



Configuration Guide

AWS Elemental Conductor Live



AWS Elemental Conductor Live: Configuration Guide

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

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

Table of Contents

About this guide	1
Rules and limits	2
Accessing the nodes	3
Working from the web interface	3
Working from the CLI	3
Key cluster setup procedures	4
Initial configuration	4
Upgrading standalone nodes	7
Adding user authentication	10
Adding Conductor redundancy	10
Designing the cluster	12
Reference: Configure connectivity	14
DNS servers	14
Clocks: NTP and PTP servers	15
Configuring NTP servers	16
Configuring PTP	17
Ethernet interfaces	18
Creating an Ethernet interface	18
Modifying an Ethernet interface	19
Creating or modifying a bond	19
Dedicating interfaces to MPTS	24
Firewalls and ports	25
Firewall recommendation	26
Enabling or disabling the product firewall	26
Working with ports on the product firewall	27
HTTPS	28
Enabling HTTPS	29
Disabling HTTPS	30
Input: Directly connected SDI inputs	31
Step A: Add the devices to each Elemental Live node	31
Step B: Import the devices into the cluster	32
Inputs: Routers for handling SDI inputs	32
Step A: Gather information	33
Step B: Run cables from the router to each node	34

Step C: Add the router	35
Step D: Complete the Router Input Mappings	36
Step E: Complete the Router Output Mappings	37
Step F: Sync the Routers	38
Step G: Use the Router Inputs	38
Mount points	38
Time zone	40
Reference: Configure worker features	42
OCR for captions	42
RTMP inputs	42
Virtual input switching	43
Reference: Configure the cluster	44
Manage nodes	44
Add (recruit) nodes	44
Remove a worker node	46
Remove a Conductor Live node	47
Redundancy groups	47
Conductor Live redundancy group	48
Worker redundancy groups	50
High availability (HA)	51
Verifying the current HA state	51
Enabling HA	51
Disabling HA	53
Reference: Configure user authentication	56
About user authentication	56
Summary of procedure	56
Types of user authentication	57
Step 1: Enable user authentication	57
Step 2: Apply authentication	61
Disabling user authentication	62
Reference: Manage users	64
Types of users	64
Users on Conductor Live nodes	65
Users on worker nodes	67
Summary	68
Adding users to Conductor Live	68

Adding users to workers	69
Role policies for PAM authentication	70
Reference: Configure notifications	72
Email notification	73
Configure sendmail relay server	76
Web callback notification	78
SNMP traps	81
SNMP polling	82
MIBs in Conductor Live	83
Reference: Backup and restore	86
Configuring for backup	86
Disabling database backups	87
Restoring a backup	87
Document History	89

About this guide

This guide describes how to configure the Conductor Live and worker nodes in a AWS Elemental Conductor Live cluster. It describes how to configure the Conductor Live manager node, AWS Elemental Live nodes, and AWS Elemental Statmux nodes (if your deployment includes that product).

This guide applies to all versions of the software that are currently available for download from AWS Elemental.

Phase 2 of installation

This guide describes how to configure the Conductor Live nodes and worker nodes (Elemental Live and optionally Elemental Statmux) in a Conductor Live cluster. The guide describes the steps to initially set up the nodes in a cluster, and describes other procedures you might want to perform after the initial deployment. The guide assumes that you have already racked the servers, and you have installed software, if necessary.

Prerequisite knowledge

We assume that you know how to:

- Connect to the Conductor Live web interface using your web browser.
- Log in to a remote terminal (Linux) session in order to work via the command line interface.

Note

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

Rules and limits

The following table provides a summary of the configuration rules and constraints that apply to an AWS Elemental Conductor Live cluster.

Feature or topic	Rule or limit
Hardware in a cluster	A Conductor Live cluster can include a maximum of 50 worker nodes.
Physical location of nodes in a cluster	Within a cluster, the Conductor Live and encoder nodes must be located in the same physical location.
	Within a cluster, the communications among cluster members shouldn't traverse the public internet.

Accessing the nodes

The procedures in this guide require you to access the AWS Elemental Conductor Live nodes and the worker nodes. You might need to work with a node using the web interface or using the CLI (command line interface).

Working from the web interface

At your workstation, open a web browser and enter the IP address or hostname of the node.

Limitations on using the web interface

You can't use the web interface to perform some configuration tasks. You must always use the CLI. The affected configuration tasks are:

- DNS server. If at least one Ethernet interface on the node uses DHCP, you can only use the CLI to work with DNS. Otherwise, you can work with DNS using the web interface or the CLI.
- Ethernet interfaces: You can't create or modify an Ethernet interface using the web interface.
- Ethernet interface bonds: You can't create or modify a bond using the web interface.

Working from the CLI

At your workstation, start a remote terminal session to the AWS Elemental Conductor Live node. For example, use SSH and connect to the node using the IP address or hostname:

```
$ ssh myConductorLive
```

Log in to the node using the credentials for an administrator that has been set up on the node. If you haven't set up users yet, use the user credentials of the default *elemental* user.

Key cluster setup procedures

This section describes the correct order of work for some key setup procedures for a AWS Elemental Conductor Live cluster. It is often important to perform tasks in a specific order.

Topics

- [Deploying a new cluster](#)
- [Upgrading standalone nodes into a cluster](#)
- [Adding user authentication to an existing cluster](#)
- [Adding Conductor redundancy to an existing cluster](#)

Deploying a new cluster

Read this section if you are setting up AWS Elemental Conductor Live, AWS Elemental Live, and (optionally) AWS Elemental Statmux for the first time.

(If you are grouping existing Elemental Live nodes into a Conductor Live cluster, see [the section called “Upgrading standalone nodes”](#).)

Warning

You must perform these steps in the specified order. Otherwise, the cluster might not get set up correctly.

Step 1: Design the cluster

As your first step, you should design the cluster. For guidelines, see [???](#).

Step 2: Install software

You might need to install the AWS Elemental software on the Conductor Live nodes and worker nodes.

- If you have obtained AWS Elemental appliances, you don't need to install software. The appliances are delivered with software already installed.

- If you have obtained qualified hardware, you must install the software. See the appropriate guide:
 - [AWS Elemental Conductor Live Install Guide](#). Keep in mind that Elemental Statmux is installed as part of Conductor Live.
 - [AWS Elemental Live Install Guide](#)

Note

Make sure that both Conductor Live nodes have the same software version installed.

Step 3: Configure connectivity features on the nodes

- On each worker node, perform the tasks that are listed in [Reference: Configure connectivity](#). Perform these tasks in any order.
- On the primary Conductor Live node, perform the tasks that are listed in [Reference: Configure connectivity](#). Perform these tasks in any order.
- On the secondary Conductor Live node, perform the following tasks. Perform these tasks in any order:
 - Configure [DNS servers](#).
 - Configure [Ethernet interfaces and bonds](#)(optional).
 - [Enable HTTPS](#) on the node.
 - Configure [NTP servers](#).

You don't need to configure as many fields on the secondary Conductor Live because the secondary Conductor Live will synchronize with the primary Conductor Live.

Step 4: Configure features on the worker

On each worker node, configure the features that apply to that node, from the list in [Reference: Configure worker features](#).

You can configure these features at any time in this setup procedure.

Step 5: Enable user authentication on the Conductor Live

We recommend that you set up the nodes so that users must log into the node. For an overview of how user authentication works, see [the section called “About user authentication”](#).

If you do decide to set up in this way, you must enable the user authentication feature on the Conductor Live, before you recruit the nodes to the cluster:

- On the primary Conductor Live, run the configuration script to [enable the user authentication feature](#).

Step 6: Recruit nodes into the cluster

Recruit the primary Conductor Live nodes and all the workers nodes into the cluster. See [the section called “Add \(recruit\) nodes”](#)

The nodes get added to the cluster, but they don't yet belong to any redundancy group.

Step 7: Configure redundancy groups in the cluster

We recommend that you set up the cluster with Conductor redundancy (a primary and a secondary Conductor Live node), and with worker node redundancy.

- Design a redundancy plan. For information, see [AWS Elemental Conductor Live User Guide](#)
- [Create the redundancy groups](#) that you identified.
- [Add the worker nodes](#) to each worker redundancy group.
- [Add the primary and secondary Conductor Live nodes](#) to the Conductor redundancy group.

Step 8: Apply user authentication on worker nodes

If you enabled user authentication on the primary Conductor Live (earlier in this procedure), you must now apply user authentication on all the worker nodes in the cluster.

- On the primary Conductor Live, apply user authentication. See [the section called “Step 2: Apply authentication”](#).

Step 9: Add users to the nodes

If you enabled user authentication, you must now add users. For information about the types of users that you can add, see [Reference: Manage users](#).

Add these types of users:

- One or two *regular administrators* on Conductor Live. For information, see [the section called “Adding users to Conductor Live”](#).
- *Operators* and *viewers* on Conductor Live. For information, see [the section called “Adding users to Conductor Live”](#).
- One or two *regular administrators* on each individual worker node. These administrators will access the worker node locally (by logging on directly on the web interface of the node) only in order to troubleshoot. For information, see [the section called “Adding users to workers”](#).

Step 10: Enable HA

[Enable HA \(high availability\)](#) on the primary Conductor Live.

Upgrading standalone nodes into a cluster

Read this section if your organization has already deployed one or more AWS Elemental Live nodes, and you now want to set them up in a AWS Elemental Conductor Live cluster.

With this task, you already have some Conductor Live nodes in deployment. You now want to put these nodes into a cluster that is controlled by a Conductor Live.

Warning

You must perform these steps in the specified order. Otherwise, the cluster might not get set up correctly.

Step 1: Design the cluster

As your first step, you should design the cluster. For guidelines, see [???](#).

Step 2: Install software

You might need to install the AWS Elemental software on the Conductor Live nodes.

- If you have obtained AWS Elemental appliances, you don't need to install software. The appliances are delivered with software already installed.

- If you have obtained qualified hardware, you must install the software. See the appropriate guide:
 - [AWS Elemental Conductor Live Install Guide](#). Keep in mind that Elemental Statmux is installed as part of Conductor Live.
 - [AWS Elemental Live Install Guide](#)

Note

Make sure that both Conductor Live nodes have the same software version installed.

Step 3: Configure connectivity features on the nodes

- On each worker node, modify the existing [NTP server configuration](#) to point to the URL of the primary Conductor Live node.
- On the primary Conductor Live node, perform the tasks that are listed in [Reference: Configure connectivity](#). Perform these tasks in any order.
- On the secondary Conductor Live node, perform the following tasks. Perform these tasks in any order:
 - Configure [DNS servers](#).
 - Configure [Ethernet interfaces and bonds](#)(optional).
 - [Enable HTTPS](#) on the node.
 - Configure [NTP servers](#).

You don't need to configure as many fields on the secondary Conductor Live because the secondary Conductor Live will synchronize with the primary Conductor Live.

Step 4: Configure user authentication on the primary Conductor Live

We recommend that you set up the nodes so that users must log into the node. For an overview of how user authentication works, see [the section called "About user authentication"](#).

If you do decide to set up in this way, you must enable the user authentication feature on the Conductor Live, before you recruit the nodes into the cluster:

- On the primary Conductor Live, run the configuration script to [enable the user authentication feature](#).

Step 5: Recruit nodes into the cluster

Recruit the secondary Conductor Live nodes and all the workers nodes into the cluster. See [the section called “Add \(recruit\) nodes”](#).

The nodes get added to the cluster, but they don't yet belong to any redundancy group.

Step 6: Configure redundancy groups in the cluster

We recommend that you set up the cluster with Conductor redundancy (a primary and a secondary Conductor Live node), and with worker node redundancy.

- Design a redundancy plan. For information, see [AWS Elemental Conductor Live User Guide](#)
- [Create the redundancy groups](#) that you identified.
- [Add the worker nodes](#) to each worker redundancy group.
- [Add the primary and secondary Conductor Live nodes](#) to the Conductor redundancy group.

Step 7: Apply user authentication on worker nodes

If you enabled user authentication on the primary Conductor Live (earlier in this procedure), you must now apply user authentication on all the worker nodes in the cluster.

- On the primary Conductor Live, apply user authentication. See [the section called “Step 2: Apply authentication”](#).

Step 8: Add users to the nodes

If you enabled user authentication, you must now add users. For information about the types of users that you can add, see [Reference: Manage users](#).

Add these types of users:

- One or two *regular administrators* on Conductor Live. For information, see [the section called “Adding users to Conductor Live”](#).
- *Operators* and *viewers* on Conductor Live. For information, see [the section called “Adding users to Conductor Live”](#).

- One or two *regular administrators* on each individual worker node. These administrators will access the worker node locally (by logging on directly on the web interface of the node) only in order to troubleshoot. For information, see [the section called “Adding users to workers”](#).

Step 9: Enable HA

[HA \(high availability\)](#) (HA) on the primary Conductor Live.

Adding user authentication to an existing cluster

The procedure to add user authentication to the AWS Elemental Conductor Livecluster includes multiple steps. You must remove the nodes from the cluster, enable user authentication, and then add the nodes back to the cluster:

- On each worker node, perform these tasks in the specified order:
 - Remove each worker node from the cluster. See [the section called “Remove a worker node”](#)
 - [Enable HTTPS](#)
 - Recruit (add) the node back into the cluster See [the section called “Add \(recruit\) nodes”](#).
- On the primary Conductor Live node, perform these tasks in the specified order:
 - [Disable HTTPS](#)
 - [Remove the primary Conductor Live node](#) from the cluster.
 - [Enable the user authentication feature](#).
 - [Recruit \(add\) the primary Conductor Live node](#) back into the cluster, then [add it](#) back into its redundancy group.
 - [Apply user authentication](#) on the cluster.
 - [Add users](#).
 - [Enable HA](#).
 - [Enable HTTPS](#).

Adding Conductor redundancy to an existing cluster

You might have originally deployed the cluster with only one AWS Elemental Conductor Live node. You might now want to add a secondary Conductor Live node, to implement node redundancy on the Conductor nodes.

Perform the following steps in the specified order:

1. As your first step, you should identify your redundancy requirements. See [???](#).
2. [Verify the firewall setup](#) is the primary Conductor Live:
 - Make sure that the firewall is enabled on both Conductor Live nodes.
 - Accept port 5432 TCP.
3. Configure the secondary Conductor Live, perform the following tasks. Perform these tasks in any order:
 - [Configure DNS servers](#).
 - [Configure Ethernet interfaces](#) and bonds (optional).
 - [Enable HTTPS](#) on the node..
 - Configure [NTP servers](#).

You don't need to configure as many fields on the secondary Conductor Live because the secondary Conductor Live will synchronize with the primary Conductor Live.

4. [Recruit \(add\)](#) the secondary Conductor Live into the existing cluster.
5. Create a redundancy group for the two Conductor Live nodes, and add the nodes to that group. See [the section called "Conductor Live redundancy group"](#).
6. [Enable HA \(high availability\)](#) on the primary Conductor Live. When you enable HA, the secondary Conductor Live synchronizes itself with the primary Conductor Live.

Designing the cluster

You should make some decisions about features you want to configure on the cluster.

Nodes in the cluster

For information about deciding about the number of AWS Elemental Conductor Live and worker nodes in the cluster, see [AWS Elemental Conductor Live User Guide](#).

For information about designing for redundancy in the cluster, see [Conductor Live User Guide](#).

Gather network information

1. We strongly recommend that, if your deployment involves several Conductor Live clusters, you set up each cluster in its own network.
2. Identify the network interfaces and devices for all the worker nodes, including the following:
 - Ethernet interfaces
 - SDI devices
 - DNS servers to connect to
 - NTP servers to connect to
 - Remote servers. For example, servers where files assets are stored that Elemental Live events will use

User authentication

Your organization might require that users present credentials in order to work with the nodes. You can implement a simple built-in user authentication, or you can implement PAM authentication.

Note

We strongly recommend that you enable user authentication, and that you implement it with HTTPS enabled.

User authentication is a mode in the cluster. If user authentication is enabled, all users must always log on to any node in the cluster. See [the section called “About user authentication”](#).

Plan for backup of databases

Conductor Live is configured by default to back up the data for the Conductor Live nodes and all workers nodes. You can configure the Conductor Live to back up to a remote server. See [Reference: Backup and restore](#).

 Note

We strongly recommend that you back up data to a remote server.

Notifications

You can configure the nodes for notifications. Some types of notifications are always enabled, but you can customize the behavior. Other types of notifications work only if you enable them. See [Reference: Configure notifications](#).

Configuring nodes for connectivity

This section describe how to configure the node so that it can communicate with Conductor Live, and with upstream systems and downstream systems.

Note

If your deployment involves several Conductor Live clusters, we strongly recommend that you set up each cluster in its own network.

Topics

- [Configure DNS servers](#)
- [Configuring NTP and PTP servers](#)
- [Configure Ethernet interfaces](#)
- [Configuring a firewall and opening ports](#)
- [Enabling and disabling HTTPS](#)
- [Adding SDI input devices](#)
- [Configuring SDI video routers](#)
- [Adding mount points to worker nodes](#)
- [Setting the web interface time zone](#)

Configure DNS servers

You can configure a node to use one or more DNS servers.

These rules apply to working with DNS on a node:

- If at least one Ethernet interface on the node uses DHCP, you can only use the CLI to work with DNS.
- If all the Ethernet interfaces have static addresses, you can work with DNS using the web interface or the CLI.

Where to perform the configuration

Make sure you perform the configuration on the correct nodes.

Node	Work on this node?
Primary Conductor Live node	Yes
Secondary Conductor Live node	Yes
Each worker node	Yes

To configure DNS using the web interface

See the information earlier on this page to determine if you can use the web interface.

1. If the node is a Conductor Live node and if HA is currently enabled, disable it now. Conductor Live redundancy (HA, or *high availability*) must be disabled before you configure Ethernet interfaces. For instructions, see [Disabling Conductor Live HA \(high availability\)](#).
2. For a Conductor Live node, on the Conductor Live web interface, go to the **Settings** page and choose **Network**. Then choose the **Domain Name Servers** tab.

For a worker node, on the worker web interface, choose **Settings** from the main menu. Choose the **Network** tab, then choose **Hostname, DNS, & Timing Server**.

3. Enter the IP address for the DNS server. Choose **Save**. Another line appears, for you to enter another DNS server, if you want.

To configure DNS servers using the CLI

1. If the node is a Conductor Live node and if HA is currently enabled, disable it now. Conductor Live redundancy (HA, or *high availability*) must be disabled before you configure Ethernet interface. For instructions, see [Disabling Conductor Live HA \(high availability\)](#).
2. To configure DNS, see the [Red Hat Networking Guide](#).

Configuring NTP and PTP servers

You must configure a clock server for the nodes to use. You can configure any of these types of clock servers:

- An NTP clock.
- A PTP clock

If you plan to set up at least one channel with a SMPTE 2110 output, you must set up the clocks as follows:

- Configure the AWS Elemental Conductor Live nodes and any AWS Elemental Statmux nodes to use NTP.
- Configure all the AWS Elemental Live nodes that include SMPTE 2110 outputs with PTP. For other Elemental Live nodes, you can configure to use PTP or NTP.

If you don't plan to set up with SMPTE 2110 outputs, you can set up all the nodes to use NTP.

Configuring NTP servers

Where to perform the configuration for NTP

Make sure you perform the configuration on the correct nodes.

Node	Work on this node?
Primary Conductor Live node	Yes
Secondary Conductor Live node	Yes
Each worker node	Yes

Recommendation

We recommend that you set up each worker node to specify the URL of the primary Conductor Live node. In this way, the worker nodes synchronize their internal clocks with Conductor Live. Setting up in this way reduces the risk that one worker clock will drift from the other clocks and cause encoding problems.

To configure NTP on Conductor Live

1. Perform this step in the remote terminal session on the primary Conductor Live node.

- On the Conductor Live node, edit the file `/etc/chrony.conf`
- Uncomment the following line, and modify the IP addresses to specify the subnet that Conductor Live and the workers are configured with.

```
#allow 192.168.0.0/16
```

- Uncomment the following line.

```
#local stratum 10
```

When you uncomment this line, the worker nodes will be able to sync with the Conductor Live node even if the Conductor Live loses communication with the NTP server. All the nodes must stay in synch all the time, in order for Conductor Live to control the worker nodes.

- Save the file.
- On the Conductor Live enter this command to restart chronyd:

```
sudo systemctl restart chronyd
```

2. On the Conductor Live web interface, go to the **Settings** page and choose **Network**.
3. On the **Network Time Protocol Servers** tab, enter the IP address for the server you want to use.

To configure NTP on Elemental Live or Elemental Statmux

1. On the Elemental Live web interface, choose **Settings** from the main menu.
2. Choose the **Network** tab, then choose **Hostname, DNS, & Timing Server**.
3. In the **NTP Servers** section, set up one line, with the IP address of the primary Conductor Live node. Choose **Save**.

Configuring PTP

PTP applies only to Elemental Live, and only starting with version 2.21.3, which is when support for SMPTE 2110 was introduced.

To configure PTP on Elemental Live

1. On the Elemental Live web interface, go to the **Settings** page and choose **Network**.
2. On the **Host, DNS & Timing Server** tab, select the Enable PTP check box and choose Save.

When you save, PTP is enabled on the Elemental Live node. If NTP was previously enabled, it is automatically disabled on the Elemental Live node. But the Conductor Live nodes and the Elemental Statmux nodes in the cluster will continue to use NTP.

Configure Ethernet interfaces

When you installed the software on the individual nodes in an AWS Elemental Conductor Live cluster, you configured eth0. If you need to set up more Ethernet interfaces (network devices), read this section. You can optionally bond Ethernet interfaces to suit your networking requirements.

Where to perform the configuration

Make sure you perform the configuration on the correct nodes.

Node	Work on this node?
Primary Conductor Live node	Yes
Secondary Conductor Live node	Yes
Each worker node	Yes

Topics

- [Creating an Ethernet interface](#)
- [Modifying an Ethernet interface](#)
- [Creating or modifying a bond](#)
- [Dedicating interfaces to MPTS](#)

Creating an Ethernet interface

You use the CLI to create Ethernet interfaces using the web interface.

- If the node is a Conductor Live node and if HA is currently enabled, disable it now. Conductor Live redundancy (HA, or *high availability*) must be disabled before you configure Ethernet interfaces. For instructions, see [Disabling Conductor Live HA \(high availability\)](#).
- To create the Ethernet interface, see the [Red Hat Networking Guide](#).

Modifying an Ethernet interface

You use the CLI to modify Ethernet interfaces using the web interface.

- If the node is a Conductor Live node and if HA is currently enabled, disable it now. Conductor Live redundancy (HA, or *high availability*) must be disabled before you configure network interfaces. For instructions, see [Disabling Conductor Live HA \(high availability\)](#).
- To modify an Ethernet interface, see the [Red Hat Networking Guide](#).

Warning

The **Devices** page on the Conductor Live web interface includes the pencil icon that lets you edit the Ethernet interface. However, you must not use the web interface to modify interfaces because you will break the configuration.

Creating or modifying a bond

If you set up more Ethernet interfaces on the Conductor Live node, you can optionally bond two Ethernet interfaces.

You can bond Ethernet interfaces to suit your networking requirements. For example, you might set up two Ethernet interfaces as an active/redundant pair.

Important

We recommend that when you set up a bond, you set up both eth0 and eth1 with static IP addresses and with eth0, eth1 and bond0 all on the same subnet.

Prerequisites

Before you begin this process, make sure that you have done the following:

- [Set up the individual Ethernet interfaces](#) that you want to bond together.
- If HA is currently enabled, disable it now. Conductor Live redundancy (HA or *high availability*) must be disabled before you configure Ethernet interfaces. For instructions, see [Disabling Conductor Live HA \(high availability\)](#).

Warning

Don't use the web interface to create or modify a bond because you will break the configuration.

Step A: Create bond configuration file

Create a configuration file for the bond interface and name it after the bond.

To create the bond configuration file

1. Create the file with the following command.

```
sudo vim /etc/sysconfig/network-scripts/ifcfg-bond0
```

2. Insert the following settings in the file:

- **DEVICE** – Type **bond0**.
- **TYPE** – Type **Bond**.
- **NAME** – Provide a name for the bond that is unique among your bonded interfaces, such as **bond0**.
- **BONDING_MASTER** – Type **yes**.
- **BOOTPROTO** – If you are using a static IP address for the bond, type **none**. If you are using DHCP, type **dhcp**.
- **ONBOOT** – Type **yes**.
- **NM_CONTROLLED** – Type **no**.
- **IPADDR** – When you are using a static IP address, complete with your networking information.
- **NETMASK** – When you are using a static IP address, complete with your networking information.

- **GATEWAY** – When you are using a static IP address, complete with your networking information.
- **BONDING_OPTS** – Type the bonding mode you are using.

Bonding modes

The following table describes the bonding modes that are available.

Bonding mode option	Mode name	Description
mode=0	Round robin	Transmissions are received and sent sequentially on each bonded interface, beginning with the first one available.
mode=1	Active backup	Transmissions are received and sent out using the first available bonded interface . The other interface is only used if the active interface fails.
mode=2	Balanced XOR	Using the exclusive-or (XOR) method, the interface matches up the incoming request's MAC address with the MAC address for one of the bonded interface NICs. When this link is established, transmissions are sent out sequentially, beginning with the first available interface.
mode=3	Broadcast	All transmissions are sent on all interfaces in the bond.
mode=4	IEEE 802.3ad dynamic link aggregation	This option creates aggregation groups that share the

Bonding mode option	Mode name	Description
		same speed and duplex settings. This transmits and receives on all interfaces in the active aggregator. Requires a switch that is 802.3ad compliant.
mode=5	Adaptive transmit load balancing	Outgoing traffic is distributed according to current load on each interface in the bond. Incoming traffic is received by the currently active interface. If the receiving interface fails, another interface takes over the MAC address of the failed interface.
mode=6	Adaptive load balancing	This option includes transmit and receive load balancing for IPV4 traffic. Receive load balancing is achieved through ARP negotiation.

Example

```

DEVICE=bond0
TYPE=Bond
NAME=bond0
BONDING_MASTER=yes
BOOTPROTO=none
ONBOOT=yes
NM_CONTROLLED=no
IPADDR=192.168.1.70
NETMASK=255.255.255.0
GATEWAY=192.168.1.1
BONDING_OPTS="mode=5 miimon=100"

```

Step B: Edit Ethernet interface configuration files

Access the configuration files for each of the interfaces that are participating in the bond. Add the following lines.

```
MASTER=bond0  
SLAVE=yes
```

For help with creating and updating bonding files through the CLI, see [Using the CLI](#) in the *Red Hat Networking Guide*.

Step C: Restart the AWS Elemental service configuration files

Restart the AWS Elemental service using the following command.

```
sudo systemctl restart network
```

Step D: Verify the bond

Enter this command to display information about a specific bond:

```
cat /proc/net/bonding/bond0
```

In this example, bond0 is correctly set up.

Example

```
[elemental@host~]$ cat /proc/net/bonding/bond0  
Ethernet Channel Bonding Driver: v3.7.1 (April 27, 2011)  
  
Bonding Mode: transmit load balancing  
Primary Slave: None  
Currently Active Slave: eth0  
MII Status: up  
MII Polling Interval (ms): 100  
Up Delay (ms): 0  
Down Delay (ms): 0  
  
Slave Interface: eth0  
MII Status: up
```

```
Speed: 10000 Mbps
Duplex: full
Link Failure Count: 0
Permanent HW addr: 00:50:56:b8:64:44
Slave queue ID: 0

Slave Interface: eth2
MII Status: up
Speed: 10000 Mbps
Duplex: full
Link Failure Count: 0
Permanent HW addr: 00:50:56:8d:8e:db
Slave queue ID: 0
[elemental@host~]$
```

Dedicating interfaces to MPTS

This section applies only if your cluster includes AWS Elemental Statmux nodes.

You must perform an extra configuration step on every Elemental Statmux node and on any Elemental Live node that will produce SPTS outputs for an Elemental Statmux MPTS:

- On each Elemental Statmux node, you must identify two interfaces that will handle MPTS communications between this node and any Elemental Live node.
- On each affected Elemental Live node, you must identify two interfaces that will handle MPTS communications between this node and any Elemental Statmux node.

To identify dedicated interfaces

1. On each Elemental Statmux node, identify two interfaces from the Ethernet interfaces [that you have created](#).

On each Elemental Live node, identify two interfaces from the Ethernet interfaces [that you have created](#).

We recommend that you dedicate two interfaces on every node, to provide network redundancy. These interfaces can be separate or they can be already bonded together.

2. On the primary Conductor Live web interface, choose the **Cluster** page, then choose **Nodes**.
3. On the **Nodes** page, select the hostname of an Elemental Statmux or Elemental Live node. Don't select the node by its IP address. The node details page appears for this node.

4. Choose the **Network** tab. On the menu across the top, choose the **MPTS Configuration** tab. (Note that this tab is the only read-write tab on this page. The other **Network** tabs are read-only.)
5. Complete the **MPTS Configuration** page as follows:
 - In the **Interface 1** and **Interface 2** fields, select the IP addresses that you identified.

If you identified only one interface, select the same value in both fields.

If you identified a bonded interface, select the same value in both fields.
 - In the **Cluster Multicast Address** field, enter a multicast address. A multicast address ensures that communications will resume if either the Elemental Statmux or the Elemental Live node fails over.
6. Repeat steps 3 to 5 on every Elemental Statmux node and on every affected Elemental Live node.

Configuring a firewall and opening ports

You can enable the firewall on each node on the cluster. You can customize which ports are open on the firewall on a node.

Where to perform the configuration

Node	Work on this node?
Primary Conductor Live node	Yes
Secondary Conductor Live node	Yes
Each worker node	Yes

Topics

- [Firewall recommendation](#)
- [Enabling or disabling the product firewall](#)
- [Working with ports on the product firewall](#)

Firewall recommendation

Your organization's firewall

We recommend that you always deploy all the nodes behind your organization's firewall, on a private network.

AWS Elemental product firewall

Each AWS Elemental product has a built-in firewall.

We recommend that you enable this *product firewall* on the entire cluster.

- If your nodes are appliances, then the software was already installed on delivery, with the product firewall enabled.
- If your nodes are on qualified hardware or on VMs, you specified whether to enable the product firewall when you installed the software. We recommend that you enable the firewall. If you didn't enable the firewall when you installed, you can enable it now. You must do this individually, on each node.

Rules for firewall configuration on the Conductor Live nodes

Both Conductor Live nodes must have the same firewall settings. If they don't, you won't be able to add the secondary node to the cluster.

Port 5432 (TCP) must be open (accepted) on both nodes.

Enabling or disabling the product firewall

Make sure that all the nodes in the cluster are configured in the same way—with the firewall enabled (recommended) or with the firewall disabled.

To enable or disable the firewall

1. If you are enabling or disabling the firewall on a Conductor Live node that has HA enabled, [disable it](#) now.
2. On the web interface for Conductor Live, go to the **Settings** page and choose **Firewall**.

Or on the web interface for the worker node, choose **Settings**, then choose the **Firewall** tab.

3. For Conductor Live, choose **Start Firewall** or **Stop Firewall**.

For a worker node, choose **Firewall On** or **Firewall Off**. Then choose **Save**.

Working with ports on the product firewall

Every node is configured by default with a list of ports that can be opened or closed. When you enable the product firewall on each node, each port is automatically configured with an open or closed state.

- Some ports are configured as open by default, and you can't change the state. These configurations are read-only because these ports must be open in order for the cluster nodes to work.
- Other ports are configured as closed by default, but you can change the state.
- You can also add custom ports and open them.

To add more incoming ports on the node firewall

1. Display the **Firewall Settings** page.
2. If necessary, choose **Firewall On** (on a worker node) or **Start Firewall** (on Conductor Live). The list of ports appears.
3. Display the dialog:
 - On Conductor Live, choose **Add Incoming Port** on the right side of the page.
 - On a worker node, go to **Add Incoming Port** at the end of the list.
4. Select **Accept**, choose the **Type** (TCP or UDP), and enter the port number. Choose **Save**.

To open or close ports on the node firewall

1. On the node web interface, go to the **Settings** page and choose **Firewall**.
2. Decide if you really want to close a port that is currently open. Look at the description, which describes the port's purpose. Some ports must be open.
3. Conductor Live: Click the edit (pencil) button. On the dialog, choose **OK** to toggle the port configuration.

A worker node: In the row for the port, choose **Accept** to open the port. Or clear the check box in **Accept** to close the port.

4. Choose **Save**.

To remove a port

You can't remove a port. Instead, clear the **Accept** field and choose **Save**.

Enabling and disabling HTTPS

Read this section only if you are using Conductor Live version 3.25 or earlier. With these versions, you must explicitly enable HTTPS if you want users or applications to have a secure connection to Conductor Live and worker nodes.

(With Conductor Live version 3.26 and later, HTTPS is enabled by default. There is no need to enable it.)

Note

The HTTPS configuration must be the same for AWS Elemental Conductor Live and all the worker nodes in the cluster. Either enable HTTPS for all nodes, or disable it for all nodes.

We recommend that you enable HTTPS.

- If your nodes are appliances, then the software was already installed on delivery, with HTTPS disabled. You can enable it now on each node.
- If your nodes are on qualified hardware or on VMs, you specified whether to enable HTTPS when you installed the software. If you didn't enable HTTPS when you installed, you can enable it now on each node.

Enabling HTTPS has the following impact:

- All the nodes use HTTPS for communications within the cluster.

- When you enter commands using the CLI, you must include the `--https` option. These commands include the following:
 - The `run` script that installs or upgrades the software.
 - The `configure` script that configures the software.

⚠ Warning

If you enter one of these commands and omit `--https`, you will *inadvertently disable HTTPS* on the node.

Where to perform the configuration

Make sure you perform the configuration on the correct nodes.

Node	Work on this node?
Primary Conductor Live node	Yes
Secondary Conductor Live node	Yes
Each worker node	Yes

Topics

- [Enabling HTTPS](#)
- [Disabling HTTPS](#)

Enabling HTTPS

i Note

This information applies only to Conductor Live 3.25 and earlier.

To enable HTTPS

1. For a worker node, if you have already added the node to the cluster, you must [remove it from the cluster](#).

For a Conductor Live node, if you have already added the node to the cluster, you must [disable HA \(high availability\)](#) and [remove the node from the cluster](#).

2. At your workstation, [start a remote terminal session](#) to the primary Conductor Live node.
3. Change to the directory where the configuration script is located and run the configuration script:

```
[elemental@hostname ~]$ cd /opt/elemental_se  
[elemental@hostname elemental_se]$ sudo ./configure --https --skip-all
```

The `skip--all` option means that the script enables HTTPS but doesn't change the configuration in any other way.

Note

If you run this command (with the `--https` option) when HTTPS is already enabled, nothing changes in the configuration. HTTPS is still enabled.

Disabling HTTPS

Note

This information applies only to Conductor Live 3.25 and earlier.

Change to the directory where the configuration script is located and run the configuration script:

```
[elemental@hostname ~]$ cd /opt/elemental_se  
[elemental@hostname elemental_se]$ sudo ./configure --skip-all
```

The `skip--all` option means that the script disables HTTPS but doesn't change the configuration in any other way.

Note

If you run the script without the `--https` option when HTTPS is already disabled, nothing changes in the configuration. HTTPS is still disabled.

Adding SDI input devices

Individual Elemental Live nodes in the AWS Elemental Conductor Live cluster might be set up with SDI cards. Each input on the card can have a direct cable connection to an SDI video input.

If your cluster deployment includes a router for handling SDI (instead of, or in addition to, direct cable connections), see [the section called “Inputs: Routers for handling SDI inputs”](#).

Where to perform the configuration

Make sure you perform the configuration on the correct nodes.

Node	Work on this node?
Primary Conductor Live node	Yes. You perform the import step on this node.
Secondary Conductor Live node	No
Each worker node	Yes. You add the devices on these nodes.

Step A: Add the devices to each Elemental Live node

You don't configure these devices manually. Each Elemental Live automatically detects its SDI cards. It creates an *input device* and *inputs* for each card, as follows:

- One *single-link input* for each input on the card (so four inputs). Each input is given a unique numerical ID.
- One *quad-link input*, if the SDI card supports quad link.

The quad-link input is used with 4K quad input. When you're creating a profile, select this quad-link input to indicate to AWS Elemental Live that the four inputs on this SDI card are the four parts of a quad-link input.

Step B: Import the devices into the cluster

After you have added the devices on each Elemental Live node, you still need to import the devices so that primary Conductor Live detects the devices. You will perform this step when you [add the worker nodes to the cluster](#).

To verify the import, go to the **Settings** page and choose **Devices**. If any devices are missing, you might have forgotten to import them.

After you have imported a device, users will be able to select **SDI Direct Input** as the input type when they create a profile in Conductor Live.

Configuring SDI video routers

If your AWS Elemental Conductor Livecluster deployment includes a router for handling SDI inputs to Elemental Live nodes, you must configure the router in the cluster.

If SDI video inputs connect directly to the Elemental Live node with a cable, see [the section called “Input: Directly connected SDI inputs”](#).

Rules for routers

- The cluster can include only one SDI router.
- That router can serve several Elemental Live nodes.

If you try to set up a second SDI router, the configuration of the second router will fail in [the section called “Step D: Complete the Router Input Mappings”](#).

Where to perform the configuration

Make sure you perform the configuration on the correct nodes.

Node	Work on this node?
Primary Conductor Live node	You <i>configure</i> the router from the primary Conductor Live.
Secondary Conductor Live node	No.

Node	Work on this node?
Worker node	You physically <i>connect</i> the router to each Elemental Live node that will use it.

Warning

If you forget to configure the router, everything looks acceptable on the event or profile, but when you run the event, you receive a no input detected error.

Topics

- [Step A: Gather information](#)
- [Step B: Run cables from the router to each node](#)
- [Step C: Add the router](#)
- [Step D: Complete the Router Input Mappings](#)
- [Step E: Complete the Router Output Mappings](#)
- [Step F: Sync the Routers](#)
- [Step G: Use the Router Inputs](#)

Step A: Gather information

To set up the router information in Conductor Live, gather the following information:

- **Information on the router:** The name you want to use for the router, its IP address, and any information required by your router's protocol, such as user, level, or matrix ID.
- The **router input** port numbers. In the diagram below, these are the green circles on the left.

These are the numbers that the router uses for its input ports, not numbers generated by AWS Elemental software. The router inputs are called **Inputs** in Conductor Live.

- The **router output** port numbers being used by the cables between the router and worker nodes. In the diagram below, these are green squares numbered 1 through 7 on the right side of the router.

The router outputs are called **Outputs** in Conductor Live.


- The **inputs to each SDI card** on each node. These correspond to the port where cables are plugged into the node from the router.

The card inputs are referred to in the **Connected to** field in Conductor Live.

Example SDI Router Configuration

In the following example, *Node 1* has two SDI cards. On Card 1, four inputs are being used. On card 2, one input is being used. *Node 2* has one card, with two inputs in use.

In total, seven inputs are in use on the node, so you need seven outputs from the router. These seven outputs are shown on the right side of the router.



Step B: Run cables from the router to each node

Perform this procedure on the Elemental Live node that is connected to the router.

1. Make sure that the cables for the router and the SDI cards are connected on the Elemental Live node.
2. On the web interface for the Elemental Live node, go to **Settings** and choose **Input Devices** screen on the node. Make sure that the inputs with cables are all listed.

Repeat for each node that uses the router.

3. Identify all SDI inputs that you want to use on the node. You are not required to use all inputs on an SDI card. Make sure that you include inputs on any reserve nodes (these are used for node failover).

4. On the router, identify the IDs of inputs that have a cable connection.
5. On the router, identify the IDs of outputs that are connected to the SDI card on the worker node. Identify the input IDs that they are connected to. You must have one router output for each SDI input that you want to use.

Step C: Add the router

Perform this procedure on the primary Conductor Live node.

1. On the web interface for the primary Conductor Live node, go to the **Settings** page and choose **Routers**.
2. On the **Routers** page, choose **Add Router** and select the type of router protocol. These are the available options:
 - Videohub Ethernet Protocol (previously BlackMagic VideoHub)
 - XY Terminal Protocol (previously Harris Panacea)
 - NV9000 Protocol (previously Miranda nVision)
 - SW-P-08 Protocol (previously Snell Aurora)
 - PassThrough Protocol
 - LRC Protocol
3. Complete the **Add New Router** fields as described in the table and choose **Add**.

Field	Description
Name	The name that appears in the Inputs field on events and profiles.
IP Address	The IP address of the router, excluding the protocol.
Level	Applies to the XY Terminal, NV9000, SW-P-08, and PassThrough.
User	Applies to the NV9000.
Matrix ID	Applies to the SW-P-08.

Step D: Complete the Router Input Mappings

Perform this procedure on the primary Conductor Live node.

You must provide Conductor Live with the port numbers for the router inputs that end up at Elemental Live nodes. When you do this, the router's inputs number are mapped to an ID that is automatically generated by the software.

Important

An *input mapping* in Conductor Live does not associate an input with an output. Instead, it specifies the inputs on the router that have cables attached, and it specifies which of those inputs you want to use. The router internally manages the mapping from router inputs to router outputs.

To complete the router input mappings

1. On the Conductor Live web interface, choose **Settings**, then choose **Routers**. Choose the router.
2. Choose **Map Inputs**. On the dialog that opens, complete the fields to identify the inputs that you're using. You must know the identification of each input on your router. Conductor Live can't detect information about the state of the input IDs. You can add inputs with the following options:
 - **Add**: Adds one or more inputs. Enter the number of inputs to be created. Conductor Live assigns input numbers to each, beginning with the first number after the largest one already created.
 - **Add inputs starting at**: Adds a range of inputs. Enter the first and last number in the range. Use this to create inputs that start at a number beyond 1 and are not consecutive with the existing input numbers.
 - **Add 4 Quadrant-4k inputs starting at**: Adds four inputs grouped together to create an HEVC input.

⚠ Important

When you enter input numbers, they must be the same as the identification of each input on your router. For example, you must know that the second input from the left on the router is *input 2*. Conductor Live can't detect information about the disposition of input IDs.

3. Choose **Add** (+ icon).
4. Repeat these steps to add all of the inputs that you need.

Step E: Complete the Router Output Mappings

Perform this procedure on the Conductor Live node.

You must map each router output to each SDI input that you plan to use. This mapping must reflect the actual cabling from the output side of the router to the input side of the SDI card.

To map the SDI outputs

1. On the Conductor Live web interface, choose **Settings**, then choose **Routers**. Choose the router.
2. Choose **Map Outputs**. Complete the first line as follows and choose **Add** (+ icon):
 - **Output:** Select an output that is one of the cabled router outputs that you plan to use. The available options have the form *Output X*, where *X* is a number that corresponds to the appropriate router output port. For example, if your cabling comes from the router's output port 20, choose **Output 20**.

i Note

The correct number for the output is determined by the router, not by Conductor Live.

- **Connected to:** Select the card and node that the router output is connected to. The node displays the cards that it has auto-detected.

If you added a 4 Quadrant-4k input in [the section called “Step D: Complete the Router Input Mappings”](#) and want to map those four inputs, choose the Quadrant 4k (HD-SDI) card. This maps all four inputs to the one output.

3. Repeat the previous steps for each line to create all necessary output mappings.

Step F: Sync the Routers

Perform this procedure on the Conductor Live node.

After you have created the router output mappings, choose **Sync Routers** (at the top right of the page) to push relevant router information down to all worker nodes.

Step G: Use the Router Inputs

When you create a profile or event, the inputs that you created are displayed in the **Input** field.

Note that you specify the input by identifying the router input, not by identifying the SDI input. So you are giving the instruction "Use the input that is coming in on Input 1 using the Videohub Ethernet Protocol." When you run the event, Conductor Live directs this input to a free SDI card. Each time you run the event, a different SDI input could be used.

Avoid Direct Inputs

Typically, all of your SDI inputs are connected to your router. Therefore, you should only ever specify the input by selecting one of the router inputs. You should *not* use any of the “direct inputs,” such as those labeled as HD-SDI.

If your inputs are all connected to your router and you select a direct input in the profile or event, when the event starts, an input not detected error occurs.

You should only use the direct inputs for inputs on the Conductor Live node that do not connect to the router but are instead direct inputs.

Adding mount points to worker nodes

Create mount points if you need to make remote assets available on the Conductor Live cluster. Remote assets include scripts, image files, and video source files.

The mount folder becomes a mount share. It's mounted to `/data/mnt/folder`.

Where to perform the configuration

Make sure you perform the configuration on the correct nodes.

Node	Work on this node?
Primary Conductor Live node	Yes
Secondary Conductor Live node	No
Each worker node	Typically no.

How mount points work

When you mount a remote folder to a local folder on the node, all of the contents of the remote folder appear as if they are actually in the local mount folder. In this way, you can view the remote folder and verify that the backup files are created. You can also copy or delete a file from the remote folder by copying or deleting it from this mount folder.

About synchronization

Mount points on the primary Conductor node automatically synchronize to the secondary Conductor Live node and worker nodes. The sync occurs within one hour for nodes that are already part of the cluster when the mount is created. For nodes added to a cluster with existing mount points, the sync occurs within three minutes.

Creating mount points on a worker node

Generally, there is no need to create mount points on a worker node. Create them here if you want only one worker node to work with the folder, perhaps for security reasons.

To create the mount point, you must use the CLI. Mount points you create using the worker web interface are overwritten the next time that the primary Conductor Live synchronizes data on all the nodes.

To create a mount

1. On the web interface of the primary Conductor Live node, go to the **Settings** page and choose **Mount Points**.

2. On the **Mount Points** page, choose **Add Mount Point**, complete the mount point fields as described in the following table, and choose **Create**.

Field	Description
Type	Choose the type of remote server: • CIFS: Choose this for a Windows CIF server or for a Windows, Linux, or Mac SMB server. • NFS: Choose this for a Linux server. • DAVFS: Choose this for a DavFS server.
Server Share	The address of the folder that you want to make available on this node. This is an address on the remote computer.
Mount Folder	The folder on the node where the remote folder is mounted. As shown, this folder must be under <code>/data/mnt</code> . You can specify a sub-subfolder; if that folder does not already exist, Conductor Live automatically creates it.
User name	If the remote server folder is protected with user credentials, enter the username here.
Password	If the remote server folder is protected with a user credentials, enter the password here.

After a few minutes, the newly mounted folder appears on the web interface.

Setting the web interface time zone

This information applies to AWS Elemental Conductor Live, AWS Elemental Live and AWS Elemental Statmux. Follow this procedure if you didn't set the time zone when you installed the software on each node, or if you want to change the time zone on any node.

The time zone set on the node is used as follows:

- The web interface shows all activity with a timestamp for the time zone that you specify.
- Activity using the Linux CLI or the REST API doesn't use this time zone.

Where to perform the configuration

Make sure you perform the configuration on the correct nodes.

Node	Work on this node?
Primary Conductor Live node	Yes
Secondary Conductor Live node	No. Instead, the primary Conductor Live pushes its configuration to the secondary Conductor Live
Each worker node	Yes

To set the time zone on the Conductor Live

Perform this procedure on the primary Conductor Live.

1. On the web interface for the primary Conductor Live, go to the **Settings** page and choose **General**.
2. In **Timezone**, choose the time zone, then choose **Update**.

To set the time zone on the worker nodes

Perform this procedure on each worker node.

1. On the worker web interface, choose **Settings** from the main menu. Choose the **Network** tab, then choose **General**.
2. Choose the time zone, then choose **Save**.

Configure worker features

There are some features of the workers that you must configure individually.

Topics

- [Configuring OCR](#)
- [Configuring for RTMP inputs](#)
- [Configuring virtual input switching](#)

Configuring OCR

Elemental Live includes a feature that lets you convert captions using OCR conversion. For information about this feature, see [Support for OCR Conversion](#) in the *AWS Elemental Live User Guide*.

If you want to use this feature, you must enable it. You might have enabled it when you installed the software on the Elemental Live node. If you didn't, you can enable it now using the configuration script (configure) instead of the install script. For example:

```
[elemental@hostname ~]$ cd /opt/elemental_se  
[elemental@hostname elemental_se]$ sudo ./configure --install-ocr --https --skip-all
```

For complete instructions, see [Install the AWS Elemental Live software](#) in the *AWS Elemental Live Install Guide*.

Configuring for RTMP inputs

The Elemental Live nodes are configured by default to support RTMP inputs. In this mode, Elemental Live is using processing resources to continually poll for input at the RTMP port. If you don't plan to support RTMP inputs, you can choose to disable these inputs, to release the processing resources.

To disable polling for RTMP inputs

If you want to enable this feature after you've enabled user authentication, you must log into the Elemental Live node as an administrator. Regular users can't log into the worker nodes.

1. On the Elemental Live web interface, go to **Settings** and choose **Advanced**.
2. Set **Enable RTMP input** to unselected.

Configuring virtual input switching

On ECL3; node, you can configure the maximum number of virtual inputs allowed with the virtual input switching feature. The default is 8 inputs on the node. For information about this feature, see [AWS Elemental Live User Guide](#).

To set the number of virtual inputs

If you want to enable this feature after you've enabled user authentication, you must log into the Elemental Live node as an administrator. Regular users can't log into the worker nodes.

1. On the Elemental Live web interface, go to **Settings** and choose **Advanced**.
2. Enter a number in **Maximum number of virtual inputs**.

Configuring the cluster

A AWS Elemental Conductor Live cluster consists of Conductor Live, Elemental Live, and Elemental Statmux nodes. To set up a cluster, you must add these nodes to the cluster. If you are implementing node redundancy, you must also create redundancy groups, then add each node to its redundancy group.

Topics

- [Managing nodes in the Conductor Live cluster](#)
- [Creating redundancy groups](#)
- [High availability \(HA\)](#)

Managing nodes in the Conductor Live cluster

Topics

- [Adding \(recruiting\) worker nodes to the cluster](#)
- [Removing a worker node from the cluster](#)
- [Removing a Conductor Live node from the cluster](#)

Adding (recruiting) worker nodes to the cluster

You can add nodes into the cluster. Here are some scenarios where you add nodes:

- When you initially configure the cluster, you must add the secondary AWS Elemental Conductor Live node and all the worker nodes into the cluster. You perform this step after you have installed the software on the primary Conductor Live node. At this point, the cluster exists but it contains only the primary node.
- You might remove a node and then want to add it back into the cluster. For example, if you decide to add user authentication to an existing cluster, one of the steps is to remove nodes. In this case, you must remove all the nodes, including the primary Conductor Live, and then must add all of them back into the cluster.
- You can add a newly obtained worker node into an existing cluster. In this case, there is no need to make any preliminary changes to the cluster. You simply add the new worker node.

Step 1: Get ready to add a secondary Conductor Live

If you are adding a secondary Conductor Live, make sure that the primary and secondary Conductor Live nodes have the same [firewall settings](#). If they don't, you won't be able to add the secondary Conductor Live.

Step 2: Get ready to add a worker node

This preliminary step applies only in the following situation:

- You are adding worker nodes to an existing cluster.
- The Conductor Live node or nodes in the cluster have Conductor Live version 3.25.2 or lower. The existing worker nodes have version 2.25.2 or lower installed.
- The new worker node or nodes you want to add have version 2.25.3 or higher.

For example, the Conductor Live nodes have version 3.23.0. You want to add workers that have version 2.25.5. You can do this, because the versions are within two major versions of each other.

The preliminary step is to set the `LEGACY_RECRUIT` environment variable to `True` on each worker node. Perform this step now.

Step 3: Add a node to the cluster

Perform the following steps on the primary Conductor Live node.

1. On the primary Conductor Live web interface, choose the **Cluster** page, then choose **Nodes**.
2. On the **Nodes** page, choose **Add Node**. The **Add Nodes to Cluster** dialog appears showing two fields:
 - **Node IP Addresses**
 - **Lookup Node IP Address**
3. Specify the addresses of the nodes to add. You can enter one address, several comma-delimited addresses, or a range of addresses. If your network has a DNS server, you can look up an address by entering a hostname.
4. When you've entered all the nodes to add, choose **Add**.
5. If an Elemental Live node has SDI cards, you must import the devices so that Conductor Live recognizes them. (You [configured these devices](#) when you configured each worker for the network.)

Choose the down arrow beside the node and select **Import Devices**.

Removing a worker node from the cluster

Generally, you remove a node only in these situations:

- To enable HTTPS in the cluster. For the complete procedure for enabling HTTPS, see [the section called "HTTPS"](#).
- To move a node to another Conductor Live cluster.
- To isolate a node, perhaps for troubleshooting purposes.
- To retire a node, perhaps because you are upgrading your hardware.

To remove a node

1. To remove an Elemental Live node, make sure that no channels are associated with the node:
 1. On the web interface of the primary Conductor Live, go to the **Channels** page. Filter the channels list so only the channels associated with this node are displayed. Make a note of these channels.
 2. Either wait for each channel to complete or manually stop a channel by choosing **Stop** beside the channel.
 3. For channels that are associated with the node that you're moving, change the node association by these steps: Choose **Edit** (pencil icon) on the stopped channel and in **Node**, select **None**.
2. To remove either an Elemental Live or AWS Elemental Statmux node, make sure there are no MPTS outputs associated with the node:
 1. On the primary Conductor Live node's web interface, go to the **MPTS** page to verify which node each MPTS output is using.
 2. If an MPTS output is using the node that you're deleting, set the MPTS output to use a different node. See the [AWS Elemental Conductor Live User Guide](#) for details.
3. Go to the **Redundancy** page and find the node that you want to remove. On the row for the node, choose the **Delete** (garbage can). On the dialog that appears, choose **OK**.
4. Go to the **Nodes** page and find the node that want to remove. On the row for the node, choose the downward triangle and select **Remove Node**. At the prompt, choose **Remove**.

Removing a Conductor Live node from the cluster

Generally, you remove a node only in these situations:

- As one of the steps when you enable HTTPS in the cluster. For the complete procedure to enable HTTPS, see [the section called “HTTPS”](#).
- To retire a node, perhaps because you are upgrading your hardware.

To remove a Conductor Live node

Follow this procedure to remove either the primary or the secondary Conductor Live. In both cases, you perform the procedure on the primary Conductor Live web interface.

1. [Disable HA](#) on the cluster.
2. Still on the **Redundancy** page, find the node that you're removing. On the row for the node, choose the **Delete** (garbage can). On the dialog that appears, choose **OK**.
3. Go to the **Nodes** page and find the node that you're removing. On the row for the node, choose the downward triangle and select **Remove Node**. At the prompt, choose **Remove**.

Creating redundancy groups

You must create redundancy groups as follows:

- If your cluster includes a primary and a secondary AWS Elemental Conductor Live node, you must create a redundancy group and add both nodes to the group. It isn't enough to just add the Conductor Live nodes to the cluster.
- If you want to set up worker nodes for failover resiliency, you must create one or more redundancy groups, then you add worker nodes to each group.

For general information about how failover resiliency works, and for detailed information about designing redundancy groups that meet your requirements, see [AWS Elemental Conductor Live User Guide](#). Then come back to this section to create the groups and add the nodes.

Topics

- [Creating a Conductor Live redundancy group](#)

- [Creating worker redundancy groups](#)

Creating a Conductor Live redundancy group

If you are implementing Conductor Live redundancy, then you should have two Conductor Live nodes — a primary node and a secondary node. You must create a redundancy group and add the nodes to this group.

Redundancy groups and a VLAN

The redundancy group that you create always runs in HA (high availability) mode. In this mode, the two Conductor Live nodes must be on the same VLAN

The VIP and the Conductor Live redundancy group

When you create the redundancy group, as described in the procedure that follows, you assign a virtual IP address to the group.

This address serves as the constant *cluster ID* for the primary and secondary Conductor Live nodes. When you enable HA (high availability), one of the two Conductor Live nodes registers as the primary node with this VIP. This node you go to enable HA, and it is the node that initially registers as the primary. Later, any time a Conductor Live node failover occurs, the node that's promoted to primary re-registers with the VIP to indicate that this node is now the primary node.

To create a Conductor Live redundancy group

1. Make sure that the secondary Conductor Live node is in the cluster [the section called “Add \(recruit\) nodes”](#).
2. On the primary Conductor Live web interface, go to the **Cluster** page and choose **Redundancy**.
3. On the **Redundancy** page, choose **New Redundancy Group** and select **Elemental Conductor Live**.
4. In the **Add New Redundancy Group** dialog, complete the fields and choose **Add**. See the table for information on each field.

Field	Description	
Redundancy Group Name	Any name that you choose.	

Field	Description	
Virtual IP Address	A valid IPv4 address. The address must meet these conditions: • It must be an address on your network that will never be allocated to any other host. • It must be on the same subnet as the Conductor Live nodes.	
Virtual Router Identifier (VRID)	The VRID must meet these conditions: • It must be an integer 1–254. • The value must not conflict with any other instance of keepalived (or any other VRRP service) that's running on the network. You must make sure that there are no conflicts. Elemental Live can't detect them.	

To add Conductor Live nodes

Follow these steps on the primary Conductor Live node.

1. On the **Redundancy** page, select the Conductor Live redundancy group. Choose **Add HA Nodes**.
2. On the dialog, select a Conductor Live node from the **Nodes** dropdown list.
3. Choose **Add**.

Creating worker redundancy groups

To set up worker nodes for failover resiliency, you create one or more redundancy groups, then you add worker nodes to each group.

For general information about how failover resiliency works, and for detailed information about design redundancy groups that meet your requirements, see [Conductor Live User Guide](#).

To create a redundancy group

1. On the primary Conductor Live web interface, go to the **Cluster** page and choose **Redundancy**.
2. On the **Redundancy** page, choose **New Redundancy Group** and select the node type.
3. Enter a name for the redundancy group and choose **Add**.

The group is added to the list on the left side of the Redundancy screen. At this point, no nodes are in the group.

To add nodes

Follow these steps on the primary Conductor Live node.

1. On the **Redundancy** page, select the group that you're adding nodes to. Two tabs appear on the right — **Active Nodes** and **Backup Nodes**.
2. Select **Active Nodes** tab, then choose **Add Active Nodes**.
3. On the dialog, select a node from the **Nodes** dropdown list. Only nodes that aren't in a redundancy group appear in this list.

If you are setting up an N+M type of group, add several nodes.

4. Choose **Add** to add the selected nodes to the group.
5. Repeat these steps to add nodes to the **Backup Nodes** tab.

The nodes are listed on the **Active Nodes** or the **Backup Nodes** tab of the redundancy group.

Make sure that you add nodes to both tabs.

High availability (HA)

If you have set up a primary and a secondary AWS Elemental Conductor Live node in a redundancy group, you must enable HA in order to set the nodes up as an active redundant pair.

Enabling HA is typically the very last step that you perform when you configure a cluster. You enable HA in order to start working with the cluster.

Verifying the current HA state

On the web interface for the primary Conductor Live node, go to the **Cluster** page and choose **Redundancy**.

Look for the button on the top right corner. If there is a green button that says **Enable**, then you know that HA is disabled.

Enabling Conductor Live HA (high availability)

If you have set up the cluster with a primary and a secondary Conductor Live node, you must enable HA before you start running channels and MPTSes in the cluster.

The results of enabling HA

When the enable HA (high availability), the following actions occurs:

The primary Conductor Live is set

The Conductor Live node where you enable HA becomes the primary Conductor Live. The other Conductor Live node becomes the secondary Conductor Live.

The virtual IP (VIP) address activates

The virtual IP (VIP) address activates. The primary Conductor Live registers with this VIP as the primary node. After this occurs:

- The primary Conductor Live node becomes inaccessible through its local IP address and only accessible through the virtual IP (VIP) address. The web page you're using automatically redirects to the VIP as you continue to access the primary Conductor Live node.
- The secondary Conductor Live node becomes completely inaccessible. If you enter the IP address of the secondary Conductor Live node, you are automatically redirected to the VIP. In

other words, while you are accessing the primary Conductor Live node, you cannot access the secondary Conductor Live node.

- Any time that a Conductor Live failover of the primary node occurs, the Conductor Live node that's promoted to become the primary Conductor Live node re-registers with the VIP, effectively indicating "I'm the primary now."

Information is copied over

Information is copied over from the primary Conductor Live node to the secondary Conductor Live node, and the secondary Conductor Live database synchronizes itself with the primary Conductor Live database. The following information that is copied over:

- Input device (SDI) configuration
- Mountpoint configuration
- Router configuration
- Time zone configuration
- Configuration of notification for alerts and SNMP
- User authentication and users
- Backup and restore configuration

To enable HA

1. Make sure that both of the Conductor Live nodes are on the same software version. If they are different, you will receive a validation error and HA won't get enabled.
2. If you're using a virtual machine (VM), take a snapshot before you enable HA. See the VMware VSphere help text for more information.
3. On the web interface for the primary Conductor Live node, go to the **Cluster** page and choose **Redundancy**. In the **High Availability** field, choose **Enable**.

To verify that HA is correctly enabled, follow these steps on each Conductor Live node.

1. At your workstation, [start a remote terminal session](#) to the Conductor Live node.
2. Enter the following command to verify that Conductor Live HA is enabled:

```
[elemental@hostname log]$ 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. Watch for the following INFO lines on the primary Conductor Live 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
```

3. Press **Ctrl - C** to exit the tail command.

4. 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

5. Confirm that you see database system is ready to accept connections on the primary Conductor Live, and database system is ready to accept read only connections on the secondary Conductor Live.

6. Press **Ctrl - C** to exit the tail command.

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

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

Disabling Conductor Live HA (high availability)

The main reason to disable HA is to make a change to the configuration of one or both Conductor Live nodes. You must disable HA to make these configuration changes:

- Configure DNS

- Configure Ethernet interfaces
- Bond Ethernet interfaces
- Configure the firewall or ports
- Enable user authentication

Important

Disabling HA occurs immediately, but enabling always involves a wait. Don't disable without a good reason!

To disable HA

1. If you're using a virtual machine (VM), take a snapshot before you disable HA. See the VMware VSphere help text for more information.
2. On the web interface for the primary Conductor Live node, go to the **Cluster** page and choose **Redundancy**.
3. Note the values in **Virtual IP Address** and **Virtual Route Identifier**. You will use these when you re-enable HA.
4. In the **High Availability** field, choose **Disable**.

To verify that HA is disabled

1. At your workstation, [start a remote terminal session](#) to each Conductor Live node.
2. In the session for each Conductor Live, enter the following command to verify that Conductor Live HA is disabled:

```
[elemental@hostname log]$ 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. Watch for the following INFO lines on the primary Conductor Live 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
```

Make sure that the secondary Conductor Live is also ready.

3. Press **Ctrl - C** to exit the tail command.
4. 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

5. Confirm that you see database system is ready to accept connections on the secondary Conductor Live.
6. Press **Ctrl - C** to exit the tail command.
7. Type the following command to exit the session as the sudo user:

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

Configuring user authentication in Conductor Live

You can enable user authentication on the cluster, so that all users must provide valid credentials to work from either the web interface or the REST API. When user authentication is enabled, users must provide the following credentials:

- For the web interface, users must enter user credentials—a user name and a password.
- For the REST API, users must include these additional HTTP headers (X-Auth-User, X-Auth-Expires, X-Auth-Key) in commands that they send.

For more information about using the API with authentication enabled, see the AWS Elemental Conductor Live REST API documentation.

Benefits

User authentication has the following benefits:

- It prevents unauthorized access to nodes.
- It lets an administrator track node activity on a per-user basis.

Topics

- [About user authentication](#)
- [Step 1: Enable the user authentication feature](#)
- [Step 2: Apply user authentication on worker nodes](#)
- [Disabling user authentication](#)

About user authentication

Summary of procedure

To set up for user authentication, you perform three steps:

- Step 1 — Enable the feature. You must enable the user authentication feature. You perform this step on the primary Conductor Live node. This step configures enables user authentication at the cluster level. See [the section called “Step 1: Enable user authentication”](#).

- **Step 2 — Apply user authentication on the nodes.** You must enable user authentication on every node. You perform this step once for all nodes, on the primary Conductor Live node. This step configures the individual nodes to require that users log in. See [the section called “Step 2: Apply authentication”](#).
- **Step 3 — Create users.** You create a user on the primary Conductor Live. These users now work from the Conductor Live to perform any work on the cluster. These users can't work on the individual worker nodes.

Typically, you also set up one or two users on the individual worker nodes, but only so that someone can perform troubleshooting tasks on these nodes.

Types of user authentication

There are two ways to implement user authentication. For both types, the first two setup steps are the same. Only the step for adding users is different.

- **Local authentication**

With this authentication, you enable authentication on the Conductor Live node.

You then create users for the entire cluster from the primary Conductor Live node. See [Reference: Manage users](#). Users are assigned a role that controls the user's permissions. These roles are built into Conductor Live. You can't modify the roles or create new roles.

- **PAM authentication**

With this authentication, you enable authentication on the Conductor Live node.

You create user credentials from an LDAP server that is external to the AWS Elemental nodes. The credentials that you assign to the users are stored on the LDAP server.

Step 1: Enable the user authentication feature

There are two steps to enabling user authentication in the cluster.

- The first step is to enable the *user authentication* feature. You perform this step on the primary Conductor Live, by running the configuration script.

- The [second step](#) is to apply user authentication to all the nodes in the cluster. To perform this step, you enable *node authentication*. You perform this step on the primary Conductor Live node, not on each worker node.

This procedure applies to both types of user authentication—local authentication and PAM authentication.

Where to perform the configuration

Make sure you perform the configuration on the correct nodes.

Node	Node where you perform this task
Primary Conductor Live node	Yes
Secondary Conductor Live node	No
Each worker node	No

To enable user authentication

To enable user authentication, follow these steps.

1. If HA redundancy is currently enabled on the Conductor Live node, [disable it](#).
2. At your workstation, [start a remote terminal session](#) to the Conductor Live node.
3. Change to the directory where the configuration script is located, then enter the configure command to enable HTTPS:

```
[elemental@hostname ~]$ cd /opt/elemental_se  
[elemental@hostname elemental_se]$ sudo ./configure --https --skip-all
```

The `--https` option enables HTTPS. When HTTPS is enabled, all user names and passwords are encrypted. When you enable user authentication, you should always enable HTTPS.

4. Enter the configure command again to enable user authentication:

```
[elemental@hostname elemental_se]$ sudo ./configure --config-auth
```

Note

Enter the configure command twice, as shown. Don't enter a command that combines the `--https` and `--config-auth` options because HTTPS won't get enabled.

5. Answer the authentication prompts as follows:

Prompt	Value to enter		
Do you wish to enable authentication?	Y		
Do you wish to enable PAM?	Y to enable PAM authentication N to enable local authentication		
Enter admin login	We recommend that you set up this default user as the API admin. Therefore, don't accept the default. Instead, assign the name <i>apiadmin</i> . For information about this user, see the section called "Types of users" .		
Enter admin email	Enter an email address.		

Prompt	Value to enter		
Enter admin password:	Create a strong password for <i>apiadmin</i> . The password must be strong: Minimum 8 characters, at least one uppercase letter, at least one lowercase letter, at least one number, and at least one symbol.		
Httpd must be restarted , which may interrupt REST commands. Restart now?	Y		

- After the configuration script has run, the following message appears. This message reminds you that users must include these additional HTTP headers in commands that they send.

Authentication has been enabled. The REST interface will require authentication as well. Please look at the REST Interface section of the Support for more information.

- When the service starts and the Conductor node is ready, [re-enable HA](#), if applicable.
- Make a note of the user name and password for *apiadmin*.

Result of this procedure

You have enabled user authentication on the primary Conductor Live node. You have also created an API admin (named *apiadmin*). This user has a specific role. For more information, see [the section called "Types of users"](#).

Step 2: Apply user authentication on worker nodes

The following steps describe how to enable *node authentication* in the cluster. Before you enable node authentication, you must [enable user authentication](#).

This procedure applies to both types of user authentication—local authentication and PAM authentication.

Where to perform the configuration

Make sure you perform the configuration on the correct nodes.

Node	Node where you perform this task
Primary Conductor Live node	Yes
Secondary Conductor Live node	No
Each worker node	No

To enable user authentication

To enable user authentication on all the worker nodes, you log onto the primary Conductor Live node and display the **Cluster Nodes** page.

1. Make sure you have followed the procedure in [Step 1: Enable the user authentication feature](#).
2. Go to the Conductor Live web interface by entering the IP address of the primary Conductor Live node in a web browser. Log into the web interface as the API admin (*apiadmin*). You created this user when you enabled user authentication on the Conductor Live node (in the [previous step](#)).
3. On the main menu, choose **Cluster**, then **Nodes**. Choose **Tasks** (in the top left corner) and select **Enable Node Authentication**.
4. On the **Select a user name** page, choose **apiadmin** and choose **Next**.

In this step, you are identifying the administrator that will serve as the reserved API user. When you use this name, you set up so that the API administrator has the same name (*apiadmin*) on all nodes. This practice reduces confusion. For more information about this user, see [the section called “Types of users”](#).

5. On the **Enter a password** page, enter the existing a password for *apiadmin*. When you enter the same password, you keep the password for *apiadmin* aligned on all the nodes in the cluster. This practice reduces confusion.
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**.

Conductor Live enables user authentication on each node. It also creates the API admin (*apiadmin*) on each worker node.

Refresh the page to track the progress of the action. When all the nodes are ready, the Nodes page displays each node with a lock icon to indicate user authentication is enabled.

Result of this procedure

You have enabled user authentication on the secondary Conductor Live node and each worker node. You have also propagated the [API admin \(*apiadmin*\)](#) to all these nodes.

Disabling user authentication

This section describes how to disable user authentication on the cluster.

To disable user authentication

You disable user authentication by running the configuration script in the same way as you ran it to enable user authentication. In other words, if authentication is enabled and you include `--config-auth`, then the script disables authentication.

1. [Check if HA redundancy](#) is currently enabled on the Conductor Live. If it is, [disable it now](#) [Disabling Conductor Live HA \(high availability\)](#).
2. At your workstation, [start a remote terminal session](#) to the primary Conductor Live node.
3. At the Linux prompt, log in with the *elemental* user credentials.
4. Change to the directory where the configuration script is located, then enter the configure command :

```
[elemental@hostname ~]$ cd /opt/elemental_se
```

```
[elemental@hostname elemental_se]$ sudo ./configure --config-auth
```

5. Answer the prompts as follows:

Prompt	Value to enter		
Httpd must be restarted , which may interrupt REST commands. Restart now?	Y		
Do you wish to update the admin user?	N		

When you return to the web interface, you are not prompted to log in, and the menu to enable or disable node authentication on the worker nodes has disappeared.

6. The configure command in the previous step disables user authentication but leaves HTTPS enabled. Therefore, if you also want to disable HTTPS, enter the configure command as follows:

```
[elemental@hostname elemental_se]$ sudo ./configure --https --skip-all
```

7. If applicable, [re-enable HA](#).

Managing users in Conductor Live

If you have enabled user authentication, you must add users to the cluster. After you've added users, you can manage existing users and add new users. If you don't enable user authentication, there is no need to create users. However, we strongly advise against deploying a cluster without user authentication enabled.

You perform user management tasks as follows:

- If you've set up with local authentication on the cluster, you manage users and user roles using AWS Elemental Conductor Live, as described in this section.
- If you've set up with PAM authentication on the cluster, you manage users on your organization's LDAP server.

Note

The username of a user is case sensitive.

Topics

- [Types of users](#)
- [Adding users to Conductor Live](#)
- [Adding users to worker nodes](#)
- [Role policies for PAM authentication](#)

Types of users

In a cluster, there are several types of users, as described in the following sections.

Topics

- [Users on Conductor Live nodes](#)
- [Users on worker nodes](#)
- [Summary](#)

Users on Conductor Live nodes

You can set up several different types of users on Conductor Live nodes.

The *elemental* user

- Purpose: When you start an SSH session to work on the software through the operating system, you log in using this user. You can't log into SSH using any other user.
- How created: This user is built into the software. You are prompted to change the password for this user either the first time you installed the software on the node, or the first time that you ran the configure script on the node.

Warning

You are prompted to change the password for this user. Make sure that you assign the same password on every node. Otherwise, you won't be able to enable user authentication on some nodes.

The default user on Conductor Live

- Purpose: This user is the default administrator.
- How created: You might have created this user when you installed the software. You might have enabled user authentication when you installed the software. In this case, you probably accepted the suggested name of *admin* for the default user.

Or you might have created this user when you followed the procedure in [the section called “Step 1: Enable user authentication”](#) to enable user authentication on Conductor Live. In this case, the default user has the role of the API admin. You should have called this user *api-admin*.

The *API admin* user on Conductor Live

- Purpose on the Conductor Live nodes: When the cluster is running (after you've configured it), Conductor Live uses the API key of the API administrator for authentication when sending commands to worker nodes. For security reasons, you should not use this administrator for regular administrative tasks.

Purpose on worker nodes: This administrator is the boot strap administrator on Conductor Live. After you've first enabled user authentication, you log into worker nodes using this administrator's credentials, and [create regular administrators](#) on the worker nodes.

- How created on the primary Conductor Live node: You should have created this user when you ran the configuration script or the installer to [enable user authentication](#).

How created on other nodes: You created the user when you [enabled node authentication](#) on all the other nodes in the cluster. Conductor Live then creates this user on every worker node.

- Username for this user: You should assign the name *api-admin*.
- Working with this user: After you've set up user authentication on the cluster, don't log in as this user. Instead, reserve it for its main role, which is to authenticate API commands between nodes in the cluster. To perform regular administration tasks, log in as a regular administrator.

Regular administrators on Conductor Live

- Purpose: Regular administrators have the same access as the default admin user. They have full read-write access, including the ability to create and manage users.
- How created: Any administrator creates these administrators on the primary Conductor Live. These users are pushed to the secondary Conductor Live when you [enable HA \(high availability\)](#) on the cluster. They aren't pushed to worker nodes.
- Username for this user: Typically, assign the person's name as the username.

Operators on Conductor Live

- Purpose: Operators have full read-write access, except that they can't create or manage users.
- How created. Any administrator creates these operators on the primary Conductor Live. These users are pushed to the secondary Conductor Live when you [enable HA \(high availability\)](#) on the cluster. They aren't pushed to worker nodes.
- Username for this user: Typically, assign the person's name as the username.

Viewers on Conductor Live

- Purpose: Operators have read-only access to all functions, except that they have no access to users.
- How created: Any administrator creates these operators on the primary Conductor Live. These users are pushed to the secondary Conductor Live when you [enable HA \(high availability\)](#) on the cluster. They aren't pushed to worker nodes.
- Username for this user: Typically, assign the person's name as the username.

Users on worker nodes

When workers are in a Conductor Live cluster, you only need to set up one or two users, to let you log on directly to the node in order to troubleshoot.

Regular administrators on worker nodes

- Purpose: Set up people as regular administrators only if they will perform troubleshooting on worker nodes. Typically, you set up the same person (for example, a manager) as a regular administrator on every worker node. Make sure that these people understand that they should log into an individual node only to troubleshoot. They should perform all regular activities on Conductor Live.
- How created: You should create these administrators on each node (Conductor Live and the worker nodes).
- Username for this user: Typically, assign the person's name as the username.

Regular users on worker nodes

Only standalone setups of Elemental Live have regular users — managers, operators, and viewers.

Note

Don't set up regular users on worker nodes when your organization is using a Conductor Live cluster.

Summary

Following is a summary of the users that you must explicitly create.

Type of user	How created
Regular administrators on Conductor Live	You manually add these users by working on the primary Conductor Live
Operators on Conductor Live	You manually add these users by working on the primary Conductor Live
Viewers on Conductor Live	You manually add these users by working on the primary Conductor Live
Regular administrators on worker nodes	You manually add these users by working on each worker node

Adding users to Conductor Live

Read this section if you've set up for [local user authentication](#). You must add regular administrators, operators, and viewers to the cluster. You add these types of users by working on the primary Conductor Live.

(If PAM authentication is enabled, you manage users on your organization's LDAP server.)

Note

When you add users to Conductor Live, keep in mind that the usernames that you assign are case sensitive. The user *Myuser* is not the same as the user *myuser*.

To add the first user

To create the first user, you must log in as the API admin (*apiadmin*). But after you've created this user, you should always log in as a regular administrator.

1. Log into the primary Conductor Live web interface as *apiadmin*. You created this user when you enabled user authentication

2. On the menu bar, choose **Settings**. Then choose **Users** from the left bar. Then choose **New User**.
3. Complete the fields as appropriate. Set up the user as **Admin**. You can leave the REST API key empty. Conductor Live will generate a key.
4. Choose **Create**. The user is created.

To add more users

Log in as a regular administrator and add more users.

1. Log into the primary Conductor Live web interface as a regular administrator.
2. On the menu bar, choose **Settings**. Then choose **Users** from the left bar. Then choose **New User**.
3. Complete the fields as appropriate. Choose any role. You can leave **API Key** empty. A key will automatically be generated.
4. Choose **Create**. The user is created with the specified role.
5. Give each user this information:
 - Give the user their user name (case sensitive) and password.
 - Advise the user to change their password. They must log onto the Conductor Live web interface. Then on the menu bar, they can select their name and choose **Account** from the dropdown menu. The **Account** page has a **Change Password** button in the top right corner.
 - If your organization uses the REST API, advise the user to make a note of their personal API key. They must log into the Conductor Live web interface. Then on the menu bar, they can select their name. The API key appears on the dropdown menu.
 - Tell the user how to log out. On the right side of the menu bar on any page, they select **Logout**.

Adding users to worker nodes

Read this section if you've set up for [local user authentication](#). You can add users to worker nodes.

(If PAM authentication is enabled, you manage users on your organization's LDAP server.)

On the worker nodes, you should only add regular administrators, and only so that they can troubleshoot problems on a node. We recommend that you create access only for one or two people. Typically you set up these people as regular administrators on a worker node.

- People who are administrators of the cluster. Make sure you create at least one administrator on each node.
- People who are managers of teams.

Note

The user names that you assign are case sensitive. The user *Myuser* is not the same as the user *myuser*.

To add users

1. Log into the worker node as *apiadmin*. If you followed the procedure in [the section called “Step 2: Apply authentication”](#), then this user is the only user that initially exists on the node.
2. Hover over **Settings** and choose **Users**, then choose **New User** (on the far right of the page).
3. Complete the fields as appropriate. You can leave **API Key** empty. A key will automatically be generated.
4. Choose **Create**. The user is created with the specified role.
5. Give each user this information:
 - Give the user their user name (case sensitive) and password.
 - Advise the user to display their user information. They must log into the worker web interface. Then on the menu bar, they can hover over **Settings** and choose **User Profile**,

Role policies for PAM authentication

If you configure the cluster to use PAM authentication, then when you create a user on the LDAP server, you assign a role policy. The role policies exist in Conductor Live, not on your LDAP server.

To view role policies

1. Log into the primary Conductor Live as an administrator.

2. On the menu bar, choose **Settings**. Then choose **Roles Policies** from the left bar.

Information about the supported role policies appears. Set up PAM authentication to use these role policies.

Configuring notifications for messages

AWS Elemental Conductor Live provides status information through alerts and messages. You can configure these notifications so you know when the node might need attention. The following table describes the differences between alerts and messages and how you can access each.

In the table, find a topic in the first column, then read across for information about alerts and about messages available for this topic.

Topic	Alerts	Messages
Ways in which you can access alerts and messages	• Web interface • REST API calls • SNMP poll • SNMP trap • Email notification • Web callback notification	• Web interface • REST API calls • SNMP poll
Information conveyed	Alerts are feedback on a problem that must be fixed. The "Channel Error" alert informs you that a channel has moved to an Error state. This can be helpful when you are receiving automatic email notifications, letting you know to check for related messages on the web interface.	There are three types of messages: • AuditMessage: Informational messages that you do not need to react to. Often, these messages are feedback to actions you performed. • WarningMessage: Messages that advise you that there is a risk that a future activity will fail unless you take action to prevent it. • ErrorMessage: Messages that indicate that a planned activity has failed or an

Topic	Alerts	Messages
		unexpected system error has occurred.
Active or inactive	Alerts are active until the underlying problem is resolved. When the cause of the alert is no longer present, the system clears the alert, and it becomes inactive.	Messages are neither active nor inactive. They are defined as <i>recent</i> when they are less than 24 hours old.
Visibility (web interface only)	You can toggle the visibility of active alerts on the web interface. Suppressing an alert this way is similar to marking an email as read. Alerts are available through the other access options, regardless of their visibility in the web interface.	You can toggle the visibility of recent error messages on the web interface. This is similar to marking an email as read. Visibility does not affect the return on SNMP and REST requests.

The following sections describe how to set up notifications. For information about viewing alerts and messages on the web interface or through the API, see the *Conductor Live User Guide* and *Conductor Live API Guide*.

Topics

- [Email notification](#)
- [Web callback notification](#)
- [SNMP traps](#)
- [SNMP polling](#)

Email notification

You can configure AWS Elemental Conductor Live to email you notifications when alerts occur.

Note

Don't configure worker nodes to send email notifications. Conductor Live sends operational status for all nodes in the cluster and all transcoding channels.

Conductor Live uses open relay to send email notifications. Before subscribing to notifications, make sure that your network allows receipt of open relay email. If your network doesn't allow open relay messages, you must also configure a Sendmail relay server with another mail server.

Important

If you subscribe to email notifications in a network that doesn't allow open relay messages and you do not relay the messages, the generated messages will collect on the Conductor Live system hard drive, eventually filling the partition and causing disk alert errors.

To set up email notifications

1. On the Conductor Live web interface, subscribe to all or some alerts using the steps described here:

Subscribe to all alerts

1. On the Conductor Live web interface, go to the **Settings** page and make sure that you're on the **General** tab.
2. Complete the **Global Alert Notification** fields as described in the following table and choose **Update**:

Field	Instructions
Email	Enter the email address of the alert recipient. Required if you don't provide a URL in the Web Callback URL field.

Field	Instructions
Web Callback URL	If you want to receive web server notifications too, enter the URL of the appropriate .php file on your web server. For instructions on how to configure your web server for notifications, see Web callback notification.
Notify	Select when you want to be notified, either when the alert is raised or when it's cleared. You can choose both options.

Subscribe to individual alerts

1. On the Conductor Live web interface, go to the **Stats** page and choose **Notifications**.
2. On the **Notifications** page, find the alert that you want to be notified on and choose the plus sign (+) to expand it.
3. Complete the fields as described in the following table and choose **Save**.

Field	Instructions
Email	Enter the email address of the alert recipient. Required if you don't provide a URL in the Web Callback URL field.
Web Callback URL	If you want to receive web server notifications too, enter the URL of the appropriate .php file on your web server. For instructions on how to configure your web server for notifications, see Web callback notification.

Field	Instructions
Notify	Select when you want to be notified, either when the alert is raised or when it's cleared. You can choose both options.

4. For each alert that you want to be notified on, find the alert, then expand and complete the fields.
2. If your network doesn't allow open relay messages, configure the sendmail server to relay the messages. For steps, see [Configure sendmail relay server](#).

Configure sendmail relay server

Use this procedure to set up a Sendmail relay server if your network doesn't accept open relay messages.

Step A: Gather the mail server information

To configure Conductor Live to relay the notification emails through a mail server, you need the following information:

- The hostname of the mail server
- If your network doesn't have DNS configured, the IP address of the mail server

Step B: Install the sendmail configuration tool

To install the configuration tool

1. Install the `sendmail.cf` configuration tool by typing the following at the command line.

```
sudo yum install sendmail-cf
```

2. When you receive a caution message asking you to confirm that you want to run the command, enter **yes**.
3. When you receive the following prompt, enter **y**.

```
Is this ok [y/N]:
```

4. When you receive the following message, move on to the next step.

```
Complete!
```

Step C: Edit the license file

To edit the file

1. With a text editor, open the `sendmail.mc` file. If you use Nano, which comes installed on all AWS Elemental systems, type the following at the command line to open the file in Nano.

```
sudo nano /etc/mail/sendmail.mc
```

2. Find the line that defines `SMART_HOST`. It's generally just past halfway down the page and should look like this.

```
dn1 define(`SMART_HOST', `smtp.your.provider')dn1
```

3. Uncomment this line by deleting the `dn1` at the beginning and end of the line.
4. Change the following text to the hostname of the mail server that is performing the relay.

```
smtp.your.provider
```

5. Save and exit the file. For Nano, press `Ctrl+O` to save and `Ctrl+X` to exit.

Step D: Check the hosts file

If your network isn't configured with DNS, add a static entry to the `hosts` file on Conductor Live.

To add an entry to the hosts file

1. With a text editor, open the `/etc/hosts` file. If you use nano, type the following at the command line.

```
sudo nano /etc/mail/hosts
```

2. Add a line to the end of the file that has the IP address of the relay server, a space, and the hostname of the relay server. The following shows an example.

10.24.34.2 ExampleMailHostname

3. Save and exit the file. For nano, press Ctrl+O to save and Ctrl+X to exit.

Step E: Apply the changes

1. To apply changes, enter the following command.

```
sudo make -C /etc/mail
```

The system responds as follows.

```
make: Entering directory `/etc/mail'
make: Leaving directory `/etc/mail'
```

2. Restart Sendmail by typing the following.

```
sudo service sendmail restart
```

Step F: Test the new configuration

Test the relay by having the system email you an alert notification.

To test the configuration

1. If you haven't already, subscribe to global alert notifications as described in [the section called "Email notification"](#). Provide an email address that you have easy access to.
2. Generate a fake alert. A simple way to do so is to create and start a channel with a simple UDP input and output with a fake input address, such as **udp://1.1.1.1:1111**.
3. Check your email for the notifications message.
4. If necessary, return the global alert notifications to your desired settings.

Web callback notification

You can configure AWS Elemental Conductor Live to send you web callback notifications when alerts occur.

To receive web callback notifications, you must have a web server that supports php scripting. Use the following steps to configure this server to receive alert notifications from Conductor Live.

To set up web callback notifications

1. Use a text editor such as Notepad on a Windows system or Nano on Linux to create a .php file containing the following text:

```
<?php
function get_raw_post(){
    $data = @file_get_contents('php://input');
    if ($data){
        return $data;
    }
    return "nothing passed";
}

$file = "../webcallback/notify";
$fh = fopen($file, "a");
$data = get_raw_post();
fwrite($fh, $data);
fclose($fh);
?>
```

2. Save the file in a directory on your web server.
3. Subscribe to all or some alerts using the steps described here:

Subscribe to all alerts

1. On the Conductor Live web interface, go to the **Settings** page and make sure that you're on the **General** tab.
2. Complete the **Global Alert Notification** fields as described in the following table and choose **Update**.

Field	Instructions
Email	Enter the email address of the alert recipient. Required if you don't provide a URL in the Web Callback URL field.

Field	Instructions
Web Callback URL	If you want to receive web server notifications too, enter the URL of the appropriate .php file on your web server.
Notify	Select when you want to be notified, either when the alert is raised or when it's cleared. You can choose both options.

Subscribe to individual alerts

1. On the Conductor Live web interface, go to the **Stats** page and choose **Notifications**.
2. On the **Notifications** page, find the alert that you want to be notified on and choose the plus sign (+) to expand it.
3. Complete the fields as described in the following table and choose **Save**.

Field	Instructions
Email	Enter the email address of the alert recipient. Required if you don't provide a URL in the Web Callback URL field.
Web Callback URL	If you want to receive web server notifications too, enter the URL of the appropriate .php file on your web server.
Notify	Select when you want to be notified, either when the alert is raised or when it's cleared. You can choose both options.

4. For each alert that you want to be notified on, find the alert, then expand and complete the fields.
4. Test your setup by typing the following at the command line of the Conductor Live node:

```
curl -X POST -d "param1=value1&param2=value2" http://yourdomain.com/webcallback/notification.php
```

5. Open your `notify.php` to check that it was updated. The text of your file should contain something like this:

```
<?xml version="1.0" encoding="UTF-8"?>
<job href="/jobs/3401">
  <node>earhart</node>
  <user_data></user_data>
  <submitted>2014-11-14 01:27:05 -0800</submitted>
  <priority>50</priority>
  <status>preprocessing</status>
  <pct_complete>0</pct_complete>
  <average_fps>0.0</average_fps>
  <elapsed>0</elapsed>
  <start_time>2014-11-14 01:27:06 -0800</start_time>
  <elapsed_time_in_words>00:00:00</elapsed_time_in_words>
</job>

param1=value&param2=value2
```

6. Enter your web callback URL into a web browser to see the HTTP post.

SNMP traps

AWS Elemental Conductor Live generates SNMP traps for activity on the cluster. You can set up to receive SNMP traps from Conductor Live. (If you prefer to poll the SNMP interface for messages, see [the section called "SNMP polling"](#).)

Conductor Live generates SNMP traps for the following events.

- Type of notification: `ELEMENTAL-MIB::alert`
- Type of event: Any alert that worker nodes in the cluster generate.
- Contents of the notification:
 - `ELEMENTAL-MIB::alertSet`. The value is 1 if the alert is being set, 0 if the alert is being cleared.
 - `ELEMENTAL-MIB::alertMessage`. Describes the alert that was set or cleared.

To set up SNMP traps

1. On the Conductor Live web interface, go to the **Settings** page and choose **SNMP**.
2. On the **SNMP** page, complete the fields. Use the instructions in the following table as a guide. Choose **Save**:

Field	Instructions
Allow external SNMP access	Choose Yes to open the SNMP port on the firewall. The port must be open if you will send an snmpwalk command.
Generate SNMP Traps for Alerts	Choose Yes to generate traps.
SNMP Management Host	Enter the IP address of the trap destination.
SNMP Management Trap Port	Enter 162 .
SNMP Management Community	Enter Public .

SNMP polling

AWS Elemental Conductor Live generates SNMP traps for activity on the cluster. You can set up to poll the SNMP interface for messages. (If you prefer to receive traps, see [the section called “SNMP polling”](#).)

You can interact with Conductor Live using a variety of network management systems. AWS Elemental products ship with the Net-SNMP (<http://www.net-snmp.org/>) command line tools. These tools let you access the SNMP interface while you are logged into the system directly or over SSH. Examples in this document are given using `net-snmp` commands.

To set up SNMP polling

1. Either disable the node firewall, or enable external access to SNMP interface.
 - For help disabling the firewall, see [the section called “Firewalls and ports”](#).

- External access to the SNMP interface is enabled by default. To check the setting, access the **Settings** page on the Conductor Live web interface and choose **SNMP**.
2. Query either individual variables, or the entire SNMP interface.

To query individual variables

Use the Net-SNMP tools to query variables as follows.

```
snmpget -c elemental_snmp -v2c -m <MIB>  
localhost MIBvariable
```

Example

```
snmpget -c elemental_snmp -v2c -m ELEMENTAL-MIB  
localhost serviceStatus
```

For a list of MIBs and their variables, see [MIBs in Conductor Live](#).

To query the entire SNMP interface

Use the Net-SNMP tools to do an snmpwalk to gather information about all of the running channels in the cluster.

```
snmpwalk -c elemental_snmp -v2c -m ELEMENTAL-MIB:ELEMENTAL-CONDUCTOR-MIB localhost  
elemental
```

MIBs in Conductor Live

AWS Elemental provides the following MIBs for use with Conductor Live:

ELEMENTAL-MIB

This is the base MIB for all AWS Elemental products.

Variable	Values
serviceStatus	• 0 if the Conductor Live isn't running. • 1 if it is running.

Variable	Values
firewallSettings	• 0 if the node firewall is off. • 1 if it is on.
networkSettings	Always 1. Required for some network management systems.
mountPoints	Number of user-mounted file systems in /mnt.
version	The version of the Conductor Live node.
httpdStatus	• 0 if the httpd service isn't running. • 1 if it is running.
databaseBackup	• 0 if writes (starting backups) is allowed. • 1 if they aren't allowed.

ELEMENTAL-CONDUCTOR-MIB

This MIB describes objects that are specific to Conductor Live.

Variable	Values
channelId	The system-assigned numerical ID of the channel.
channelName	The user-defined name of the channel.
channelRunning	• 0 if the channel isn't running. • 1 if it is running.
channelError	• 0 if the channel isn't in an error state. • 1 if it is in an error state.

Variable	Values
<code>channelLiveEventId</code>	The system-assigned ID of the event associated with the channel.
<code>channelStartTime</code>	Start time of the channel which is provided if the channel is currently running only.
<code>channelDuration</code>	The duration of time that the channel has been running which is provided if the channel is currently running only.
<code>channelAlerts</code>	The text bodies of any active alerts related to the channel, including the time the alert was last set. Each alert is separated by semicolons.
<code>channelMessages</code>	The text bodies of any messages generated in the last 24 hours related to the channel, including the time the message was last set. Each message is separated by semicolons.
<code>nodeId</code>	The numerical ID of the node on which the channel is running.
<code>nodeHostname</code>	Hostname of the node that the channel is running on.

Both the ELEMENTAL-MIB and ELEMENTAL-LIVE-MIB come installed on Conductor Live. They are located in `/opt/elemental_se/web/public/mib/`.

Configuring backup and restore on Conductor Live

AWS Elemental Conductor Live is configured by default to create database backups to a directory on the node. We recommend that you modify the configuration to back up to a remote server. This section describes how to modify the configuration.

The Conductor Live backup command copies the following data to a backup server: profiles, channels, MPTS outputs, nodes, and redundancy groups. Backup files are named in the following format:

```
elemental-db-backup_yyyy-mm-dd_hh-mm-ss.tar.bz2
```

Backup when a Conductor Live node fails

If the primary Conductor Live fails, the other Conductor Live node (the new primary) takes over backups. The new primary stores the backups in the same location as the failed primary. You don't have to manage two backup files.

Topics

- [Configuring for backup](#)
- [Disabling database backups](#)
- [Restoring a backup](#)

Configuring for backup

This section describes how to modify the backup configuration so that AWS Elemental Conductor Live creates database backups on a remote server. (The default configuration is to back up to a directory on the node.)

You only need to change the configuration on the primary Conductor Live node. The secondary node will copy the configuration information from the primary node.

To configure for backups

1. Identify a server and directory on your network for backups. Make a note of the path.
2. Mount the server to the Conductor Live nodes, as described in [Adding mount points to worker nodes](#).

3. On the Conductor Live web interface, go to the **Settings** page and choose **General**.
4. In the **Cluster Tasks** section, change these fields as desired:
 - **Minutes between management database backups**: Change if you want.
 - **Management database backups to keep**: Change if you want.
 - **Path to store management database backups**: Specify the path on the remote server.
5. Choose **Save**.

Disabling database backups

Follow these steps to disable automatic backups in a AWS Elemental Conductor Live cluster.

1. On the Conductor Live web interface, go to the **Settings** page and choose **General**.
2. In the **Cluster Tasks** section, change the value in **Minutes between management database backups** to **0**.
3. Choose **Save**.

Restoring a backup

To restore a backup version of the database on a Conductor Live, Elemental Live, or Elemental Statmux node, follow this procedure.

If you're restoring the Conductor Live database, restore to the primary Conductor Live node.

To restore the database

1. At your workstation, [start a remote terminal session](#) to the Conductor Live node.
2. Enter the following command to identify the version of Conductor Live that is currently installed.

```
[elemental@hostname ~]$ cat /opt/elemental_se/versions.txt
```

Several lines of information appear, including the version number. For example: Conductor Live (2.20.1.12345).

3. Run the install script with the restore option.

```
[elemental@hostname ~]$ sudo sh <product> --restore-db-backup <path> <backup-file> --https
```

Where:

- `product` is the product installer, including the version number that you obtained in the previous step. For example, if you are restoring on Conductor Live, the product installed might be:

```
elemental_production_conductor_live_2.20.1.12345.run.
```

- `path` is the path to the backup file.
 - `backup-file` is the file that you want to restore. This is a zipped file. Don't unzip the file; the restore command automatically unzips the file for you.
 - `--https` keeps HTTPS enabled. If you currently have HTTP enabled, include this option. If you don't currently have HTTPS enabled, omit this option.
4. The file is unzipped and copied to the appropriate folder.

Document History for Configuration Guide

The following table describes the documentation for this release of AWS Elemental Conductor Live.

- **API version:** 3.22.0 and later
- **Release notes:** [current Release Notes](#)

The following table describes the documentation for this release of AWS Elemental Conductor Live. For notification about updates to this documentation, you can subscribe to an RSS feed.

Change	Description	Date
Complete revision of the guide	The guide has been completely revised to correct missing and inaccurate information and procedures.	November 4, 2022
Enabling user authentication	The procedure to enable user authentication has been updated. Don't combine the <code>--config-auth</code> option and the <code>--https</code> option in one command because the configuration script will ignore the request to enable HTTPS.	June 13, 2022
Strong passwords	The guide has been updated to include a recommendation to always set a strong password.	December 21, 2021
Cross-version release of the guide	This guide has been modified so that it isn't for a specific version of AWS Elemental Live. The configuration	November 11, 2021

procedure doesn't change
from version to version.