



Migration Guide

AWS Elemental Conductor File



AWS Elemental Conductor File: Migration Guide

Copyright © 2026 Amazon Web Services, Inc. and/or its affiliates. All rights reserved.

Amazon's trademarks and trade dress may not be used in connection with any product or service that is not Amazon's, in any manner that is likely to cause confusion among customers, or in any manner that disparages or discredits Amazon. All other trademarks not owned by Amazon are the property of their respective owners, who may or may not be affiliated with, connected to, or sponsored by Amazon.

Table of Contents

About this guide	1
Standard cluster migration	2
Plan maintenance windows	3
Step A: Get ready	4
Read the essential notes	4
Modify your automation system for HTTPS	4
Verify installer type	4
Create a boot USB drive	5
Verify space on each node	5
Step B: Prepare each node	5
Upgrade to the latest 2.17 minor version	5
Verify access to the BMC on the appliances	6
Note the network adapter for the management interface	6
Update firmware	6
Move custom files	6
Step C: Tear down an AWS Elemental Conductor File cluster	7
Step D: Create backups	8
Step E: Upgrade nodes	8
Step F: Rebuild the cluster	9
Tasks	12
Boot mode — UEFI	12
Switching to UEFI on a Dell	13
Switching to UEFI on a SuperMicro	15
Boot mode — Legacy	17
Switching to Legacy on a Dell	17
Switching to Legacy on a SuperMicro	19
Boot USB drive — create	21
Cluster — enable HA or disable HA	21
Database — back up	23
About the backup process	24
Step 1: Download the lifeboat script	25
Create the backup	25
Database — restore	27
Perform the restore	28

Result of the restore	29
Conductor File — install	31
AWS Elemental Server — install	33
Firmware — update	34
Step 1: Update the firmware	34
Step 2: Reboot the chassis	34
RHEL 9 — install	35
Installing on a Dell	35
Installing on an SuperMicro	37
RPM repository	38
Document history	39

About this guide

This guide describes how to upgrade or downgrade a Conductor cluster consisting of AWS Elemental Conductor File nodes and AWS Elemental Server worker nodes. It describes how to perform an upgrade to version 2.18 or higher of the AWS Elemental software, and how to downgrade to a version below 2.18.

This special guide exists because this software upgrade requires that you install the RHEL 9 version of the Linux operating system. A software downgrade requires that you re-install RHEL 7 or CentOS 7.

After you have migrated to a 2.18 version, you don't need to read this guide to perform further upgrades. Instead, read the regular upgrade guide.

Important

We strongly recommend that you test the entire migration procedure in your lab. This strategy lets you test the migration process itself, and test the entire workflow on the new software.

Note

For assistance with your AWS Elemental appliances and software products, see the [AWS Elemental Support Center](#).

Performing a standard cluster migration on an AWS Elemental Conductor File cluster

This procedure describes how to take the nodes in a AWS Elemental Conductor File cluster from a version below 2.18 and migrate them to version 2.18 or higher. The nodes in the cluster might be the following:

- A primary Conductor node and, optionally, a secondary Conductor node.
- AWS Elemental Server nodes.

Important

We strongly recommend that you test the entire migration procedure in your lab. This strategy lets you test the migration process itself, and test the entire workflow on the new software.

In this procedure, we show how to upgrade the cluster from Conductor File version 2.17.5 (worker nodes version 2.17.5) to version 2.18.5. Modify the commands you enter to match your versions.

Important

You must upgrade all the nodes in the cluster to a 2.18 version. You can't, for example, set up the cluster so that Conductor File is running 2.17 and the workers (or some of the workers) are running 2.18.

Topics

- [Plan maintenance windows for migrating an AWS Elemental Conductor File cluster](#)
- [Step A: Get ready to migrate an AWS Elemental Conductor File cluster](#)
- [Step B: Prepare each AWS Elemental Conductor File node for migration](#)
- [Step C: Tear down an AWS Elemental Conductor File cluster](#)
- [Step D: Create backups](#)
- [Step E: Upgrade nodes](#)

- [Step F: Rebuild the cluster](#)

Plan maintenance windows for migrating an AWS Elemental Conductor File cluster

You should plan to perform the cluster migration in several phases:

First phase

You can perform the tasks in [Step A: Get ready to migrate an AWS Elemental Conductor File cluster](#) outside of a maintenance window.

Second phase

Perform the following tasks in one or more maintenance windows. The number of windows depends on the number of nodes you can complete in one maintenance window.

- [the section called "Step B: Prepare each node"](#)

Third phase

Perform all the following tasks on every node, all in one maintenance window.

- [the section called "Step C: Tear down an AWS Elemental Conductor File cluster"](#)
- [the section called "Step D: Create backups"](#)
- [the section called "Step F: Rebuild the cluster"](#)

These steps upgrade all the nodes at one time. You must perform the upgrade in this way because you can't have a cluster where some nodes are on the previous version of the AWS Elemental software and some are on the new version.

Step A: Get ready to migrate an AWS Elemental Conductor File cluster

Read the essential notes

Refer to the essential notes in the [AWS Elemental Conductor File Release Notes](#) to identify key changes to the behavior of the Conductor File and worker nodes.

Important

Make sure that you have the latest version of the Release Notes. If you downloaded the file more than a few days ago, we recommend that you download it again.

Modify your automation system for HTTPS

After a node has been migrated, it uses HTTPS. By default, the nodes are set up with self-signed certificates. Make sure of the following points:

- You might need to change your automation system to use HTTPS.

Verify installer type

The software installer that you use depends on whether you have GPU-accelerated software type, or CPU-only.

To determine whether you have a GPU-accelerated or CPU-only system, run the following command:

lspci | grep NVIDIA

If the output is empty, your system is CPU-only.

If the output is similar to the following, your system is GPU-accelerated:

```
VGA compatible controller: NVIDIA Corporation GM204GL [Tesla M60] (rev a1)
VGA compatible controller: NVIDIA Corporation GM204GL [Tesla M60] (rev a1)
```

```
VGA compatible controller: NVIDIA Corporation GM204GL [Tesla M60] (rev a1)
VGA compatible controller: NVIDIA Corporation GM204GL [Tesla M60] (rev a1)
```

Create a boot USB drive

On Dell hardware, you have the option to install RHEL 9 by using a boot USB drive or by using or iDRAC (for Dell). (Note that SuperMicro hardware, you can only install RHEL 9 by using IPMI.)

If you want to use a boot USB drive with Dell, you should make the drive now. You might want to make several drives, depending on how many people will be performing the migration tasks. For instructions, see [the section called “Boot USB drive — create”](#).

Verify space on each node

As part of the upgrade, you create a backup of the data on each node. You must make sure that you have enough free space for the backup. Follow these guidelines:

- The backup for a freshly kickstarted and licensed appliance generates a small backup directory andB zipped version of the backup (<hostname>_lifeboat-archive.zip).
- Your configuration will generate larger files because of the data that you create, so review available space before starting.
- Check the contents of the /home partition, and clear out old files, unnecessary files, and old installers.

Step B: Prepare each AWS Elemental Conductor File node for migration

Prepare the nodes during one or more maintenance windows. The number of windows depends on the number of nodes you can complete in one maintenance window.

Upgrade to the latest 2.17 minor version

To upgrade to version 2.18 or higher, the software currently installed on the node must be version 2.17.0 or higher.

- If the Conductor nodes are on one version and the worker nodes are on a different version, upgrade all nodes to the same version.

Verify access to the BMC on the appliances

Make sure that you have access to the BMC on each appliance:

- On a Dell server, make sure that iDRAC is installed and that you can start it.
- On an SMC server, make sure that IPMI is installed and that you can start it.

You can install iDRAC or IPMI even when the node is active — when AWS Elemental Server is running events or Conductor is controlling the cluster.

Note the network adapter for the management interface

Make a note of the management network device listed in the web UI under **Settings, Network, Network Settings, Current Settings, Network Devices**. By default, eth0 is the management network device, but this may differ on your system. You'll need to know this adapter later during the migration process.

Update firmware

Both the BIOS firmware and the BMC firmware (IPMI for SuperMicro, iDRAC for Dell) must be at the latest versions available from the manufacturer. They must be at the latest versions before you can set the boot mode to UEFI.

We recommend that you update the firmware on all your nodes at the same time. We also recommend that you perform this update during a maintenance window. If you need to upgrade to the latest 2.17 version of the AWS Elemental software, you might want to perform both tasks during the same maintenance window.

After you install the firmware, you must reboot each node. For more information, see [the section called “Firmware — update”](#).

Move custom files

You might have custom files in `/opt/elemental_se/scripts` on the node. These are files that you created. They aren't part of the installation of the Conductor File or AWS Elemental Server software, and they aren't backed up and restored.

Copy these files to storage off the node, so that you can copy the files back to the node after you've upgraded it.

Step C: Tear down an AWS Elemental Conductor File cluster

Before you can install RHEL 9 and the new software version, you must remove all the nodes from the cluster.

1. Remove every worker node from the cluster. You perform this action from the primary Conductor node.
 - a. On the primary Conductor web interface, choose **Nodes**.
 - b. On the **Nodes** screen, scroll down to the list of nodes.
 - c. Choose **View All Nodes** to display a list of all the conductor and worker nodes in the network.
 - d. Select the nodes to remove and choose **Remove from Cluster**.
2. If you have a secondary conductor in the cluster, you need to remove the secondary conductor node from the cluster.
 - a. If HA mode is enabled, you need to disable HA mode on the **secondary** conductor.
 - b. Remove the secondary conductor node from the cluster
 - i. From the Linux prompt, log in to the conductor node with the *elemental* user credentials.
 - ii. Run the following commands:

```
[elemental@hostname ~]$ cd /opt/elemental_se  
[elemental@hostname ~]$ sudo ./configure -s -c -t -n -z -xeula
```

- iii. You will see the following prompts:

```
Remove this node from the Conductor system? [N]  
Y  
Configure this node as a secondary Elemental Conductor in an existing cluster?  
[N]  
N
```

3. If HA mode is enabled, you need to disable HA mode on the **primary** conductor.

After you remove the last worker node, the cluster still exists but it doesn't contain any worker nodes or a secondary Conductor. The single Conductor exists, but it isn't controlling any worker nodes.

Step D: Create backups

Create a backup of the data on every node — the primary Conductor node, the secondary Conductor node, and all the workers.

Important

After you make a backup of the first node in the cluster, don't make any changes to any worker node or Conductor node or to cluster until you've finished this migration process. Don't change the setup of the Conductor node.

To create database backups, see [the section called “Database — back up”](#).

Step E: Upgrade nodes

Perform these steps on each node in the cluster, after you've removed all the nodes from the cluster. Before you start, make sure that you've performed the steps in [Step B: Prepare each AWS Elemental Conductor File node for migration](#).

1. Set the boot mode on the node to UEFI.
2. Perform a kickstart to upgrade the operating system to RHEL 9.
3. From the Linux command line, log in to the node. Use the **elemental** user credentials.
4. Check routing table by running the following command in the command line:

```
ip r show
```

The system returns something similar to the following:

```
default via 10.x.x.x dev eth0 proto dhcp src 10.x.x.x metric 103
10.x.x.x/x dev eth1 proto kernel scope link src 10.12.107.43 metric 102
...
```

To proceed with the upgrade, your management interface must be listed for the first route. To find which network interface is your system's management interface, see: [Note the network adapter for the management interface](#).

5. If your management interface isn't listed for the first route, you must update the default route.
6. Run the installer. Use the appropriate command:
 - For GPU versions of the software (for Conductor File only):

```
[elemental@hostname ~]$ sudo sh ./elemental_production_server_2.18.n.nnnn.run --skip-all --start --xeula
```

- For CPU-only versions of the software:

```
[elemental@hostname ~]$ sudo sh ./elemental_production_server_cpu_2.18.n.nnnn.run --skip-all --start --xeula
```

Where:

`--skip-all` skips all the prompts. There is no need to view prompts about configuration because when you restore the database to the node, all the configuration data is copied over and overwrites any configuration data already on the node.

`--start` restarts the software after installation.

`--xeula` skips the display of the license agreement. There is no need to view this prompt because you have previously accepted the agreement.

7. When the installation is complete, restart the node:

```
[elemental@hostname ~] sudo reboot
```

Step F: Rebuild the cluster

1. Restore the database on primary conductor file node.
2. If you have a secondary conductor in the cluster, you need to restore the database on secondary conductor file node.
3. If you have a secondary conductor in the cluster, you need to add the secondary conductor to the cluster.
 - a. From the Linux prompt, log in to the secondary conductor node with the *elemental* user credentials.

b. Run the following commands:

```
[elemental@hostname ~]$ cd /opt/elemental_se  
[elemental@hostname ~]$ sudo ./configure -s -c -t -n -z -xeula
```

c. You will see the following prompts:

```
[elemental@hostname ~]$ Configure this node as a secondary Elemental Conductor in  
an existing cluster? [N]  
Y  
[elemental@hostname ~]$ What is the hostname or IP address of the cluster  
management node?  
<hostname of primary conductor>  
...  
[elemental@hostname ~]$ Are you sure you want to continue connecting (yes/no/  
[fingerprint])?  
yes  
[elemental@hostname ~]$ elemental@<hostname> password:  
<hostname of primary conductor>  
...  
[elemental@hostname ~]$ Trust this public key? [Y]  
Y
```

4. If HA mode is enabled before, you need to re-enable it again.

5. Restore the database on all the worker nodes.

6. Add all the worker nodes back to the cluster

- a. From the Linux prompt, log in to the worker node with the *elemental* user credentials.
- b. Run the following commands:

```
[elemental@hostname ~]$ cd /opt/elemental_se  
[elemental@hostname ~]$ sudo ./configure -s -c -t -n -z -xeula
```

c. After being prompted to add the worker node to the Conductor File cluster, you are asked to trust the public key from the conductor node(s). You must accept this for the Conductor File node to control the worker node.

d. After trusting the public key, continue through the configuration prompts as normal.

7. Use the configure script to add worker nodes to the cluster.

⚠ Important

If this is your first time adding worker nodes to a cluster you must add worker nodes to the Conductor File cluster via the CLI. Using the Conductor File web interface to discover the worker node and add to the cluster causes the worker node to go into a failed state.

- a. On the worker node, enter the following command to change the directory:

```
cd /opt/elemental_se
```

- b. Enter the following command to start the configurations script on the worker node:

```
sudo ./configure
```

- c. After being prompted to add the worker node to the Conductor File cluster, you are asked to trust the public key from the conductor node(s). You must accept this for the Conductor File node to control the worker node.
- d. After trusting the public key, continue through the configuration prompts as normal.

Tasks for migrating an AWS Elemental Conductor File cluster

This section lists the tasks that are part of migrating the nodes in a cluster. The tasks are listed alphabetically.

For information about the correct order for following these tasks, see the procedure for your setup:

- [Standard cluster migration](#)

Topics

- [Switching boot mode to UEFI](#)
- [Switching boot back to Legacy](#)
- [Create a boot USB drive](#)
- [Enabling or disabling high availability \(HA\)](#)
- [Backing up data](#)
- [Restoring the database](#)
- [Installing Conductor File on nodes](#)
- [Installing AWS Elemental Server on a worker node](#)
- [Updating firmware](#)
- [Installing RHEL 9](#)
- [Working with RPM repository](#)

Switching boot mode to UEFI

RHEL 9 requires that the boot mode for the appliance is UEFI. You can change the boot mode from BIOS (or Legacy mode) to UEFI.

Topics

- [Switching to UEFI on a Dell](#)
- [Switching to UEFI on a SuperMicro](#)

Switching to UEFI on a Dell

There are three ways to switch the boot mode from Legacy mode to UEFI.

Topics

- [Switch using the iDRAC user interface](#)
- [Switch using RACADM](#)
- [Switch using the F2 boot menu](#)

Switch using the iDRAC user interface

iDRAC is a system for controlling Dell servers remotely. It is already installed and enabled on the Dell server. However, you might need to configure it. For more information about configuring iDRAC, see the official [Dell iDRAC User Guide](#).

This procedure is identical to the procedure for switching to BIOS, except that you choose **UEFI** instead of **BIOS**.

1. Log into the iDRAC user interface as an administrative user.
2. On the iDRAC menu, choose **Configuration**, then **BIOS Settings**, then **Boot Settings**.
3. On the **Boot Mode** line, change the **Current Value** from **BIOS** to **UEFI**.
4. Scroll down to the **Apply** button and choose that button. The **Pending Value** changes to **UEFI**.
5. Scroll down to the bottom of the page and choose **Apply And Reboot**.

The system reboots. UEFI is now enabled.

Switch using RACADM

You can switch to UEFI mode by logging into RACADM, which is the iDRAC command line interface.

This procedure is identical to the procedure for switching to BIOS, except that you specify **UEFI** instead of **BIOS**.

1. Start a Linux session and log into the iDRAC command line interface as a Linux Admin user. For example:

```
ssh ADMIN@<iDRAC hostname or IP>
```

The IDRAC command line interface appears, with the **racadm>>** prompt.

2. To verify that the current boot environment is BIOS, enter this command:

```
get BIOS.biosBootSettings.BootMode
```

If the environment is BIOS, a message similar to the following appears:

```
[Key=BIOS.Setup.1-1#biosBootSettings]  
BootMode=Bios
```

3. Set the **BIOS settings to UEFI**:

```
set BIOS.BiosBootSettings.BootMode Uefi
```

4. Apply and reboot:

```
jobqueue create BIOS.Setup.1-1 -r Forced
```

The system reboots. UEFI is now enabled.

Switch using the F2 boot menu

You can use the boot menu from a direct connection to the server, or through the IDRAC virtual console.

This procedure is identical to the procedure for switching to BIOS, except that you specify **UEFI** instead of **BIOS**.

1. This step applies only if you want to use the virtual console: log into the iDRAC user interface and launch the Virtual Console.
2. Reboot the appliance.

```
sudo reboot
```

3. The appliance starts to reboot using BIOS, which is currently enabled.
4. As soon as the reboot starts, repeatedly press **F2** on the keyboard, until the message **Entering System Setup** appears. Then wait for the **System Setup** screen to appear.
5. Choose **System BIOS**, then choose **Boot Settings**.

6. On the **Boot Mode** line, choose **UEFI**.
7. Choose the **Exit** option and follow the prompts to save. At the success message, choose **OK**.

The system reboots. UEFI is now enabled.

Switching to UEFI on a SuperMicro

To switch the boot mode from BIOS (Legacy mode) to UEFI, you can use the IPMI interface, or you can work when directly connected to the server.

Step 1: Install Java applet

Decide if you want to use the IPMI management console, or if you plan to connect directly to the server. If you want to use the console, decide if you want to use the Java remote console applet to access the console, or if you want to use HTML5.

If you want to use the IPMI management console and you want to use the Java remote console applet, you must install the applet.

1. Make sure you have the IP address of the IPMI. If you don't have it, connect to the appliance using SSH, then type the following command:

```
sudo ipmiutil lan | grep Param\{3\}
```

The IP address appears in the response. For example:

```
Lan Param(3) IP address: 10 4 130 12
```

2. Log in to the IPMI management console via a web browser. Use the ADMIN credentials, with the user name entered in uppercase.
3. From the menu bar, choose **Console Redirection**, then **Launch Console**. The download of a JNLP file starts.
4. When the download is complete, open the applet. The applet is self signed. Typically, this file is already associated with Java so you should just be able to open it directly.
5. Change the security level in the Java control panel in order for the applet to run:
 - a. In Windows, open **Control Panel**, **Programs**, and then **Java**.
 - b. Click the **Security** tab. Move the slider to the lowest setting: **Medium**.

- c. Click **OK**.

You can now open the remote console window.

Step 2: Change the mode to UEFI

This procedure is nearly identical to the procedure for switching to BIOS. You change the same fields on the **Setup Utility** screen.

1. From the IPMI management console, sign in to the server as the *elemental* user.
2. Reboot the system:

```
[elemental@hostname]$ sudo reboot
```

The system starts to reboot. The window size might change as the system is rebooting.

3. While the system is rebooting, repeatedly press the **Delete** key on the keyboard (or the **del** button on the virtual keyboard). The **Setup Utility** screen appears.

You can use these keys to work on the screen:

- The arrow keys
 - Enter to select
 - ESC to return to the previous screen.
4. On the main menu, choose **Advanced**.
 5. In **sSATA Configuration**, look for fields that have one of these values:
 - **BIOS**
 - **DUAL**
 - **Legacy**
 - **Legacy BIOS**

Change the value to **EFI**. If there are no fields with these values, go to the next step.

6. In **PCIe/PCI/PnP Configuration**, find every field that has one of these values:
 - **BIOS**
 - **DUAL**
 - **Legacy**
 - **Legacy BIOS**

In each of these fields, change the value to **EFI**.

7. On the main menu, choose **Boot**. In **Boot Mode Select**, change the value from **DUAL** to **UEFI**.
8. Select **F4**. On the **Save & Exit** dialog, choose **Yes**.

Switching boot back to Legacy

RHEL 7 and CentOS 7 require that the boot mode for the appliance is BIOS. If you have changed the boot mode to UEFI to support RHEL 9, you can change the boot mode from UEFI back to BIOS (Legacy mode).

Topics

- [Switching to Legacy on a Dell](#)
- [Switching to Legacy on a SuperMicro](#)

Switching to Legacy on a Dell

There are three ways to switch the boot mode from UEFI back to BIOS (Legacy mode).

Topics

- [Switch using the iDRAC user interface](#)
- [Switch using RACADM](#)
- [Switch using the F2 boot menu](#)

Switch using the iDRAC user interface

iDRAC is a system for controlling Dell servers remotely. It is already installed and enabled on the Dell server.

This procedure is identical to the procedure for switching to UEFI, except that you specify BIOS instead of UEFI.

1. Log into the iDRAC user interface as an administrative user.
2. On the iDRAC menu, choose **Configuration**, then **BIOS Settings**, then **Boot Settings**.
3. On the **Boot Mode** line, change the **Current Value** from **UEFI** to **BIOS**.

4. Scroll down to the **Apply** button and choose that button. The **Pending Value** changes to **BIOS**.
5. Scroll down to the bottom of the page and choose **Apply And Reboot**.

The system reboots. Legacy mode is now enabled.

Switch using RACADM

You can revert to Legacy mode by logging into RACADM, which is the iDRAC command line interface.

This procedure is identical to the procedure for switching to UEFI, except that you specify BIOS instead of UEFI.

1. Start a Linux session and log into the iDRAC command line interface as a Linux Admin user. For example:

```
ssh ADMIN@<iDRAC hostname or IP>
```

The iDRAC command line interface appears, with the **racadm>>** prompt.

2. To verify that the current boot environment is UEFI, enter this command:

```
get BIOS.biosBootSettings.BootMode
```

If the environment is UEFI, a message similar to the following appears:

```
[Key=BIOS.Setup.1-1#biosBootSettings]  
BootMode=Uefi
```

3. Set the **BIOS settings** to **BIOS**:

```
set BIOS.BiosBootSettings.BootMode Bios
```

4. Apply and reboot:

```
jobqueue create BIOS.Setup.1-1 -r forced -s TIME_NOW
```

The system reboots. Legacy mode is now enabled.

Switch using the F2 boot menu

You can use the boot menu from a direct connection to the server, or through the iDRAC virtual console.

This procedure is identical to the procedure for switching to UEFI, except that you specify BIOS instead of UEFI.

1. This step applies only if you want to use the virtual console: log into the iDRAC user interface and launch the Virtual Console.
2. Reboot the appliance.

```
sudo reboot
```

3. The appliance starts to reboot using UEFI, which is currently enabled.
4. As soon as the reboot starts, repeatedly press **F2** on the keyboard, until the message **Entering System Setup** appears. Then wait for the **System Setup** screen to appear.
5. Choose **System BIOS**, then choose **Boot Settings**.
6. On the **Boot Mode** line, choose **BIOS**.
7. Choose the **Exit** option and follow the prompts to save. At the success message, choose **OK**.

The system reboots. Legacy mode is now enabled.

Switching to Legacy on a SuperMicro

To switch the boot mode from UEFI back to BIOS (Legacy mode), you can use the IPMI interface, or you can work when directly connected to the server.

Install Java applet

Decide if you want to use the IPMI management console, or if you plan to connect directly to the server. If you want to use the console, decide if you want to use the Java remote console applet to access the console, or if you want to use HTML5.

If you want to use the Java remote console applet, you might need to install it. See [the section called “Step 1: Install Java applet”](#).

Change the mode to BIOS

This procedure is nearly identical to the procedure for [switching to UEFI](#). You change the same fields on the **Setup Utility** screen, but you specify either Legacy or Disabled.

1. From the IPMI management console, sign in to the server as the *elemental* user.
2. Reboot the system:

```
[elemental@hostname sudo reboot
```

The system starts to reboot. The window size might change as the system is rebooting.

3. While the system is rebooting, repeatedly press the **Delete** key on the keyboard (or the **del** button on the virtual keyboard). The **Setup Utility** screen appears.

You can use these keys to work on the screen:

- The arrow keys
 - Enter to select
 - ESC to return to the previous screen.
4. On the main menu, choose **Advanced**.
 5. In **sSATA Configuration**, change the following line to **Legacy**:
 - **sSATA RAID Option ROM/UEFI Driver**
 6. In **PCIe/PCI/PnP Configuration**, change the following lines to **Legacy**:
 - **AOC-URN2-14GXS-SLOT1 PCI-E 3.0 X8 OPROM**
 - **RSC-RIUW-EBR SLOT1 PCI-E X8 OPROM**
 - **RSC-RIUW-2E16 SLOT1 PCI-E X16 OPROM**
 - **RSC-RIUW-2E16 SLOT2 PCI-E X16 OPROM**
 - **Onboard LAN OPROM Type**
 - **Onboard Video OPROM**
 7. Still in **PCIe/PCI/PnP Configuration**, change the following lines to **Disabled**:
 - **Onboard LAN NVMe1 OPROM**
 - **Onboard LAN NVMe2 OPROM**

8. Select **F4**. On the **Save & Exit** dialog, choose **Yes**.

Create a boot USB drive

1. Obtain the RHEL 9 .iso file from [AWS Elemental Software Download page](#).

Find the AWS Elemental product and version they are planning to use. The appropriate ISO file appears beside that version.

2. At your workstation, use a third-party utility (such as PowerISO or ISO2USB) to create a bootable USB drive from your .iso file. For help, see the knowledge base article [Creating Bootable Recovery \(kickstart\) Media](#).

Enabling or disabling high availability (HA)

Disable high availability prior to performing any changes on the Conductor nodes.

To disable high availability on a conductor:

1. From the Linux prompt, log in to the conductor node with the *elemental* user credentials.
2. Enter the following command to configure the primary Conductor node for HA.

```
[elemental@hostname ~]$ sudo /opt/elemental_se/.support_utils/dbrepl disable
```

3. Restart the service using the following commands.

```
[elemental@hostname ~]$ sudo systemctl restart postgresql-15  
[elemental@hostname ~]$ sudo systemctl restart elemental_se
```

4. Enter the following command to verify that Conductor high availability is disabled.

```
[elemental@hostname ~]$ tail -F /opt/elemental_se/web/log/conductor.output
```

The conductor.output log starts to scroll on the screen and shows messages as they are occurring. Watch for the following INFO lines on the primary Conductor node.

```
CONDUCTOR: Initializing environment  
I, [2024-06-04T23:07:54.439807 #131824] INFO -- : HA environment not enabled  
[2024-06-04 23:08:00 UTC SERVICE]: Elemental Conductor File 2.18.x.x
```

5. Enter **Ctrl+C** to exit the tail command.

6. Enter the following command, where `<day>` is today (the day you are upgrading), typed with an initial capital letter: Mon, Tue, Wed, Thu, Fri, Sat, Sun:

```
[elemental@hostname ~]$ sudo tail -F /data/pgsql/logs/postgresql-<day>.log
```

7. Confirm that you see this line.

```
database system is ready to accept connections.
```

If the `elemental_se` or `postgres` process has already started when you starting tailing the logs, you might not see the **ready to accept connections** message. Instead, you could see **rejects connection for host messages** until you upgrade the worker nodes.

8. Enter Ctrl+C to exit the tail command.

You must re-enable high availability on the Conductor nodes, to put the nodes back into a redundant configuration. If you have only one Conductor File node, skip this step.

To enable high availability on the primary Conductor

1. From the Linux prompt, log in to the primary Conductor node with the *elemental* user credentials.
2. Enter the following command to configure the primary Conductor node for high availability. Where `<dbrepl_config_file_name>` is your `dbrepl_config.yml` file.

```
[elemental@hostname ~]$ sudo /opt/elemental_se/.support_utils/dbrepl configure  
<dbrepl_config_file_name> primary
```

3. Restart the `elemental_se` service.

```
[elemental@hostname ~] sudo systemctl restart elemental_se
```

4. Enter the following command to verify that the service is running.

```
[elemental@hostname ~] tail -F /opt/elemental_se/web/log/conductor.output
```

The `conductor.output` log starts to scroll on the screen and shows messages as they are occurring. Watch for the following three INFO lines on the primary Conductor node.

```
CONDUCTOR: Initializing environment
I, [2024-06-04T20:58:01.073149 #9844] INFO -- : Configuring the HA environment
I, [2024-06-04T20:58:03.170375 #9844] INFO -- : Preparing database as replication
master
...
[2024-06-04 20:58:09 UTC SERVICE]: Elemental Conductor File 2.18.x.x
```

5. Enter Ctrl+C to exit the tail command.
6. Enter the following command, where <day> is today (the day you are upgrading), typed with an initial capital letter: Mon, Tue, Wed, Thu, Fri, Sat, Suns.

```
[elemental@hostname ~]$ sudo -s
[elemental@hostname ~]$ cd /data/pgsql/logs
[elemental@hostname ~]$ tail -F postgresql-<day>.log
```

7. Confirm that you see this line.

```
database system is ready to accept connections.
```

If the elemental_se or postgres process has already started when you starting tailing the logs, you might not see the **ready to accept connections message**. Instead, you could see **rejects connection for host messages** until you upgrade the worker nodes.

8. Enter Ctrl+C to exit the tail command.
9. Type the following command to exit the session as the sudo user.

```
[elemental@hostname ~]$ exit
```

On the secondary Conductor

Repeat high availability steps on the secondary Conductor but use the following command instead.

```
[elemental@hostname ~]$ sudo /opt/elemental_se/.support_utils/dbrepl configure
<dbrepl_config_file_name> secondary
```

Backing up data

You back up data using the special lifeboat script.

⚠ Warning

This warning applies if you are migrating to Conductor File version 3.26.1 or 3.26.2 and if you use the channel schedules feature of Conductor File.

There is an issue with the schedules feature that affects the lifeboat script that you use to back up and restore the database. If the database includes schedule data, the backup and restore will fail.

Do not migrate your Conductor File cluster until AWS Elemental releases a version of the software that fixes this issue.

⚠ Important

The lifeboat script creates a backup of multiple files that are relevant to the AWS Elemental software. These files might include credentials and other sensitive system information. Handle the backup according to your organization's best practices for handling sensitive data.

About the backup process

The script backs up the following data:

- Licenses.
- Network settings for the node, including Ethernet configurations, DNS information, and host addresses.
- Timecode configuraton such as NTP, PTP, and chronyd.
- Firewall settings.
- SSL certificates.
- Optionally, the user credentials used in various components on the cluster. It is convenient to include these credentials, if your organization's policies allow them to be handled in this way.
- Configuration files for features of the AWS Elemental software.
- Remote storage mounts. The data is included only in the database for the primary and secondary Conductor nodes.
- Cluster data. Data relating to the cluster, including data about the channels, MPTSeS, channel and MPTS node assignments, users setup, redundancy groups, cluster members. The data is

included only in the database for the primary Conductor. The primary Conductor pushes data down to the secondary Conductor and to the appropriate worker nodes.

Step 1: Download the lifeboat script

Perform this procedure on every node in the cluster, to copy the lifeboat script onto every node.

1. Download the latest version of the lifeboat script from <https://a.co/ElementalRHEL9Lifeboat> to your laptop. The lifeboat file is called `elemental_lifeboat_e1.tar`.

Important

Download the script just before you are ready to create the backup. AWS Elemental is continually making improvements to the script, therefore you want to make sure that you always have the latest version.

2. Copy the lifeboat file to the `/home/elemental` directory on every node in the cluster.
3. From the Linux prompt, use the *elemental* user to start a remote terminal session with the node. Don't log in as `sudo`.
4. Untar the lifeboat file.

```
[elemental@hostname ~]$ cd /home/elemental && tar xvf elemental_lifeboat_e19.tar
```

Create the backup

Important

Make sure that you have stopped the node. We recommend that you don't run the script on an active node. The script temporarily stops `elemental_se` and `httpd` services.

Enter the backup command as follows:

```
[elemental@hostname ~]$ ./lifeboat.sh --backup --include-creds
```

Where `--include-creds` (optional) includes the following credentials in the backup: SSH, AWS, SMB/CIFs.

Results of the backup

The script creates the following assets:

- Asset 1. One version of the data that is compatible with 2.26.1 or later. When you restore the backup after you've installed RHEL 9, the lifeboat script will automatically select and copy over this version.
- Asset 2. One version of the data that is compatible with 2.25.x and earlier. You might later decide to downgrade a node back to a version below 2.26.0. When you restore the backup after you've installed RHEL 7 or CentOS 7, the lifeboat script will automatically select and copy over this version.
- Asset 3. An MD5 checksum of the contents of asset 3.
- Asset 4. A SHA1 checksum of the content of asset 3.

The script also creates the following files:

- File 1. A file that contains assets 1 and 2. The file has this name, where `hostname` is the name of the current node:

```
<hostname>_lifeboat-archive.zip
```

- File 2. A file that contains assets 3 and 4. The file has this name, where `hostname` is the name of the current node:

```
<hostname>_lifeboat-archive_export-checksum.txt
```

- File 3. A file that contains assets 1, 2, 3 and 4. The file is stored on the current node at this location:

```
/opt/upgrade-backups/system-backup.tar.gz
```

Verify the backup

Verify the integrity of the backup archive. This step is optional but we strongly recommend that you follow it because the [restore operation](#) that you later perform might fail if the backup file is corrupted.

You verify the integrity by comparing the checksum that the backup script creates to the checksum that you perform on the `<hostname>_lifeboat-archive_export-checksum.txt` file. You can compare an MD5 or a SHA1 checksum.

1. Enter the `cat` command to view the checksums currently listed in the checksum file.:

```
~]$ cat <hostname>_lifeboat-archive-export-checksum.txt
```

The `cat` command simply displays the file contents on your screen. For example:

```
md5sum
    d41d8cd98f00b204e9800998ecf8427e
sha1sum
    e3b0c44298fc1c149afbf4c8996fb92427ae41e4649b934ca495991b7852b855
```

2. Now run a checksum command on the `lifeboat-archive.zip` file. For example:

```
~]$ md5sum /home/elemental/<hostname>_lifeboat-archive.zip
```

Or

```
~]$ sha1sum /home/elemental/<hostname>_lifeboat-archive.zip
```

3. Compare the results from step 1 to the results from step 2. If the checksums don't match, copy the archive file again.

Store the backup archive

Copy the `<hostname>_lifeboat-archive.zip` file to storage off the node, so that you can copy it back to the node when you want to perform the restore operation.

Important

The lifeboat script creates a backup of multiple files that are relevant to the AWS Elemental software. These files might include credentials and other sensitive system information. Handle the backup according to your organization's best practices for handling sensitive data.

Restoring the database

You restore data using the same lifeboat script that you used to create the backup.

⚠ Important

Make sure that you have stopped the node. Don't run the script on an active node.

Perform the restore

1. Download the lifeboat script, following the procedure you followed when you [created the backup](#).

2.

⚠ Important

Make sure that you have latest version of the script. AWS Elemental is continually making improvements to the script.

Enter the restore command.

- On a worker node or the secondary Conductor node, enter this command:

```
[elemental@hostname ~]$ ./lifeboat.sh --restore
```

Don't include any options.

- On the primary Conductor node, enter this command:

```
[elemental@hostname ~]$ ./lifeboat.sh --restore --import-database
```

The script tries to extract the version of the backup that is stored in this folder:

```
/opt/upgrade-backups/system-backup.tar.gz
```

This file was created when you created the backup. The script automatically copied it to this directory. The installation of RHEL 9 should not have deleted this file. Therefore, it should be in this location.

If this file doesn't exist or if there is a problem with it, the scripts stops. See the recovery steps below to continue.

The script tries to extract the version of the backup that is stored in this folder: `/opt/upgrade-backups/system-backup.tar.gz`. This file was created when you created the backup. The script automatically copied it to this directory. The installation of RHEL 9 should not have deleted this file. Therefore, it should be in this location. If this file doesn't exist or if there is a problem with it, the script stops. See the recovery steps below to continue.

Important

If you updated the routing table in earlier, add the `--exclude netscripts-rest` flag when you run the `lifeboat.sh` script:

```
[elemental@hostname ~]$ ./lifeboat.sh --restore --import-database --exclude  
netscripts-rest
```

3. After the restore has succeeded, reboot the node:

```
[elemental@hostname ~]$ sudo reboot
```

Recovery steps

1. Locate the other copies of the backup and of the checksum files that you should have copied to storage off the node. The files to locate are:
 - `<hostname>_lifeboat-archive.zip`
 - `<hostname>_lifeboat-archive_export-checksum.txt`
2. Copy the files to `/home/elemental`
3. Enter the restore command again:

```
[elemental@hostname ~]$ ./lifeboat.sh --restore
```

This time the script looks for the files that are in `/home/elemental`, and restores those files.

Result of the restore

Restored data

As a result of the restore command, the following data from the backup is restored on the nodes:

Node	Worker nodes	Secondary Conductor	Primary Conductor
Licenses	Yes	Yes	Yes
Network settings for the node, including Ethernet configurations, DNS information, and host addresses	Yes	Yes	Yes
Timecode configuration such as NTP, PTP, and chronyd	Yes	Yes	Yes
Firewall settings	Yes	Yes	Yes
The user credentials used in various components on the cluster (if you included them in the backup)	Yes	Yes	Yes
Configuration files for features of the AWS Elemental software	Yes	Yes	Yes
Remote storage mounts.		Yes	Yes
Cluster data. Data relating to the cluster, including data about the	Yes		

Node	Worker nodes	Secondary Conductor	Primary Conductor
channels, MPTSES, channel and MPTS node assignments, users setup, redundancy groups, cluster members.			

Notes

- The cluster data is only ever stored on the Conductor nodes. It is restored only to the primary Conductor because when you enable HA later in this migration procedure, the primary Conductor pushes the data to the secondary Conductor and to the appropriate worker nodes.
- The remote storage mounts is only ever stored on the Conductor nodes. The data specific to the node is restored to that node.

Installing Conductor File on nodes

1. From the Linux command line, log in to the Conductor node where you want to install Conductor File software. Use the *elemental* user credentials.
2. Run the installer:
 - For the **primary** (or only) Conductor File, enter this command:

```
[elemental@hostname ~]$ sudo sh ./  
elemental_production_conductor_file_2.18.0.12345.run --skip-all --start -xeula
```

Where:

`--skip-all` skips all the prompts, which means you won't change anything about the configuration.

Don't change anything about the configuration as part of the upgrade either of the Conductor nodes. Don't change the hostname. If you want to change anything about the configuration, you can do so after you've completed the migration.

--start restarts the software after installation.

--xeula skips the display of the license agreement. There is no need to view this prompt because you have previously accepted the agreement.

- For the **secondary** Conductor File, enter this command:

```
[elemental@hostname ~]$ sudo sh ./
elemental_production_conductor_file_2.18.0.12345.run --cleandb --skip-all --start -
xeula
```

Where:

--cleandb deletes the application database. This option is required on the secondary Conductor node. You don't need the application database because when you add the secondary Conductor node back into the cluster, the secondary Conductor node will synchronize with the database of the primary Conductor node.

Note that this option doesn't clean operating system configuration data.

--skip-all skips all the prompts, which means you won't change anything about the configuration.

Don't change anything about the configuration as part of the upgrade either of the Conductor nodes. Don't change the hostname. If you want to change anything about the configuration, you can do so after you've completed the migration.

--start restarts the software after installation.

--xeula skips the display of the license agreement. There is no need to view this prompt because you have previously accepted the agreement.

3. Make sure that the elemental_se service restarts. Look for this prompt on the primary Conductor command line:

```
Starting elemental_se: [OK]
```

Installing AWS Elemental Server on a worker node

This install procedure isn't the same as the install procedure on a newly obtained appliance. You don't have to configure the node.

This install procedure is very similar to the upgrade procedure, but there are significant differences in the options you include.

1. From the Linux command line, log in to the worker node. Use the **elemental** user credentials.
2. Run the installer. Use the appropriate command:

- For GPU versions of the software:

```
[elemental@hostname ~]$ sudo sh ./elemental_production_server_2.18.0.12345.run --skip-all --start -xeula
```

- For CPU-only versions of the software:

```
[elemental@hostname ~]$ sudo sh ./elemental_production_server_cpu_2.18.0.12345.run --skip-all --start -xeula
```

Where:

--skip-all skips all the prompts. There is no need to view prompts about configuration because when you restore the database to the node, all the configuration data is copied over and overwrites any configuration data already on the node.

--start restarts the software after installation.

--xeula skips the display of the license agreement. There is no need to view this prompt because you have previously accepted the agreement.

3. When the installation is complete, restart the node:

```
[elemental@hostname ~] sudo reboot
```

Updating firmware

Step 1: Update the firmware

To update the BIOS firmware and the BMC firmware (IPMI for SuperMicro, iDRAC for Dell), check with the firmware manufacturers for the following information:

- The latest firmware versions.
- Information about which BIOS firmware and BMC firmware versions are compatible with each other.

To update the firmware, follow the manufacturer's instructions.

Step 2: Reboot the chassis

After you update the firmware, you must perform a cold reboot.

- You can physically power the appliance on and off.
- Or if you don't have physical access to the appliance, you can use `ipmiutil` to perform a cold reboot. This method resets the System Management processor. It doesn't reset the chassis.

This method is not considered to be a clean restart. This cold power cycle cuts the power quickly and discharges all remnant power in its components. Don't use this method as a standard way of rebooting.

To reset using `ipmiutil`

This procedure takes about 5 minutes. Don't be tempted to skip this reset because if you do, the node might not run properly.

1. Stop the node from the web interface or command line of the AWS Elemental software.
2. From the Linux prompt, use the **elemental** user to start a remote terminal session to the AWS Elemental Server node.
3. Run the following command:

```
[elemental@hostname ~]$ sudo ipmiutil power -k
```

4. Wait 30 seconds. Then run the following command continually until the output shows a state that includes the codes S0 or 2a:

```
[elemental@hostname ~]$ watch -n 5 "sudo ipmiutil health | grep 'Power State'"
```

5. Press **Ctrl-C** to exit watch.
6. Run the following command:

```
[elemental@hostname ~]$ sudo sync
```

7. Run the following power cycle command:

```
[elemental@hostname ~]$ sudo ipmiutil power -c
```

This command turns off and turns on the appliance and terminates the SSH connection. The RAID status might show as *Verify* during the boot sequence. This is normal.

8. You now restart the node, if you want the node to resume activity.

Installing RHEL 9

This section provides instructions to install RHEL 9 on a Dell chassis and on an SuperMicro chassis.

Topics

- [Installing on a Dell](#)
- [Installing on an SuperMicro](#)

Installing on a Dell

You can install RHEL 9 on a Dell chassis either from the iDRAC interface or using a USB stick.

Install using the iDRAC interface

Get Ready

1. Make sure that there are no physical USB drives plugged into the system.
2. Make sure that you are at a workstation that has direct access to the network that the iDRAC interface is on. (So don't use a VPN connection.)

3. Log into iDRAC through the web interface. Use an administrative username and password.
4. Launch the Virtual Console. On the main menu, select **Virtual Media**. On the next screen, select **Connect Virtual Media**. The **Virtual Media** screen appears.
5. In the **Map CD/DVD** section, in **Image File**, click **Choose File**. In the window that appears, navigate to the kickstart .iso file, select it, and click **Open**. The **Image File** field in the **Virtual Media** screen now specifies the image file.
6. Click **Map Device**. Then at the bottom of the screen, click **Close**.

The kickstart .ISO image file is now mapped to the virtual CD/DVD drive.

1. On the main menu of the Virtual Console, click **Boot**. On the **Boot Controls** list, click **Virtual CD/DVD/ISO**. Then at the **Confirm Boot Action** prompt, click **Yes**.
2. On the main menu of the Virtual Console, click **Power**, then click **Reset System (warm boot)**, and at the **Confirm** prompt, click **Yes**.

The system reboots into the kickstart .iso. Lines of text appear, and finally the prompt **Enter the server complete hostname** appears.

Install the operating system

1. At the **Enter the server complete hostname** prompt, enter the hostname that already applies to this node, then press **Enter**. The installation starts.
2. When the installation is complete, press **Enter** to quit and reboot.
3. You can now install any third-party packages. To obtain these packages, see [the section called "RPM repository"](#).

USB stick

1. Make sure that you have created a boot USB drive. See [the section called "Boot USB drive — create"](#).
2. Insert the USB drive into an available USB port. You might need to press **F2** while booting in order to select the boot device. The recovery (kickstart) screen appears.
3. Enter the hostname that already applies to this node, then press **Enter**. The installation starts.
4. When the installation is complete, remove the USB drive from the system and store it in a secure location.

5. Then on the screen, press the reboot button shown or press the **Enter** key.
6. You can now install any third-party packages. To obtain these packages, see [the section called "RPM repository"](#).

Installing on an SuperMicro

You install RHEL 9 on a SuperMicro chassis from the IPMI interface.

1. Install the Java applet and change the security level. See [the section called "Step 1: Install Java applet"](#).
2. Make sure that there are no physical USB drives plugged into the system.
3. Make sure that you are at a workstation that has direct access to the network that the IPMI interface is on.

Note

Don't use a VPN connection.

4. Copy the ISO file for RHEL 9 to your laptop.
5. Open the IPMI remote console viewer. On the main menu, choose **Virtual Media** or **Media**, then choose **Virtual Storage/Virtual Media Wizard**.
6. Choose **CD/ISO media** and browse to the ISO that you want to use. Choose **Connect/Plug in**.
7. Reboot the system. The image should start to boot.

If the image does not start to boot, click the **F11** key while the splash screen is displaying. Then when the **Please select boot device** prompt appear, choose **UEFI: Virtual CDROM**. Move this item to the top of the list by pressing the **+** key repeatedly.

8. The installer starts. At the prompt, enter the hostname of the appliance and press **Enter**. The installation starts and takes 20 to 30 minutes.
9. When the installation completes, press the **Enter** key to reboot.
10. Plug Out the ISO before it reboots or it takes you back into the kickstart menu.
11. You can now install any third-party packages. To obtain these packages, see [the section called "RPM repository"](#).

Working with RPM repository

AWS Elemental maintains an RPM repository for use with RHEL 9. The repository contains the following types of third-party packages:

- Packages that are stored in the Red Hat BaseOS repository, and that are required to run AWS Elemental software.
- Packages that are stored in the Red Hat AppStream repository, that aren't required but that you want to include.

For more information about the packages that you must obtain from the AWS Elemental RPM repository, and for instructions about configuring the repository, see the knowledge base article [Advisory](#).

Document history for migration guide

The following table describes the main changes to this guide.

Change	Description	Date
New guide	First release of this guide	June 7, 2024