



Perform the prerequisite tasks for installation

Element Software

NetApp
June 11, 2021

Table of Contents

- Perform the prerequisite tasks for installation 1
- Find more information 6
- Considerations for network configuration 6
- Configure the host networking 7

Perform the prerequisite tasks for installation

Ensure that you perform the necessary checks and verify that your environment meets the configuration, IP addressing, and networking requirements before you install SolidFire eSDS.

Install the required hardware

- Install the supported server. See [NetApp Interoperability Matrix \(login required\)](#) for more information.
- Ensure that your hardware configuration is balanced and all channels are populated. For more information about maximizing bandwidth, see the [KB article](#) (login required).

Configure the host (node)

- Install RHEL based on the supported versions listed in the [NetApp Interoperability Matrix \(login required\)](#).
- Configure a Network Time Protocol (NTP) server to use with all the hosts in your network.
- When selecting the installation destination, select the radio button to manually configure file system partitioning. On the **Manual Partitioning** page, use the + and - buttons to remove any existing partitions, and create new partitions and size them following the recommendations listed here. Using the default LVM partitioning scheme enables you to easily resize later, if needed.



By default, RHEL chooses `xf`s as the default file system for the partitions that you manually create. You should change it to `ext4`, except for the `/boot` and `swap` partitions. Your `/boot` partition should use `ext2`.

Partition	Size
/boot	1GB
/opt	50GB
/var	50GB
swap	4GB
/home	5GB
/ and /usr	Split remaining space

The minimum required partition layout is as follows:

Partition	Size
/opt	40GB
/var	40GB
/	10GB (RedHat recommendation)



The `/dev/sdb` disk is not used by any process.

- Disable RAID for `/boot`.
- On the Software Selection screen, where you select specific packages to install, select **Server** or **Infrastructure Server** based on your RHEL version.
- After the first boot, do the following:
 - Install Red Hat Subscription Manager, and enable the following repositories:

```
rhel-7-server-ansible-2.9-rpms
rhel-7-server-optional-rpms
rhel-7-server-extras-rpms
```

- Enable SSH on your nodes.
- If you want to disable IPv6, follow the steps detailed in this [KB article \(login required\)](#).

Install the required software

- Install Ansible, Git, and Python 3.0.

Verify that your configuration matches NetApp's requirements for installing SolidFire eSDS

- Use the SolidFire eSDS configuration listed in the [NetApp Interoperability Matrix Tool \(IMT\)](#) as a reference.



If you contact NetApp Support for assistance with issues relating to SolidFire eSDS, Support will first verify that your platform complies with the reference configuration for SolidFire eSDS listed in the IMT. If Support determines that your underlying platform does not comply with the reference configuration, Support will guide you in aligning the non-compliant firmware, software, and/or hardware components with the correct versions in the IMT.

- Run a compliance check for SolidFire eSDS.
 - i. Run the `ansible-galaxy install` command to install the `nar_solidfire_sds_compliance` role.

```
ansible-galaxy install git+https://github.com/NetApp-
Automation/nar_solidfire_sds_compliance.git
```

You can also manually install the role by copying it from the [NetApp GitHub repository](#) and placing the role in the `~/.ansible/roles` directory. NetApp provides a README file, which includes information about how to run a role.



Ensure that you always download the latest versions of the roles.

- ii. Move the roles that you downloaded up one directory from where they were installed.

```
$ mv ~/.ansible/roles/ansible/nar_solidfire_sds_* ~/.ansible/roles/
```

- iii. Run the `ansible-galaxy role list` command to ensure that Ansible is configured to utilize the new roles.

```
$ ansible-galaxy role list
# ~/.ansible/roles
- nar_solidfire_sds_install, (unknown version)
- nar_solidfire_sds_upgrade, (unknown version)
- ansible, (unknown version)
- nar_solidfire_sds_compliance, (unknown version)
```

- iv. Create the playbook to use for the compliance check.
- v. Run the compliance check playbook as shown in the following example:

```
$ ansible-playbook -i yourinventory.yml yourplaybook.yml
```




Even after you start using the SolidFire eSDS system, you should run the compliance check regularly to ensure that your system is in compliance. In some cases, NetApp Support will ask you to run the compliance check to help diagnose and troubleshoot issues.

Understand network and IP address requirements

- Familiarize yourself with how to configure and manage networks and network interfaces in RHEL. See the [RedHat documentation](#).
- Configure your network following the IP requirements detailed here:

Component	Storage network IP address	Management network IP address	Total # of IP addresses
Storage node	1	1	2 per node
Management node	(Optional) 1	1	1 per cluster on the storage network + 1 per cluster on the management network + 1 FQDN per cluster for the management node
Storage cluster	1 storage IP (SVIP)	1 management IP (MVIP)	2 per storage cluster

- Configure the storage network on 25GbE Ethernet switches and the management network on 10GbE switches. See the following cabling illustration:



Item	Description
1	Ports for storage network
2	Port for IPMI
3	Ports for management network



The illustration given here is intended to be an example. Your actual hardware might be different based on the server you have.

- Change the switch port MTU to 9216 bytes.

Allow specific ports through your datacenter's firewall

- If `firewalld` is enabled on the storage node running RHEL, ensure that you have the following ports open, so that you can manage the system remotely, allow clients outside of your datacenter to connect to resources, and ensure that internal services can function properly:

Source	Destination	Port	Description
Storage node MIP	Management node	80 TCP/UDP	Cluster upgrades
SNMP server	Storage node MIP	161 UDP	SNMP polling
System administrator PC	Management node	442 TCP	HTTPS UI access to management node
System administrator PC	Storage node MIP	442 TCP	HTTPS UI access to storage node
iSCSI clients	Storage cluster MVIP	443 TCP	(Optional) UI and API access
Management node	monitoring.solidfire.com	443 TCP	Storage cluster reporting to Active IQ

Source	Destination	Port	Description
Storage node MIP	Remote storage cluster MVIP	443 TCP	Remote replication cluster pairing communication
Storage node MIP	Remote storage node MIP	443 TCP	Remote replication cluster pairing communication
SolidFire eSDSsfapp	Per-node UI and API access to create a cluster	2010 UDP	Cluster beacon (to discover nodes to add to a cluster)
iSCSI clients	Storage cluster SVIP	3260 TCP	Client iSCSI communications
iSCSI clients	Storage cluster SIP	3260 TCP	Client iSCSI communications
SOAP server	SolidFire eSDSsfapp	7627 TCP	SOAP web services
System administrator PC	N/A	8080 TCP	System administrator communications
vCenter Server	Management node	8443 TCP	vCenter Plug-in QoSSIOC service



Ports 2181, 2182, and 2183 are needed for are needed for the Element distributed database, and will be dynamically opened from the Element container when you install SolidFire eSDS.

- Use the following commands to open the ports mentioned above:

```
systemctl start firewalld
firewall-cmd --permanent --add-service=snmp
firewall-cmd --permanent --add-port=80/tcp
firewall-cmd --permanent --add-port=80/udp
firewall-cmd --permanent --add-port=442-443/tcp
firewall-cmd --permanent --add-port=442-443/udp
firewall-cmd --permanent --add-port=2010/udp
firewall-cmd --permanent --add-source-port=2010/udp
firewall-cmd --permanent --add-port=3260/tcp
firewall-cmd --permanent --add-port=7627/tcp
firewall-cmd --permanent --add-port=8080/tcp
firewall-cmd --permanent --add-port=8443/tcp
firewall-cmd --reload
```

Configure your host network

- Configure your host network using the [best practices](#) provided.



You should complete the steps to configure your host network to ensure a successful installation of SolidFire eSDS.

Complete additional requirements

- Install One Collect, which will be used by NetApp Support for host log collection. You can install One Collect from [here](#). You need a NetApp account to access the download. You can also find the One Collect Installation Guide and Release Notes at the same location.



You must download and install One Collect in order to receive an optimal support experience.

- Install the management node for log collection and to enable NetApp Support access for troubleshooting. For information about management node and installation steps, see [here](#).

Find more information

- [NetApp SolidFire Resources Page](#)
- [Documentation for earlier versions of NetApp SolidFire and Element products](#)

Considerations for network configuration

Before you install SolidFire eSDS, you should set up the required networks on the storage nodes running RHEL. You are responsible for network routing in your environment. You can use the best practices provided as a framework.



- Use bonded or teamed interfaces.
- Use the same interface names for all the nodes in the cluster (example: team-mgt for each node's management interface and team-stg for each node's storage interface).
- Ensure that NetworkManager is running.
- Ensure that the NetworkManager-dispatcher-routing-rules package is installed on all storage nodes for policy-based routing.
- See [Best Practices for Networking and Network Maintenance on NetApp SolidFire Storage Systems](#).
- Configure the management and storage networks on each node to use multiple, redundant physical interfaces via bond or team configurations.

For information about network teaming, see [Configure network teaming](#). By default, all storage node 10GbE interfaces are enabled with a maximum transmission unit (MTU) of 9000 bytes. For optimal performance, configure all server-side storage interfaces with the same MTU as the NetApp SolidFire storage nodes. You should configure network switches to support an MTU of at least 9016 bytes or more to account for jumbo frame overhead and for proper forwarding through the network. If you want to change this configuration to support a lower MTU setting, you should contact NetApp Support.

See the following table for information about the storage and management networks that SolidFire eSDS requires for the different types of traffic:

Type of network	Description
Storage network	<ul style="list-style-type: none">• Includes all storage/iSCSI traffic.• Can be routed if you want to mount from hosts located on a different layer 3 network or if you plan to replicate data between clusters.• Should be configured with network interfaces on the same layer 2 broadcast domain.
Management network	<ul style="list-style-type: none">• Includes all management traffic.• Can be routed if you want to access the cluster API or UI from a different layer 3 network.• Should be configured with network interfaces on the same layer 2 broadcast domain.



For examples and tips to configure the host network, see [here](#).

Find more information

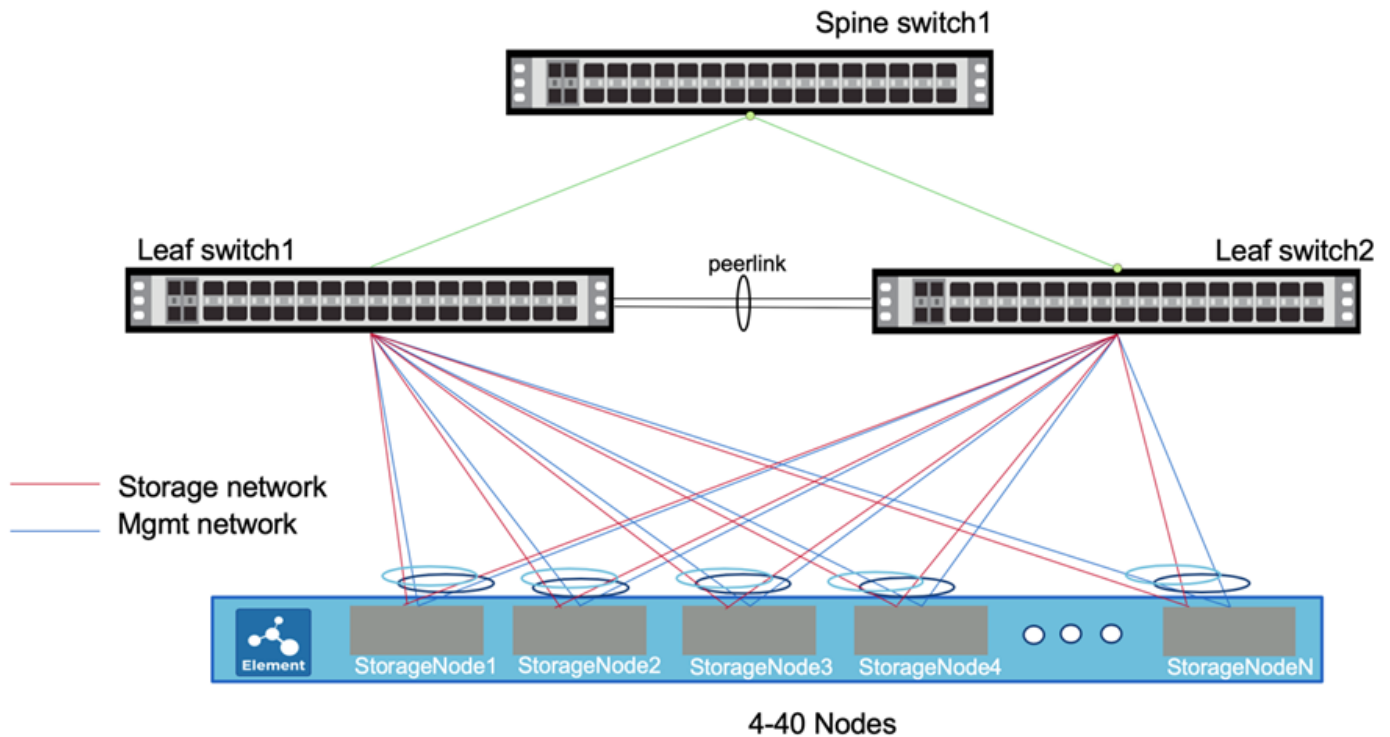
- [NetApp SolidFire Resources Page](#)
- [Documentation for earlier versions of NetApp SolidFire and Element products](#)

Configure the host networking

Use the examples and tips provided to configure the host networking before you install

SolidFire eSDS.

Here is a sample network configuration:



In this example, two interfaces on the storage node are network [teamed](#) and connected to the management network for redundancy purposes. Likewise, two additional interfaces are teamed and connected to the storage network.



Each interface has a configuration file named `ifcfg-<interface-name>X`, where X is the number of the interface, starting with zero or 1 depending upon the naming convention used. The configuration files are created when you first create the interfaces. One configuration file should already exist for each of the two physical interfaces connected to the storage network. One configuration file should also exist for each of the two physical interfaces connected to the management network. The interface configuration files are placed in the `/etc/sysconfig/network-scripts` directory. See [Interface configuration files](#).



The examples provided here have the storage and management interface names for HPE servers. If you have a Dell server, the interface names are different. Storage interface names for the Dell server are `em1` and `em2`. Management interface names for the Dell server are `p3p1` and `p3p2`.

Steps

1. Install `NetworkManager-dispatcher-routing-rules` package, and ensure that the appropriate repositories are configured.
2. Configure your network switch using your switch vendor's documentation. For specific instructions about configuring Multi-chassis Link Aggregation Group (MLAG) protocol and Link Aggregation Control Protocol (LACP), see your switch vendor's documentation.



It is recommended that you configure LACP fallback and disable LACP individual port suspension behavior by running `no lacp suspend-individual`. This enables the AccessPoint link to come up even without LACP packets being broadcast in case of misconfiguration.

3. Edit the two configuration files for the physical interfaces connected to the storage network using the following example. Jumbo frame setting on the storage network is highly recommended, but not required. In this example below, the storage interface name is `ens2f0` and the storage team name is `team10G`:



In all the example configurations listed here, `NAME` and `DEVICE` use the same values. You can use different values if you choose to do so.

```
# cat /etc/sysconfig/network-scripts/ifcfg-ens2f0
# 10G Team Physical Port to Storage Network
NAME=ens2f0
DEVICE=ens2f0
ONBOOT=yes
TEAM_MASTER=team10G
DEVICETYPE=TeamPort
MTU=9000
```

4. Edit the two configuration files for the interfaces connected to the management network using the following example. In this example, the management interface name is `eno5` and the management team name is `team1G`:

```
# cat ifcfg-eno5
# 1G Team Physical Port to Management Network
NAME=eno5
DEVICE=eno5
ONBOOT=yes
TEAM_MASTER=team1G
DEVICETYPE=TeamPort
```

5. Create the team interface file for the storage team using the following example. In this example, the team is called `team10G`. It is on the storage network running the network teaming `lacp runner`.



The active/active configuration is recommended for storage interfaces. This configuration requires extra active/active Multi-chassis Link Aggregation Group (MLAG) protocol and Link Aggregation Control Protocol (LACP) to be configured on the switches. This configuration requires the [network teaming lacp runner](#).

```

# cat /etc/sysconfig/network-scripts/ifcfg-team10G
# IPADDR= "SIP"
# GATEWAY= "SIP_GATEWAY"
# Pick one TEAM_CONFIG, activebackup or lacp
# note that lacp require changing switch port to lacp as well

TEAM_CONFIG="{\"runner\": {\"name\": \"lacp\"}, \"link_watch\":
{\"name\": \"ethtool\"}}"
PROXY_METHOD=none
BROWSER_ONLY=no
BOOTPROTO=none
DEFROUTE=no
IPV4_FAILURE_FATAL=no
IPV6INIT=no
NAME=team10G
DEVICE=team10G
ONBOOT=yes
DEVICETYPE=Team
IPADDR=192.0.2.2
PREFIX=24
GATEWAY=192.0.2.1
NM_CONTROLLED=yes
MTU=9000

```

6. Create the team interface file for the management team using the following example. In this example, the team is called team1G. It is on the management network running the network teaming activebackup runner.



The active/passive configuration is recommended for management interfaces, though you can also use the active/active configuration. This does not require extra configurations on the leaf switches. This configuration uses the [network teaming activebackup runner](#).

```

# cat /etc/sysconfig/network-scripts/ifcfg-team1G
# IPADDR= "MIP"
# GATEWAY= "MIP_GATEWAY"
# DNS1= "DNS"
# Pick one TEAM_CONFIG, activebackup or lacp
# note that lacp require changing switch port to lacp as well

TEAM_CONFIG="{\"runner\": {\"name\": \"activebackup\"}, \"link_watch\":
{\"name\": \"ethtool\"}}"
#TEAM_CONFIG="{ \"runner\": {\"name\": \"lacp\", \"active\": true,
\"fast_rate\": true } }"
PROXY_METHOD=none
BROWSER_ONLY=no
BOOTPROTO=none
DEFROUTE=yes
IPV4_FAILURE_FATAL=no
IPV6INIT=no
IPV6_AUTOCONF=yes
IPV6_DEFROUTE=yes
IPV6_FAILURE_FATAL=no
IPV6_ADDR_GEN_MODE=stable-privacy
NAME=team1G
DEVICE=team1G
ONBOOT=yes
DEVICETYPE=Team
IPADDR=198.51.100.2
PREFIX=24
GATEWAY=198.51.100.1
DNS1=198.51.100.250
NM_CONTROLLED=yes

```

7. Edit the `/etc/iproute2/rt_tables` file to enable a new routing table using the following sample. This file defines the mappings to use the routing table names instead of index numbers to refer to a specific table. In the following example, the new storage routing table called team10G can be called by its index (20) or its name (team10G):

```
# cat /etc/iproute2/rt_tables
#
# reserved values
#
255local
254main
253default
0unspec

20    team10G
```

8. Add routes to the routing table for storage traffic using the following example. This routing table points to the storage network as a default gateway and must be used for iSCSI traffic. In the following example, the teamed interface name is team10G.



You should replace `$storage_network`, `$storage_if_name src`, `$SIP table`, `$routing_table_name`, `$storage_default_gw dev`, `$storage_if_name src`, `$SIP table`, and `$routing_table_name` with your own values.

```
# cat /etc/sysconfig/network-scripts/route-team10G
$storage_network/24 dev $storage_if_name src $SIP table
$routing_table_name
default via $storage_default_gw dev $storage_if_name src $SIP table \
$routing_table_name
```

9. Add policy-based routing to use the new routing table that you created, if the traffic originates from the SIP or SVIP. Use the following example and substitute with your own values:

```
# cat /etc/sysconfig/network-scripts/rule-team10G
from $SIP table
$routing_table_name
```

10. Restart networking for all the changes to be applied.

```
# systemctl restart network.service
```

11. To check the policy-based routing rules, run the `ip rule show` command.
12. To check the routing table, run the `ip route show table` command.

Find more information

- [NetApp SolidFire Resources Page](#)

- [Documentation for earlier versions of NetApp SolidFire and Element products](#)

Copyright Information

Copyright © 2021 NetApp, Inc. All rights reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means-graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system- without prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP "AS IS" AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

RESTRICTED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.277-7103 (October 1988) and FAR 52-227-19 (June 1987).

Trademark Information

NETAPP, the NETAPP logo, and the marks listed at <http://www.netapp.com/TM> are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.