by the LOCKSS software. We recommend hitting **Enter** to accept the suggested value, which will be derived from the first storage area (e.g. `/cache0`).
- 20. Enter the username of the main user of the LOCKSS node's administrative user interface (UI), `#{UIUSER}`, and hit **Enter**. Using the recommended value in Table 1 you would enter `lockss`.

21. Enter the password for the main user of the LOCKSS node's administrative user interface (UI), and hit **Enter**. Then confirm the password, again ending with **Enter**.
22. Verify that the configuration values you have entered are correct, and if so, hit **Enter** to accept the default response **Y** (for "yes"). If a value is incorrect, type **N** (for "no") and hit **Enter** to go back to the beginning.
23. The configuration script will then create files and directories and perform other necessary tasks, asking for confirmation at each step. The response **Y** (for "yes") is the default for each, so you can simply hit **Enter**.

Representative screenshots of this interactive process are shown in Figure 6, Figure 7, Figure 8 and Figure 9.

Figure 6. Screenshot of hostconfig (1)

```
[jsmith@localhost ~]$ sudo /etc/lockss/hostconfig
root is configuring
eth0: error fetching interface information: Device not found
LOCKSS host configuration for Linux.
For more information see /etc/lockss/README
Configuring for user lockss
Fully qualified hostname (FQDN) of this machine: [localhost.localdomain] lockss.
myuniversity.edu
IP address of this machine: [ ] 192.168.1.123
Is this machine behind NAT?: [N]
Initial subnet for admin UI access: [192.168.1.0/24]
LDAP V3 protocol port: [9729]
PROXY port: [8080]
Admin UI port: [8081]
Mail relay for this machine: [localhost] smtp.myuniversity.edu
Does mail relay smtp.myuniversity.edu need user & password: [N]
E-mail address for administrator: [ ] jsmith@myuniversity.edu
Path to java: [/bin/java]
Java switches: [ ]
Configuration URL: [http://props.lockss.org:8001/daemon/lockss.xml] http://infra.
mybigpln.org/mybigpln/lockss.xml
Configuration proxy (host:port): [NONE]
Enable config failover: [Y]
Config failover max age: [ ]
Preservation group(s): [prod] _
```

Figure 7. Screenshot of hostconfig (2)

```
Mail relay for this machine: [localhost] smtp.myuniversity.edu
Does mail relay smtp.myuniversity.edu need user & password: [N]
E-mail address for administrator: [ ] jsmith@myuniversity.edu
Path to java: [/bin/java]
Java switches: [ ]
Configuration URL: [http://props.lockss.org:8001/daemon/lockss.xml] http://infra.
mybigpln.org/mybigpln/lockss.xml
Configuration proxy (host:port): [NONE]
Enable config failover: [Y]
Config failover max age: [ ]
Preservation group(s): [prod] mybigpln
Content storage directories: [ ] /cache0/gamma:/cache1/gamma
Temporary storage directory: [/cache0/gamma/tmp]
User name for web UI administration: [ ] lockss
Password for web UI administration user lockss: [ ]
Password for web UI administration (again): [ ]

Configuration:
LOCKSS_CONFIG_VERSION=1
LOCKSS_USER="lockss"
LOCKSS_HOSTNAME=lockss.myuniversity.edu
LOCKSS_IPADDR=192.168.1.123
LOCKSS_EXTERNAL_IPADDR=
LOCKSS_V3_PORT=9729
LOCKSS_ACCESS_SUBNET="192.168.1.0/24"
```

Figure 8. Screenshot of `hostconfig` (3)

```

Temporary storage directory: [/cache0/gamma/tmp]
User name for web UI administration: [] lockss
Password for web UI administration user lockss: []
Password for web UI administration (again): []

Configuration:
LOCKSS_CONFIG_VERSION=1
LOCKSS_USER="lockss"
LOCKSS_HOSTNAME=lockss.myuniversity.edu
LOCKSS_IPADDR=192.168.1.123
LOCKSS_EXTERNAL_IPADDR=
LOCKSS_V3_PORT=9729
LOCKSS_ACCESS_SUBNET="192.168.1.0/24"
LOCKSS_MAILHUB=smtp.myuniversity.edu
LOCKSS_MAILHUB_USER=
LOCKSS_MAILHUB_PASSWORD=
LOCKSS_EMAIL=jsmith@myuniversity.edu
LOCKSS_JAVA_CMD=/bin/java
LOCKSS_JAVA_SWITCHES=
LOCKSS_JAVA_HEAP=
LOCKSS_PROPS_URL="http://infra.mybigpln.org/mybigpln/lockss.xml"
LOCKSS_PROPS_PROXY="NONE"
LOCKSS_PROPS_SERVER_AUTHENTICATE_KEYSTORE=""
LOCKSS_CONFIG_FAILOVER_ENABLE="Y"
LOCKSS_CONFIG_FAILOVER_MAX_AGE=""

```

Figure 9. Screenshot of `hostconfig` (4)

```

LOCKSS_ACCESS_SUBNET="192.168.1.0/24"
LOCKSS_MAILHUB=smtp.myuniversity.edu
LOCKSS_MAILHUB_USER=
LOCKSS_MAILHUB_PASSWORD=
LOCKSS_EMAIL=jsmith@myuniversity.edu
LOCKSS_JAVA_CMD=/bin/java
LOCKSS_JAVA_SWITCHES=
LOCKSS_JAVA_HEAP=
LOCKSS_PROPS_URL="http://infra.mybigpln.org/mybigpln/lockss.xml"
LOCKSS_PROPS_PROXY="NONE"
LOCKSS_PROPS_SERVER_AUTHENTICATE_KEYSTORE=""
LOCKSS_CONFIG_FAILOVER_ENABLE="Y"
LOCKSS_CONFIG_FAILOVER_MAX_AGE=""
LOCKSS_TEST_GROUP="mybigpln"
LOCKSS_DISK_PATHS="/cache0/gamma:/cache1/gamma"
LOCKSS_ADMIN_USER=lockss
LOCKSS_ADMIN_PASSWD=SHA-256:bc0b3ecbe9a9564b64680ce5d2fad19d61f67ab656d2056c30732a5da2006cef
LOCKSS_PROXY_PORT=8080
LOCKSS_UI_PORT=8081
LOCKSS_TMPDIR=/cache0/gamma/tmp
LOCKSS_CLEAR_TMPDIR=yes
LOCKSS_RELEASE=1.72.3-1
LOCKSS_HOME is
OK to store this configuration: [Y] Y_

```

If you encounter any difficulty during this process, contact lockss-support@lockss.org for assistance.

6. Verifying the Installation

The `hostconfig` script sets up the LOCKSS daemon to start up when the system boots. You can either reboot the machine, or to cause the LOCKSS daemon to start right away, type this command:

```
sudo /etc/init.d/lockss start
```

From an IP address on the machine's authorized subnet (`${USERNET}`), use a Web browser to access `http://${NODEHOST}:8081` or `http://${NODEIP}:8081`.

A. Installing CentOS 7

Section last updated: 2017-05-10

CentOS 7 is a popular, enterprise-grade Linux distribution based on Red Hat Enterprise Linux (RHEL). Visit the CentOS Web site at <https://www.centos.org/> for downloads, documentation, support, and more.

Figure A.1. CentOS Web site front page



This section presents a guided overview of the CentOS 7 installation process.

As of this writing, the latest CentOS 7 release is CentOS 7 Release 1611 (codename for November 2016).



Note

Although CentOS 6 will receive maintenance updates until 2020, it will only receive full updates until mid-2017, so we no longer recommend it for new installations.

Several variants are offered. The "Minimal" variant fits on a CD and does not install a graphical desktop environment. The "DVD" and "Everything" variants do, but only the "DVD" variant fits on a DVD. The "Netinstall" variant is the smallest to download upfront, and downloads software packages over the network. Since a typical installation does not require many available packages, this variant is the most efficient in total time spent downloading from the network, so we recommend it. The remainder of this section is written in terms of the **"Netinstall" variant**.

Procedure A.1. Overview of CentOS 7 installation

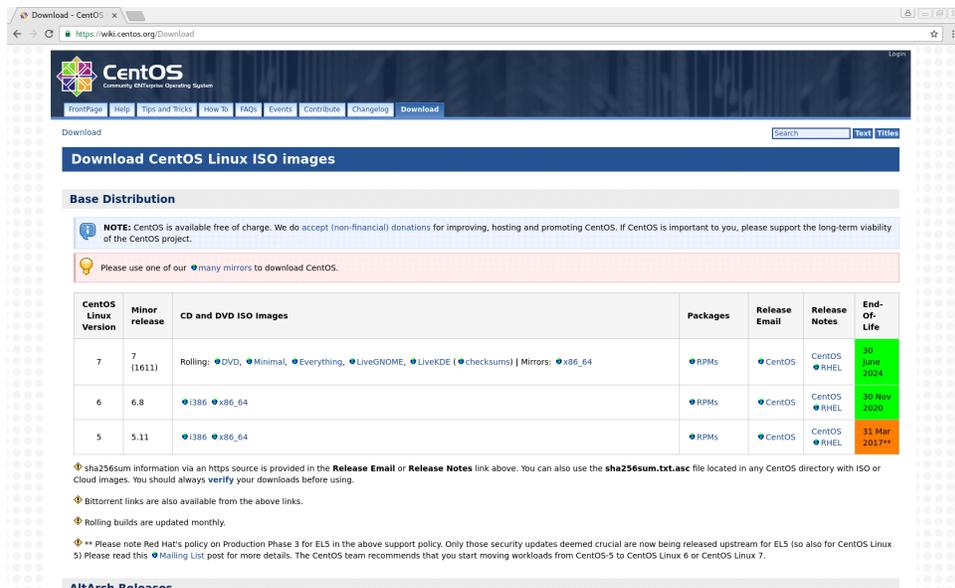
1. Visit <https://www.centos.org/download/>.

Figure A.2. CentOS Web site download page

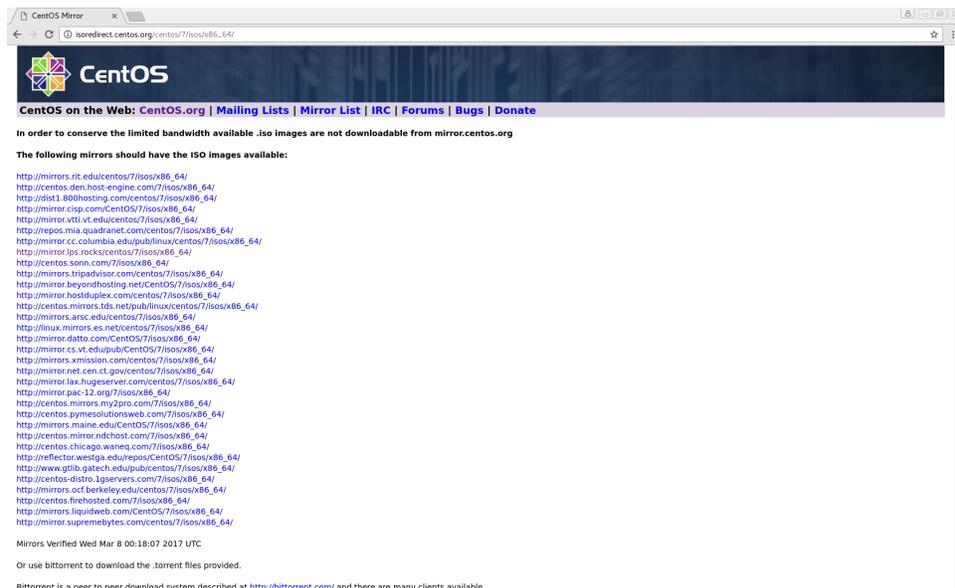


2. Click on “alternative downloads” or “More download choices”, which takes you to <https://wiki.centos.org/Download>.

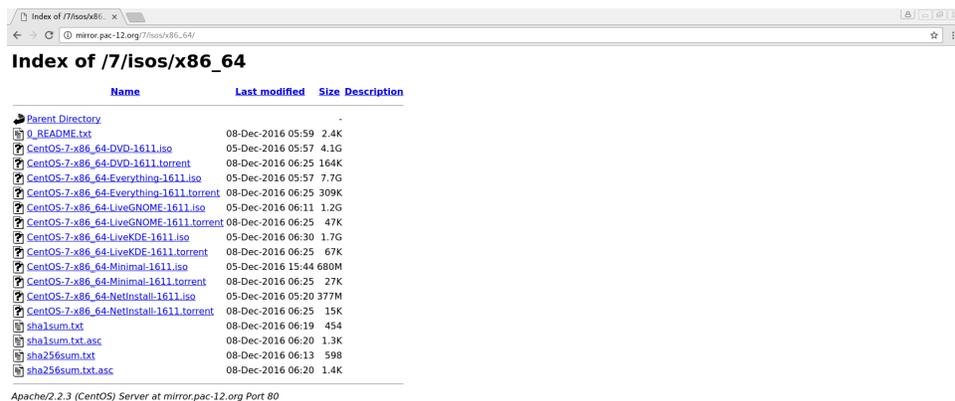
Figure A.3. CentOS Web site alternative downloads page



3. In the row for CentOS version 7, click on “Mirrors: x86_64”, which takes you to http://isoredirect.centos.org/centos/7/isos/x86_64/.

Figure A.4. CentOS Web site download mirror selection

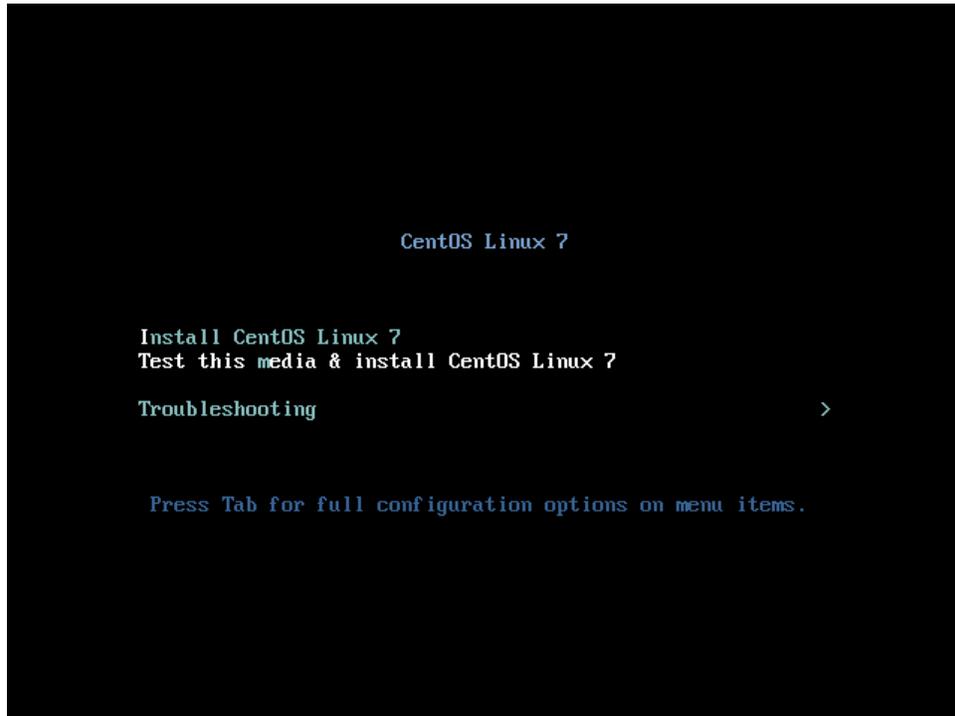
4. Select a mirror.

Figure A.5. CentOS Web site download mirror

5. Save the “Netinstall” ISO image to your computer. There is also an option to download the image as a torrent.
6. Verify the integrity of the downloaded ISO image using the SHA-256 checksum provided on the mirror. General instructions for how to do this on various platforms are provided at <https://wiki.centos.org/TipsAndTricks/sha256sum>.
7. Burn the ISO image to a CD or DVD. Alternatively, you can perform the installation from a USB key; see <https://wiki.centos.org/HowTos/InstallFromUSBkey> for details.

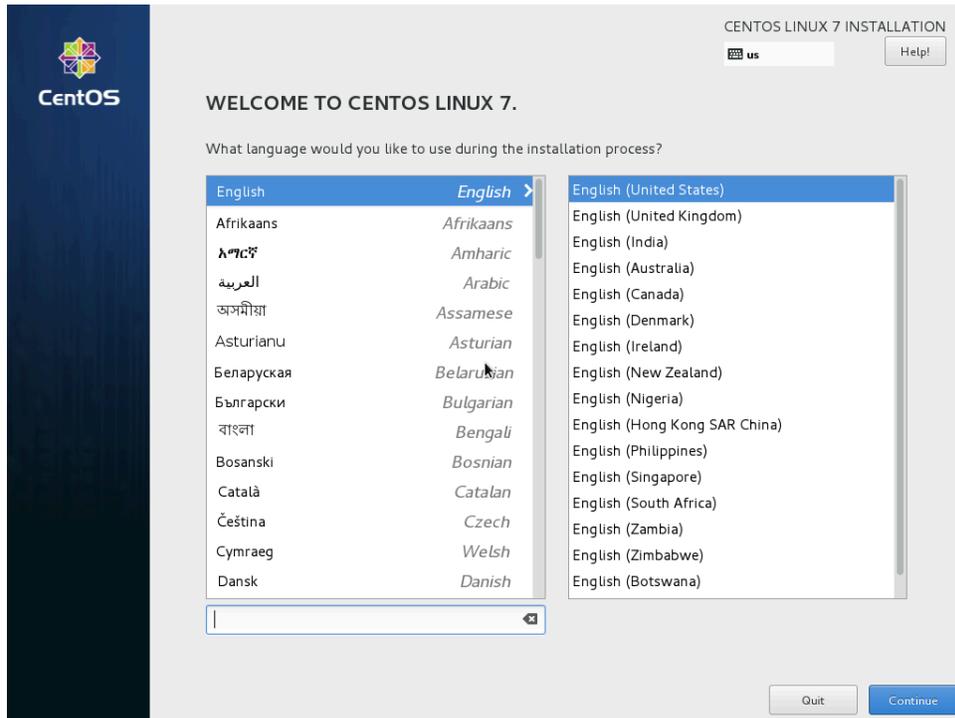
8. Power the target machine on, such that it boots from the installation CD or DVD (or USB key).
9. Use the keyboard arrows to select Install CentOS Linux 7 (I), then hit **Enter**. The graphical installer will then load.

Figure A.6. CentOS installation boot screen



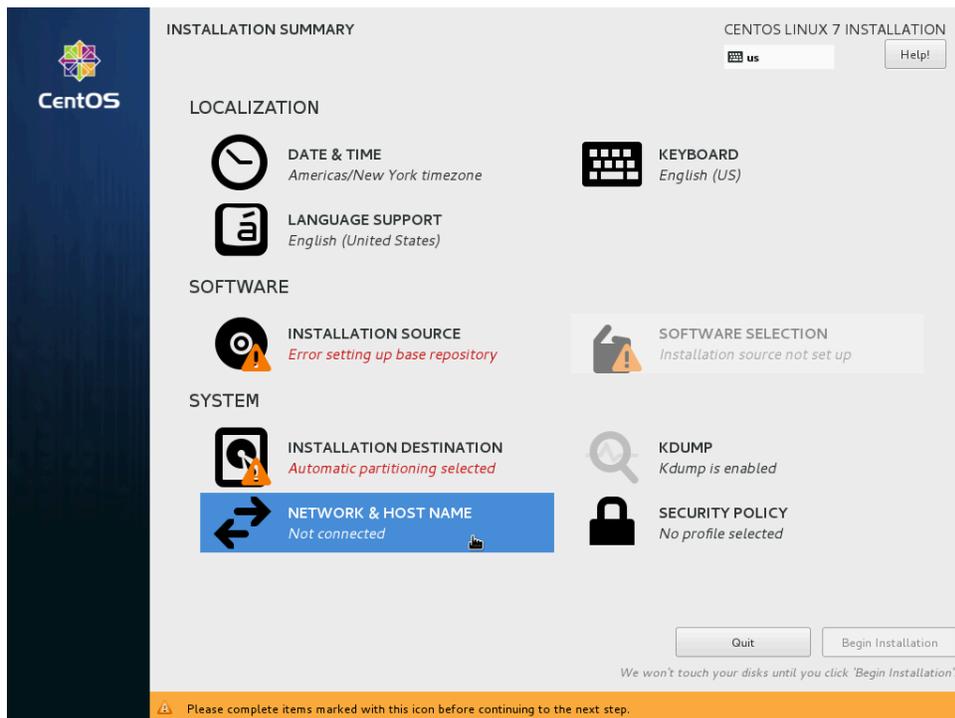
10. Select the installation language, for instance English → English (United States), then click Continue.

Figure A.7. CentOS installation language



11. Click on Network & Host Name.

Figure A.8. CentOS installation network and host name

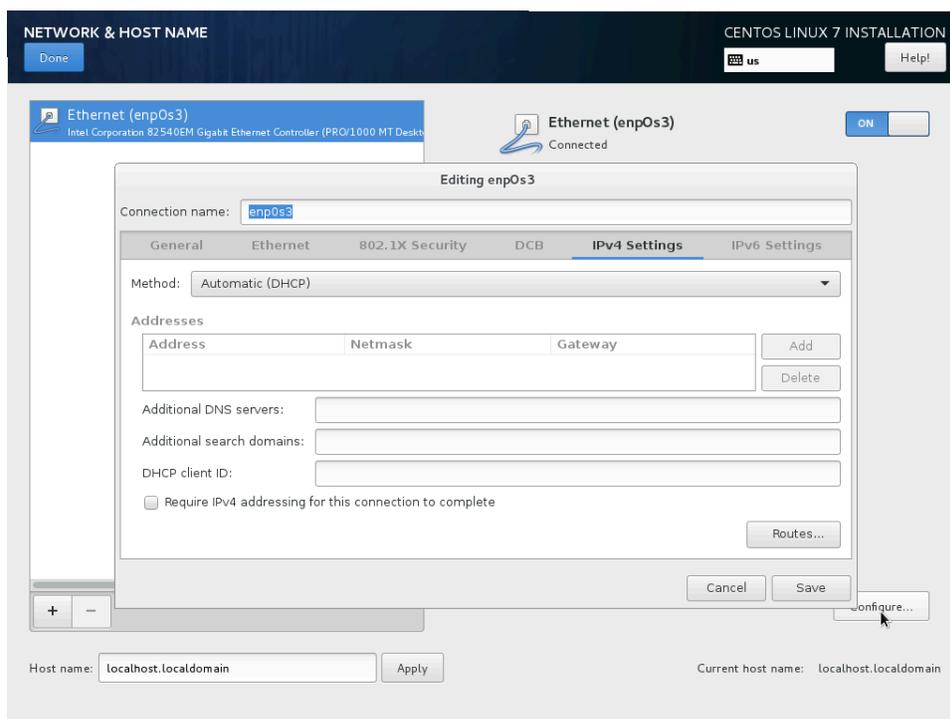


- Configure network interfaces as necessary for the machine to have access to the Internet. Enter the machine's fully-qualified domain name in the Host name box. You can then turn an interface on and off with the On/Off toggle, add interfaces with the + button, and configure a selected interface with the Configure button.

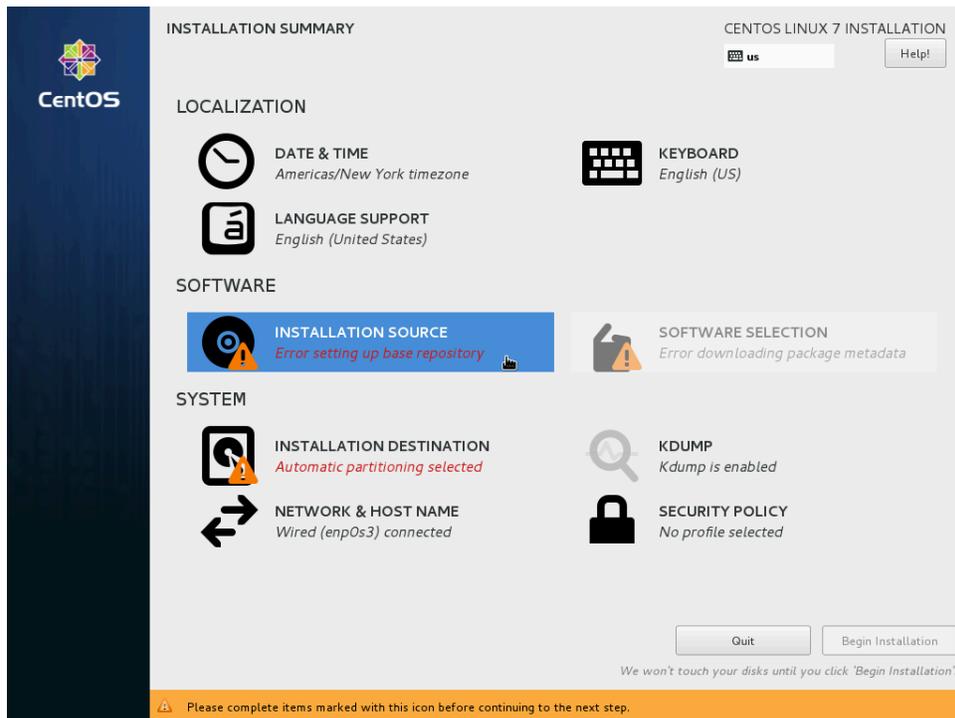
The IPv4 Settings (or IPv6 Settings, if applicable) tab in the interface configuration dialog (shown below) enables you to enter networking information, such as the IP address, netmask, gateway, DNS servers, and more.

When the machine is successfully connected, click Done in the top left corner.

Figure A.9. CentOS installation network and host name screen



- Click on Installation Source.

Figure A.10. CentOS installation source

- Click the On the network radio button, select `http://` from the drop-down menu, and enter the following into the adjacent text box: `mirror.centos.org/centos/7/os/x86_64/`, so that the result is `http://mirror.centos.org/centos/7/os/x86_64/`. Then click Done in the top left corner.

Figure A.11. CentOS installation source screen

INSTALLATION SOURCE

CENTOS LINUX 7 INSTALLATION

Done

Help!

Which installation source would you like to use?

On the network:

http:// mirror.centos.org/centos/7/os/x86_64/ Proxy setup...

This URL refers to a mirror list.

Additional repositories

Enabled	Name
---------	------

+ - ↻

Name:

http://

This URL refers to a mirror list.

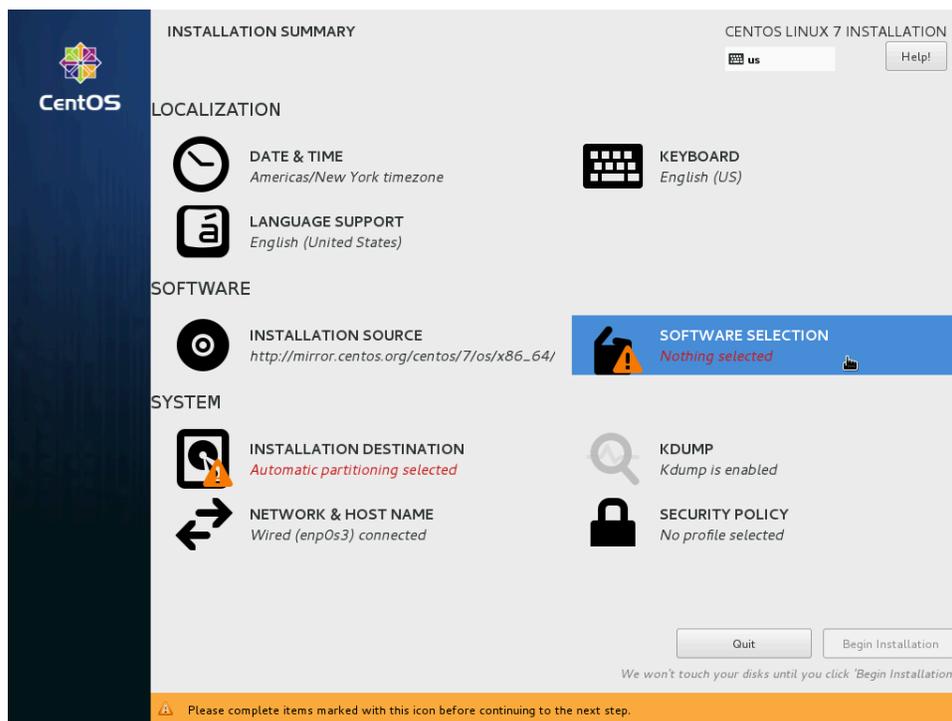
Proxy URL:

User name:

Password:

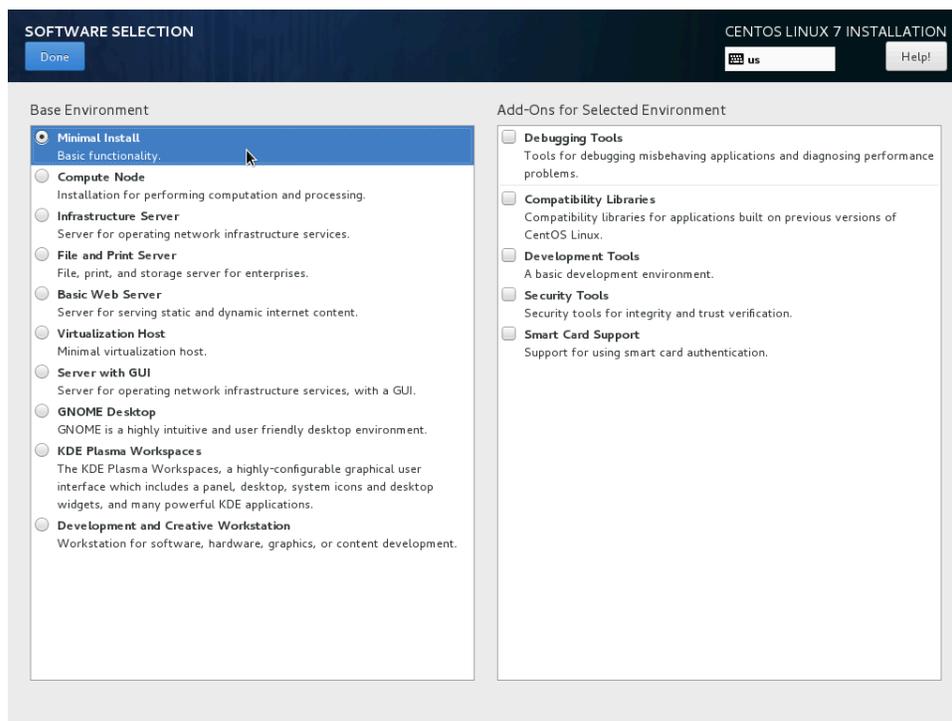
15. If networking is set up properly and the installation source URL is entered correctly, the installer will download package information from the network, and eventually the installation URL from above (`http://mirror.centos.org/centos/7/os/x86_64/`) will be displayed under the label Installation Source, and the warning icon will disappear.
 - If this process does not succeed, go back to the networking and installation source steps above to establish a network connection and point to the installation mirror.
 - If this process succeeds, click on Software Selection.

Figure A.12. CentOS installation software selection



16. Select an installation profile.

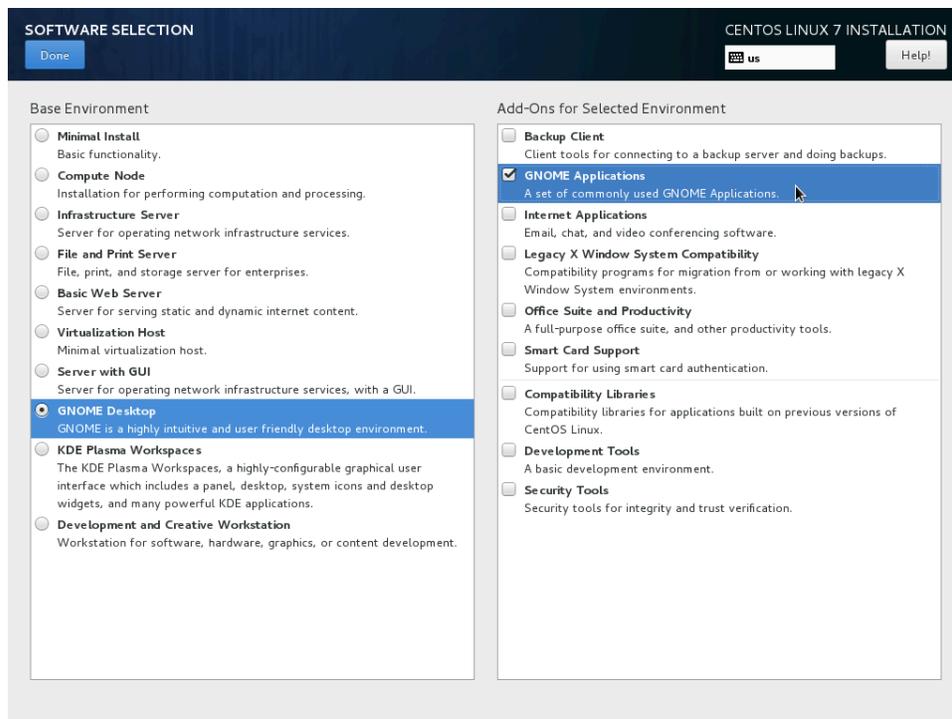
- If you do not need a graphical desktop environment, select Base Environment → Minimal Install.

Figure A.13. CentOS installation with minimal software

- If you need a graphical desktop environment, select Base Environment → GNOME Desktop and Add-Ons for Selected Environment → GNOME Applications (recommended), or Base

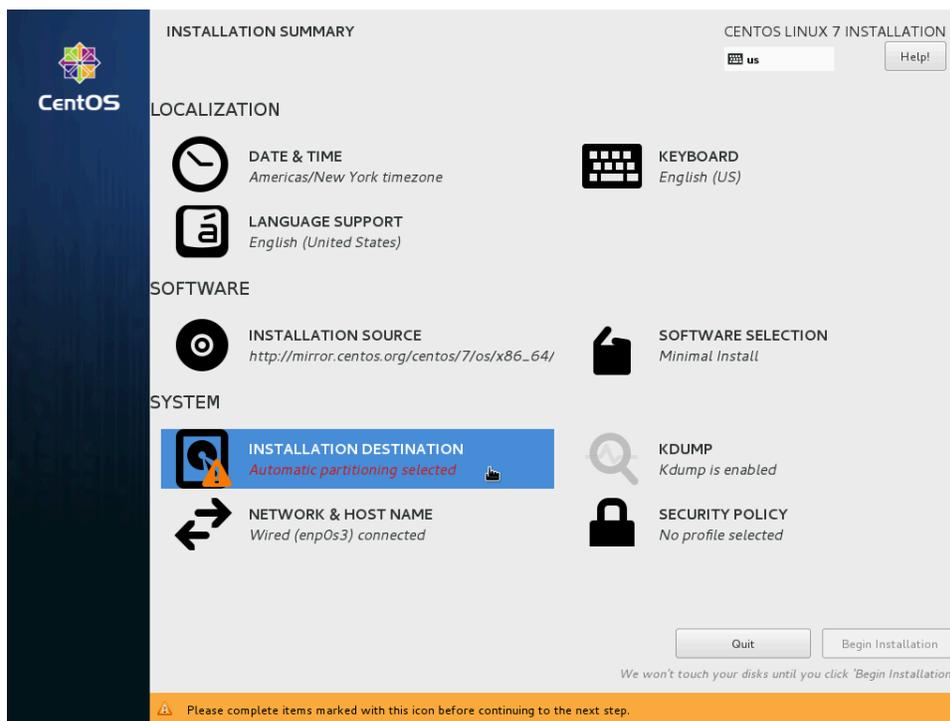
Environment → KDE Plasma Workspaces and Add-Ons for Selected Environment → KDE Applications.

Figure A.14. CentOS installation with graphical desktop



Then click Done in the top left corner.

17. Click on Installation Destination.

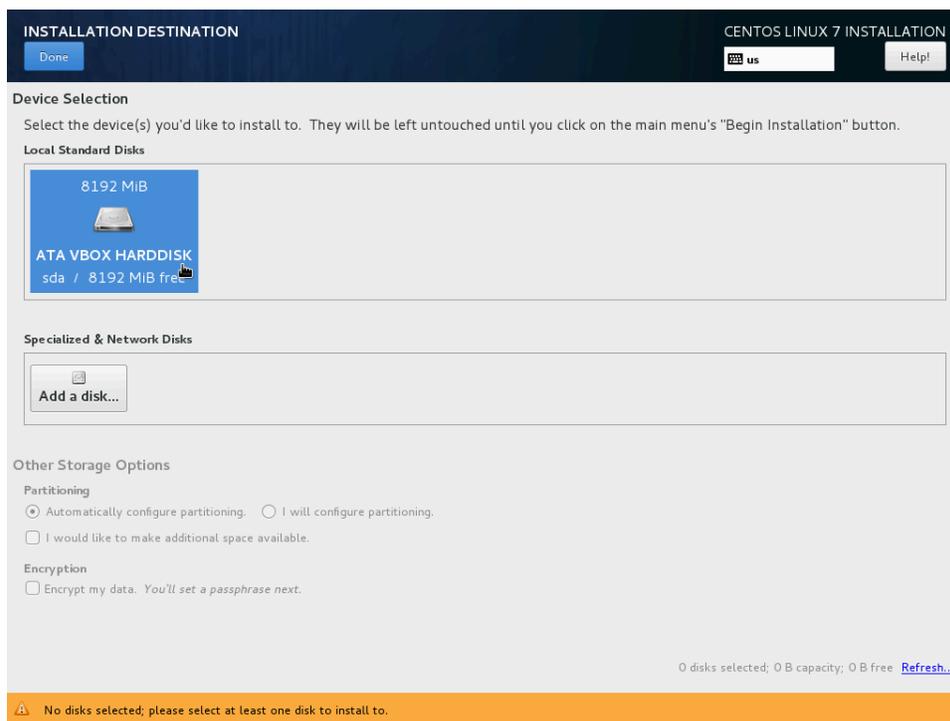
Figure A.15. CentOS installation destination

18. The Installation Destination screen is the starting point for setting up disk partitions and storage.

For a straightforward installation of a simple system onto a single disk:

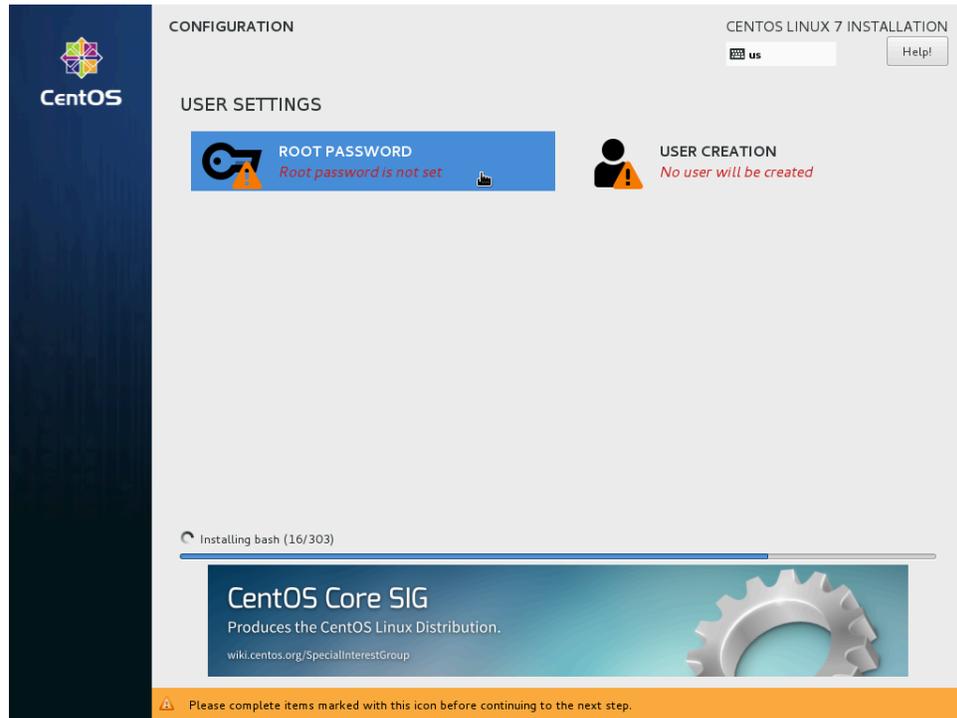
- Select the single disk in the Local Standard Disks section. The selected disk gets a check mark.
- In the Partitioning section, select Automatically configure partitioning.
- Click the Done button in the top-left corner.

If you are installing a more complex system, refer to the documented procedure for how to use this screen to customize partitioning, add network storage devices like iSCSI targets, encrypt, and more.

Figure A.16. CentOS installation disk partitioning and selection

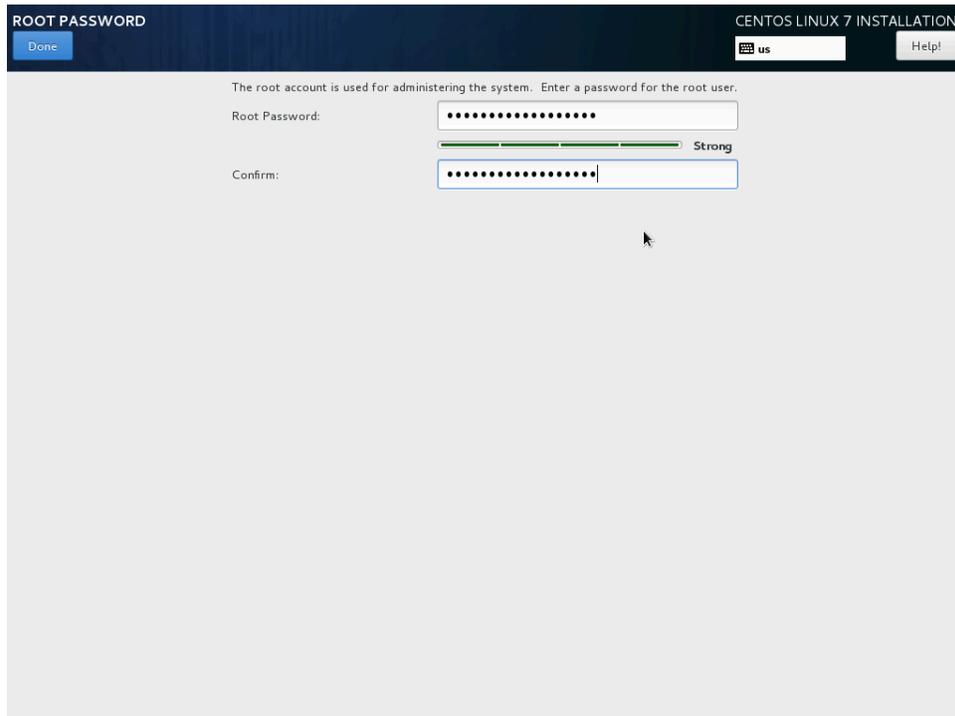
19. If desired, click on Date & Time, Keyboard, Language Support, Kdump or Security Policy, to review and change settings as needed.
20. Click on Begin Installation in the bottom right corner to begin the installation process.
21. While packages are downloading and installing, click on Root Password.

Figure A.17. CentOS installation root password



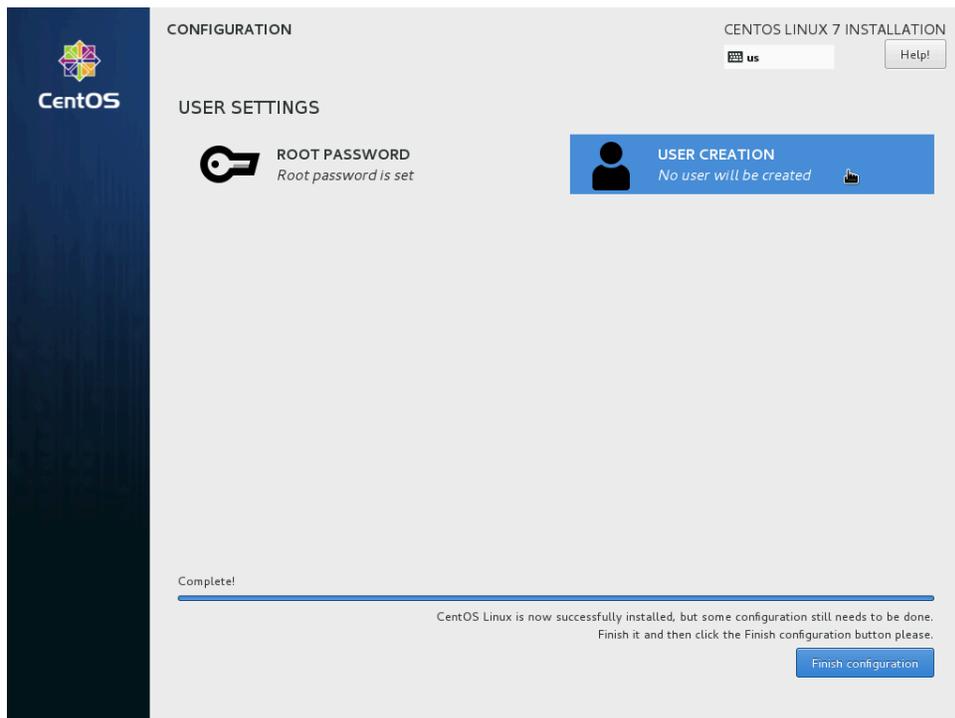
22. Enter and re-enter a root (administrator) password for the infrastructure server, then click Done in the top left corner.

Figure A.18. CentOS installation root password screen



23. To create a first user account right now (recommended), click on User Creation.

Figure A.19. CentOS installation user creation



24. Enter the user's full name and login name, then enter and re-enter a password.

Check the Make this user administrator box to add this user to the `wheel` user group, which in CentOS has the effect of granting privileges via **sudo**.



Note

The policies and procedures at your institution for managing users, granting administrator privileges, administering systems, etc. may vary. The remainder of this document assumes that users with administrator privileges cannot log in as root directly or only do so infrequently, and use **sudo** to issue commands as root instead.

Finally, click on Done in the top left corner.

Figure A.20. CentOS installation user creation screen

The screenshot shows the 'CREATE USER' screen during the CentOS Linux 7 installation. The interface includes a 'Done' button in the top left and a 'Help!' button in the top right. The main form fields are: 'Full name' (filled with 'Fake User'), 'User name' (filled with 'fakeuser'), a tip 'Tip: Keep your user name shorter than 32 characters and do not use spaces.', two checked checkboxes 'Make this user administrator' and 'Require a password to use this account', a 'Password' field with a strength indicator showing 'Strong', and a 'Confirm password' field. An 'Advanced...' button is located at the bottom of the form.

25. Wait for the installation process to finish. The progress bar will eventually say Complete! Click on Finish Configuration.
26. Wait for the last installation steps to complete. The label below the progress bar will eventually say CentOS Linux is now successfully installed and ready for you to use. Click on Reboot to reboot the server.
27. Before anything else, bring your freshly installed server completely up to date with the latest software and security patches.

In a terminal as the user with administrator privileges created above, enter the command **sudo yum update**. You will be prompted for your user password to authenticate with **sudo**, and since this is the first time this user uses **sudo**, the system will display a security warning.

Figure A.21. Initial CentOS update (1)

```
CentOS Linux 7 (Core)
Kernel 3.10.0-514.el7.x86_64 on an x86_64

Hint: Num Lock on

localhost login: fakeuser
Password:
lfakeuser@localhost ~1$ sudo yum update

We trust you have received the usual lecture from the local System
Administrator. It usually boils down to these three things:

    #1) Respect the privacy of others.
    #2) Think before you type.
    #3) With great power comes great responsibility.

[sudo] password for fakeuser: _
```

Accept the change set by entering **Y** and hitting **Enter**.

Figure A.22. Initial CentOS update (2)

```
openssl                x86_64      1:1.0.1e-60.el7_3.1    updates      713 k
openssl-libs           x86_64      1:1.0.1e-60.el7_3.1    updates      959 k
policycoreutils        x86_64      2.5-11.el7_3           updates      841 k
polkit                  x86_64      0.112-11.el7_3         updates      167 k
python-firewall        noarch     0.4.3.2-8.1.el7_3.2    updates      301 k
python-perf            x86_64      3.10.0-514.10.2.el7    updates      4.0 M
selinux-policy         noarch     3.13.1-102.el7_3.15   updates      414 k
selinux-policy-targeted noarch     3.13.1-102.el7_3.15   updates      6.4 M
sudo                    x86_64      1.8.6p7-21.el7_3       updates      735 k
systemd                 x86_64      219-30.el7_3.7         updates      5.2 M
systemd-libs           x86_64      219-30.el7_3.7         updates      369 k
systemd-sysv           x86_64      219-30.el7_3.7         updates      63 k
tuned                   noarch     2.7.1-3.el7_3.1        updates      210 k
tzdata                  noarch     2016j-1.el7            updates      442 k
vim-minimal             x86_64      2:7.4.160-1.el7_3.1    updates      436 k
wpa_supplicant         x86_64      1:2.0-21.el7_3         updates      788 k
xfsprogs                x86_64      4.5.0-9.el7_3          updates      895 k

Transaction Summary
-----
Install  1 Package
Upgrade  57 Packages

Total download size: 98 M
Is this ok [y/d/N]: y_
```

Because you are performing updates for the first time, you may be asked to confirm whether to download and install the CentOS software signing key. Accept by entering **Y** and hitting **Enter**.

Figure A.23. Initial CentOS update (3)

```

(44/58): polycoreutils-2.5-11.el7_3.x86_64.rpm | 841 kB 00:00
(45/58): python-perf-3.10.0-514.10.2.el7.x86_64.rpm | 4.0 MB 00:01
(46/58): selinux-policy-3.13.1-102.el7_3.15.noarch.rpm | 414 kB 00:00
(47/58): NetworkManager-libnm-1.4.0-17.el7_3.x86_64.rpm | 443 kB 00:10
(48/58): sudo-1.8.6p7-21.el7_3.x86_64.rpm | 735 kB 00:00
(49/58): kernel-3.10.0-514.10.2.el7.x86_64.rpm | 37 MB 00:09
(50/58): systemd-libs-219-30.el7_3.7.x86_64.rpm | 369 kB 00:00
(51/58): systemd-sysv-219-30.el7_3.7.x86_64.rpm | 63 kB 00:00
(52/58): tuned-2.7.1-3.el7_3.1.noarch.rpm | 210 kB 00:00
(53/58): tzdata-2016j-1.el7.noarch.rpm | 442 kB 00:00
(54/58): vim-minimal-7.4.160-1.el7_3.1.x86_64.rpm | 436 kB 00:00
(55/58): wpa_supplicant-2.0-21.el7_3.x86_64.rpm | 788 kB 00:00
(56/58): systemd-219-30.el7_3.7.x86_64.rpm | 5.2 MB 00:01
(57/58): xfsprogs-4.5.0-9.el7_3.x86_64.rpm | 895 kB 00:00
(58/58): selinux-policy-targeted-3.13.1-102.el7_3.15.noarc | 6.4 MB 00:05
-----
Total 5.9 MB/s | 98 MB 00:16
Retrieving key from file:///etc/pki/rpm-gpg/RPM-GPG-KEY-CentOS-7
Importing GPG key 0xF4A80EB5:
  Userid      : "CentOS-7 Key (CentOS 7 Official Signing Key) <security@centos.org>"
  Fingerprint: 6341 ab27 53d7 8a78 a7c2 7bb1 24c6 a8a7 f4a8 0eb5
  Package     : centos-release-7-3.1611.el7.centos.x86_64 (@anaconda)
  From        : /etc/pki/rpm-gpg/RPM-GPG-KEY-CentOS-7
Is this ok [y/N]: y_

```

When complete, reboot the server again using the `sudo reboot` command (or using your graphical desktop environment's logout/reboot function).

B. CentOS 7 Storage Array Configuration

Section last updated: 2018-09-09

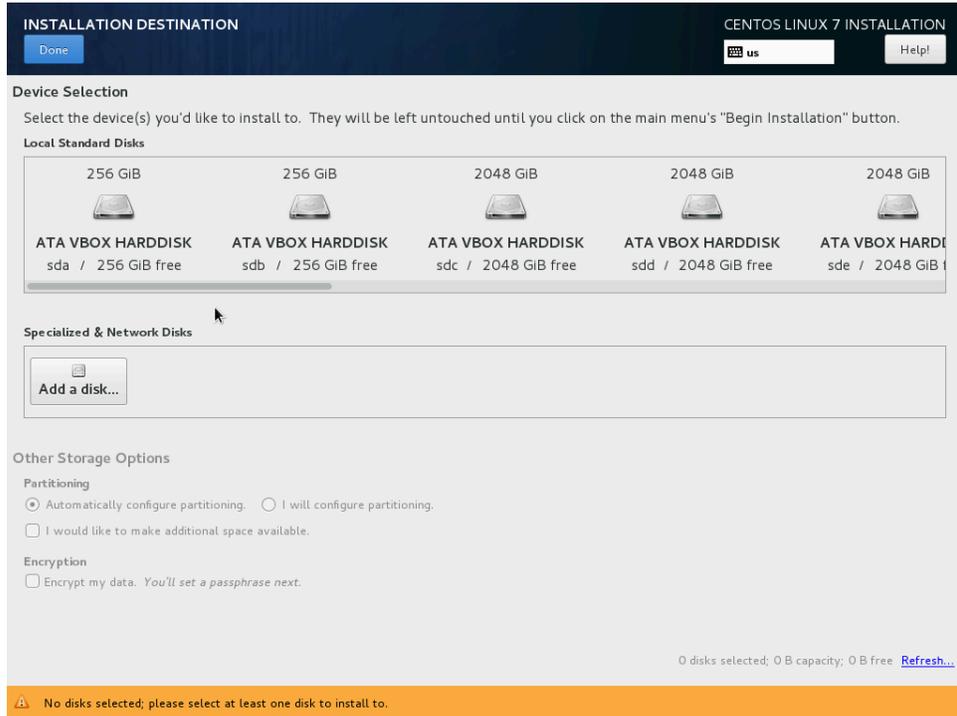
We present the following procedure to implement the designs of Section 3.3 for CentOS 7. (For other operating systems, refer to the usage manual for disk partitioning instructions.)

Unless otherwise specified, the machine used as an example has two 256GB solid-state disks as the first and second disks (labeled `sda` and `sdb`), and twelve 2TB disks (labeled `sdc` through `sdl`), intended for a 2-disk dedicated system array and two 6-disk RAID6 storage arrays respectively; and further assume the machine has 4GB of memory, meaning we intend to devote 8GB to the swap partition.

Procedure B.1. Storage array configuration (Procedure A.1 Step 18)

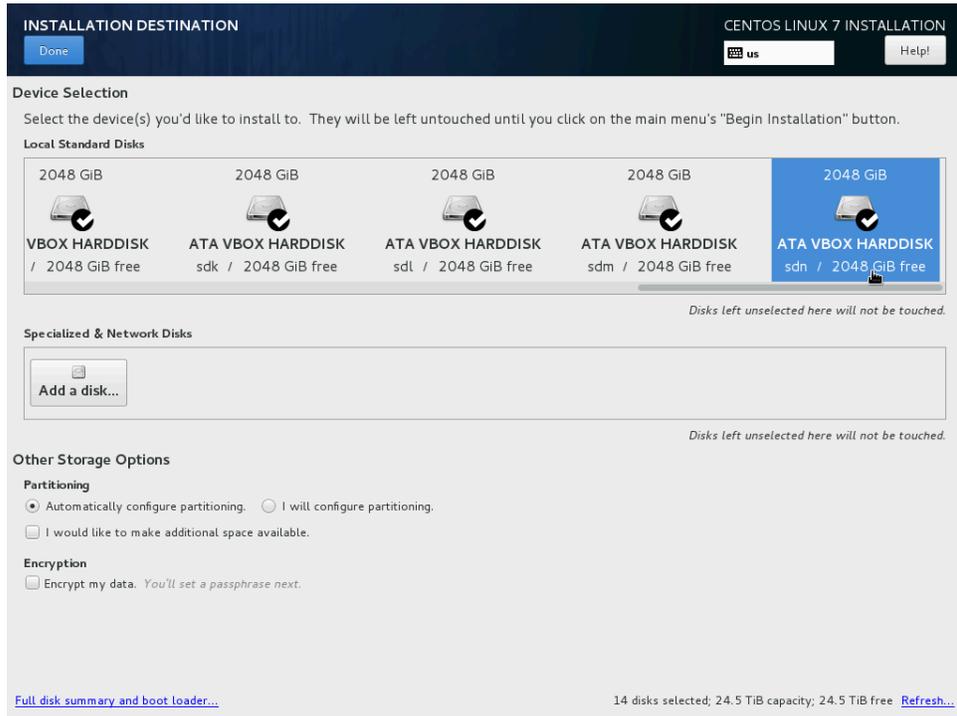
1. In this step, you will designate all the disks in the machines for manual partitioning:
 - a. In the Installation Destination screen (Procedure A.1 Step 18), you should see the machine's disks listed. See Figure B.1.

Figure B.1. Installation Destination screen



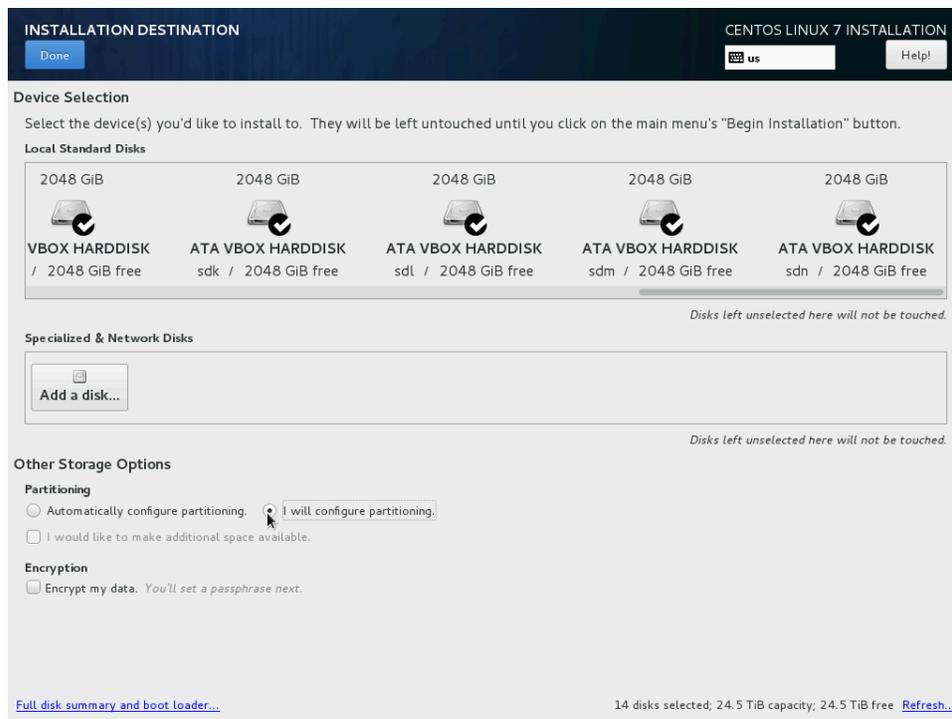
- b. Select the machine's disks one by one until all have a check mark. See Figure B.2.

Figure B.2. Selecting all the disks



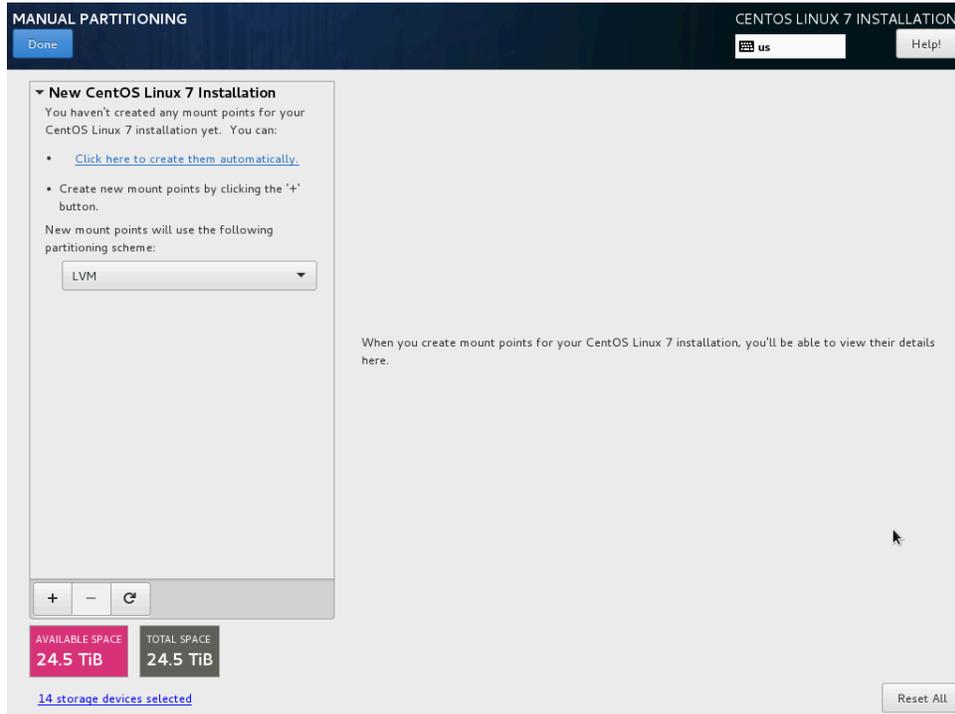
- c. Under Other Storage Options, in the Partitioning category, select I will configure partitioning. See Figure B.3.

Figure B.3. Selecting manual partitioning



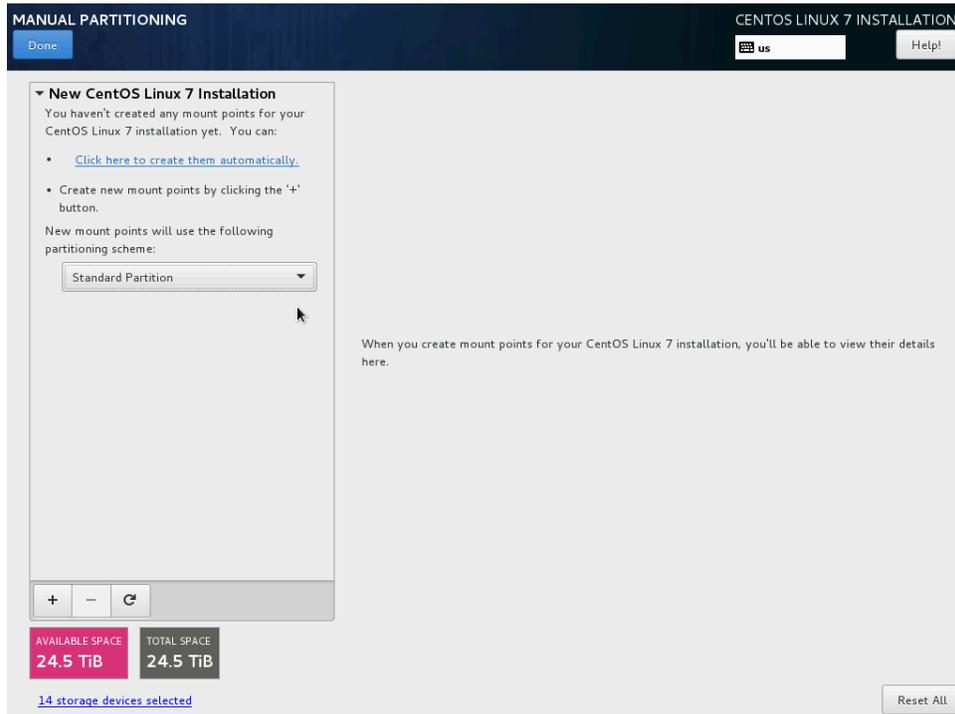
- d. Click on the Done button in the top-left corner of the Installation Destination screen.
2. After clicking on the Done button in the top-left corner of the Installation Destination screen, you will be taken to the Manual Partitioning screen. See Figure B.4.

Figure B.4. Manual partitioning screen



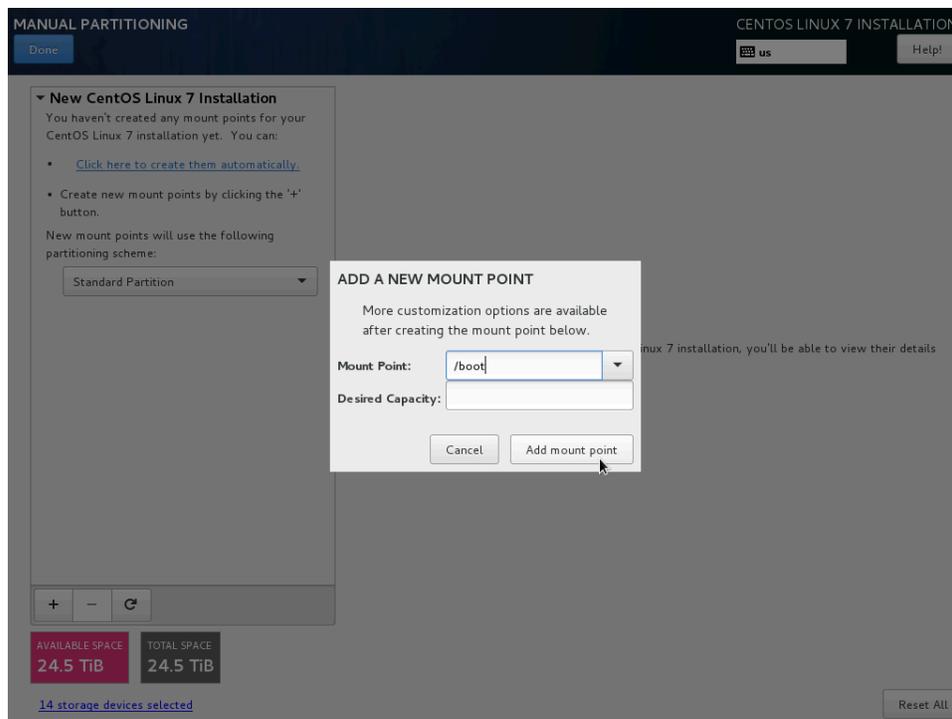
In the left pane, select Standard partition. See Figure B.5.

Figure B.5. Manual partitioning screen with standard partitions



3. In this step, you will create a partition. The first partition you will create is the boot partition. (This step will then be repeated for other partitions.)
 - a. At the bottom of the left pane, click the + sign to create a new partition.
 - b. In the Mount Point text field, type the mount point. For the boot partition, the value is `/boot`.
 - c. Leave the Desired Capacity field blank. The result so far is shown in Figure B.6.

Figure B.6. Creating the boot partition



- d. Click the Add mount point button.
4. In this step, you will set the device type and file system type of the newly created partition, and its size. (Again here, the first partition you will configure is the boot partition, then this step will be repeated for other partitions.)
 - a. In the Desired Capacity field, enter the partition size. In the case of the boot partition, the value is **512 MiB**.



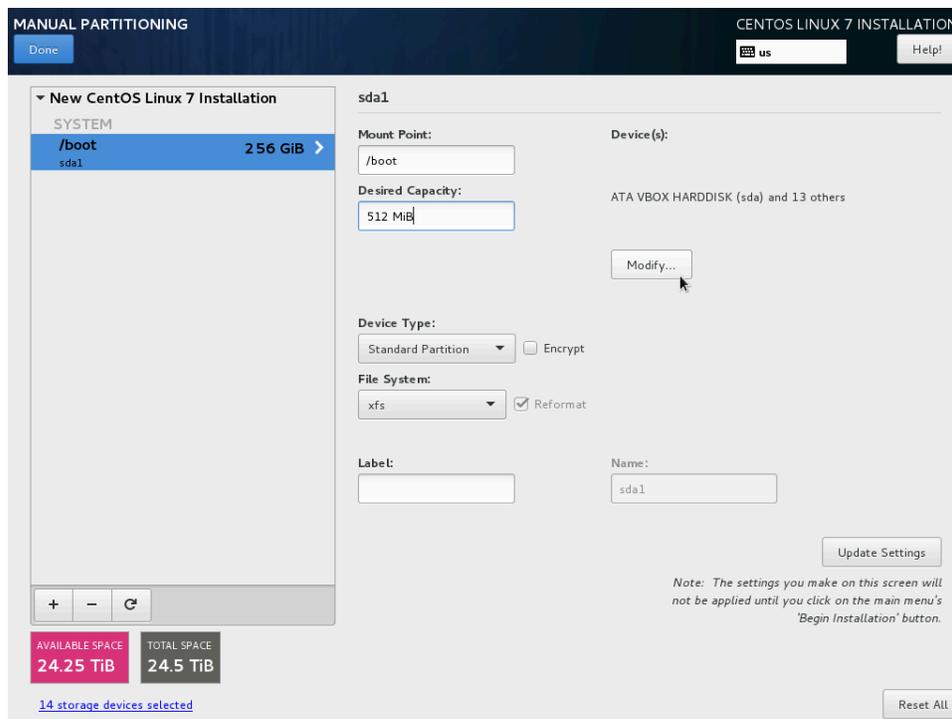
Important

Note the "i" in "MiB". Under some definitions, one megabyte is 1,024 kilobytes, and under others, 1,000 kilobytes. CentOS uses the notation with an extra "i", which makes it explicit that the non-decimal definition of 1,024 is meant. In this step and all similar steps, use "MiB" for megabyte, "GiB" for gigabyte, and "TiB" for terabyte -- not the traditional but ambiguous "MB", "GB" and "TB".

- b. In the Device Type drop down box, select the correct device type for the partition, which is often Standard Partition or RAID, or the special type Swap for the swap partition. In the case of the boot partition, the device type is Standard Partition.

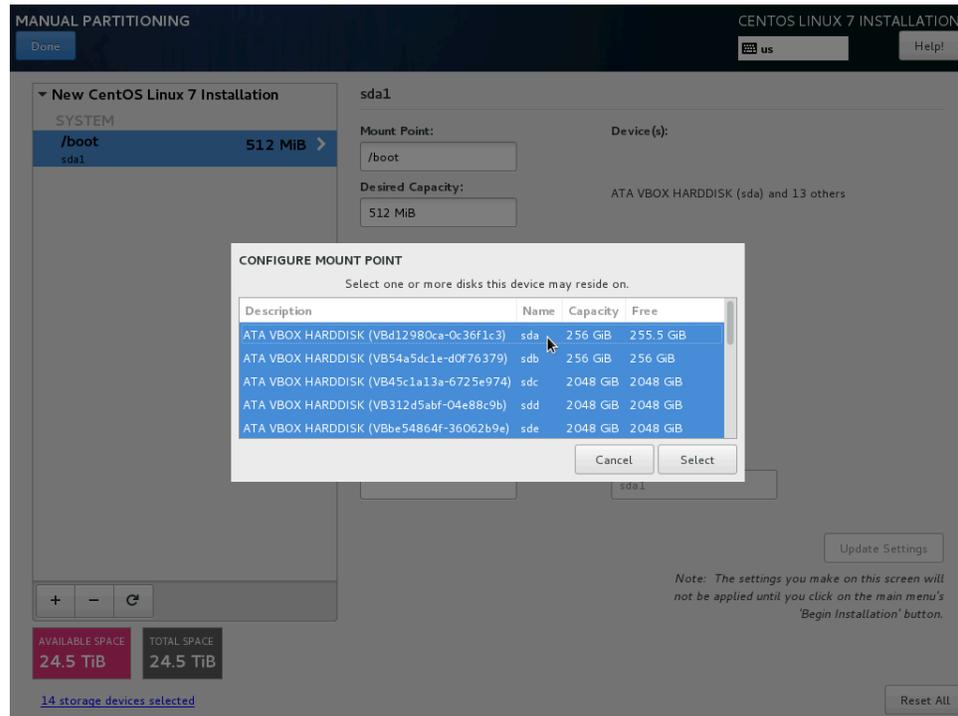
- c. (Optional) If the device type is RAID, a new drop down box labeled RAID Level will appear so that you can further select the RAID scheme for a RAID partition. (This does not apply to the boot partition.)
- d. In the File System drop down box, select the desired file system type for the partition, which is often xfs, or the special types EFI System Partition for the EFI system partition, or swap for the swap partition. In the case of the boot partition, the correct value is xfs. The result so far is illustrated in Figure B.7.

Figure B.7. Modifying the boot partition devices

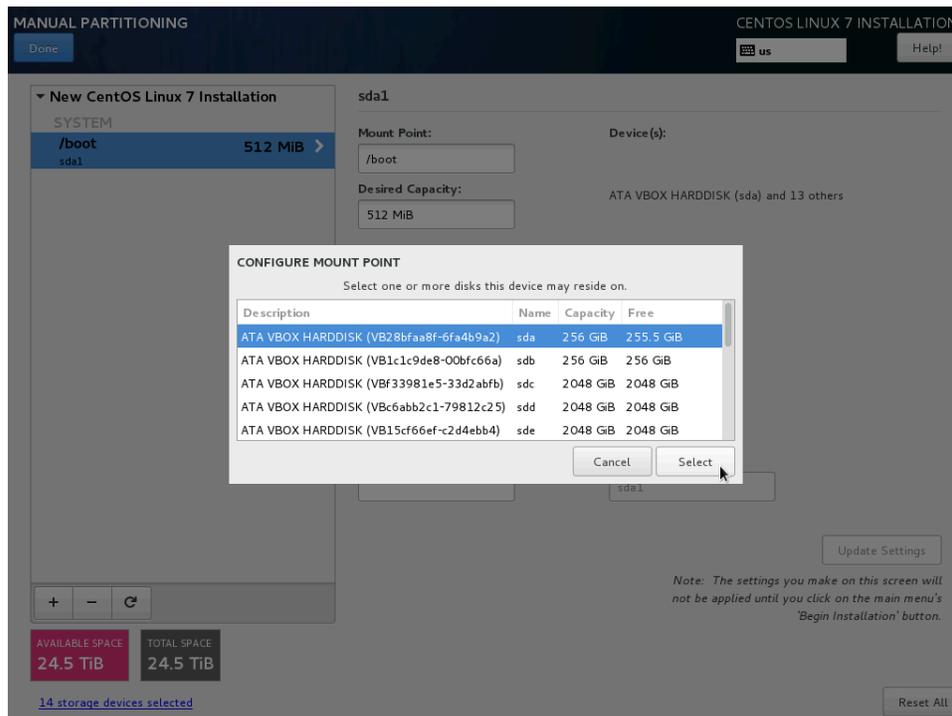


- e. In the top-right Device(s) category, click the Modify... button.
5. In this step, you will select the disks the partition applies to. (The first partition you will do this for is the boot partition, then later this step will be repeated for other partitions.)
 - a. After you click the Modify... button, the Configure Mount Point dialog to select the partition's devices will pop up. By default, every disk with remaining free space is selected. See Figure B.8.

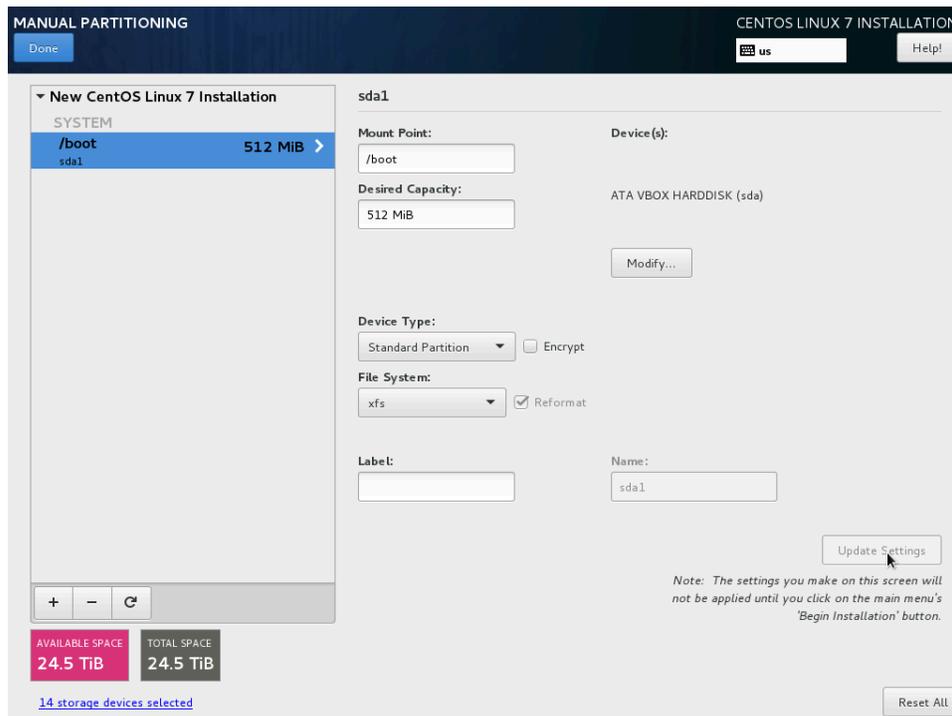
Figure B.8. Partition device selection dialog



- b. Select only the disks corresponding to the partition. In the case of the boot partition, this is only the first disk in the machine, sda. See Figure B.9.

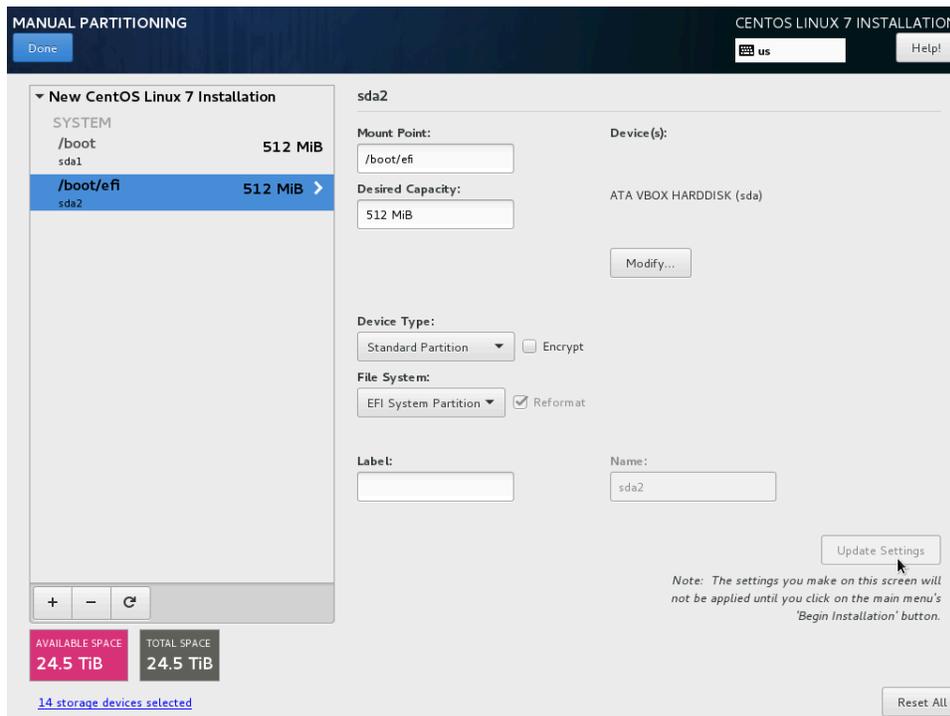
Figure B.9. Selecting the boot partition devices

- c. Click the Select button to close the Configure Mount Point dialog.
- d. In the bottom-right corner, click the Update Settings button and allow the system to finalize the characteristics of the partition. The result for the boot partition is shown in Figure B.10.

Figure B.10. Finalized boot partition

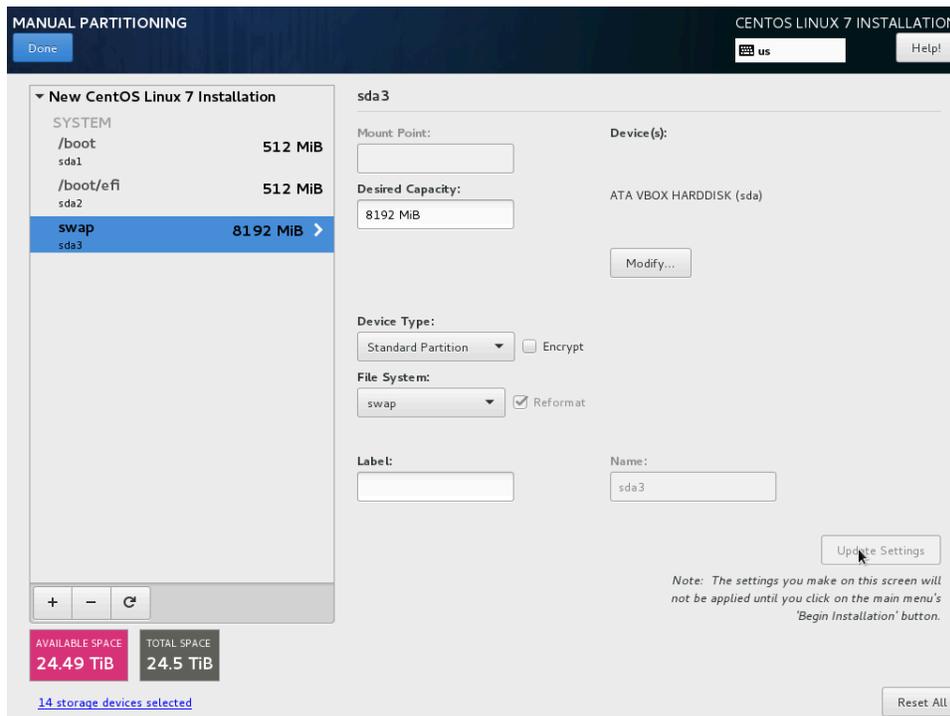
6. Repeat Step 3, Step 4 and Step 5 for the EFI system partition:
 - In Step 3, enter `/boot/efi` in the Mount Point text field.
 - In Step 4, enter the desired partition size (in this example **512 MiB**), select **Standard Partition** in the Device Type drop down box, and select **EFI System Partition** in the File System drop down box.
 - In Step 5, select only the first disk (in this example `sda`).

The result of configuring the EFI system partition is illustrated in Figure B.11.

Figure B.11. Configuring the EFI system partition

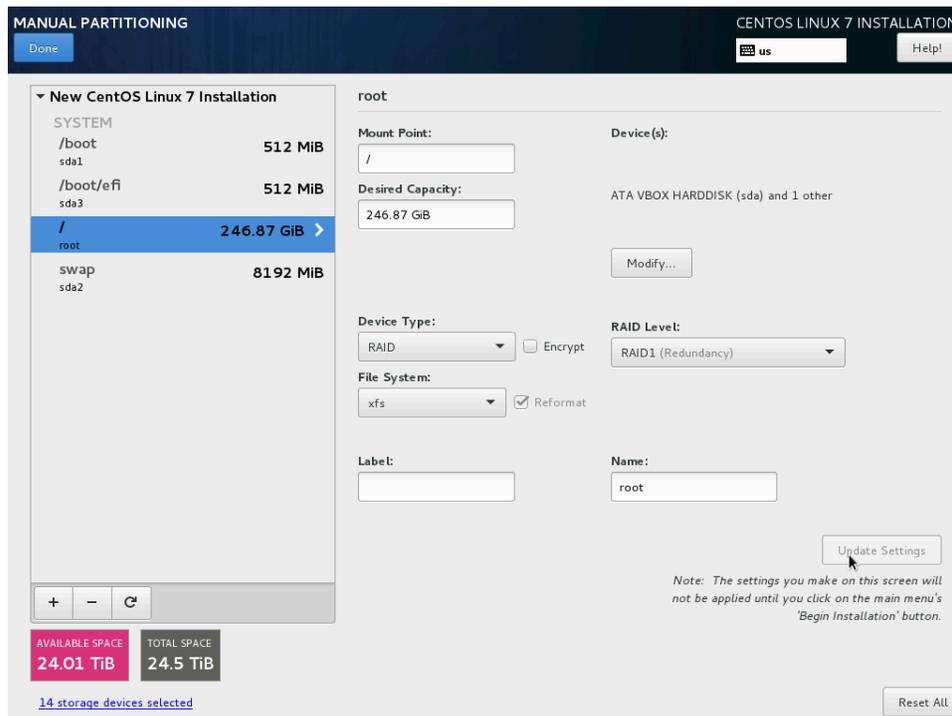
7. Repeat Step 3, Step 4 and Step 5 for the swap partition:
 - In Step 3, enter `swap` in the Mount Point text field.
 - In Step 4, enter the desired swap capacity (recommended: twice as much as there is memory, which in our example means `8 GiB`), select `Standard Partition` in the Device Type drop down box, and select `swap` in the File System drop down box.
 - In Step 5, select only the first disk (in this example `sda`).

The result of configuring the swap partition is illustrated in Figure B.12.

Figure B.12. Configuring the swap partition

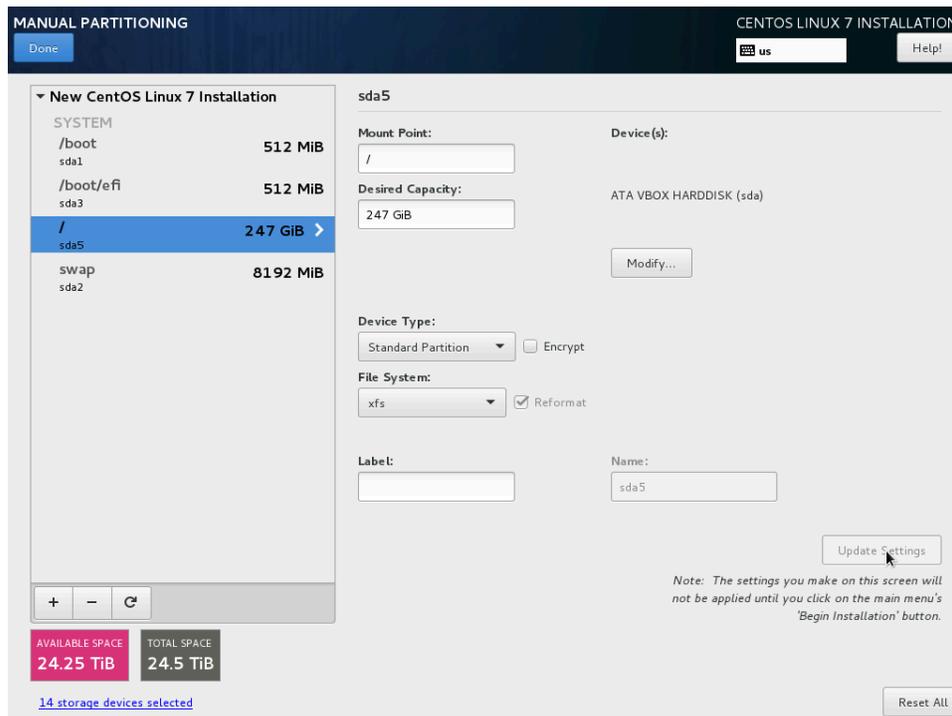
8. Repeat Step 3, Step 4 and Step 5 for the root partition, according to one of these three alternatives:
 - If you are using a dedicated system array or a shared system array:
 - In Step 3, enter / in the Mount Point text field.
 - In Step 4, enter the special keyword **max** into the Desired Capacity field, select RAID in the Device Type drop down box, select RAID1 (Redundancy) in the RAID Level drop down box, and select xfs in the File System drop down box.
 - In Step 5, select the two disks in the system array (in our example the first two disks in the machine, `sda` and `sdb`).

The result of this alternative is shown in Figure B.13.

Figure B.13. Configuring the root partition (dedicated system array)

- If you are using a dedicated system disk:
 - In Step 3, enter / in the Mount Point text field.
 - In Step 4, enter the special keyword **max** into the Desired Capacity field, select Standard Partition in the Device Type drop down box, and select xfs in the File System drop down box.
 - In Step 5, select the first disk in the machine, sda.

The result of this alternative is shown in Figure B.14.

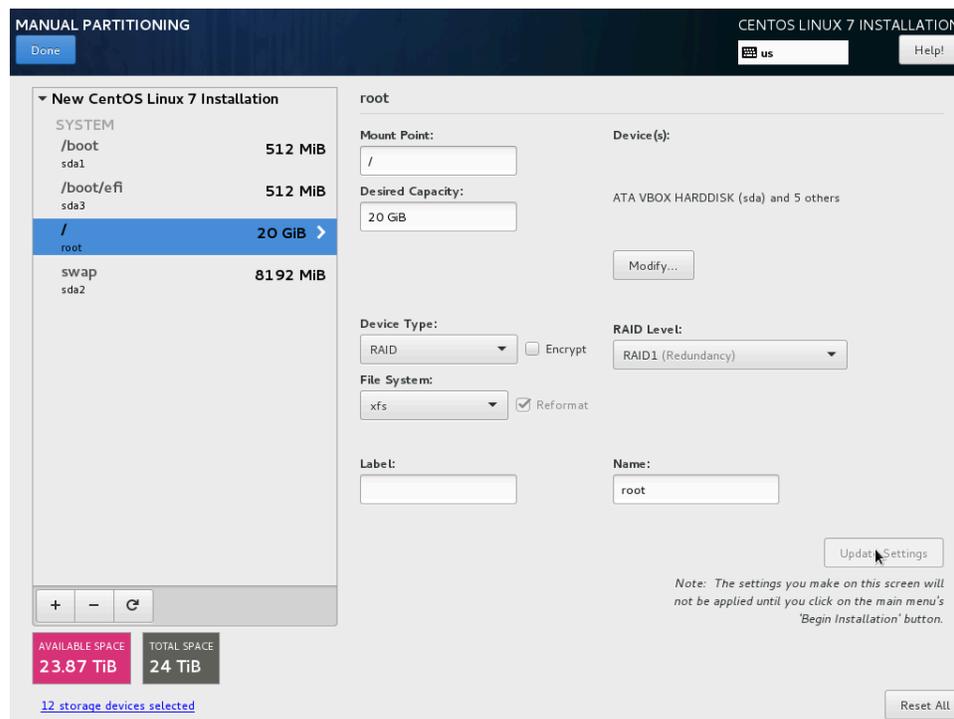
Figure B.14. Configuring the root partition (dedicated system disk)

- If you are using a shared system array:
 - In Step 3, enter / in the Mount Point text field.
 - In Step 4, enter the intended usable size of the root partition into the Desired Capacity field, select RAID in the Device Type drop down box, select RAID1 (Redundancy) in the RAID Level drop down box, and select xfs in the File System drop down box. (In Figure B.15, the usable size of 20 GiB is used as an example. In RAID1, it means that 20GB are used on each of the disks in the array, with the root partition appearing to the user as having 20GB of usable space.)
 - In Step 5, select all the disks in the shared system array, that is, all the disks in the first storage array. (In all the examples thus far, we were using a hypothetical machine with two small disks for a dedicated system array and twelve large disks for two 6-disk storage arrays. In Figure B.15, the depiction is of a machine with only twelve large disks named sda through

sd1 for two 6-disk arrays. The disks selected in this case would then be the first six, sda through sdf.)

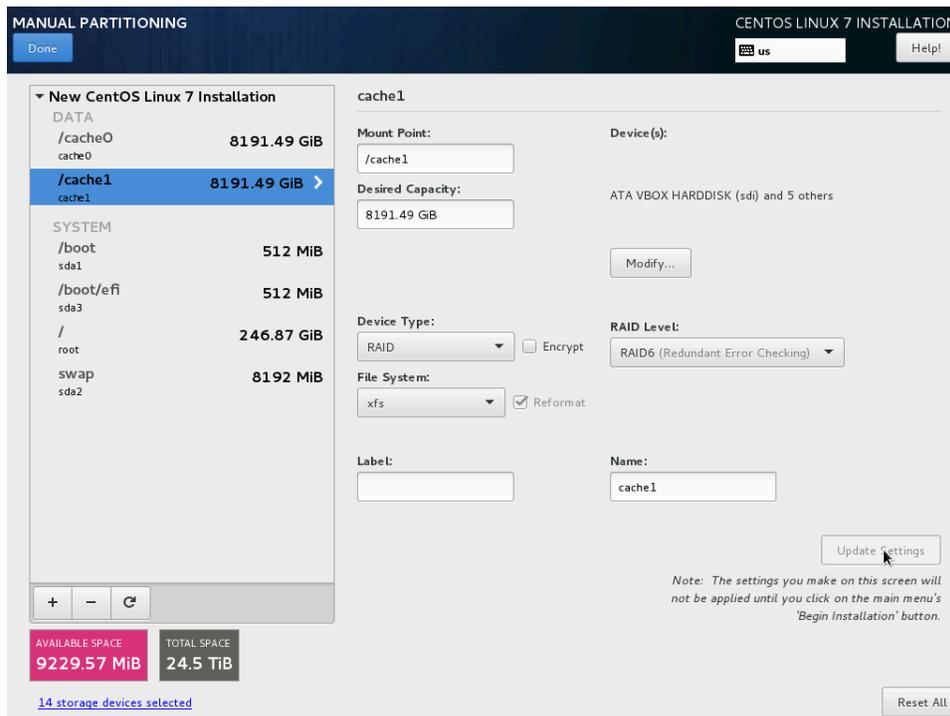
The result of this alternative is shown in Figure B.15.

Figure B.15. Configuring the root partition (shared system array)



9. The next step is to configure the storage arrays. Repeat Step 3, Step 4 and Step 5 for **each** storage array:
 - In Step 3, enter the mount point for the storage array in the Mount Point text field: `/cache0` for the first storage array, `/cache1` for the second, `/cache2` for the third, etc.
 - In Step 4, enter the special keyword `max` into the Desired Capacity field, select RAID in the Device Type drop down box, select RAID6 (Redundant Error Checking) in the RAID Level drop down box, and select xfs in the File System drop down box.
 - In Step 5, select only the first disk (in this example `sda`).

Figure B.16 shows the result of creating storage arrays on our hypothetical machine. The first storage array (mount point `/cache0`) spans the six disks `sdc` through `sdh`, and the second storage array (mount point `/cache1`) spans the six disks `sd1` through `sdn`.

Figure B.16. Configuring the storage arrays

10. In the top-left corner of the Manual Partitioning screen, click the Done button.
11. A Summary of Changes dialog appears (see Figure B.17), allowing you to review the low-level disk partitioning and formatting operations that will take place once the CentOS installation proceeds. Click on the Accept Changes button to return to the main CentOS installation screen.

Figure B.17. Confirming disk partitioning and formatting actions

