



Migration Guide

AWS Elemental Conductor Live



AWS Elemental Conductor Live: Migration Guide

Copyright © 2025 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
Important notes	2
Standard cluster migration	3
Plan maintenance windows	4
Step A: Get ready	5
Read the essential notes	5
Modify your automation system for HTTPS	5
Verify installer type	5
Create a boot USB drive	6
Obtain the default element password	6
Verify space on each node	6
Step B: Prepare each node	7
Upgrade to the latest 3.25 minor version	7
Verify access to the BMC on the appliances	7
Update firmware	7
Make a note of node assignments	8
Make a note of router information	8
Move custom files	8
Step C: Tear down the cluster	9
Step D: Create backups	9
Step E: Upgrade nodes	10
Upgrade the worker nodes	10
Upgrade the Conductor nodes	10
Step F: Rebuild the cluster	11
Step F1: Add and configure the worker nodes	11
Step F2: Reconfigure routers	12
Step F3: Add Conductor node	12
Step F4: Final steps	12
Split cluster migration	13
Plan maintenance windows	14
Step A: Get ready	15
Read the essential notes	15
Modify your automation system	16
Verify installer type	17

Create a boot USB drive	17
Obtain the default elemental password	17
Verify space on each node	17
Step B: Prepare each node	18
Upgrade to the latest 3.25 minor version	18
Verify access to the BMC on the appliances	19
Update firmware	19
Make a note of node assignments	19
Make a note of router information	20
Move custom files	20
Step C: Split the cluster	20
Step D: Upgrade node X	22
Step E: Upgrade the worker nodes	23
Step E1: Remove the worker node	23
Step E2: Upgrade the worker node	24
Step F: Upgrade node Y	26
Step G: Add node Y to cluster	26
Tasks	28
Boot mode — UEFI	29
Switching to UEFI on a Dell	29
Switching to UEFI on a SuperMicro	31
Boot mode — Legacy	33
Switching to Legacy on a Dell	33
Switching to Legacy on a SuperMicro	35
Boot USB drive — create	37
Cluster — add Conductor	37
Cluster — add worker node	38
Step A: Add worker nodes to the cluster	38
Step B: Add worker nodes to redundancy groups	39
Step C: Add node assignments	39
Cluster — enable HA or disable HA	40
Cluster — enable user authentication	41
Cluster — remove Conductor node	42
Cluster — remove worker node	43
Step A: Stop running channels	43
Step B: Remove node assignments	43

Step C: Remove node from redundancy group	44
Step D: Remove node from cluster	44
Cluster — restart channels	44
Database — back up	45
About the backup process	45
Step A: Verify hostnames	46
Step B: Download the lifeboat script	46
Step C: Create the backup	47
Database — restore	50
Step A: Perform the restore	50
Step B: Perform manual restore tasks	52
Result of the restore	53
Conductor Live — install	56
Elemental Live — install	57
Firmware — update	59
Step A: Update the firmware	59
Step B: Reboot the chassis	59
RHEL 9 — install	60
Installing on a Dell	60
Installing on a SuperMicro	62
RPM repository	64
Document history	65

About this guide

This guide describes how to upgrade a Conductor cluster consisting of AWS Elemental Conductor Live nodes and worker nodes (AWS Elemental Live and AWS Elemental Statmux). It describes how to perform an upgrade to a version in the range of 3.26.1 (2.26.1) to 3.26.5 (2.26.5) of the AWS Elemental software.

This special guide exists because an upgrade to a version in that range requires that you install the RHEL 9 version of the Linux operating system.

See [AWS Elemental Conductor Live Upgrade Guide](#) if the following situations apply:

- You are upgrading to a version earlier than 3.26.1.
- You have migrated to a 3.26 version, and you now want to upgrade to a higher 3.26.x version to 3.27.x or higher.

Note

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

Important notes

Following is some key advice about migration.

Monitor ongoing migration issues

Before you start the migration, read the [current Release Notes](#). The Essential Notes section of the release notes includes topics about recently discovered issues.

For each issue, the release notes initially describe the issue and specify the conditions for continuing with the migration. These issues might affect assets such as the lifeboat script that you use to create a database backup. We can often fix issues without waiting for a new release of the AWS Elemental software.

When an issue is fixed, we either revise the current release notes or include new information in the release notes for the new software release.

Test in a lab

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.

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

This procedure describes how to take the nodes in a AWS Elemental Conductor Live cluster from a version below x.26.0 and migrate them to version x.26.1 or higher. The nodes in the cluster might be the following:

- A primary Conductor node and, optionally, a secondary Conductor node.
- AWS Elemental Live nodes.
- AWS Elemental Statmux 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 *standard* procedure, you take all nodes offline for a while and tear down the cluster. Therefore there is a period when no channels or MPTSes are running on any of the worker nodes. If you need to decrease the downtime, you might want to perform the [split Conductor migration](#) instead.

In this procedure, we show how to upgrade the cluster from Conductor Live version 3.25.5 (worker nodes version 2.25.5) to version 3.26.x (2.26.x). Modify the commands you enter to match your versions.

Important

You must upgrade all the nodes in the cluster to an x.26 version. You can't, for example, set up the cluster so that Conductor Live is running 3.25.x and the workers (or some of the workers) are running 2.26.x.

Topics

- [Plan maintenance windows for migrating an AWS Elemental Conductor Live cluster](#)

- [Step A: Get ready to migrate an AWS Elemental Conductor Live cluster](#)
- [Step B: Prepare each AWS Elemental Conductor Live node for migration](#)
- [Step C: Tear down the cluster](#)
- [Step D: Create backups](#)
- [Step E: Upgrade nodes](#)
- [Step F: Rebuild the cluster](#)

Plan maintenance windows for migrating an AWS Elemental Conductor Live cluster

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

First phase

You can perform the tasks in [the section called “Step A: Get ready”](#) 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 the cluster”](#)
- [the section called “Step D: Create backups”](#)
- [the section called “Step E: Upgrade nodes”](#)
- [the section called “Step F: Rebuild the cluster”](#)

These steps upgrade all the nodes in one maintenance window. 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 Live cluster

Read the essential notes

Refer to the essential notes in the [current Release Notes](#) to identify key changes to the behavior of the Conductor Live 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.

Follow the instructions in the Release Notes to prepare for these changes. Make a note of all the decisions that you make, particularly:

- If you have Elemental Live nodes that have a Mellanox NIC, you must decide whether you will include the `-skip-mellanox` option in the install command on the node.
- If you have Elemental Live nodes that handle SMPTE 2110 inputs or outputs, you must a new license for each. The procedure for obtaining a new license is described in the essential notes.

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.
- In addition, if you decide not to provide custom certificates, use HTTPS with the `-k` option.

Verify installer type

For Elemental Live, the software installer that you use depends on whether you have GPU-accelerated software type, or CPU-only. To determine the type of software, look at any page on the web interface of each worker node. At the top, look for these icons:

- CPU and GPU icons: the software is *GPU-accelerated*.

- CPU icon only: the software is *CPU-only*.

Make a note of the installer type for every Elemental Live node in the cluster.

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”](#).

Obtain the default element password

When you install RHEL 9, the *elemental* password will be reset to the global default value. When you install the AWS Elemental software, you will need to enter this default password.

If your organization had previously kickstarted an AWS Elemental appliance, someone in your organization might have the global default password. Note that this global default password isn't the same as the unique default password that applied when you first obtained the appliance.

If you need help finding the password, go to the [AWS Elemental Support Center](#), and read the Knowledge article [Unique passwords for AWS Elemental appliances](#) or open a case.

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 5 MB backup directory and a 2 MB 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 Live 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 3.25 minor version

To upgrade to version 3.26.x or higher (version 2.26.x or higher for workers nodes), the software currently installed on the node must be version 3.25.5 or higher (or 2.25.5).

- If the Conductor nodes are currently on version 3.25.5 or higher (or 2.25.5 or higher), you don't have to upgrade to a higher patch.
- If the Conductor nodes are on one major version and the worker nodes are on a different major version, upgrade all nodes to the same major version. Normally, there is a rule that Conductor Live can be a lower major version than the workers. This rule doesn't apply when upgrading to 3.26.1 or higher. All the nodes in the cluster must be running version 3.25.5 (or 2.25.5) or higher.
- If the Conductor nodes are on a version below 3.25.5, upgrade to the current highest patch version.

To upgrade to 3.25.5 or higher (or 2.25.5), see the [AWS Elemental Conductor Live Upgrade Guide](#).

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 install it.

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

Update firmware

Both the BIOS firmware and the BMC firmware (IPMI for SuperMicro, iDRAC for Dell) must be at the latest versions that have been qualified by AWS Elemental. They must be at the latest versions

before you can set the boot mode to UEFI, as part of [upgrading each node](#). To obtain the versions, go to the [AWS Elemental Support Center](#), and read the Knowledge article [Latest AWS Elemental Qualified Remote Management and BIOS Firmware](#) or open a case.

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 3.25.x (and 2.25.x) 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”](#).

Make a note of node assignments

Before you upgrade any worker node, you must make a note of the channels that are assigned to this node. You will use this information to restart the channels, after you've completed the migration.

1. On the web interface for the primary Conductor node, access the **Channels** screen.
2. Filter the information on the screen to show one node. Then make a note of all the channels that are assigned to that node.
3. Repeat for each node.

Make a note of router information

This information applies if the cluster includes nodes that connected to an SDI input using a router. After you upgrade, the cluster will still have information about the SDI inputs and about the router, but it will be missing the mapping from the inputs to the router.

In order to reconfigure the information accurately, make a note of the current configuration. For more information see the information about configuring routers in the *Reference: Configure connectivity* section of the [AWS Elemental Conductor Live Configuration Guide](#).

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 Live or Elemental Live 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 the cluster

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

1. Disable high availability (HA) on the cluster. You must disable HA before you can remove the secondary Conductor node. See [the section called "Cluster — enable HA or disable HA"](#). After you disable HA, only the primary Conductor can control the cluster.

If you don't have HA enabled, skip this step.

2. Remove the secondary Conductor node from the cluster. You must remove the secondary node so that when you shut down the primary Conductor node, control doesn't fail over to the secondary Conductor node. See [the section called "Cluster — remove Conductor node"](#).

If you have only one Conductor node, skip this step.

- 3.
4. Remove the workers from the cluster. You perform this action from the primary Conductor node. See [the section called "Cluster — remove worker node"](#).

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, don't create channels, don't create new node assignments for any channel, and so on.

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

Step E: Upgrade nodes

Upgrade the worker nodes

Perform these steps on each worker node in the cluster, after you've removed all the nodes from the cluster.

Before you start, make sure that you have performed the tasks in [the section called “Step B: Prepare each node”](#).

1. Create a backup of the database on the node. See [the section called “Database — back up”](#).
2. Set boot mode on the node to UEFI. See [the section called “Boot mode — UEFI”](#).
3. Perform a kickstart to upgrade the operating system to RHEL 9. See [the section called “RHEL 9 — install”](#).
4. Install Elemental Live version 2.26.x on the node. See [the section called “Elemental Live — install”](#).
5. Restore the database onto the node. See [the section called “Database — restore”](#).
6. Install new licenses.

If a specific node handles SMPTE 2110 inputs or outputs, you should have obtained a new license that includes the SMPTE 2110 add-on package. (The procedure for obtaining a new license is described in the essential notes in the [current Release Notes](#).) To deploy the license, see the section about configuring licenses in the [AWS Elemental Live Configuration Guide](#).

Each worker node now has the new operating system and software installed, and it restored to its former database, including most of its configuration data.

Upgrade the Conductor nodes

Perform these steps on both Conductor nodes.

1. Set boot mode on the node to UEFI. See [the section called “Boot mode — UEFI”](#).
2. Perform a kickstart to upgrade the operating system to RHEL 9. See [the section called “Database — back up”](#).

3. Install Conductor Live version 3.26.1 on the node. See [the section called “Conductor Live — install”](#).
4. Restore the database onto each node. See [the section called “Database — restore”](#).

The primary Conductor node now has all the configuration information about the cluster. This means that when you add worker nodes back into the cluster, all the information about redundancy groups, for example, is present on the primary Conductor node. You don't have to set it up again.

5. If you moved custom files to a safe location as part of your preparation, you can now copy these files back to their original location.

Step F: Rebuild the cluster

Step F1: Add and configure the worker nodes

Follow these steps on each worker node:

1. Add the worker node to the cluster. Then add the node to its redundancy group. Finally, assign channels to the node, using the [list of channel assignments](#)) that you created. See [the section called “Cluster — add worker node”](#).
2. Set up worker features.

Configuration information about some features isn't included in the database, so you must set them up again. The features are:

- Enabling OCR conversion to handle captions
- Disabling RTMP inputs in order to release processing resources
- Setting the maximum for virtual input switching

See the section about features in the [AWS Elemental Live Configuration Guide](#).

Note

You might want to voluntarily change the configuration of one or more of the worker nodes.

We strongly recommend that you don't make any voluntary changes to the configuration until you have tested your workflows in the new setup.

Step F2: Reconfigure routers

This section applies if the cluster previously included nodes that connected to an SDI input using a router. After you upgrade, the cluster still has information about the SDI inputs and about the router, but it is missing the mapping from the inputs to the router. You must reconfigure this information. You should have [made a note of the configuration](#).

For more information see the information about configuring routers in the *Reference: Configure connectivity* section of the [AWS Elemental Conductor Live Configuration Guide](#).

See the information about configuring routers in the *Reference: Configure connectivity* section of the [AWS Elemental Conductor Live Configuration Guide](#). Specifically, start at the *Complete the Router Output Mappings* step in that procedure.

Step F3: Add Conductor node

If the cluster included a secondary Conductor, add it to the cluster and to the Conductor Live redundancy group. See [the section called "Cluster — add Conductor"](#).

Step F4: Final steps

1. Start channels. You can start the channels that were previously running. See [the section called "Cluster — restart channels"](#).

If you have only one Conductor, the upgrade process is now complete.

2. If you have two Conductors, re-enable HA. The secondary Conductor synchs itself to the primary Conductor. The upgrade process is now complete.

Performing a split cluster migration

This procedure describes how to take the nodes in a AWS Elemental Conductor Live cluster from a version below x.26.0 and migrate them to version x.26.1 or higher. The nodes in the cluster might be the following:

- Both a primary Conductor node and a secondary Conductor node.
- AWS Elemental Live nodes.
- AWS Elemental Statmux 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 Live version 3.25.5 (worker nodes version 2.25.5) to version 3.26.x (2.26.x). Modify the commands you enter to match your versions.

Important

You must upgrade all the nodes in the cluster to an x.26 version. You can't, for example, set up the cluster so that Conductor Live is running 3.25.x and the workers (or some of the workers) are running 2.26.x.

How a split cluster migration works

In this procedure, which you typically perform over several maintenance windows, you remove the primary Conductor node from the cluster. The secondary Conductor node takes control of this cluster (the original cluster). You upgrade the primary Conductor to the new operating system and Conductor Live software. This node is then ready to control a new cluster (the new cluster).

You then remove the worker nodes from the original cluster, upgrade them, add them to the new cluster, and restart their channels. Eventually, the secondary Conductor has no worker nodes. At

that point, you upgrade that node and add it to the new cluster, then enable HA. You now have the cluster running again, with all the nodes running the new version of the AWS Elemental software on the new operating system.

If this procedure seems unnecessarily complicated and you don't mind losing all processing for a few hours, then you might want to perform the [standard migration](#) instead.

Topics

- [Plan maintenance windows](#)
- [Step A: Get ready](#)
- [Step B: Prepare each node](#)
- [Step C: Split the cluster](#)
- [Step D: Upgrade node X](#)
- [Step E: Upgrade the worker nodes](#)
- [Step F: Upgrade node Y](#)
- [Step G: Add node Y to cluster](#)

Plan maintenance windows

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

First phase

You can perform the tasks in [the section called "Step A: Get ready"](#) 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"](#)

Remaining phases

Perform all the following tasks on several maintenance windows.

- [the section called "Step C: Split the cluster"](#)
- [the section called "Step D: Upgrade node X"](#)
- [the section called "Step E: Upgrade the worker nodes"](#)
- [the section called "Step F: Upgrade node Y"](#)
- [the section called "Step G: Add node Y to cluster"](#)

Follow these rules:

- Perform step C in one window, and perform step D in the next window. Or combine steps C and D in one window.

You could also perform step D outside of a maintenance window (but before you perform step E). You can do this because node X is no longer active — it's not part of the working cluster.

- Perform step E on as many nodes as you can in one maintenance window. In each window, you remove one or more workers nodes from the original cluster and put them in the new cluster. Eventually, all the worker nodes will be in the new cluster.

A consideration when you decide which worker nodes to migrate at the same time is whether you have routers in the cluster. Identify all worker nodes that use the SDI inputs attached to the router. Then plan to migrate those worker nodes together. After you've migrated all those nodes, you will be able to reconfigure the new cluster for that router.

- Then perform step F in one window, and step G in the next window. Or combine steps F and G in one window.

You could also perform step F outside of a maintenance window (but before you perform step G). You can do this because node Y is no longer active — it's not part of the working cluster.

Step A: Get ready

Read the essential notes

Refer to the essential notes in the [current Release Notes](#) to identify key changes to the behavior of the Conductor Live 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.

Follow the instructions in the Release Notes to prepare for these changes. Make a note of all the decisions that you make, particularly:

- If you have Elemental Live nodes that have a Mellanox NIC, you must decide whether you will include the `-skip-mellanox` option in the install command on the node.
- If you have Elemental Live nodes that handle SMPTE 2110 inputs or outputs, you must a new license for each. The procedure for obtaining a new license is described in the essential notes.

Modify your automation system

In order for processing to continue after you split the cluster, you must make changes to the way your automation system and your web interface users connect to the cluster.

Set up automation system to use physical IP addresses

After you split the cluster, the primary Conductor is accessible only via its physical IP address, while the secondary Conductor is accessible via the VIP (and also via its physical IP address). By comparison, when the cluster is set up in the normal way, the primary Conductor is accessible via the VIP, while the secondary Conductor isn't accessible.

Therefore you must make sure that your automation system can connect to the primary Conductor via that node's physical IP address.

Plan the new connection for the primary Conductor. Then switch over to the new way of connecting right after you [split the cluster](#).

Set up users to use physical IP addresses

The users of the Conductor Live web interface must enter the physical IP address of the appropriate Conductor node. Remember to provide them with that address at the same time as you switch over the automation system.

Set up automation system to use permalinks

You must make sure that your automation system isn't referencing resources using an ID. For example, if you use profile IDs to reference profiles, you should instead use permalinks. Perform this setup before you start the first maintenance window.

Note that when you make this change, the cluster will be in line with existing AWS Elemental best practice.

Set up automation system to use 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.
- In addition, if you decide not to provide custom certificates, use HTTPS with the `-k` option.

Verify installer type

For Elemental Live, the software installer that you use depends on whether you have GPU-accelerated software type, or CPU-only. To determine the type of software, look at any page on the web interface of each worker node. At the top, look for these icons:

- CPU and GPU icons: the software is *GPU-accelerated*.
- CPU icon only: the software is *CPU-only*.

Make a note of the installer type for every Elemental Live node in the cluster.

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"](#).

Obtain the default elemental password

When you install RHEL 9, the *elemental* password will be reset to the global default value. When you install the AWS Elemental software, you will need to enter this default password.

If your organization had previously kickstarted an AWS Elemental appliance, someone in your organization might have the global default password. Note that this global default password isn't the same as the unique default password that applied when you first obtained the appliance.

If you need help finding the password, go to the [AWS Elemental Support Center](#), and read the Knowledge article [Unique passwords for AWS Elemental appliances](#) or open a case.

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 5 MB backup directory and a 2 MB 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 node

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 3.25 minor version

To upgrade to version 3.26.1 or higher (version 2.26.x or higher for workers nodes), the software currently installed on the node must be version 3.25.5 or higher (or 2.25.5).

- If the Conductor nodes are currently on version 3.25.5 or higher (or 2.25.5 or higher), you don't have to upgrade to a higher patch.
- If the Conductor nodes are on one major version and the worker nodes are on a different major version, upgrade all nodes to the same major version. Normally, there is a rule that Conductor Live can be a lower major version than the workers. This rule doesn't apply when upgrading to 3.26.1 or higher. All the nodes in the cluster must be running version 3.25.5 (or 2.25.5) or higher.
- If the Conductor nodes are on a version below 3.25.5, upgrade to the current highest patch version.

To upgrade to 3.25.5 or higher (or 2.25.5), see the [AWS Elemental Conductor Live Upgrade Guide](#).

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 install it.

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

Update firmware

Both the BIOS firmware and the BMC firmware (IPMI for SuperMicro, iDRAC for Dell) must be at the latest versions that have been qualified by AWS Elemental. They must be at the latest versions before you can set the boot mode to UEFI, as part of [upgrading each node](#). To obtain these versions, go to the [AWS Elemental Support Center](#), and read the Knowledge article [Latest AWS Elemental Qualified Remote Management and BIOS Firmware](#) or open a case.

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 3.25.5 (2.26.x) 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”](#).

Make a note of node assignments

Before you upgrade any worker node, you must make a note of the channels that are assigned to this node. You will use this information to restart the channels, after you've completed the migration.

1. On the web interface for the primary Conductor node, access the **Channels** screen.
2. Filter the information on the screen to show one node. Then make a note of all the channels that are assigned to that node.
3. Repeat for each node.

Make a note of router information

This information applies if the cluster includes nodes that connected to an SDI input using a router. After you upgrade, the cluster will still have information about the SDI inputs and about the router, but it will be missing the mapping from the inputs to the router.

In order to reconfigure the information accurately, make a note of the current configuration. For more information see the information about configuring routers in the *Reference: Configure connectivity* section of the [AWS Elemental Conductor Live Configuration Guide](#).

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 Live or Elemental Live 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.

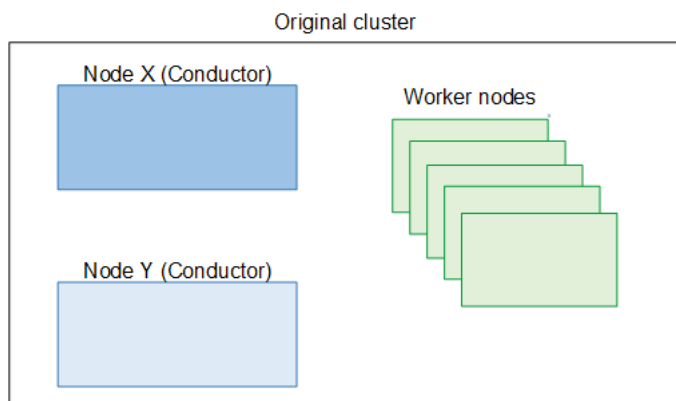
Step C: Split the cluster

Note

In this procedure, the Conductor nodes change their roles. Therefore, the node that is originally the primary node is called node X. The node that is originally the secondary node is called node Y.

To split the cluster into two clusters, you first fail over control to the secondary Conductor node (*node Y*), then you remove the primary Conductor node (*node X*).

You start with the original deployment of one cluster, which looks like the following diagram.



You end up with:

- The original cluster controlled by the second node. This cluster is using the older version of the software. For example, Conductor Live 3.25.5 and Elemental Live 2.25.5.
- A new cluster that right now consists only of the primary Conductor node. The Conductor is using 3.26.x.

During this time, the worker nodes (on the original cluster) continue processing without interruption.

1. Disable high availability (HA) on the cluster. See [the section called “Cluster — enable HA or disable HA”](#).
2. Create a backup of the data on node X. See [the section called “Database — back up”](#).

⚠ Important

Now that you have made a backup of this Conductor node, don't make any changes to this Conductor node or to cluster until you've finished this migration process. Don't change the setup of the Conductor node, don't create channels, don't create new node assignments for any channel, and so on.

3. On the web interface for node X, re-enable HA. See [the section called “Cluster — enable HA or disable HA”](#).
4. Perform the commands in the next few steps from the command line.
 - a. From the Linux prompt, use the *elemental* user to start a remote terminal session with node X.
 - b. Fail over to the secondary Conductor node. To fail over, enter these commands from the CLI on node X:

```
sudo systemctl stop elemental_se  
sudo systemctl stop serf
```

Node Y is now controlling the cluster.

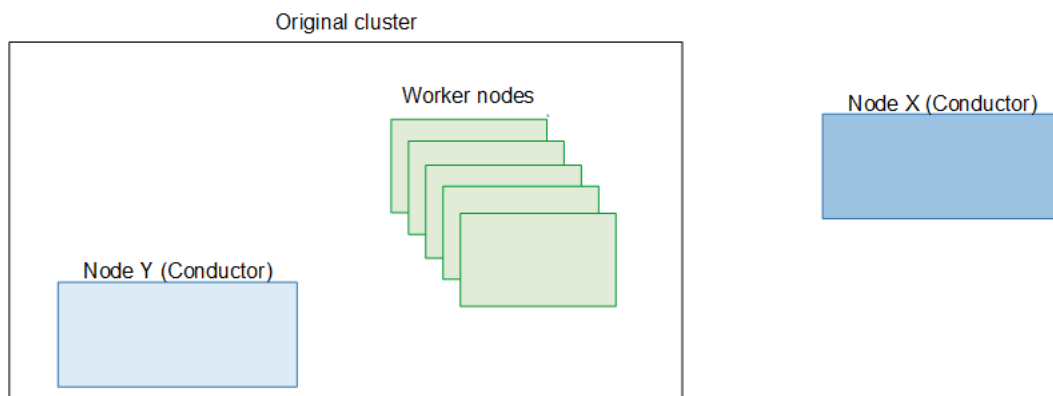
5. To enable user authentication on node Y, enter this command from the CLI:

```
sudo systemctl restart httpd
```

6. On the web interface for node Y, disable HA. See [the section called “Cluster — enable HA or disable HA”](#).
7. Now that HA is disabled, you can remove node X from the cluster. See [the section called “Cluster — remove Conductor node”](#).
8. On the CLI, enter the following commands to push the VIP into node Y:

```
sudo systemctl start keepalived  
sudo systemctl enable keepalived
```

Node Y is now the Conductor that is controlling the cluster. Note that you no longer have Conductor redundancy in the cluster.



Step D: Upgrade node X

Now that node X is not controlling the cluster, you can upgrade it.

1. Set boot mode on the node to UEFI. See [the section called “Database — back up”](#).
2. Perform a kickstart to upgrade the operating system to RHEL 9. See [the section called “RHEL 9 — install”](#).

3. Install Conductor Live version 3.26.1 on the node. See [the section called “Conductor Live — install”](#).
4. In [the section called “Step C: Split the cluster”](#), you created a backup of the database on node X. You can now restore the backup onto the node. See [the section called “Database — restore”](#).

As a result of restoring the database, this Conductor node is now configured as it was before you removed it from the cluster. Specifically, it has the data relating to the channels, MPTs, node assignments for channels and MPTs, user setup, redundancy groups, and cluster members.

This means that when you add worker nodes back into the cluster (in the next section), there is less worker node configuration required compared to when you set up the cluster for the very first time.

5. If you moved custom files to a safe location as part of your preparation, you can now copy these files back to their original location.

Step E: Upgrade the worker nodes

You must now remove each worker from the original cluster (that is controlled by node Y), upgrade it, and add it to the new cluster (that is controlled by node X).

Note

Perform the following steps on one node at a time.

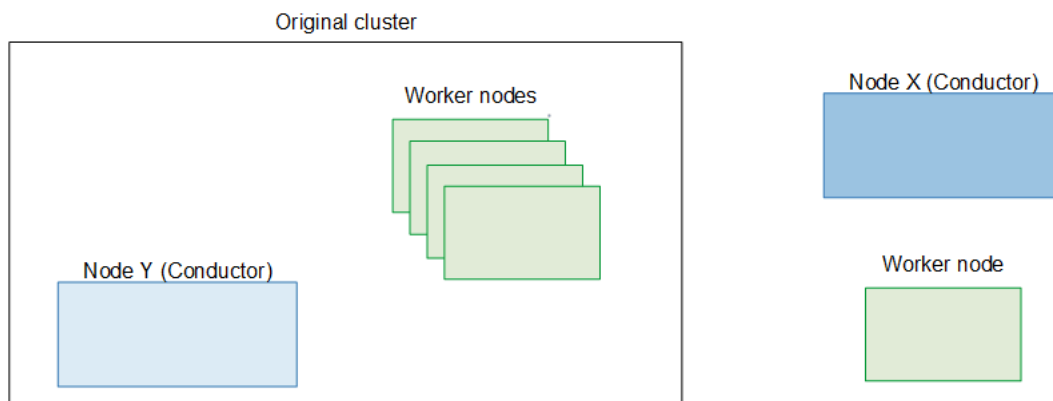
Typically, you move over as many worker nodes as you can in one maintenance window. Keep in mind that the workers nodes that you don't move over will still be processing.

Step E1: Remove the worker node

Before you start, make sure that you have performed the tasks in [the section called “Step B: Prepare each node”](#).

Remove the worker from the cluster. You perform this action from node Y (the Conductor node that is controlling the original cluster). See [the section called “Cluster — remove worker node”](#).

The deployment now looks like the following diagram. Note that the worker node is standalone. It isn't being controlled by any Conductor.



Step E2: Upgrade the worker node

1. Create a backup of the database on the node. See [the section called “Database — back up”](#).
2. Set boot mode on the node to UEFI. See [the section called “Database — back up”](#).
3. Perform a kickstart to upgrade the operating system to RHEL 9. See [the section called “RHEL 9 — install”](#).
4. Install Elemental Live version 2.26.1 on the node. See [the section called “Elemental Live — install”](#).
5. Restore the database onto each node. See [the section called “Database — restore”](#).
6. Install new licenses.

If this worker node handles SMPTE 2110 inputs or outputs, you should have obtained a new license that includes the SMPTE 2110 add-on package. (The procedure for obtaining a new license is described in the essential notes in the [current Release Notes](#).) To deploy the license, see the section about configuring licenses in the [AWS Elemental Live Configuration Guide](#).

Step E3: Add the worker node to the new cluster

1. Add the worker node to the cluster. Then add the node to its redundancy group. Finally, assign channels to the node, using the [list of channel assignments](#)) that you created. See [the section called “Cluster — add worker node”](#).
2. If the original cluster is configured for user authentication, then log onto node X and enable node authentication, in order to let operators work on the newly added worker node. Note that you must enable node authorization each time you add a worker node. For more information, see [the section called “Cluster — enable user authentication”](#).

3. This step applies if the original cluster design included nodes that connected to an SDI input using a router.

In the section about planning maintenance windows, we recommended that you migrate all worker nodes that rely on the router at the same time. If you have followed this advice, and you are now moving those nodes, you need to reconfigure the router.

You need to reconfigure the router because the new cluster has information about the SDI inputs and about the router, but it is missing the mapping from the inputs to the router. You must reconfigure this information now. You should have [made a note of the previous mappings](#).

You perform this work on the primary Conductor (node X). See the information about configuring routers in the *Reference: Configure connectivity* section of the [AWS Elemental Conductor Live Configuration Guide](#). Specifically, start at Step E in that procedure.

4. Set up worker features.

Configuration information about some features isn't included in the database. If one of these features applies to the current node, you must set up the feature again. The features are:

- Enabling OCR conversion to handle captions
- Disabling RTMP inputs in order to release processing resources
- Setting the maximum for virtual input switching

You perform this work on the primary Conductor (node X). See the section about features in the [AWS Elemental Live Configuration Guide](#).

 **Note**

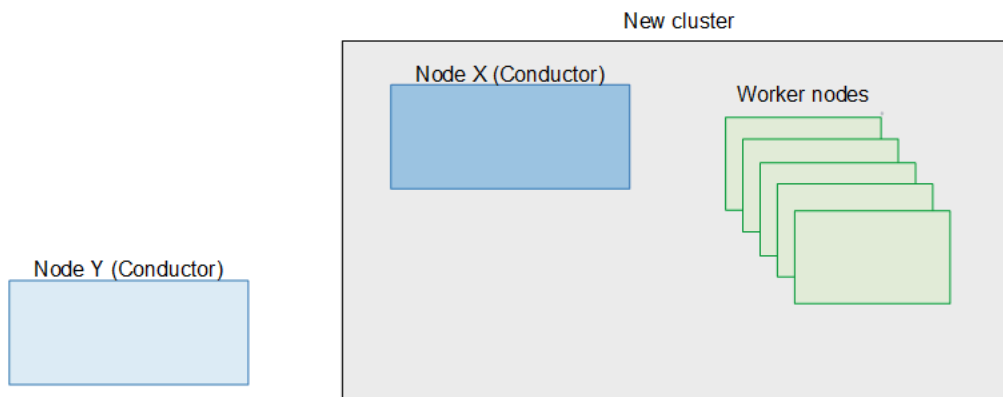
You might want to take this opportunity to make other configuration changes on one or more worker nodes.

We strongly recommend that you don't make these changes to the configuration until you have tested your workflows in the new setup.

5. Start channels on this node. You assigned the channels to this node when you added the node to the cluster. Therefore, the channels that were running on this node in the original cluster will restart on this node in the new cluster. See [the section called “Cluster — restart channels”](#).

Step F: Upgrade node Y

After you remove the last worker node, the original cluster no longer exists. Node Y is no longer controlling any worker nodes. The deployment now looks like the following diagram.



You can now upgrade node Y. Perform all the following steps on node Y.

1. On node Y, which will soon be the secondary Conductor in the new cluster, you must clean the database. Enter the configure commands as follows:

```
cd /opt/elemental_se; sudo ./configure -xeula --skip-all --cleandb --start
```

2. Create a backup of the data on node Y. See [the section called “Database — back up”](#).
3. Set boot mode on the node to UEFI. See [the section called “Boot mode — UEFI”](#).
4. Perform a kickstart to upgrade the operating system to RHEL 9. See [the section called “RHEL 9 — install”](#).
5. Install Conductor Live version 3.26.x on the node. See [the section called “Conductor Live — install”](#).
6. Restore the backup onto the node. See [the section called “Database — restore”](#).
7. If you moved custom files to a safe location as part of your preparation, you can now copy these files back to their original location.

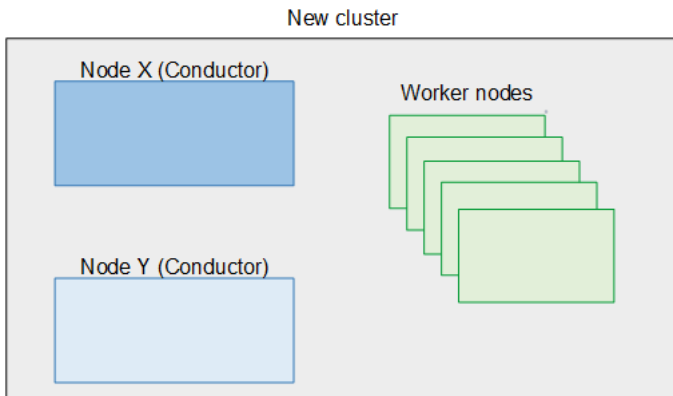
Step G: Add node Y to cluster

1. Add node Y to the new cluster, and add it to the HA group in that cluster. Node Y will assume the role of secondary Conductor. See [the section called “Cluster — add Conductor”](#).
2. Re-enable HA on node X. See [the section called “Cluster — enable HA or disable HA”](#).

The secondary Conductor (node Y) synchs itself to the primary Conductor (node X). See [the section called “Cluster — enable HA or disable HA”](#).

The upgrade process is now complete. Node X is acting as the primary Conductor, node Y is acting as the secondary Conductor, and all the nodes are running the new software version.

The deployment now looks like the following diagram.



Tasks for migrating an AWS Elemental Conductor Live 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*](#)
- [*Split cluster migration*](#)

Topics

- [Switching boot mode to UEFI](#)
- [Switching boot back to Legacy](#)
- [Create a boot USB drive](#)
- [Adding the secondary Conductor node to the cluster](#)
- [Adding worker nodes](#)
- [Enabling or disabling high availability \(HA\)](#)
- [Enabling user authentication](#)
- [Removing a Conductor node from the cluster](#)
- [Removing a worker node from the cluster](#)
- [Restarting channels](#)
- [Backing up data](#)
- [Restoring the database](#)
- [Installing Conductor Live on nodes](#)
- [Installing Elemental Live 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 A: 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 B: 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 A: 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

Only Dell servers support the ability to install RHEL 9 from a boot USB drive. SuperMicro servers don't support this ability.

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, read the Knowledge article [Creating Bootable Recovery \(kickstart\) Media](#).

Adding the secondary Conductor node to the cluster

Add the secondary node back to the cluster, and then to the redundancy group.

To avoid errors when you're adding the secondary Conductor back to the cluster, wait approximately three minutes after the upgrade before performing these steps. This wait ensures that the elemental_se service has restarted and is running.

Add the secondary node to the cluster

1. On the web interface for the primary Conductor node, choose **Cluster > Nodes**, then choose **Add Node**.
2. In the **Add Nodes to Cluster** pop-up, complete the node information for the secondary Conductor and choose **Add**.

Add the secondary node to the redundancy group

1. On the web interface for the primary Conductor node, choose **Cluster > Redundancy**, then select the Conductor Live redundancy group.
2. Choose **Add HA Nodes**.
3. In the **Add** pop-up, use the **Node** drop-down to select the secondary Conductor. Choose **Add**.

Adding worker nodes

Topics

- [Step A: Add worker nodes to the cluster](#)
- [Step B: Add worker nodes to redundancy groups](#)
- [Step C: Add node assignments](#)

Step A: Add worker nodes to the cluster

1. On the web interface for the primary Conductor node, choose the **Cluster** page, then choose **Nodes**.
2. On the **Nodes** page, choose **Add Node**.
3. In the **Add Nodes to Cluster** dialog, do one of the following:
 - In **Node IP Address/s**, enter the IP address or range of IP addresses for multiple nodes and choose **Add**.
 - If your network has a DNS server, search for the node by its hostname:
 - a. In **Lookup Node IP Address**, enter the hostname of the node that you're adding. You set the hostname during installation of the node.
 - b. Choose the **search** icon.
 - c. When Conductor Live displays the IP address that corresponds to the hostname, choose the plus sign beside the address to add it to the Node IP address list.
 - d. Choose **Add**.
4. Still in the **Add Nodes to Cluster** dialog, add each node that will be a part of the cluster.
5. Verify that the nodes have been added to the list on the **Nodes** page and the correct information is shown:
 - The **Status** is **Online**.
 - The **Elemental Product** is the correct type of node, either Conductor Live or a worker.

6. If any worker nodes contain SDI cards, import the devices so that the Conductor node knows about them. Do the following:
 - a. Still on the primary Conductor, go to **Nodes** page and find the first worker node that has an SDI card. Beside that node, choose the downward triangle and select **Import Devices**.

Conductor Live detects the device and adds the device's configuration to the Conductor Live database.

- b. Repeat for every worker node that has an SDI card.

If you don't import the devices, they won't appear in the Conductor Live web interface. You won't be able to specify these devices as video sources in a channel.

Step B: Add worker nodes to redundancy groups

1. On the web interface for the primary Conductor node, choose the **Cluster** page, then choose **Redundancy**. In the navigation bar, choose the Elemental Live redundancy group.
2. On the **Active Nodes** tab, choose **Add to Active**. In the **Nodes** field, select the nodes to add to the group.
3. Choose **Add**.
4. On the **Backup Nodes** tab, choose **Add to Backup Nodes**. In the **Nodes** field, select the nodes to add to the group. Make sure that you add nodes in order, starting with the reserve node with the highest priority.
5. Choose **Add**.
6. If you have multiple Elemental Live redundancy groups, repeat this procedure on each group.

Step C: Add node assignments

To assign channels back to a node, use the [list of channel assignments](#) that you created for the node.

1. On the web interface for the primary Conductor node, display the **Channels** page, then choose **Tasks**, then choose **Change channel node assignments**.
2. Select the channels that should all be assigned to one node, according to your list. Choose **Next**.
3. On the **Select a new node** page, in **New Node**, choose the node for these channel.
4. Choose **Next**, then choose **Process Now**.

Enabling or disabling high availability (HA)

To enable or disable HA in a cluster, follow these steps.

1. If you're using a virtual machine (VM), take a snapshot before disabling high availability. See the VMware VSphere help text for more information.
2. On the web interface for the primary Conductor node, choose the **Cluster** page, then choose **Redundancy**. In the **High Availability** field, select **Enable** or **Disable**.
3. Verify that high availability is enabled or disabled. From the Linux prompt, use the **elemental** user start a remote terminal session with the primary and with the secondary Conductor nodes.

Then, in the remote terminal session for each Conductor, enter the following command:

```
[elemental@hostname ~]$ tail -F /opt/elemental_se/web/log/conductor_live247.output
```

The conductor_live247.output log starts to scroll on the screen and shows messages as they are occurring.

If you enabled HA, watch for the following INFO lines on the primary Conductor node:

```
CONDUCTOR: Initializing environment
I, [2015-11-13T04:37:54.491204 #4978] INFO -- : Configuring the HA environment
I, [2015-11-13T04:37:54.660644 #4978] INFO -- : configuring keepalived
.
.
.
I, [2015-11-13T04:38:03.905069 #4978] INFO -- : Elemental Conductor is ready
```

If you disabled HA, watch for the following INFO lines on the primary Conductor node:

```
WARN -- : Disabling HA, elemental_se restarting...
.
.
.
I, [2015-11-13T04:37:54.491204 #4978] INFO -- : HA environment not enabled
.
.
.
I, [2015-11-13T04:38:03.905069 #4978] INFO -- : Elemental Conductor is ready
```

4. Press Ctrl+C to exit the tail command.

5. Enter the following commands:

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

where *<day>* is today (the day you are upgrading), typed with an initial capital letter: Mon, Tue, Wed, Thu, Fri, Sat, Sun

6. Confirm the message that appears.

If you enabled HA, look for database system is ready to accept connections on the primary Conductor, and database system is ready to accept read only connections on the secondary Conductor.

If you disabled HA, look for database system is ready to accept connections on the secondary Conductor.

7. Press Ctrl+C to exit the tail command.

8. Type the following command to exit the session as the sudo user:

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

Enabling user authentication

This section applies only if you had previously enabled user authentication on the cluster.

User authentication is already set up in the cluster, because the configuration information was included when you restored the database. But you must apply it (enable it) again, to push user authentication to every worker node.

We assume that you are familiar with the user authentication process, described in the [AWS Elemental Conductor Live Configuration Guide](#).

We also assume that you created the recommended administrator users, including the *api-admin* user). For information about the role of this user, see the information about types of users in the Reference: Manage users section of the [AWS Elemental Conductor Live Configuration Guide](#).

To enable user authentication

This procedure involves working with two users (that should already exist in the cluster). When you are logged in as the *api-admin* user on the primary Conductor, you copy the *api-admin* user to each worker node. As part of the copy action, you must specify the credentials for a user that has SSH access. We recommend that you specify the *elemental* user,

1. Log into the web interface on the primary Conductor node as the API admin (*api-admin*).
2. On the main menu, choose **Cluster**, then choose **Nodes**. Choose **Tasks** (in the top left corner) and select **Enable Node Authentication**.
3. On the **Select a user name** page, choose **api-admin**.
4. Choose **Next**.
5. On the **Enter a password** page, enter the existing password for the *api-admin* user. When you enter the same password, you are setting up *api-admin* with the same password on every node in the cluster. Setting up with the same password reduces effort with password management.
6. Choose **Next**.
7. On the **Enter the SSH credentials to access nodes page**, enter the default user (*elemental*) and its password. Then choose **Next**.
8. Choose **Configure Now**.

Refresh the page to track the progress of the action. When all the action has finished, the **Nodes** page displays each node with a lock icon.

9. Verify that enabling has succeeded. Enabling succeeds only if the *elemental* user on the primary Conductor and on the worker node have the same password. The lock icon might not be a valid indicator that user authentication succeeded on a worker node.

If your organization has the policy of setting a different *elemental* password on every node, you must repeat this process. Each time that you display the **Enter the SSH credentials to access nodes page**, enter the password for another worker node, until you have set up all the worker nodes.

Removing a Conductor node from the cluster

You can remove a Conductor node that is not acting as the primary Conductor node — that isn't controlling the cluster. To remove a Conductor node, you first remove the node from the redundancy group, and then remove the node from the cluster.

1. Disable HA. On the web interface for the Conductor node, choose **Cluster** then choose **Redundancy**. Make sure that the Conductor Live redundancy group is selected. In the High Availability field, choose Disable. To verify that high availability is disabled, see the instructions in [the section called “Cluster — enable HA or disable HA”](#).
2. Locate the Conductor to remove and click **Delete** (trash icon) to delete it from the redundancy group.
3. On the web interface for the primary Conductor node, choose **Cluster**, then choose **Nodes**.
4. Locate the Conductor node and display the options by choosing the down arrow. Select **Remove Node**.

Removing a worker node from the cluster

To remove a worker node from the cluster, first you remove channel assignments from the node. Then you remove the worker node from its redundancy group, if it is in one. Finally, you remove the node from the cluster.

You perform all these steps on the primary Conductor node, using the web interface.

Step A: Stop running channels

To stop one channel

Choose the **Channels** page, then select the stop button for the channel to stop.

To stop several channels or all channels

1. On the web interface for the primary Conductor node, choose the **Channels** page.
2. Toward the top of the page, choose **Tasks**, then choose **Stop Channels**. Or if you want to stop only some channels, select the box next to each channel you want to stop.
3. Choose **Next**, then choose **Process Now**.

Wait for all the channels to stop.

Step B: Remove node assignments

1. On the web interface for the primary Conductor node, choose the **Channels** page.

2. Toward the top of the page, choose **Tasks**, then choose **Change Channel Node Assignments**.
The list of channels on all nodes in the cluster appears.
3. Select the channels that you want to de-assign from the worker node. There are several ways to select the channels.
 - If you want to de-assign all channels from all nodes, choose **Select all channels**.
 - To de-assign the channels from one node, filter the list by node then select all the channels.
4. Choose **Next**.
5. On the **Select a new node** page, in **New Node**, choose **None**.
6. Choose **Process Now**.

Step C: Remove node from redundancy group

1. On the web interface for the primary Conductor node, choose the **Cluster** page.
2. On the **Cluster** page, choose **Redundancy**. In the navigation bar, choose the Elemental Live redundancy group.
3. On the **Active Nodes** tab, choose **Delete** (trash icon) for the node.

Step D: Remove node from cluster

1. On the **Cluster** page, choose **Nodes**.
2. Find the worker node, choose the downward triangle and select **Remove Node**.

Restarting channels

To restart one channel

1. To decide which channels to restart, use the list of running channels that you made when you stopped the channel. This list identifies the node that each channel is assigned to.
2. Choose the **Channels** page, then select the start button beside the channel to start.

To restart several or all channels

1. On the web interface for the primary Conductor node, choose the **Channels** page.
2. Toward the top of the page, choose **Tasks**, then choose **Start Channels**.

3. Choose **Select all channels**.

Or select individual channels. If you are starting channels on a specific node, use the list of running channels that you made when you stopped the channel. This list identifies the node that each channel is assigned to.

4. Choose **Next**, then choose **Process Now**.

Backing up data

You back up data using the special lifeboat script.

Topics

- [About the backup process](#)
- [Step A: Verify hostnames](#)
- [Step B: Download the lifeboat script](#)
- [Step C: Create the backup](#)

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 configuration such as NTP, PTP, and chronyd.
- Firewall settings.
- SSL certificates that are in the following directories:

- `/etc/httpd/conf`
- `/home/elemental/cert`
- 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 A: Verify hostnames

RHEL 9 (specifically `systemd`) doesn't support underscores in hostnames. If any of the nodes in the cluster contain underscores, there are two ways to proceed:

- Continue with this procedure. When you run the lifeboat script, a prompt will appear that will force you to change any hostname that includes an underscore.
- Change the hostnames before you run the lifeboat script. To change a hostname, see the Red Hat documentation.

If your migration process means that you don't run the lifeboat script for any reason, make sure that you change the hostnames before you boot the node after installing RHEL 9.

Step B: 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_el.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_el9.tar
```

5. Change to the `elemental_lifeboat_el9` directory:

```
[elemental@hostname ~]$ cd elemental_lifeboat_el9
```

Step C: 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.

Step C1: Run the backup command

Enter the backup command as follows.

On the primary Conductor:

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

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

On the secondary Conductor and on any worker node:

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

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
```

Step C2: Store the backup archive

Copy two files to storage off the node, so that you can copy them back to the node when you want to perform the restore operation. The files to store off the node are the following:

- `<hostname>_lifeboat-archive.zip` (File 1)

- `<hostname>_lifeboat-archive_export-checksum.txt` (File 2)

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.

Step C3: 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.

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

On a Linux system, enter this command:

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

On a macOS system, enter this command:

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

On a Windows system, enter this command:

```
~]$ certutil -hashfile <hostname>_lifeboat-archive.zip MD5
```

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

Restoring the database

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

Topics

- [Step A: Perform the restore](#)
- [Step B: Perform manual restore tasks](#)
- [Result of the restore](#)

Important

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

Step A: Perform the restore

1. Download the lifeboat script, following the procedure that 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 directory:

```
/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.

The script tries to extract the version of the backup that is stored in this directory: `/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.

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

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

Recovery steps

Read this information if there was a problem with the file, as mentioned in step 2 earlier on this page.

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.

Step B: Perform manual restore tasks

After you have run the lifeboat script, you might need to perform some manual steps.

SSL certificates

Read this section if your organization uses custom SSL certificates.

The lifeboat script backs up and restores the SSL certificates, both the default certificates and custom certificates you previously set up. However, the script doesn't update the Apache configuration file to point to the certificates. You must update the Apache configuration to include these paths.

Follow these steps on each node.

1. Determine where the custom certificates are located on the node. If you originally followed the recommendation from AWS Elemental, then the certificates are in `/home/elemental/cert`.
2. Enter this command:

```
cd /opt/elemental_se
```

3. Enter the configure command:

```
sudo ./configure -xeula --skip-all <certificate options>
```

There are three certificate options:

```
--https-crt <path to Apache SSLCertificateFile
```

```
--https-key <path to Apache SSLCertificateKeyFile
```

```
--https-chain <path to Apache SSLCertificateChainFile
```

Typically you need to enter either the first two options (to specify the location of the certificate and the key file) or all three options.

For example, to specify the location of the certificate and the key:

```
sudo ./configure -xeula --skip-all --https --https-cert /home/elemental/cert --https-key /home/elemental/cert
```

Firewall configuration

Read this section if the cluster was previously configured to use HTTP. Starting with versions 2.26.0 (and 3.26.0), HTTPS is always enabled by default. For more information, see the [current Release Notes](#).

You must open port 443 on every node to allow access to the web interfaces for the AWS Elemental software.

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	Yes	Yes	Yes

Node	Worker nodes	Secondary Conductor	Primary Conductor
components on the cluster (if you included them in the backup)			
Configuration files for features of the AWS Elemental software	Yes	Yes	Yes
Remote storage mounts. The remote storage mounts is only ever stored on the Conductor nodes. The data specific to the node is restored to that node.		Yes	Yes

Node	Worker nodes	Secondary Conductor	Primary Conductor
<p>Cluster data. Data relating to the cluster, including data about the channels, MPTSeS, channel and MPTS node assignments, users setup, redundancy groups, cluster members.</p> <p>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 is this migration procedure, the primary Conductor pushes the data to the secondary Conductor and to the appropriate worker nodes.</p>	Yes		

Directory structure for restored data

On every node, the directory structure after backup-and-restore is identical to the structure before backup-and-restore.

Cluster cleanup

On the primary Conductor node only, the restore command also cleans up the cluster configuration data. It performs the following actions

- Removes data about nodes that are assigned to each channel.
- Deletes data about nodes in the cluster.
- Deletes data about nodes in Conductor and worker redundancy groups. But keep the redundancy groups themselves.
- Changes the state of channels and MPTSes to Idle, and delete any data about channel activity.
- Clears all active Alerts.

This cleanup ensures you can rebuild the cluster without running into configuration conflicts.

Installing Conductor Live on nodes

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

```
[elemental@hostname ~]$ sudo sh ./
elemental_production_conductor_live247_3.26.x.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 Live, enter this command:

```
[elemental@hostname ~]$ sudo sh ./
elemental_production_conductor_live247_3.26.x.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 Elemental Live on a worker node

This install procedure isn't the same as the install procedure on a newly obtained appliance (as described in [AWS Elemental Live Installation Guide](#)). You don't have to configure the node.

This install procedure is very similar to the upgrade procedure (as described in [AWS Elemental Live Upgrade Guide](#)), 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 and CPU versions of the software:

```
[elemental@hostname ~]$ sudo sh ./elemental_production_live_2.26.x.12345.run --cleandb --skip-mellanox --skip-all --start -xeula
```

- For CPU-only versions of the software:

```
[elemental@hostname ~]$ sudo sh ./elemental_production_live_cpu_2.26.x.12345.run --cleandb --skip-mellanox --skip-all --start -xeula
```

Where:

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

`--skip-mellanox`. Optional. Skips installation of the Mellanox driver, even if the script detects that a Mellanox NIC is installed in the appliance. For more information, see the [current Release Notes](#).

`--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 A: Update the firmware

To update the BIOS firmware and the BMC firmware (IPMI for SuperMicro, iDRAC for Dell), read the Knowledge article [Latest AWS Elemental Qualified Remote Management and BIOS Firmware](#). This article provides instructions and includes links to downloads.

Step B: 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 Elemental Live 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 a 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.

Finishing steps

1. Installation of the RHEL 9 operating system will remove all network configurations. You must manually configure at least one interface before you can install the Elemental Live software. For information, go to the [AWS Elemental Support Center](#), and read the Knowledge article [What to do if kickstart removes on-disk network configuration](#) or open a case.
2. Sometimes, a kickstart or upgrade generates incorrect mappings for Ethernet ports. For information about how to fix this issue, go to the [AWS Elemental Support Center](#), and read the Knowledge article [How to restore default ethernet port order on an Elemental appliance](#) or open a case.
3. You can now install any third-party packages. To obtain these packages, see [the section called "RPM repository"](#).

USB stick

Install the operating system

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.

Finishing steps

1. Installation of the RHEL 9 operating system will remove all network configurations. You must manually configure at least one interface before you can install the Elemental Live software. For information, go to the [AWS Elemental Support Center](#), and read the Knowledge article [What to do if kickstart removes on-disk network configuration](#) or open a case.
2. Sometimes, a kickstart or upgrade generates incorrect mappings for Ethernet ports. For information about how to fix this issue, go to the [AWS Elemental Support Center](#), and read the Knowledge article [How to restore default ethernet port order on an Elemental appliance](#) or open a case.
3. You can now install any third-party packages. To obtain these packages, see [the section called “RPM repository”](#).

Installing on a SuperMicro

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

Install the operating system

1. Install the Java applet and change the security level. See [the section called “Step A: 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, otherwise you return to the kickstart menu.

Finishing steps

1. Installation of the RHEL 9 operating system will remove all network configurations. You must manually configure at least one interface before you can install the Elemental Live software. For information, go to the [AWS Elemental Support Center](#), and read the Knowledge article [What to do if kickstart removes on-disk network configuration](#) or open a case.
2. Sometimes, a kickstart or upgrade generates incorrect mappings for Ethernet ports. For information about how to fix this issue, go to the [AWS Elemental Support Center](#), and read the Knowledge article [How to restore default ethernet port order on an Elemental appliance](#) or open a case.
3. 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, go to the [AWS Elemental Support Center](#) and read the Knowledge article [about the introduction of RHEL 9](#).

Document history for migration guide

The following table describes the main changes to this guide.

Change	Description	Date
Installing on worker node	The procedure to install software on a worker node has been updated to include the --cleandb in the command.	June 13, 2025
Corrections	General corrections and additions have been made to the guide.	July 29, 2024
New guide	First release of this guide	March 25, 2024